

Scale-Resolving Simulations of a Supersonic Retro-Propulsion Concept For Mars Entry, Descent, and Landing

Francois Cadieux*, Scott Neuhoﬀ†, Michael F. Barad‡ and Cetin C. Kiris§
NASA Ames Research Center, Moﬀett Field, CA, 94063

Supersonic retro-propulsion (SRP) is a deceleration technology that could enable larger payloads and potentially humans to be brought from space to the surface of Mars safely. Accurate and reliable CFD predictions are needed to help design future Mars entry, descent, and landing (EDL) vehicles due to the cost and limitations of wind tunnel tests. Typical CFD approaches struggle at points along the Mars EDL trajectory with high enough altitude (low enough ambient pressure) that the SRP motor plumes are over-expanded. This work demonstrates how high-order scale-resolving simulations on Cartesian AMR grids can provide accurate and reliable time-averaged aerodynamic loads on the vehicle for this challenging flow regime. Greater than second order convergence in the time-averaged integrated axial force C_A is obtained across the three finest mesh resolutions simulated, and the estimated uncertainty is smaller than the unsteady standard deviation – demonstrating the accuracy and reliability of the method. Although costly, this approach can yield higher confidence in C_A than previously obtained in a reasonable turnaround time and is within reach for a few points in the trajectory (i.e. Mach 2.4, 2.52, and 2.78). Further computational performance and numerical methods improvements are needed to reduce the cost enough to bring this method into the EDL engineering design cycle.

I. Nomenclature

C_p	=	pressure coefficient
C_A	=	axial force coefficient
C_y	=	side force coefficient in the y direction
C_{M_y}	=	moment coefficient in the y direction
$\langle q \rangle$	=	time-averaged value of any time-varying quantity $q(t)$
$\sigma(q)$	=	standard deviation of any time-varying quantity $q(t)$
D	=	diameter
Δt	=	time step
Δx	=	distance between cell centers along a Cartesian direction

II. Introduction

SUPERSONIC retro-propulsion (SRP) is a deceleration technology that could enable larger payloads and potentially humans to be brought from space to the surface of Mars safely [1]. There have been a number of wind tunnel tests dating back to the beginning of the space race to explore the physics of SRP and its potential for planetary entry, descent, and landing (EDL) [2–5]. More recently, the focus has shifted toward experiments focusing on the dynamics of SRP [6–8] and to provide validation benchmarks for computational fluid dynamics (CFD) predictions [9–11]. Unfortunately, these experiments do not reflect the different composition of Mars’ atmosphere, do not simulate the retrorockets’ exhaust, and do not quite match the Reynolds number or Mach number of flight vehicles being considered for Mars EDL. Although these deficiencies could be addressed in future wind tunnel tests, the cost of such tests are orders of magnitude higher than CFD simulations. There is an urgent need to develop best practices so that CFD can be

*Aerospace Engineer, Computational Aerosciences Branch, N-258, AIAA Member.

†Research Scientist, Computational Aerosciences Branch, N-258, AIAA Member.

‡Aerospace Engineer, Computational Aerosciences Branch, N-258.

§Branch Chief, Computational Aerosciences Branch, N-258, AIAA Senior Member.

relied upon to design the EDL system and reduce the need for comparatively costly wind tunnel tests. Multi-year efforts have yielded significant advancements toward validating steady Euler CFD predictions as well as steady and unsteady Reynolds-Averaged Navier-Stokes CFD predictions against these benchmark wind tunnel experiments [12–15]. Best practices developed through these efforts have shown good agreement between CFD and experiment when considering under-expanded plume conditions, where the nozzle exit pressure p_e is lower than the surrounding pressure (which can be different than p_∞ due to the presence of a bow shock). These under-expanded plume conditions typically give rise to stable and relatively static (or at worst periodic) bow shock and plume structures. However, at altitudes where SRP would be first initiated, this pressure ratio is likely to be lower, and the plumes are expected to be over-expanded (p_e is larger than the surrounding pressure) – yielding the so-called long-jet penetration mode causing the chaotic interactions between the plume and the bow shock. Predicting the effects of these interactions on the vehicle aerodynamic loads with CFD then becomes costly due to the need for a time-accurate methodology, and a finer spatial resolution over a larger volume upstream of the vehicle. Nevertheless, if CFD is to drive the design of future Mars EDL systems, CFD predictions must be accompanied with a measure of confidence. This is most often achieved with a mixture of convergence studies in both temporal and spatial resolution, and code-to-code verification, where different CFD solvers and mesh paradigms (e.g. structured overset vs unstructured) are used to solve the same problem. Used in tandem, grid convergence studies and code-to-code verification can provide both confidence in the predictions and the level of uncertainty associated with each prediction.




