

# Toward Adaptive Mixed-Element Unstructured Grids for Simulations of Viscous Flows

Gabriel Nastac\*, Forrest Miller†, Dominic Schneider‡, Kevin Jacobson§, and William T. Jones¶  
*NASA Langley Research Center, Hampton, Virginia, 23681*

Matthew Opgenorth||  
*Sierra Space, Louisville, Colorado, 80027*

**Accurate prediction of surface skin friction and heat transfer for viscous flow applications using computational fluid dynamics tools necessitates thin wall-normal elements in the boundary layer. A tetrahedral adaptation process is extended to incorporate the advancing front local reconnection procedure. Thin boundary-layer grids are adapted from the solution, which governs the boundary-layer grid parameters such as first cell height off the wall. The process enables grid-converged steady-state results from a CAD file. Results are presented across the speed range from low subsonic flow over an airfoil to high-speed flows over blunt bodies and a rocket.**

## I. Introduction

HIGH Reynolds number flows commonly encountered in the aerospace field involve thin boundary layers. Skin friction and heat transfer are often key drivers for vehicle design. Computational Fluid Dynamics (CFD) is commonly used to predict these flows. To handle complex geometries, unstructured grids are typically used for quick turnaround of simulations. To accurately capture boundary layers, hybrid mixed-element grids involving prisms or hexahedra have historically been used due to their improved accuracy when using finite-volume solvers to predict skin friction and heat transfer [1]. While unstructured grids are not particularly difficult to generate, it is desirable to automate the process further to enable even more rapid design cycles. Grid adaptation is becoming increasingly common in engineering simulations [2]. The NASA CFD Vision 2030 study [3] also emphasizes grid refinement as one of its key areas of development.

There have been many studies in recent years investigating adaptation using mixed-element grids. These range from fixed-grid studies investigating topologies [4] to adaptation techniques ranging from aligning prisms with shock fronts [5, 6] to using overset with a dual-solver paradigm [7]. Ref. [8] demonstrated the capability of a frozen thin prismatic boundary layer with off-body anisotropic tetrahedral adaptation to predict surface skin friction and heat transfer. Ref. [9] explored various adaptation techniques for supersonic flow over a capsule geometry. The study explored fully tetrahedral adapted grids and frozen mixed-element boundary-layer grids with tetrahedral adaptation outside the boundary layer. In terms of overall forces and moments on the forebody, the approaches predicted similar coefficients due to the pressure-dominated flow. The fully tetrahedral grid, however, has large wall-normal spacing greater than a  $y^+$  of 10, leading to increased error in wall-normal gradients and poor prediction of the boundary layer near the vehicle. For cases where this is important, such as prediction of skin friction and heat transfer, the latter of which is commonly the primary driver for design of high-speed entry vehicles, it is desirable to more reliably enforce wall-normal spacing.

The present effort expands on Ref. [8] by working towards automating the boundary-layer generation process. This paper will first describe the general governing equations and numerical methods for the flow solver (Section II) followed by the grid generation and adaptation techniques (Section III). Results using the proposed approach are demonstrated on various problems of interest across the speed range (Section IV).

---

\*Research Aerospace Engineer, AIAA Member

†Intern

‡Intern

§Research Aerospace Engineer, AIAA Member

¶Computer Engineer, AIAA Associate Fellow

||Principal Systems Engineer

## II. Governing Equations and Numerical Methods

FUN3D, an unstructured-grid node-based finite-volume CFD solver developed at the NASA Langley Research Center [10], is the target solver for this work. The solver is used across the speed range which is demonstrated in this work for flows ranging from subsonic incompressible speeds to hypersonic speeds involving chemical nonequilibrium. The governing equations for compressible reacting flow are the conservation of species, mixture momentum, and total energy:

$$\frac{\partial}{\partial t} (\rho y_s) + \frac{\partial}{\partial x_j} (\rho y_s u_j) - \frac{\partial}{\partial x_j} (J_{sj}) = \dot{\omega}_s, \quad (1)$$

$$\frac{\partial}{\partial t} (\rho u_i) + \frac{\partial}{\partial x_j} (\rho u_i u_j + p \delta_{ij}) - \frac{\partial}{\partial x_j} (\tau_{ij}) = 0, \quad (2)$$

$$\frac{\partial}{\partial t} (\rho E) + \frac{\partial}{\partial x_j} ((\rho E + p) u_j) - \frac{\partial}{\partial x_j} \left( u_k \tau_{kj} + \dot{q}_j + \sum_{s=1}^{N_s} h_s J_{sj} \right) = 0, \quad (3)$$

where  $y_s$  is the mass fraction of species  $s$ ,  $\rho = \sum_{s=1}^{N_s} \rho y_s$  is the mixture density,  $\rho_s = \rho y_s$  is the species density of species  $s$ ,  $u_i$  is the  $i$ th component of the velocity, and  $E$  is the total energy.  $J_{sj}$  is the  $j$ th component of the diffusive flux,  $\dot{\omega}_s$  is the chemical source term of species  $s$ ,  $p$  is the pressure,  $\tau_{ij}$  is the shear-stress tensor,  $\dot{q}_j$  is the  $j$ th component of the heat flux, and  $h_s$  is the enthalpy of species  $s$ . Pressure is modeled as an ideal gas, heat transfer is modeled with Fourier's law, species diffusion is modeled using Fick's law, and the shear-stress is modeled with a Newtonian model. NASA polynomials [11] are used to compute thermodynamic properties on a per-species basis. The transport properties (diffusivity, viscosity, and thermal conductivity) are computed using collision integrals [12]. The compressible perfect gas equations are obtained by assuming a single perfect gas species with thermally perfect thermodynamic properties and simplified transport properties (e.g., constant Prandtl number and Sutherland's law). Turbulence, when modeled, is closed with the SA-neg model in this work [13]. More details on the implementation can be found in Ref. [14].

A median-dual finite-volume approach is used for spatial discretization on general unstructured grids containing tetrahedra, pyramids, prisms, and hexahedra. A dual-grid is generated by bisecting every edge and connecting them to cell centroids to generate polyhedra. Inviscid fluxes are computed at each dual-face using an approximate Riemann solver, Roe's scheme [15] or HLLC++ [8] in this work, with second-order accuracy obtained using an unstructured Monotonic Upstream-centered Scheme for Conservation Laws (UMUSCL) reconstruction [16] with unweighted least-squares gradients. In this work, the van Albada flux limiter is utilized for UMUSCL [17]. Viscous fluxes are computed using an edge-based approach [18], equivalent to a Galerkin-based approach. The equations are integrated in time implicitly. Many Jacobians for the implicit time integration are computed using automatic differentiation [19]. The turbulence model is solved loosely coupled to the meanflow equations.

The simulations in this work are primarily run on GPUs and utilize the FUN3D Library for Universal Device Acceleration (FLUDA) described in Ref. [20]. FLUDA is written in a thin abstraction over CUDA C++ that enables primarily single-source software capable of running efficiently on NVIDIA, AMD, and Intel GPUs and multicore CPUs. Performance is commensurate with memory bandwidth, which enables modern GPU nodes consisting of a few GPUs to be equivalent to thousands of CPU cores of performance, substantially reducing power and space requirements versus CPU-based systems.

## III. Grid Adaptation Process

This work extends *refine* [21, 22], a 3D grid adaptation tool developed at the NASA Langley Research Center, to support prismatic boundary layers through the advancing-front local-reconnection (AFLR) method [23]. This enables an accurate CAD-to-solution (C2S) [24] capability for high Reynolds number viscous flows. By tying the solution into the boundary-layer grid generation process through a user specified quantity such as a target  $y^+$  or surface cell Reynolds number based on the local speed of sound and wall properties,  $Re_C = \rho_w c_w \Delta n / \mu_w$ , grid-converged simulations (with sufficient grid points) can be performed with simply a CAD file with the user not needing to know the details of grid generation.

The C2S+Hybrid process is as follows. First, a clean CAD file along with other domains (e.g., a farfield) must be specified using Engineering Sketch Pad [25]. An initial grid is then generated using either AFLR3 [23] or TetGen [26]; this is commonly referred to as bootstrapping. This is followed by an initial C2S pure tetrahedral adaptation phase which is run to some target grid size. The final surface grid is then extracted from the final C2S grid and used to

generate a mixed-element boundary-layer grid using an open advancing-layer approach. AFLR3 [23], DOD HPCMP Capstone [27], and Fidelity Pointwise [28] have all been successfully employed to generate the initial hybrid grid, although AFLR3 is the primary mesher used in this work due to ease of scripting and direct support for the UGRID grid format which is the current default grid format for FUN3D. An additional adaptation phase is run on this hybrid grid, where only the off-body tetrahedra are adapted. This process is similar to that described in Ref. [8]. While the open advancing-layer approach can be repeated, it is performed only one time for the current results; the boundary-layer grid is fixed for the hybrid phase.

The process is scripted in NASA *pyrefine* [29], an open source Python module for driving the C2S *refine* and FUN3D workflow. Modules for bootstrapping the initial grid, calling AFLR3, and splitting hardware and MPI ranks were added during this work. *pyrefine* generates and runs job scheduling scripts for clusters or can alternatively be called inside a job scheduling script if the same hardware is desired for all the simulations. For the former, hardware can be dynamically changed based on problem size and different hardware can be allocated for *refine* and FUN3D, e.g., GPU nodes for FUN3D and CPU nodes for *refine* (which is not GPU-capable currently).

The AFLR3 module in *pyrefine* has three different extrusion options: 'manual', 'yplus', and 're\_cell'. The first is a manually prescribed boundary-layer grid which can be set with a specified first layer height and number of layers or using a built-in  $y^+$  calculator based on turbulent flat plate theory and first layer height and total boundary-layer grid height in terms of  $y^+$  units (defaulting to a total height of  $y^+ = 1000$ ). The last two approaches require the maximum surface  $y^+$  and surface cell Reynolds number,  $Re_C$ , along with corresponding wall spacing at those locations. These are dynamic and will automatically generate the boundary-layer grid based on the solution.

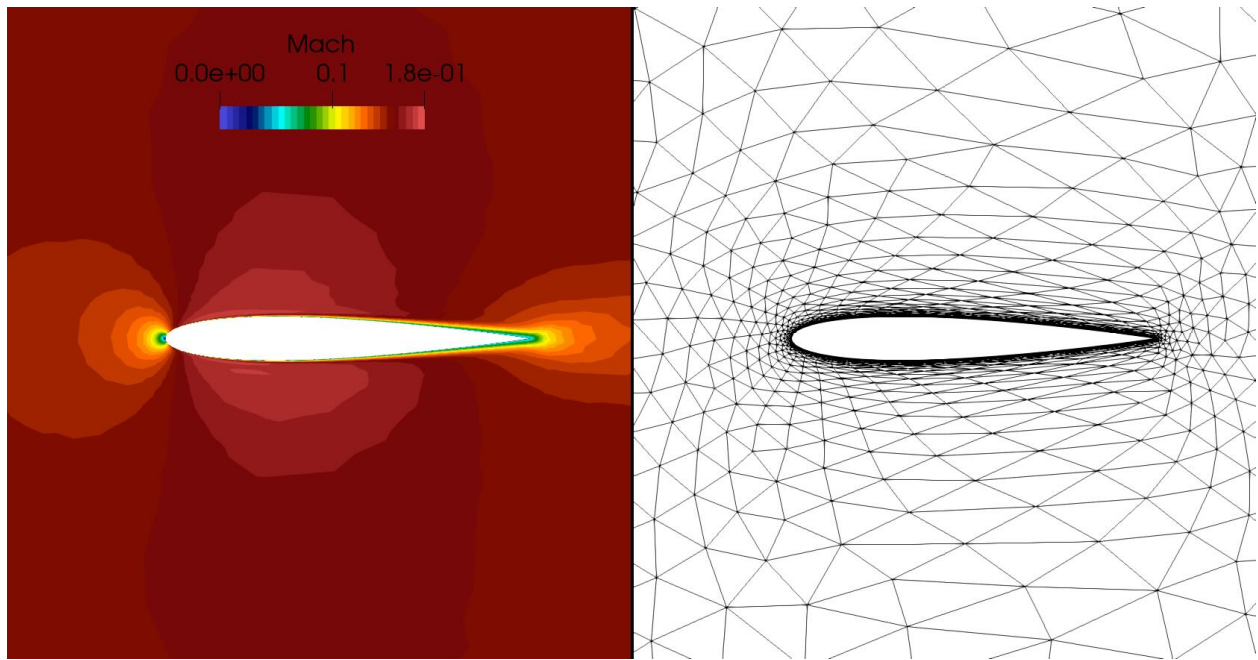
## IV. Results

The following section details the results for the adaptation approaches. The first flow examined is subsonic flow over the NACA 0012 airfoil to demonstrate the capability of simplex elements to capture turbulent boundary layers in a limited 2D and shock-free setting. The second case examined is high enthalpy hypersonic flow over a capsule to demonstrate the capability for entry conditions. This case examines C2S and C2S+Hybrid approaches with both prescribed and dynamic boundary layers. The third flow examined is high Reynolds number transonic flow over the ONERA M6 wing as a representative example for transonic flight. This case demonstrates the limitations of pure tetrahedral adaptation in capturing turbulent boundary layers and shocks with reasonable grid resolutions and the benefit of the hybrid approach in ensuring adequate boundary-layer spacing. The fourth case examined is supersonic flow over a four-engine rocket configuration. This case is used to demonstrate the capability on rocket engine base flow prediction. Lastly, hypersonic flow over Dream Chaser is considered as a more realistic hypersonic example with complex curvatures and features. Results are compared to those obtained using structured grids and experimental data when available.


### A. Subsonic Flow over the NACA 0012 Airfoil

The tripped experimental condition investigated by Ref. [30] is considered as a 2D example to investigate how well triangle elements can capture a turbulent boundary layer. The freestream Mach number, Reynolds number per chord, angle of attack, and temperature are 0.15,  $6 \times 10^6$ ,  $0^\circ$ , and 300 K, respectively. The trailing edge is considered sharp. The gas is modeled as perfect gas air. The wall is modeled as no slip with the wall temperature set to the flat plate adiabatic wall temperature at the freestream Mach number. Roe's flux scheme is employed [15]. Turbulence is modeled using the SA-neg model [13]. The solution is advanced using the hierarchical adaptive nonlinear iteration method (HANIM) [31] with a max local CFL of 1000. Two-hundred iterations are run per CFD cycle. Overall, the residual norms are reduced by 7 orders of magnitude.

The multiscale [32] metric grid adaptation process uses Mach number as the target metric. The grid size increases from an initial 5,000 grid points to 166,000 grid points. Twenty-three cycles are run using pure triangle adaptation. The initial and final grid and Mach number contours are plotted in Figures 1 and 2, respectively. Overall  $y^+$  values are less than one on the airfoil surface for the final grid. In this scenario, due to the 2D grid and subsonic flow without shocks, the metric-based adaptation fully captures the turbulent boundary layer without the need for a structured boundary-layer grid. Results are compared to structured grid NASA CFL3D results on the Turbulence Modeling Resource website [33] which utilize the SA turbulence model [34] and have been shown to be nearly grid converged. Surface results are plotted in Figure 3. Surface pressure agrees extremely well to the structured CFL3D and experimental results. Skin friction results for the adapted grid are smooth and agree well with CFL3D results.




**Fig. 1 Subsonic flow over the NACA0012 airfoil initial Mach contour and grid.**



**Fig. 2 Subsonic flow over the NACA0012 airfoil final Mach contour and grid.**





**Fig. 3** Subsonic flow over the NACA0012 airfoil surface results. Left: pressure coefficient. Right: skin friction.

### B. Hypersonic Flow over the Crew Exploration Vehicle


Hypersonic flow over the Crew Exploration Vehicle (CEV) is considered. The high-enthalpy test performed in the CUBRC LENS-I facility in air [35] is investigated using a CEV model with a radius,  $R$ , of 0.073025 meters. In addition to experimental pressure and heat transfer on the surface, NASA Data-Parallel Line Relaxation (DPLR) [36] data are also available from Ref. [37] for additional comparisons. Run 10, the highest enthalpy condition ( $H_0 = 12.4$  MJ/kg), is simulated. All inputs for this study, including geometry, are publicly available as a NASA *pyrefine* example [29]. Other simulations in this work generally employed the same template file with minor changes (e.g., different geometry, flow solver inputs, and grid point targets). A 5-species, one-temperature gas model is used. The surface is modeled as a noncatalytic cold ( $T = 300$  K) no-slip wall. The freestream temperature, density, and velocity are 631 K,  $1.08 \times 10^{-3}$  kg/m<sup>3</sup>, and 4601 m/s, respectively. The freestream mass fractions are  $\tilde{Y} = Y_{N_2}, Y_{O_2}, Y_{NO}, Y_O, Y_N = 0.7377, 0.1387, 0.0590, 0.0646, 0.0000$ . The angle of attack is  $28^\circ$ . The HLLE++ inviscid flux scheme [8] is employed with the van Albada limiter [17]. Laminar flow is assumed. Edge-based viscous terms are utilized [18]. The solution is advanced using steady-state pseudo-time stepping with a max local CFL of 5. Two-thousand iterations are run per cycle. Overall, the residual norms are reduced by 4 orders of magnitude without limiter freezing. Heat transfer magnitudes were observed unchanged by subsequent iterations and cycles.