Fig. 1 Conceptual vehicle used for this study of SRP for Mars EDL

A recent steady and unsteady Reynolds-averaged Navier-Stokes (RANS and URANS) simulation campaign performed code-to-code verification of aerodynamic loads for a full-scale low L/D vehicle concept (see Figure 1) in Mars atmosphere firing eight liquid oxygen-methane SRP motors (two species: Mars atmosphere and exhaust mixture, frozen chemistry) [16]. The authors of that study saw excellent code-to-code agreement in unpowered simulations, but noted some disagreements for powered simulations in the over-expanded plume regime spanning the altitudes and Mach numbers (2.4 to 2.78) associated with points in the trajectory where SRP could be first initiated. Different CFD solvers predicted different trends for integrated axial force coefficient C_A as altitude and Mach number increased. C_A consists of the integrated drag force on the vehicle, ignoring the momentum generated by the SRP motors, non-dimensionalized by the force associated with the oncoming freestream flow $\rho_\infty U_\infty^2 \pi (D/2)^2$. Given the 8 SRP rocket engines each provide 80 kN of thrust, the C_A metric defined in this way provides a much more sensitive figure of merit for comparing CFD results across mesh resolutions, solvers, and methodologies (URANS vs scale-resolving methods). In other recent work, the same vehicle was adapted to be simulated with air-only using the detached-eddy simulation (DES) methodology – a flavor of hybrid RANS and large-eddy simulation (LES). This approach was applied across a wide range of mesh resolutions and Mach numbers, including the aforementioned Mach 2.4 over-expanded plume case, in an effort to investigate sensitivity of predicted time-averaged C_A to grid resolution [17]. Although the authors conclude that the “Mach 2.4 cases show the smallest change with increasing spatial resolution”, the reported time-averaged C_A did not converge for the over-expanded case despite reaching resolutions over 1 billion cells [17].

The overarching goal of this work is to provide an alternate path to reliable CFD predictions for SRP in the over-expanded plume regime using scale-resolving simulations. To this end, we leverage the Launch, Ascent, and Vehicle Aerodynamics (LAVA) framework’s Cartesian adaptive mesh refinement (AMR) module to simulate the same

Table 1 Freestream conditions in Mars atmosphere

Case	M_∞	V_∞ (m/s)	T_∞ (K)	p_∞ (Pa)	q_∞ (Pa)	ρ_∞ (kg/m ³)
1	2.52	587.63	216.19	273.37	1155.67	0.0067
2	2.40	560.60	216.53	284.73	1093.82	0.0070
6	2.78	642.78	214.95	235.35	1212.83	0.0058

conceptual flight vehicle – see Figure 1. We focus solely on the points in the trajectory for which disagreement among CFD predictions was previously observed [16], namely the three Mach numbers and associated freestream conditions in the over-expanded plume regime – see Table 1. Section III describes the conceptual vehicle, SRP motors, and freestream conditions at the Mach numbers of interest. The specifics of the methodology are explained in Section IV. LAVA’s predicted trends in C_A as the Mach number and altitude decrease are explained in Section V.B. We then investigate grid convergence of time-averaged C_A for the Mach 2.4 case in Section V.C and explain its impact on the uncertainty associated with predictions at other Mach numbers. The computational cost of each of these simulations is detailed in Section VI. To conclude, a summary of our findings is presented along with notes on future directions.

III. Problem Description

The conceptual SRP vehicle is the low lift-to-drag ratio vehicle detailed in previous works [16, 17]. It has a 16.4 meter diameter (D) Hypersonic Inflatable Aerodynamic Decelerator (HIAD) heatshield with a 9.1 m diameter rigid connecting body and an aft-mounted surface habitat payload as shown in Figure 1. Despite the heatshield being inflatable, it is considered as a rigid, non-deforming body for the purposes of CFD analysis.

The conditions considered for this analysis are described in Table 1 and are representative of Mars atmosphere conditions at relevant altitudes.