The multiscale metric grid adaptation process uses temperature as the target metric. The grid size increases from an initial 0.4 million grid points to 1.6 million grid points using pure tetrahedral adaptation across 20 cycles. This phase utilizes a metric-space gradation stretching ratio of 10 to reduce surface grid anisotropy to improve grid quality for boundary-layer grid extrusion. This process is followed by 10 cycles at fixed complexity (targeting 5 million grid points) using a hybrid grid with the boundary-layer grid generated targeting a surface cell Reynolds number based on speed of sound ( $Re_C$ ) of unity from the final tetrahedral solution (denoted as dynamic C2S+Hybrid). A prescribed boundary-layer grid simulation is also run with the first layer set to  $2.5 \times 10^{-7}$  m (corresponding to a  $Re_C = 1.0$ ) with 35 layers at a growth rate of 1.2. The aspect ratios near the wall are over one thousand. Visual results are shown for the dynamic C2S+Hybrid simulation. For additional comparison, a pure tetrahedral C2S simulation is run without metric-space gradation to a final grid size of 6.4 million points.


The adapted surface grid is shown in Figure 4. The adaptation process primarily focuses on the forebody which is the primary region of interest. Figure 5 shows centerline grid slices for both the dynamic and prescribed C2S+Hybrid simulations. Figures 6 and 7 depict contours of the flow around the vehicle. The temperature peaks at 8100 K at the shock and is smoothly captured due to grid adaptation. At this enthalpy, molecular oxygen dissociates completely in the post-shock region as well as in the wake. Molecular nitrogen is partially dissociated, as noted in the contours. The bow shock and wake are captured well with anisotropic tetrahedra using only 5 million total points in the entire domain. Figure 8 plots surface quantities of interest. All quantities are smoothly captured on the unstructured grid. The cell

Reynolds number peaks at 2.3 near the stagnation point. For the prescribed C2S+Hybrid simulation, the peak cell Reynolds number is 1. The peak surface cell Reynolds number for the initial tetrahedral phase was over 500 at 1.6 million grid points; the spacing is a nonlinear function of the solution and in this case linear extrapolation led to some error. If the first phase was run with more grid points to reduce the peak cell Reynolds number, or if the hybrid phase was repeated with these hybrid results, the specified target would be better matched. As seen in Figure 5, the dynamic mismatch causes a slightly larger boundary-layer grid to be generated for the dynamic approach, although still well below the shock standoff distance.

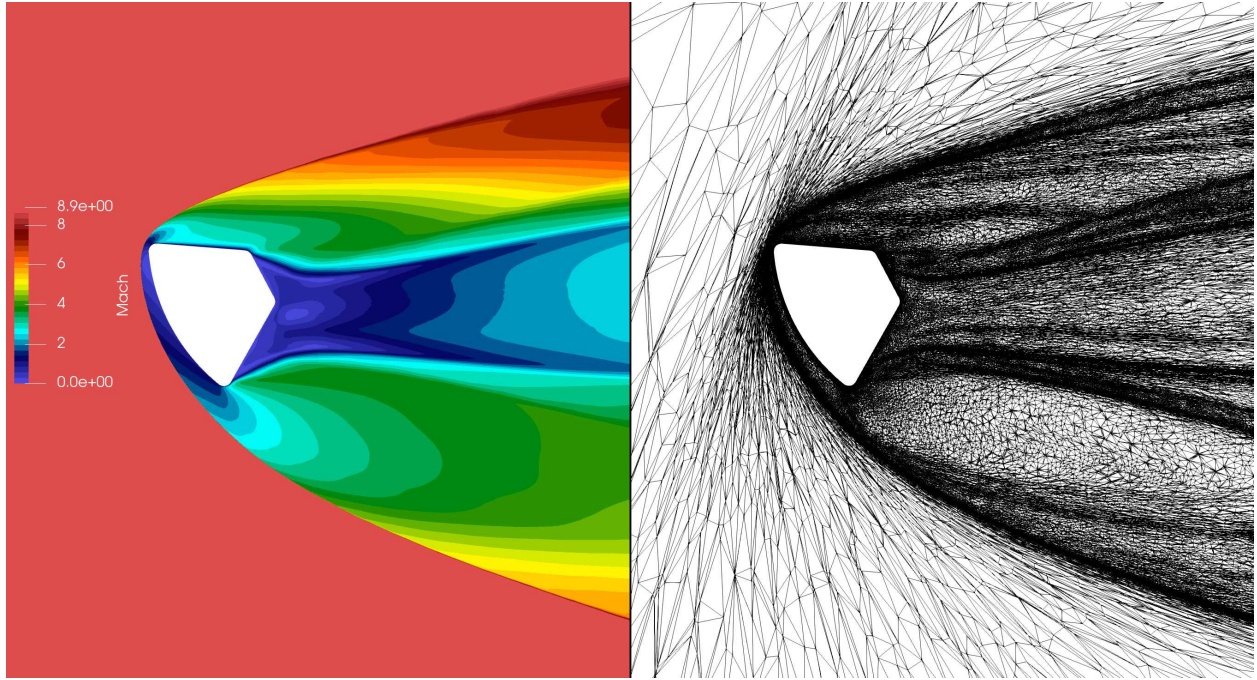
Comparisons to experiment and DPLR are shown in Fig. 9. Surface properties are smooth along the heatshield and backshell of the capsule. The surface pressures agree extremely well with DPLR results, but overpredict experiment slightly at the center and rear shoulder of the heatshield. Heat transfer predictions by both FUN3D and DPLR are somewhat higher than the experiment. FUN3D predicts slightly larger heat transfer than DPLR. The dynamic and prescribed hybrid results are indistinguishable from each other. The hybrid FUN3D approaches predict smooth heat transfer results whereas the pure tetrahedral simulation is noisier and overpredicts the shoulder versus the other grids. It is expected the pure tetrahedral simulation to eventually predict smooth heat transfer in the limit of a very fine grid, but the current maximum surface cell Reynolds number is roughly 100 at 6.4 million grid points versus  $O(1)$  for the hybrid grids, which only have 5 million grid points.



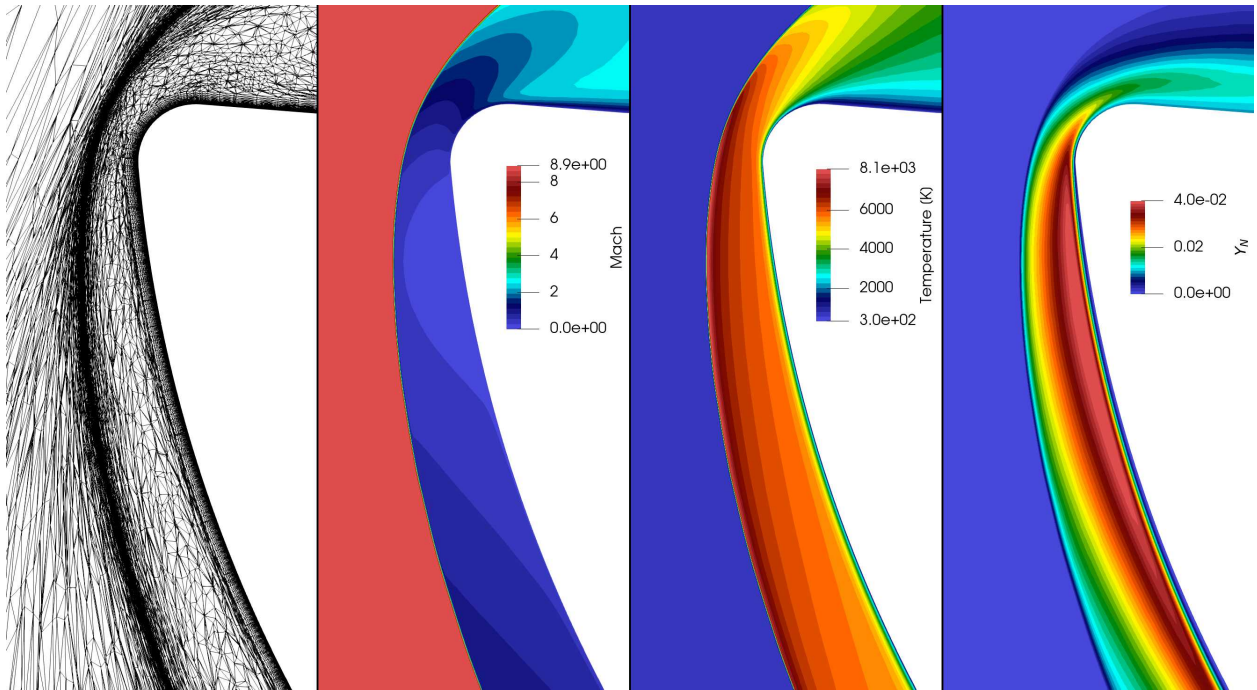
**Fig. 4 Hypersonic Flow over the CEV adapted surface grid. Left: forebody. Right: aftbody.**




**Fig. 5 Hypersonic Flow over the CEV centerline slices through the final adapted grids. Left: dynamic boundary-layer grid adapted on surface cell Reynolds number case. Right: prescribed boundary-layer grid case.**




**Fig. 6** Hypersonic Flow over the CEV centerline slice. Left: Mach contour. Right: grid.



**Fig. 7** Hypersonic Flow over the CEV centerline slice near the surface. From left to right: grid, Mach contour, temperature, and atomic nitrogen mass fraction.



**Fig. 8** Hypersonic Flow over the CEV surface results for the dynamic C2S+Hybrid simulation. Top Left: Surface pressure coefficient. Top Right: Heat transfer. Bottom Left: Surface cell Reynolds number. Bottom Right: Surface grid.



**Fig. 9** Hypersonic Flow over the CEV comparisons to DPLR and experiment. Left: Surface pressure along the centerline. Right: Heat transfer along the centerline.

### 1. Computational Performance

This case is run on various hardware at the NASA Advanced Supercomputing (NAS) facility. The NAS Cabeus cluster is composed of GPU nodes (denoted as MIL\_A100) consisting of 4 NVIDIA A100 80 GB SXM4 GPUs and one single-socket 64-core AMD EPYC 7763 CPU. Some CPU nodes (denoted as ROM\_AIT) from the NAS Aitken cluster are used for CPU runs; each node has dual-socket 64-core AMD EPYC 7742 CPUs (128 total cores). The whole simulation process is repeated on 1, 2, and 4 MIL\_A100 nodes and 10 ROM\_AIT nodes and are summarized in Table 1. The initial bootstrap and grid process and AFLR3 are serial and take a few minutes total. For the GPU nodes, refinement is run using all the CPU cores (64 cores per node) and CFD is run using 16 cores per node (4 ranks per



GPU) using the CUDA Multi-Process Service [38]. GPU-aware MPI was not used for this work, but would improve scaling results. In addition, the reported times include total executable time including pre- and post-processing. As CFD and refinement are run 30 times during the simulation, these times are not negligible. For GPU runs, pre- and post-processing takes approximately 25% of the CFD run time. For all the GPU node runs, CFD and refinement each take approximately 50% of the run time. For the CPU run, CFD is 83% of the run time and refinement is 16% of the run time. It is important to note that the GPU nodes utilize newer generation AMD CPUs, but even taking that into account, refinement strong scaling is poor and slows down with more cores. For the single GPU node result, each A100 corresponds to approximately 850 EPYC 7742 cores, ignoring IO (which is roughly 25% for GPU simulations and 10% for CPU simulations). Four MIL\_A100 nodes can complete the simulation in under 2 hours.

NASA *pyrefine* does have the capability to launch individual compute jobs per process (each CFD and *refine* run are individual jobs), but overall throughput due to queue times on NASA systems limit this capability. For dedicated hardware, it is recommend to utilize this feature and run *refine* separately on CPU nodes and CFD on GPU nodes. Using separate hardware, a simulation can be performed in under 2 hours using approximately 16 A100-hours and 4 64-core EPYC 7763 CPU-hours. The same simulation can be performed in around 4 hours for about half the cost using the one GPU node configuration using approximately 8 A100-hours and 2 64-core EPYC 7763 CPU-hours.

**Table 1 Hypersonic Flow over the CEV C2S+Hybrid computational performance.**

Hardware	MIL_A100	MIL_A100	MIL_A100	ROM_AIT
Nodes	1	2	4	10
Bootstrap [min]	2	2	1	3
AFLR3 [min]	1	1	1	1
C2S [min]	89	64	49	190
Hybrid [min]	179	108	66	184
Total [min]	271	174	118	378
CFD [min]	141	99	71	313
<i>refine</i> [min]	127	72	43	60
CFD [%]	52	57	60	83
<i>refine</i> [%]	47	41	37	16

### C. Transonic Flow over the ONERA M6 Wing


Transonic turbulent flow over the ONERA M6 wing [39] is considered as a representative 3D geometry and condition for transonic flow. The condition matches that described on the Turbulence Modeling Resource website [33]. The trailing edge is considered sharp. The freestream Mach number, Reynolds number per root chord, angle of attack, and temperature are 1.04,  $14.6 \times 10^6$ ,  $3.06^\circ$ , and 300 K, respectively. The full span wing is considered and simulated. The wall is modeled as no slip with the wall temperature set to the flat plate adiabatic wall temperature at the freestream Mach number. Roe's flux scheme is employed [15] with the van Albada limiter [17]. Turbulence is modeled using the SA-neg model [13]. The solution is advanced using HANIM with a max CFL of 100 [31]. Five-hundred iterations are run per cycle. Overall, the residual norms are reduced by 4 orders of magnitude without limiter freezing.

The multiscale metric grid adaptation process uses Mach number as the target metric. The grid size increases from an initial 0.2 million grid points to 1.6 million grid points using pure tetrahedral adaptation. This first phase utilizes a metric-space gradation stretching ratio of 10 to reduce surface grid anisotropy to improve grid quality for boundary-layer grid extrusion. This phase is also run as inviscid instead of turbulent due to surface grid quality issues described in the following paragraph. The process is followed by 10 cycles at fixed complexity (targeting 8 million grid points) using a hybrid grid. The hybrid mixed-element boundary-layer grid is prescribed from the freestream Reynolds number using a  $y^+$  of 1 from turbulent flat plate theory (with the root chord as the reference length) and a boundary-layer grid height of  $y^+ = 1000$ . The final C2S+hybrid grid has 7.7 million grid points and 4.9 million prisms (approximately 2.4 million prism points). The final surface grid has 0.17 million surface triangles.

One limitation of the current approach is that the surface grid that is used for boundary-layer grid extrusion must have reasonable surface element aspect ratios or the boundary-layer grid generation will fail. Running this example

with modeled turbulence and pure tetrahedral adaptation (C2S) and a grid size of 12.8 million points (for a half span simulation) leads to results shown in Figures 10 and 11. The surface grid is highly anisotropic, with some surface triangle aspect ratios over 100. Even so, the surface  $y^+$  is still roughly 5 along the surface causing underprediction of skin friction. AFLR3 [23], DOD HPCMP Capstone [27], and Fidelity Pointwise [28] were not able to generate a boundary-layer grid from this highly anisotropic surface grid. Employing metric-space gradation did not improve the surface quality in this case, but was sufficient to successfully run the other examples in this work. This issue is primarily observed for transonic flows. Faster flows typically have stronger shocks which lead to the adaptation primarily focusing on the shocks rather than the surface. While NASA *refine* does have aspect ratio limits, they are not strongly enforced. Strongly enforcing the aspect ratio leads to reasonable aspect ratio surface grids but with millions of points which is not desirable (as the boundary layer extrusion process will typically generate about 30 layers by default). Future work should investigate ways to generate reasonably sized high quality surface grids from turbulent subsonic and transonic solutions.

The final surface contours are plotted in Figure 12. The lambda shock is well captured on the surface for the turbulent solution despite using an inviscid simulation to generate the surface grid. The skin friction is smoothly predicted. The surface  $y^+$  maximum is 1.4, although it is generally less than 1 on the surface. Results at  $\eta = 65\%$  span are plotted in Figure 13. The shocks are well captured from the adaptation along with the upper part of the boundary layer. The wake is also well captured all the way to the exit 100 chords away (not explicitly shown). Results are compared to predominantly structured grid results from NASA CFL3D, NASA USM3D, and NASA FUN3D on the Turbulence Modeling Resource website [33]. Surface results along a slice at 65% span are plotted in Figure 14. Surface pressure agrees extremely well with the structured grid CFD and experimental results. Skin friction results for the full tetrahedral adapted grid are noisier and are underpredicted versus the C2S+Hybrid and structured grid results. The C2S+Hybrid approach is able to match, and potentially better predict, results versus the structured grids with over an order of magnitude fewer points. Integrated force results are plotted in Figure 15. The forces for the C2S+Hybrid approach are slightly larger than the other solutions.