The SRP engines are each throttled to 80% of their stated maximum thrust of 100 kN specifically to allow operating margin for dispersions, mass growth, and differential throttling (to maintain attitude control). Total conditions at the plenum upstream of the throat are given as $p_{0,j} = 1200$ pounds per square inch (psi), $T_{0,j} = 3581.6$ K. An axisymmetric chemical equilibrium reacting flow simulation was used to obtain the combustion exhaust gas mixture thermodynamic properties given the oxidizer-to-fuel ratio of 3.5, and the unscarfed nozzle area ratio of 177:1. The thermodynamic properties of the exhaust mixture are as follows: $\gamma = 1.2$, molecular weight is 24.07 kg/kmol. The expected mass flow (\dot{m}) at the nozzle exit plane is 19.8388 kg/s, with a thrust of 73.8715 kN, and an area-averaged Mach number of 4.869. Only the doublet-A nozzle configuration, as seen in Figure 1 is investigated here [16]. It is important to note that each of the SRP nozzles is canted by 5 degrees radially from the vehicle axis – the nozzles point slightly outward. For the exhaust gas transport properties, we assumed a Prandtl number of 0.7592, and a viscosity equivalent to that of carbon dioxide gas given by Sutherland’s law with reference viscosity 2.018e-5 Pa s at reference temperature of 292.25 K and Sutherland temperature of 127 K. Given the high temperature of the exhaust and the high Reynolds number nature of the flow, it can be shown that the numerical viscosity required to capture the shocks and strong gradients in the plume dwarf the physical viscosity experienced by the flow, so no attempt was made to obtain a more representative Prandtl number or viscosity. It is also assumed that the exhaust gas mixture behaves as a single, non-reacting ideal gas. Such an assumption used in CFD drastically reduces the complexity of the equation set that needs to be solved to the multi-species Navier-Stokes formulation at the cost of accuracy in predictions of temperatures in the plume due to the absence of after-burning, soot, and heat radiation. Given that the goal is to obtain time-averaged loads on the vehicle, as opposed to surface heating for example, this assumption is justified.

IV. Methodology

The LAVA Cartesian solver is chosen to simulate SRP based on previous successes in scale-resolving simulations of launch vehicles like the Space Launch System and the Orion Launch Abort System. It solves the non-reacting multi-species formulation of the Navier-Stokes equations where species’ mass fractions are solved for in addition to conservation of mass, momentum and energy: a gas mixture representative of the Martian atmosphere (mostly carbon dioxide) and another gas mixture representative of the liquid oxygen methane SRP motor exhaust. The solver makes use of a fifth-order weighted essentially non-oscillatory shock-capturing scheme with incremental stencil (WENO5-IS) to reconstruct convective fluxes at cell faces using Rusanov flux-vector splitting and characteristic transformation [18, 19].

This scheme provides the necessary robustness at strong shocks and already provides more numerical dissipation in turbulent flows than a typical modern subgrid-scale (SGS) model would [20], so no SGS closure model is used. Viscous fluxes are computed using traditional second order centered fluxes. Explicit time-integration is used with the strong stability preserving third order Runge-Kutta method. Each nested Cartesian refinement level has half the Δx and Δt of its parent, and is worked on separately by all available processors thanks to our subcycling algorithm for time advancement. Coarse-fine interfaces are treated with conservative second order coarsening, and the prolongation operator uses linear interpolation in time and space to fill fine values under coarse.

The farfield boundary conditions consist of a supersonic inlet specified by Mach number, pressure, temperature and mass fraction at the $x = -163.84$ m plane or $x \approx -10D$, and all other outer boundaries use a simple supersonic outlet boundary condition with all variables extrapolated from the interior values. The domain spans 327.68 m or $20D$ in each direction and is centered on the center of mass of the vehicle.




Fig. 2 Overview of domain and Cartesian grid used for unpowered simulations. Each box contains 32^3 cells. The logarithm of pressure shows the location of the bow shock and the wake structure.