**Fig. 10** Transonic flow over the ONERA M6 wing C2S final surface grid. Right: Surface grid around leading edge shock.






Fig. 11 Transonic flow over the ONERA M6 wing C2S final surface skin friction and  $y^+$ .

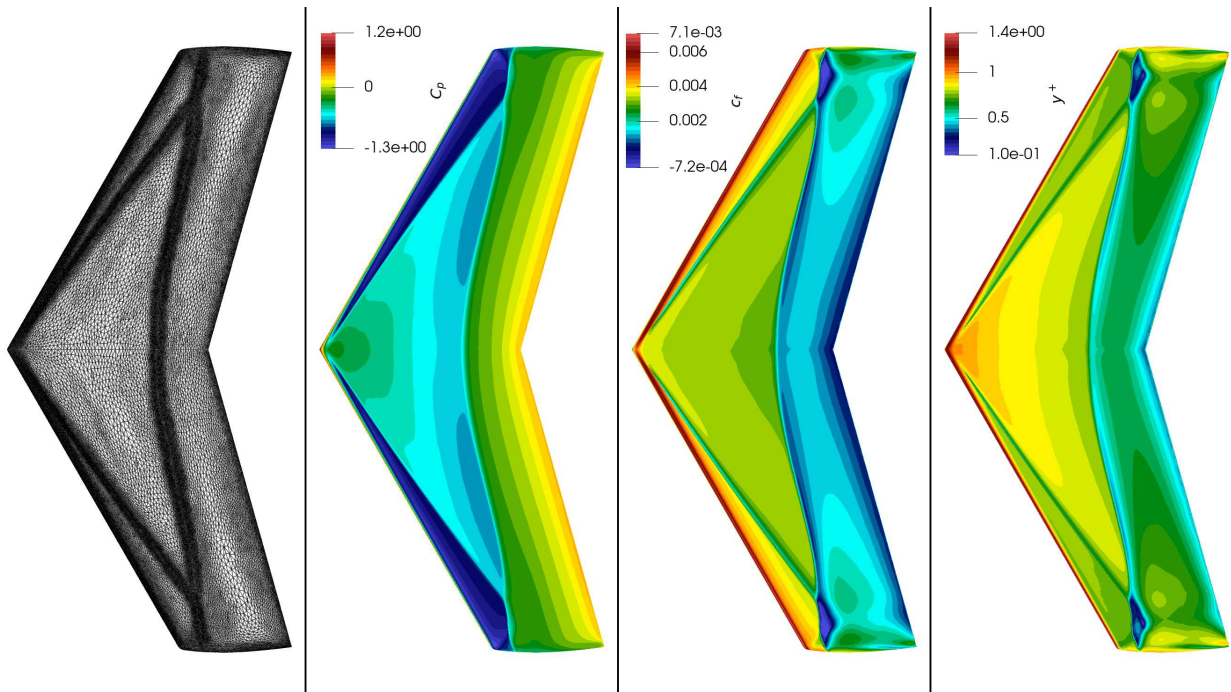





Fig. 12 Transonic flow over the ONERA M6 wing C2S+Hybrid surface grid, pressure coefficient, skin friction, and  $y^+$ .



**Fig. 13** Transonic flow over the ONERA M6 wing C2S+Hybrid slice at  $\eta = 65\%$  span Mach contour and grid. Bottom: Zoomed-in view near the terminating shock around the mid-chord.



**Fig. 14** Transonic flow over the ONERA M6 wing surface pressure and skin friction along a slice at  $\eta = 65\%$  span comparing various flow solver solutions and experimental data. The denoted grid points are for half span simulations.



**Fig. 15** Transonic flow over the ONERA M6 wing lift and drag coefficient versus grid spacing ( $h$ ). The grid spacing is for half span simulations.

#### D. Supersonic Flow over a Four-Engine Rocket Configuration

Flow over a rocket configuration is considered to demonstrate the approach for rocket engine base flow prediction. The NASA experiments of Musial and Ward from 1961 are investigated [40]. These experiments have been examined numerically before [41, 42]. A four-engine rocket configuration was investigated in the then Lewis Research Center 10×10 foot supersonic wind tunnel at freestream Mach numbers from 2.0 to 3.5 over a range of pressure altitudes from 14.33 to 30.48 kilometers. Specifically, a high altitude condition (27.65 kilometers) at Mach 2.75 is simulated in this work. The freestream velocity, density, and temperature are 952 m/s, 0.0196 kg/m<sup>3</sup>, and 298 K, respectively.

The rocket has a cone-cylinder body with a maximum diameter of 0.3048 meters. There are four engines mounted on the base. The bell-shaped nozzles exit-to-throat area ratios are 12 and have exit half-angles of 3°. The exit diameter of each engine is 0.07468 meters. The liquid rocket engines use JP-4 and LOX as propellants with an oxidizer to fuel ratio of 2.2. The total pressure of each engine is 4.137 MPa. The equilibrium temperature using NASA CEA [43] is 3460 K, which is used as the engine total temperature in the simulations. The computed equilibrium mass fractions, which are used as the engine plenum mass fractions, are presented in Table 2. Each engine produces a nominal thrust of 2.22 kN. The engines are cooled with water-cooled jackets. Injectors are not simulated in this study. The experiments are run for 10 seconds. More details including detailed schematics of geometry can be found Ref. [40].

**Table 2** Rocket engine plenum species mass fractions.

Species	Mass Fraction
H <sub>2</sub> O	0.25091
CO <sub>2</sub>	0.20602
H <sub>2</sub>	0.01290
CO	0.49517
OH	0.02567
O <sub>2</sub>	0.00493
O	0.00295
H	0.00145

The experimental details are not fully documented; the internal engine wall temperature distribution, sting, rocket nose geometry, and injectors are notably missing. Specifically, the internal engine wall temperature distribution is essential in predicting base heat transfer and temperature profiles. Reference [41] performed a 1D analysis based on the coolant and combustion gas data and estimated an internal wall temperature of approximately 700 K. The study also simulated adiabatic walls to bound the uncertainty of heat transfer. In the present effort, the external vehicle walls are set to the freestream temperature and the internal engine walls are modeled as both isothermal (at 700 K) and adiabatic in two separate analyses. Visual results in this section are of the isothermal boundary condition. All walls are modeled as noncatalytic.


A 10-species, 19-reaction mechanism is employed, identical to past studies investigating afterburning effects of hydrocarbon rocket engines [44–48]. Turbulence is modeled with the SA-neg model [13]. The HLLE++ inviscid flux [8] is employed with the van Albada flux limiter [17]. Maximum local CFL numbers for meanflow and turbulence are 5 and 2, respectively. Two-thousand iterations are run per cycle except for the final cycle which is run longer to ensure convergence. Overall, the final residual norms are reduced by 6 orders of magnitude without limiter freezing. Viscous forces were monitored and determined to be converged.

An initial C2S phase is performed by using multiscale metric adaptation with temperature as the target metric. The grid size increases from an initial 0.4 million grid points to 3.2 million grid points. A prescribed boundary-layer grid is extruded from the final C2S surface grid using past studies values [41], ensuring a surface  $y^+$  of below one on the base. The first point off the wall is  $2.54 \times 10^{-6}$  m and is grown for 30 layers at a growth rate of 1.2 leading to approximately 20 million prisms (about 10 million grid points) in the boundary-layer grid. The hybrid phase is then run for 5 cycles, targeting a final grid size of approximately 30 million points. No attempt was made to reduce the grid count or domain to reduce problem size. In practice, the farfield domain can be much smaller, which would reduce the grid size.


The final surface and volume grids are shown in Figure 16. The multiscale metric adaptation captures both off-body features such as plumes and shocks as well as near-body features such as boundary layers and base flow. The extruded boundary-layer grid is visible with a total height of approximately 0.3 cm. Flow features are well captured over seven vehicle lengths downstream as shown in Figure 17. Flow results near the base are plotted in Figure 18. The engine shocks are well captured with the adaptation metric. The larger temperature at the base is clearly visible along with the specific heat ratio changes; the base gas is primarily engine exhaust. The plume results are plotted in Figure 19. The plume length, denoted by the terminal Mach disk, is approximately 75% of the vehicle length away from the base. Surface results are plotted in Figure 20. The results are mostly symmetric, although not completely due to the general unstructured nature of the grid. The surface  $y^+$  is below one on the base as specified. The surface pressure coefficient and heat transfer peak at the stagnation point. The grid points are also clustered on the base, which is desirable. Figure 21 depicts the results of the external engine wall facing the stagnation point. All engine results are visually symmetric. The peak heat transfer on the engine external walls is approximately the stagnation point heat transfer on the base. There are three peaks: the base point, before the throat, and before the exit. The surface pressure coefficients on the engine walls on these peaks are slightly lower than the stagnation point value.