Given the nozzles' rather extreme area ratio of 177:1 and our explicit time integration approach, it would be very expensive to resolve the flow in the nozzle throat – regardless of the mesh paradigm – within the larger vehicle simulation due to the two orders of magnitude smaller cell size (and thus time step Δt) it would require locally. This expense is avoided by instead performing a precursor simulation of the axisymmetric part of the unscarfed nozzle by itself and using the resulting radially-varying state vector at the exit as a supersonic inlet boundary condition face inside the larger vehicle simulation. The precursor truncated nozzle is simulated using the LAVA unstructured steady state RANS capability with the Spalart-Allmaras turbulence model. It takes full advantage of the axisymmetric nature of the nozzle to drastically reduce grid requirements. The unstructured grid is wall-resolved with a near-wall spacing in viscous units $\Delta y^+ < 1$ across the entire nozzle surface, and grows to a cell volume no greater than our best practice Cartesian mesh inside the nozzle diameter $\Delta x/(40D)$. The simulation uses the plenum total conditions and the exhaust gas mixture described in Section III and the far-field conditions are commensurate with the Mars atmosphere corresponding to freestream conditions for Mach 2.4. Since the flow in the nozzle is supersonic from the throat to the nozzle exit, the freestream conditions are irrelevant to the information we will be extracting from the simulation. The simulation is iterated until the residuals have dropped by at least 5 orders of magnitude. The initial results yielded a 3% higher mass flow rate and 13% higher thrust than the target values predicted by the chemical equilibrium nozzle code – see Section

III. Given the RANS simulation does not take into account any aspects of combustion or heat radiation, this overshoot is primarily due to an overshoot in area-averaged temperature at the exit, as expected. To address this problem, we use the best practices that were developed for simulations of launch vehicles like SLS: we use the isentropic nozzle equations to find a more appropriate throat area, total pressure and total temperature to match the target area-averaged values given by the chemically reacting nozzle simulation tool. The throat area was scaled up by 1.12804, the total temperature increased by a factor of 1.26741, and the total pressure scaled down by 0.91155. With these changes, the precursor simulation is performed again and the mass flow rate and thrust obtained are now only 0.05% higher, and the area averaged Mach number is 0.11% higher than the target values.

We extract the full radially-varying state vector at the truncated nozzle exit face, which we use as a supersonic inlet boundary condition in the Cartesian simulation at the same location, which falls just upstream of the full scarfed nozzle exit plane. The rest of the vehicle is treated with a ghost-cell immersed-boundary method that imposes adiabatic walls, and makes use of a mixture of slip and no-slip boundary conditions where appropriate. For example, the nozzle wall boundary layer is extremely thin and cannot be modeled efficiently with Cartesian isotropic cells, so the slip boundary condition is used there. Similarly, on the surface of the unpowered vehicle's heat shield, the Reynolds number becomes extremely high and so slip is again utilized. To avoid a strong Prandtl-Meyer expansion which could lead to negative densities and pressures where the nozzle walls intersect the heat shield, no-slip is used on the heat shield surface surrounding the nozzle exit and spans 1 extra nozzle exit diameter (roughly 1 m) for powered simulations. On the shoulder and on the aft-end of the vehicle, the flow is driven by large turbulent recirculating regions, and so it is treated with no-slip boundary condition for both powered and unpowered simulations. Such engineering simplifications unfortunately continue to be necessary until robust turbulent wall-layer modeling becomes more mature for flows of this complexity and for Cartesian grids that cannot afford to have 10 to 20 points within the boundary layer over the entire vehicle and are thus necessarily under-resolved in certain regions.

Two sets of automatically generated meshes with user-specified refinement regions are used: one to obtain a converged initial condition for the unpowered, and one for the powered vehicle. The refinement regions are defined using flared cylinders and are illustrated in Figures 2 and 3 for the unpowered mesh, and Figure 4 for the powered mesh. The unpowered meshes are designed to capture the strong gradients from the bow shock all the way to the vehicle and the vehicle wake on the second finest mesh level (4 cm), and the finest mesh (2 cm) is concentrated solely on high curvature regions like the nozzles and the heat shield shoulder. This is done to minimize the cost of the simulations as the finest level necessarily is updated every time step, whereas coarser mesh levels are worked on only every 2^{N-l} time steps in our subcycling algorithm, where N is the total number of levels, and l is the level number and $l = N$ on the finest mesh level.