Line plots along the base are plotted in Figure 22. Isothermal and adiabatic simulations in addition to experimental data are compared. Overall, results agree favorably with experimental data [40]. There are no reported error bars on the experimental data. Heat transfer is slightly underpredicted for the isothermal computation and slightly overpredicted for the adiabatic computation. Base pressures are predicted to be slightly larger for the adiabatic condition as well. Base pressures are slightly underpredicted away from the centerline versus experimental data. Turbulence model variations have not been considered in this study to further bound simulation uncertainty. Overall, these simulations demonstrate the potential of using a fully automated C2S+Hybrid approach from geometry to simulate rocket base flows which are important for rocket design and analysis.

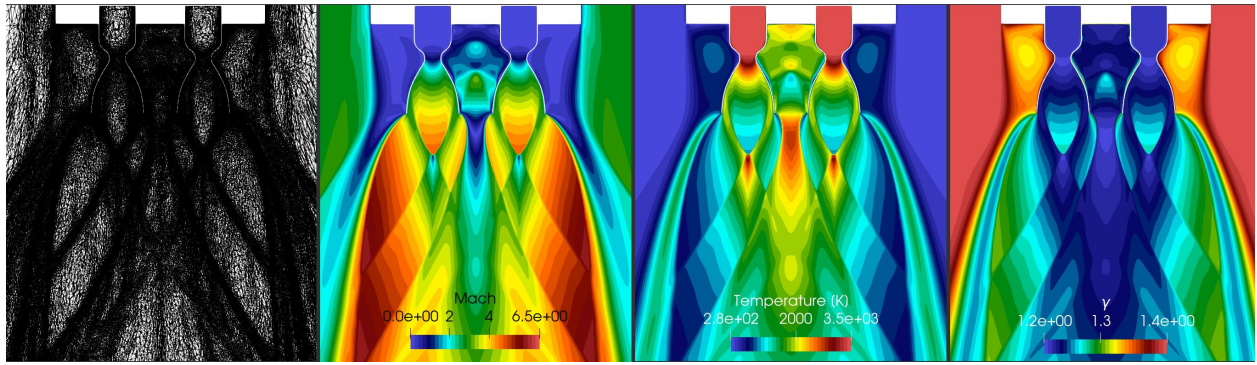





**Fig. 16** Supersonic flow over a four-engine rocket configuration final grid. Top Left: Farfield centerline slice. Top Right: Centerline slice through an engine. Bottom Left: Centerline slice closer to the vehicle. Bottom Right: Base surface grid.



**Fig. 17** Supersonic flow over a four-engine rocket configuration farfield centerline slice Mach contour.




**Fig. 18** Supersonic flow over a four-engine rocket configuration results near the base. From left to right: centerline grid slice, Mach contours, temperature contours, and specific heat ratio contours.

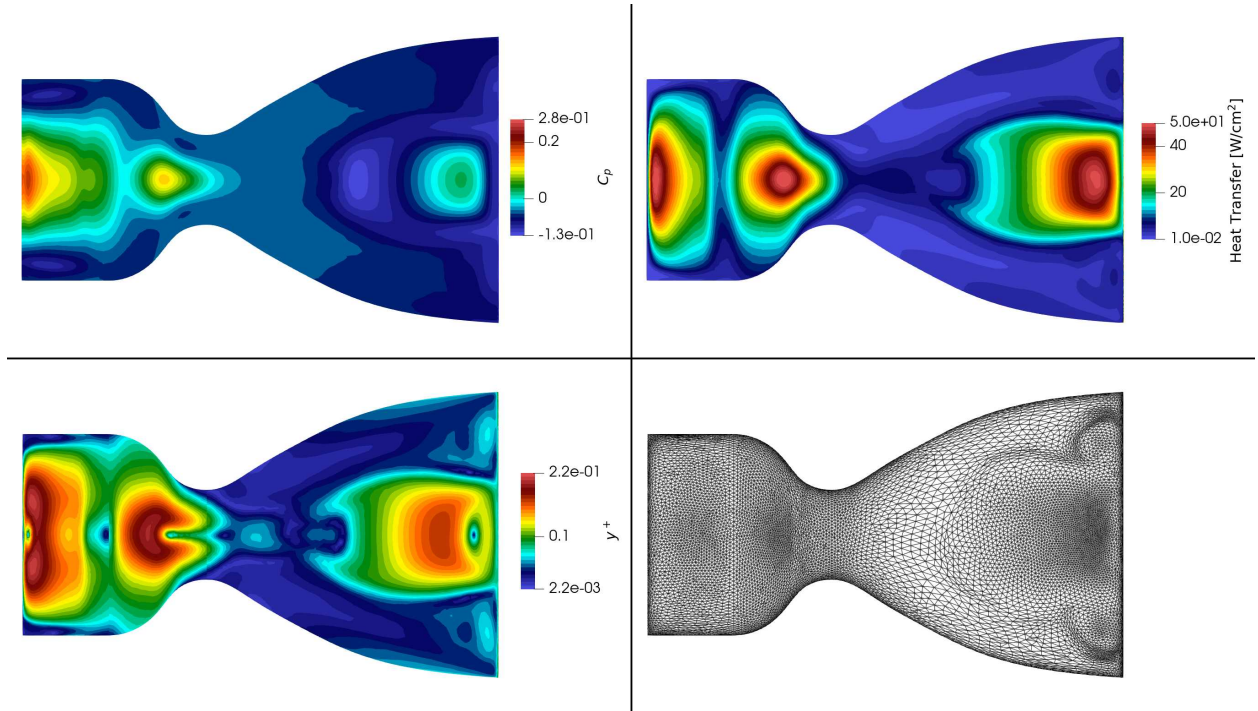


**Fig. 19** Supersonic flow over a four-engine rocket configuration plume results. From left to right: centerline grid slice, Mach contours, temperature contours, and specific heat ratio contours.






**Fig. 20** Supersonic flow over a four-engine rocket configuration base surface results. Top Left: Surface pressure coefficient. Top Right: Heat transfer. Bottom Left:  $y^+$ . Bottom Right: Surface grid.



**Fig. 21** Supersonic flow over a four-engine rocket configuration external engine wall facing stagnation point results. Top Left: Surface pressure coefficient. Top Right: Heat transfer. Bottom Left:  $y^+$ . Bottom Right: Surface grid.



**Fig. 22** Supersonic flow over a four-engine rocket configuration base surface line plots. Isothermal internal engine wall results are plotted in blue and adiabatic internal engine wall results are plotted in red. Left: Base pressure. Right: Base heat transfer.

### E. Hypersonic Flow over Dream Chaser®


To further examine the applicability and automation of this approach, a realistic example with complex surface curvatures and features is explored using Sierra Space's Dream Chaser Spaceplane. For this high speed test case, a 5 km/s freestream velocity is enforced at an angle of attack of  $40^\circ$  with zero side slip. This velocity equates to Mach 15.8. The C2S+Hybrid approach is employed for adaptation.

The wall boundaries are modeled as reaction cured glass (RCG) coating and assumed to be in radiative equilibrium. Laminar flow is modeled using a 5-species, one-temperature gas model. The HLLE++ inviscid flux scheme [8] is employed with the van Albada limiter [17]. Edge-based viscous terms are utilized [18]. A conservative local CFL of 2 is used due to the complexity of Dream Chaser's outer mold line. The local CFL can likely be increased to further reduce the solution time. One-thousand iterations are used for each C2S phase run and two-thousand iterations are used for each hybrid phase run. In general, the residuals are reduced by 3 to 4 orders of magnitude without limiter freezing. The heat flux is observed after each cycle. For the C2S simulation, the forward part of the vehicle saw little change after 5 adaptations, whereas the aft portion takes 12 cycles to converge. During the hybrid grid simulation, the whole vehicle takes 12 cycles to converge.


The grid size increases from an initial 0.1 million grid points to 5 million grid points using pure tetrahedral adaptation. As mentioned above, full vehicle heat flux convergence is achieved after 12 cycles. This is followed by extracting the adapted surface grid and generating 30 prism layers with a volume grid using Fidelity Pointwise [28]. The first cell height is prescribed to ensure the cell Reynolds number is unity. The resulting hybrid grid has 2 million grid points, which is used for the first cycle. After 5 cycles, the grid reaches 7 million points and remains at a fixed complexity for the subsequent cycles.