The same approach is used for the powered meshes where only the heat shield shoulder and four large flared cylinders to capture the plumes are refined to the finest mesh spacing (2 cm). The plumes span multiple vehicle diameters in length and so cannot affordably be entirely covered with fine mesh. That's where the second finest mesh level steps in, encompassing the entire length of the plume all the way to the bow shock, and downstream to capture the interactions of the plume turbulence as it convects downstream and interacts with the vehicle wake. The third finest mesh level is a longer and larger version of the 4 cm level designed to capture any run-away events where the bow shock might be pushed further ahead of the vehicle as was observed in previous DES results [17]. The other mesh levels are used to provide some minimal level of resolution in the larger footprint of the vehicle's plumes, bow shock, and wake. They use simple cylinder refinement regions because they are not updated as often and thus cheaper to compute. These meshes were designed with our best practices derived from previous work, namely to target a resolution of 40-50 cells across the nozzle exit and in most of the potential core of the plumes.

The typical workflow consists of simulating the unpowered vehicle, where the truncated nozzle face is treated with slip until mean aerodynamic drag converges, which typically occurs after 15-20 convective time units (D/V_∞), or roughly 0.5 seconds. Then in a separate simulation, we restart from the last checkpoint file in the unpowered simulation directory, reduce the time step, and re-grid, which automatically interpolates the solution onto the new much larger grid with plume refinements. From that moment on, the supersonic inlet boundary condition is applied at the truncated nozzle face, but is ramped from current conditions at that face to the full thrust conditions linearly over 0.04 seconds, which corresponds to less than 2 convective time units obtained from the oncoming flow speed. This was found to be necessary to avoid an extremely strong ignition over-pressure wave that would otherwise have severely restricted our time step due to the CFL condition. This is justified because the time-averaged loads are of interest for this work, and no engine can turn on and instantly reach its peak pressure – in previous work we have noted rockets often take up to 0.1 seconds to reach peak pressure. However, should realistic time-dependent total conditions become available at the nozzle plenum face, we could easily simulate the engine startup as designed. It is important to note that once




Fig. 3 Zoomed in view of Cartesian grid used for unpowered simulations. Each box contains 32^3 cells. The logarithm of pressure shows the location of the bow shock and the wake structure.




Fig. 4 Zoomed in view of Cartesian grid used for powered simulations. Each box contains 32^3 cells. The logarithm of temperature is shown to indicate the location of the plumes, bow shock, and the wake structure in this snapshot taken from the Mach 2.4 fine mesh simulation.

the engines are turned on, the turbulent time scales that the time-integration must capture accurately are driven by the

hot exhaust plume velocity, which measures on average over 3500 m/s, which is on the order of 5 to 6.25 times faster than the time-scales associated with the freestream. However, the time-scale for a fluid particle in the plume to reach the bow shock and be convected back downstream to pass by the vehicle is much longer, on the order of 0.08 to 0.09 seconds, which is what drives the need for long integration times to obtain confidence in time-averaged aerodynamic loads. We aim to simulate enough convective time units to obtain at least 0.5 seconds of integration after the simulation has reached a statistically stationary state (periodic or chaotic oscillations about a mean value).

V. Results

A. Unpowered Vehicle




Fig. 5 Time history of instantaneous axial force coefficient C_A for both *unpowered* and *powered* vehicle simulations at Mach 2.52 on fine mesh.

Using our best practice mesh (188 million cells) described in Figure 3, the three cases described in Table 1 were simulated using a time step $\Delta t = 1.7578125 \times 10^{-5}$ on our finest cells, corresponding to a CFL number of approximately 0.5 for all cases based on the freestream velocity. The flow accelerates considerably over the shoulder of the heat shield and reaches a maximum CFL of 0.8 to 0.95 there during the simulation, which is what drives this choice of time step using our explicit time integration approach. Figure 5 shows a typical time history of the integrated axial force coefficient C_A for unpowered in blue, and powered in orange. The final value that the unpowered LAVA solutions converge to for the three Mach numbers of interest is tabulated in Figure 6. The maximum percent difference in C_A across previous CFD predictions from literature [16] and across these three Mach numbers is 4.4% at Mach 2.52. The root-mean-square (RMS) percent difference is 1.7% across the same three cases and previously published results when we include LAVA predictions. Given the range of mesh paradigms (unstructured, structured overset, and Cartesian AMR) and CFD methods (steady RANS vs scale-resolving), this level of code-to-code agreement is excellent.