Figure 23 provides a zoomed-in view of the adapted grid and contours near the nose of Dream Chaser. At this velocity and altitude, molecular nitrogen partially dissociates and is adequately captured. At the final grid size, the bow shock is well-resolved which can be seen in both the temperature and Mach contours. The grid size could likely be reduced without sacrificing results. Figure 24 depicts the full flow field around Dream Chaser. Multiple shock structures are captured along with a detailed wake region. Figure 25 shows surface contours of pressure and heat flux. The surface grid adaptation has captured smooth gradients for each variable. The adapted surface grid is also shown.


Results are also compared to Data-Parallel Line Relaxation (DPLR) [36], which is considered a gold standard hypersonic structured CFD code. Details of the DPLR simulation have previously been published [20]. Line plots of the centerline and a fixed x-station location of the surface pressure and heating are provided in Figures 26 and 27. A hybrid grid approach with a relatively uniform surface grid [20] is included in these comparisons. All three cases are in agreement. It is worth noting that the benefit of C2S+Hybrid approach is that the analyst is not required to manually generate a grid. Also, the resulting adapted surface grid requires a third of the grid points compared to manually generating a surface grid while still capturing detailed surface gradients.




**Fig. 23** Sierra Space's Dream Chaser centerline slice near the nose surface showing the grid and contours of atomic nitrogen mass fraction, temperature, and Mach number, respectively.  
*Credit: Sierra Space Corporation*



**Fig. 24** Sierra Space's Dream Chaser centerline slice showing the grid and Mach number contour.  
*Credit: Sierra Space Corporation*




**Fig. 25** Sierra Space's Dream Chaser surface contours for pressure coefficient, heat transfer, and grid.  
*Credit: Sierra Space Corporation*



**Fig. 26** Sierra Space's Dream Chaser C2S+Hybrid centerline comparisons to Hybrid approach and DPLR.  
*Credit: Sierra Space Corporation*





**Fig. 27** Sierra Space’s Dream Chaser C2S+Hybrid x-station location slice comparisons to Hybrid approach and DPLR.

*Credit: Sierra Space Corporation*

### 1. Computational Performance

The *refine* and FUN3D processes are split between CPU and GPU nodes. *refine* uses five CPU nodes, each consisting of an AMD 64-core EPYC 7742 CPU, while FUN3D runs on a single GPU node with eight NVIDIA A100 SXM4 GPUs. A total of 12,000 CFD iterations are run during the C2S phase and 24,000 CFD iterations during the hybrid phase. More CFD iterations were needed for the hybrid grid to converge the prism layer cells. Table 3 provides the breakdown of run times. It can be seen that *refine* is two thirds of the total run time. A fully converged solution is achieved in less than 5.5 hours using approximately 16 A100-hours and 18 64-core EPYC 7742 CPU-hours.

**Table 3** Hypersonic Flow over Dream Chaser C2S+Hybrid computational performance. *refine* is run on 5 CPU nodes, where each node consists of an AMD 64-core EPYC 7742 CPU, and FUN3D is run on 1 GPU node, which consists of 8 NVIDIA A100 SXM4 GPUs.

*Credit: Sierra Space Corporation*

Process	Value
Bootstrap [min]	2
Pointwise [min]	7
C2S [min]	155
Hybrid [min]	162
Total [min]	326
CFD [min]	107
<i>refine</i> [min]	210
CFD [%]	33
<i>refine</i> [%]	64

## V. Summary and Future Work

An adaptive mixed-element open advancing-layer capability has been developed and demonstrated across the speed range for viscous flow applications. The mixed-element grid results better predict skin friction and heat transfer compared with experiments and structured grid results over fully tetrahedral grid results with significantly fewer grid points. The boundary-layer grid first cell height can be determined based on the flow solution or prescribed manually (e.g., using traditional turbulent flat plate  $y^+$  relations). The approach has also been demonstrated using GPU hardware for CFD, leading to total run times of less than 2 hours for a simulation of hypersonic flow over a capsule. These improvements are aimed toward eventual full automation of the grid generation process from clean CAD to enable significantly faster design cycles for next-generation aerospace vehicles.

Future work involves further integration of the advancing-layer process in the refinement process by extruding the boundary-layer grid for every adaptation cycle similar to the closed-advancing-layer approach in Ref. [49]. This would reduce any grid adaptation errors associated with shocks moving based on the better-resolved boundary layer; it is assumed the first tetrahedral adaptation process leads to a sufficiently refined surface and boundary-layer grid and that any shocks are well-captured. If a better-resolved boundary layer leads to a different separation zone, for example, the process described here will not properly refine the separation point, whereas, more frequent integration eliminates this deficiency. Improved computational performance (with GPU-capability) and surface grid quality for *refine* is also desired. Asynchronous refinement could further improve performance. While a CFD simulation is in progress, an intermediate flow solution can be used for the refinement process to generate a new grid concurrently; the final flow solution can be interpolated onto this new grid. For steady state analyses, this should converge to the same answer.

## Acknowledgments

This research was sponsored by the NASA Transformational Tools and Technologies (TTT) Project of the Transformative Aeronautics Concepts Program under the Aeronautics Research Mission Directorate. Resources supporting this work were provided by the NASA Langley K cluster and by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center.

## References

- [1] Mavriplis, D., “Adaptive Meshing Techniques for Viscous Flow Calculations on Mixed Element Unstructured Meshes,” *International Journal for Numerical Methods in Fluids*, Vol. 34, No. 2, 2000, pp. 93–111.
- [2] Park, M. A., Loseille, A., Krakos, J., Michal, T. R., and Alonso, J. J., “Unstructured Grid Adaptation: Status, Potential Impacts, and Recommended Investments Towards CFD 2030,” *AIAA Paper 2016-3323*, 2016.
- [3] Slotnick, J., Khodadoust, A., Alonso, J., Darmofal, D., Gropp, W., Lurie, E., and Mavriplis, D., “CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences,” *NASA CR-2014-218178*, 2014.
- [4] Candler, G., Barnhardt, M., Drayna, T., Nompelis, I., Peterson, D., and Subbareddy, P., “Unstructured Grid Approaches for Accurate Aeroheating Simulations,” *AIAA Paper 2007-3959*, 2007.
- [5] McCloud, P. L., “Best Practices for Unstructured Grid Shock Fitting,” *AIAA Paper 2017-1149*, 2017.
- [6] Tang, C., “Rapid Hypersonic Simulations using US3D and Pointwise,” *International Conference on Computational Fluid Dynamics (ICCFD11-3205)*, 2022.
- [7] McQuaid, J. A., and Brehm, C., “Heat Flux Predictions for Hypersonic Flows with an Overset Near Body Solver on an Adaptive Block-Structured Cartesian Off-Body Grid,” *Computers & Fluids*, Vol. 269, 2024, p. 106121.
- [8] Nastac, G., Nielsen, E., and Tramel, R., “Improved Heat Transfer Prediction for High-Speed Flows over Blunt Bodies using Adaptive Mixed-Element Unstructured Grids,” *AIAA Paper 2022-0111*, 2022.
- [9] Pederson, C., and Schoenenberger, M., “Exploring the Accuracy of RANS Simulations for Mars Entry Vehicles,” *AIAA Paper 2023-3695*, 2023.
- [10] Anderson, W. K., Biedron, R. T., Carlson, J.-R., Derlaga, J. M., Druyor Jr, C. T., Gnoffo, P. A., Hammond, D. P., Jacobson, K. E., Jones, W. T., Kleb, B., Lee-Rausch, E. M., Liu, Y., Nastac, G. C., Rumsey, C. L., Thomas, J. L., Thompson, K. B., Walden, A. C., Wang, L., Wood, S. L., Wood, W. A., and Zhang, X., *FUN3D Manual: 14.1*, NASA TM 20240006306, 2024.