Figure 6 shows that the variation in our predicted C_A across the three Mach numbers is less than 0.46%. LAVA and previous CFD results all predict that the C_A is insensitive to the Mach number for the cases considered, at least to within associated uncertainty (grid convergence, or unsteady standard deviation). Although such good code-to-code agreement is to be celebrated, it is expected for this case because loads induced by supersonic flow over a blunt vehicle like this one are largely driven by the pressure distribution induced by the bow shock, which can be obtained to a fairly




Fig. 6 Trends in time-averaged integrated axial force coefficient $\langle C_A \rangle$ for *unpowered* vehicle on fine mesh across different freestream dynamic pressure conditions and their associated Mach numbers with full range in (a) and zoomed in to distinguish small differences in (b).

high degree of accuracy from normal shock relations.

B. Powered Vehicle




Fig. 7 Time history of instantaneous C_A for *powered* vehicle at 3 Mach numbers simulated with the fine mesh from SRP initiation in (a) and zoomed in over stationary state interval in (b). Dashed lines represent time-average $\langle C_A \rangle$ and shaded region shows $\langle C_A \rangle \pm \sigma(C_A)$, where σ is the unsteady standard deviation.

Starting from their respective statistically stationary state in the unpowered simulation, all three Mach numbers are simulated using the fine mesh (856 million cells) with the 8 nozzles at 80% of maximum thrust. The time step is reduced to $\Delta t = 1 \times 10^{-6}$ seconds as the SRP motors reach full thrust to ensure $CFL \leq 1$. The time history of the instantaneous axial force coefficient for each case is shown in Figure 7. A stationary state is reached for all cases after roughly 0.3 seconds. This corresponds to 3 convective time units associated with passive particles leaving the nozzle and whizzing by the vehicle once the motors are ramped up. Once the plumes settle to an equilibrium with the oncoming supersonic flow, the location of the bow shock chaotically oscillates about a new time-averaged distance

away from the vehicle. This ratio of momentum of the plumes to the oncoming flow dictates the bow shock location and in turn largely determines the pressure distribution over the vehicle. In order to obtain a time-averaged C_A , we numerically integrate its time history using the mid-point rule from $t = 0.3s$ until the end of available data and divide by the integral of time over the same interval. These values are tabulated in Figure 8 to highlight the relationship between C_A and Mach number for fixed thrust. The thrust provided by the motors remains constant for all cases. So as the Mach number increases, the dynamic pressure of the oncoming flow also increases, despite the associated change in altitude – see Table 1. Dynamic pressure is the relevant independent variable because it is a measure of the momentum of the incoming flow: of how much it pushes against the momentum generated by the SRP motors. The predicted time-averaged C_A from these simulations follows a linear trend with dynamic pressure: $\langle C_A \rangle$ increases linearly with q_∞ . Figure 9 illustrates how the increase in $\langle C_A \rangle$ is directly linked to the bow shock stand-off distance: the bow shock moves closer to the vehicle with increase in q_∞ , and causes a larger region of peak $\langle C_p \rangle$ between diametrically opposed nozzles. In reality, such a clean linear trend will rapidly break down as the freestream pressure increases (with lower altitude). At lower altitude, the plumes transition from over-expanded to under-expanded, and thus have a larger impact on the pressure on the heat shield in the vicinity of the nozzles, preventing pressure recovery and lowering the $\langle C_A \rangle$. However, at the Mach numbers and altitudes considered here, the plumes do not merge to block pressure recovery on the heat shield in part because of the nozzles' 5 degree outward cant angle and in part because of the fact that the plumes are over-expanded at these altitudes and Mach numbers. Thus the linear increase in $\langle C_A \rangle$ with q_∞ aligns with expectations because the time-averaged bow shock strength should increase, and its distance from the vehicle decrease, leading to increased pressure on the heat shield. Again, the expectation would be different should the plumes be under-expanded because it would increase their effective cross-sectional area near the vehicle, thereby reducing pressure recovery on the heat shield.




Fig. 8 Trends in time-averaged C_A for *powered* vehicle on fine mesh versus freestream dynamic pressure q_∞ . Unsteady standard deviation $\sigma(C_A)$ is shown as error bars. Predicted best fit linear relationship follows $\langle C_A \rangle = 0.0008q_\infty - 0.5209$ with a high confidence indicated by $R^2 = 0.9981$.