- [11] McBride, B. J., Gordon, S., and Reno, M. A., “Coefficients for Calculating Thermodynamic and Transport Properties of Individual Species,” *NASA TM 4513*, 1993.
- [12] Gnoffo, P. A., Gupta, R. N., and Shinn, J. L., “Conservation Equations and Physical Models for Hypersonic Air Flows in Thermal and Chemical Nonequilibrium,” *NASA TP-2867*, 1989.
- [13] Diskin, B., Pandya, M. J., Pomeroy, B. W., and Boyett, T., “USM3D-ME Revised SA-neg Implementation and Assessment for Seventh Drag Prediction Workshop Solutions,” *AIAA Paper 2023-3252*, 2023.
- [14] Nastac, G., Walden, A., Nielsen, E. J., and Frendi, A., “Implicit Thermochemical Nonequilibrium Flow Simulations on Unstructured Grids using GPUs,” *AIAA Paper 2021-0159*, 2021.
- [15] Roe, P. L., “Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes,” *Journal of Computational Physics*, Vol. 43, No. 2, 1981, pp. 357–372.
- [16] Burg, C., “Higher Order Variable Extrapolation For Unstructured Finite Volume RANS Flow Solvers,” *AIAA Paper 2005-4999*, 2005.
- [17] Van Leer, B., “Towards the Ultimate Conservative Difference Scheme. V. A Second-Order Sequel to Godunov’s Method,” *Journal of Computational Physics*, Vol. 32, No. 1, 1979, pp. 101–136.
- [18] Liu, Y., Diskin, B., Nishikawa, H., Anderson, W. K., Nastac, G., Nielsen, E. J., and Wang, L., “Edge-Based Viscous Method for Mixed-Element Node-Centered Finite-Volume Formulation,” *AIAA Paper 2022-4083*, 2022.
- [19] Zubair, M., Ranjan, D., Walden, A., Nastac, G., Nielsen, E., Diskin, B., Paterno, M., Jung, S., and Davis, J. H., “Efficient GPU Implementation of Automatic Differentiation for Computational Fluid Dynamics,” *2023 IEEE 30th International Conference on High Performance Computing (HiPC)*, 2023.
- [20] Nastac, G., Walden, A., Wang, L., Nielsen, E. J., Liu, Y., Opgenorth, M., Orender, J., and Zubair, M., “A Multi-Architecture Approach for Implicit Computational Fluid Dynamics on Unstructured Grids,” *AIAA Paper 2023-1226*, 2023.
- [21] Park, M. A., “Anisotropic Output-Based Adaptation with Tetrahedral Cut Cells for Compressible Flows,” Ph.D. thesis, Massachusetts Institute of Technology, Sep. 2008.
- [22] Park, M. A., and Darmofal, D., “Parallel Anisotropic Tetrahedral Adaptation,” *AIAA Paper 2008-917*, 2008.
- [23] Marcum, D., and Gaither, J., “Mixed Element Type Unstructured Grid Generation for Viscous Flow Applications,” *AIAA Paper 1999-3252*, 1999.
- [24] Kleb, W. L., Park, M. A., Wood, W. A., Bibb, K. L., Thompson, K. B., and Gomez, R. J., “Sketch-to-Solution: An Exploration of Viscous CFD with Automatic Grids,” *AIAA Paper 2019-2948*, 2019.
- [25] Haimes, R., and Dannenhoffer, J., “The Engineering Sketch Pad: A Solid-Modeling, Feature-Based, Web-Enabled System for Building Parametric Geometry,” *AIAA Paper 2013-3073*, 2013.
- [26] Hang, S., “TetGen, a Delaunay-Based Quality Tetrahedral Mesh Generator,” *ACM Trans. Math. Softw.*, Vol. 41, No. 2, 2015, p. 11.
- [27] Dey, S., Aubry, R. M., Karamete, B. K., and Mestreau, E. L., “Capstone: A Geometry-Centric Platform to Enable Physics-Based Simulation and System Design,” *Computing in Science & Engineering*, Vol. 18, No. 1, 2015, pp. 32–39.
- [28] Cadence Design Systems, “Fidelity Pointwise,” 2024. URL [https://www.cadence.com/en\\_US/home/tools/system-analysis/computational-fluid-dynamics/fidelity.html#fidelity-pointwise](https://www.cadence.com/en_US/home/tools/system-analysis/computational-fluid-dynamics/fidelity.html#fidelity-pointwise).
- [29] NASA, “pyrefine,” 2024. URL <https://github.com/nasa/pyrefine>.
- [30] Ladson, C. L., “Effects of Independent Variation of Mach and Reynolds Numbers on the Low-Speed Aerodynamic Characteristics of the NACA 0012 Airfoil Section,” *NASA TM-1988-4074*, 1988.
- [31] Wang, L., Diskin, B., Nielsen, E., and Liu, Y., “Improvements in Iterative Convergence of FUN3D Solutions,” *AIAA Paper 2021-0857*, 2021.
- [32] Loseille, A., Dervieux, A., Frey, P., and Alauzet, F., “Achievement of Global Second Order Mesh Convergence for Discontinuous Flows with Adapted Unstructured Meshes,” *AIAA Paper 2007-4186*, 2007.

- [33] Rumsey, C. L., “NASA Langley Turbulence Modeling Resource,” <https://turbmodels.larc.nasa.gov>, 2024. Last Accessed November 1, 2024.
- [34] Spalart, P., and Allmaras, S., “A One-Equation Turbulence Model for Aerodynamic Flows,” *30th Aerospace Sciences Meeting and Exhibit 1992-439*, 1992.
- [35] MacLean, M., Mundy, E., Wadhams, T., Holden, M., and Parker, R., “Analysis and Ground test of Aerothermal Effects on Spherical Capsule Geometries,” *AIAA Paper 2008-4273*, 2008.
- [36] Wright, M. J., Candler, G. V., and Bose, D., “Data-Parallel Line Relaxation Method for the Navier-Stokes Equations,” *AIAA Journal*, Vol. 36, No. 9, 1998, pp. 1603–1609. <https://doi.org/10.2514/2.586>.
- [37] Hollis, B. R., and Borrelli, S., “Aerothermodynamics of Blunt Body Entry Vehicles,” *Progress in Aerospace Sciences*, Vol. 48, 2012, pp. 42–56.
- [38] NVIDIA, “Multi-Process Service,” [https://docs.nvidia.com/deploy/pdf/CUDA\\_Multi\\_Process\\_Service\\_Overview.pdf](https://docs.nvidia.com/deploy/pdf/CUDA_Multi_Process_Service_Overview.pdf), 2024. Last Accessed November 1, 2024.
- [39] Schmitt, V., “Pressure Distributions on the ONERA M6-wing at Transonic Mach Numbers, Experimental Data Base for Computer Program Assessment,” *AGARD AR-138*, 1979.
- [40] Musial, N. T., and Ward, J. J., *Base Flow Characteristics for Several Four-Clustered Rocket Configurations at Mach Numbers from 2.0 to 3.5*, NASA TN-D-1093, 1961.
- [41] Mehta, M., Canabal, F., Tashakkor, S. B., and Smith, S. D., “Numerical Base Heating Sensitivity Study for a Four-Rocket Engine Core Configuration,” *Journal of Spacecraft and Rockets*, Vol. 50, No. 3, 2013, pp. 509–526.
- [42] Zhou, Z., Bao, Y., Sun, P., and Le, G., “Afterburning Effect on Thermal Environment of Four-Engine Liquid Rockets at Different Altitudes,” *Engineering Applications of Computational Fluid Mechanics*, Vol. 15, No. 1, 2021, pp. 1134–1146.
- [43] McBride, B. J., and Gordon, S., *Computer Program for Calculation of Complex Chemical Equilibrium Compositions and Applications*, Vol. 2, NASA Lewis Research Center, 1996.
- [44] Nastac, G., Korzun, A., Walden, A., Nielsen, E. J., Jones, W., and Moran, P., “Computational Investigation of the Effect of Chemistry on Mars Supersonic Retropropulsion Environments,” *AIAA Paper 2022-2299*, 2022.
- [45] Korzun, A., Nastac, G., Walden, A., Nielsen, E. J., Jones, W., and Moran, P., “Application of a Detached Eddy Simulation Approach with Finite-Rate Chemistry to Mars-Relevant Retropropulsion Operating Environments,” *AIAA Paper 2022-2298*, 2022.
- [46] Nastac, G., Ernst, Z., Hickey, A., Walden, A., Jacobson, K., Jones, W., Nielsen, E. J., Diskin, B., Wang, L., Korzun, A., et al., “Closed-Loop Simulations of Human-Scale Mars Lander Descent Trajectories on Frontier,” *AIAA Paper 2024-3535*, 2024.
- [47] Nielsen, E. J., Walden, A., Nastac, G., Wang, L., Jones, W., Lohry, M., Anderson, W. K., Diskin, B., Liu, Y., Rumsey, C. L., et al., “Large-Scale Computational Fluid Dynamics Simulations of Aerospace Configurations on the Frontier Exascale System,” *AIAA Paper 2024-3866*, 2024.
- [48] Nastac, G., and Frendi, A., “Numerical Investigation of Gas Models for Retropropulsion Flows for Future Mars Missions,” *Journal of Spacecraft and Rockets*, 2024, pp. 1–9.
- [49] Alauzet, F., Loseille, A., and Marcum, D., “On a Robust Boundary Layer Mesh Generation Process,” *AIAA Paper 2017-0585*, 2017.