LAVA predictions for powered $\langle C_A \rangle$ fall within the range of predicted values across other CFD codes [16], and show fairly good agreement at the two higher Mach numbers with two out of the three previous unsteady RANS predictions. The trend in $\langle C_A \rangle$ predicted across the three Mach numbers is consistent with that of one of the previous CFD results, albeit the drop in drag as Mach number decreases is less pronounced in the LAVA results. It is important to reiterate here that these cases showed the most code-to-code disagreement in unsteady RANS results,




Fig. 9 Time-averaged pressure coefficient $\langle C_p \rangle$ on the vehicle surface and on cut plane through center-left nozzles taken from fine mesh results. Time-averaged bow shock standoff location measured as peak $\langle C_p \rangle$ yields $x \approx -27$ m, $x \approx -25$ m, and $x \approx -23$ m respectively.

and there is no wind tunnel data for this configuration or an accepted benchmark CFD solution, so a different metric for success is needed. Given that the time history of C_A demonstrates a convincing statistically stationary state for each Mach number for $t > 0.3$ s, we can look at each instantaneous value of C_A as an independent measurement and thus regard the unsteady standard deviation as a measure of uncertainty to our time-averaged C_A value. Although the histogram of C_A does not fit a Gaussian curve perfectly due to some low amplitude and low frequency oscillations, it is still representative of the spread, and thus of the range of possible values of C_A . For each of the three cases, the unsteady standard deviation is approximately 0.02 as seen in Table 2, which corresponds to at most 5.6% of the mean C_A value at Mach 2.4 where the C_A value is lowest. These unsteady standard deviation values are nearly one order of magnitude smaller than the differences between previous unsteady RANS results at any of these three Mach numbers [16]. It is clear that unsteady standard deviation alone is insufficient to provide high confidence in the expected value of C_A . Rather, it indicates that the dynamics of the interactions of the plumes with the bow shock do not influence the pressure recovery on the heat shield as much as previously gleaned from DES results [17]. Indeed, LAVA predicts oscillations in C_A with relatively small amplitudes and higher frequencies – as seen in its time history – with near zero asymmetry as shown in the low value for time-averaged side force coefficient $\langle C_y \rangle$ in Table 2. Unsteady side load and moment amplitudes are on the order of 10% of $\langle C_A \rangle$ in magnitude centered around a nearly zero mean. Given that the total force coefficient in the axial direction $\langle C_{A_{total}} \rangle$ is on the order of 585 when we include the momentum from the SRP motors, variations in moment of 0.05 are unlikely to negatively impact the vehicle’s attitude or trajectory during descent.

C. Grid Convergence Study

Grid convergence of mean integrated aerodynamic loads (like $\langle C_A \rangle$) is the most reliable means by which we can increase confidence in scale-resolving simulation results. Due to the cost associated with scale-resolving simulations, we limit our attention to the Mach 2.4 case where previous results are available [17]. Starting from our 856-million cell “fine” mesh shown in Figure 4 that was derived from our best practices, a medium mesh was created by stretching

Table 2 Time-average and standard deviation of forces and moments on *powered* vehicle for fine mesh results. Results in z-direction are commensurate in magnitude with y-direction and as such are omitted for brevity.

Mach Number	$\langle C_A \rangle$	$\sigma(C_A)$	$\langle C_y \rangle$	$\sigma(C_y)$	$\langle C_{M_y} \rangle$	$\sigma(C_{M_y})$
2.40	0.3601552	0.0203117	0.0003209	0.0065782	0.0121689	0.0409772
2.52	0.4136464	0.0176187	0.0001552	0.0049874	0.0075731	0.0433312
2.78	0.4560513	0.0251202	0.0006066	0.0045459	-0.0037890	0.0487426






Fig. 10 Time history of instantaneous C_A for *powered* vehicle at Mach 2.4 simulated across different mesh resolutions shown from SRP initiation in (a) and zoomed in over stationary state interval in (b). Dashed lines represent time-average $\langle C_A \rangle$ and shaded region shows $\langle C_A \rangle \pm \sigma(C_A)$, where σ is the unsteady standard deviation.

the entire domain by a factor of $2^{\frac{1}{3}}$ in each direction, but keeping the number of cells across the coarsest level the same, yielding 575 M cells, and an extra fine mesh was created by compressing the domain by the same factor resulting in 1.6B cells. This is different than a typical mesh refinement study where the domain is kept constant and only the cell sizes change. However, because our specified refinement regions are the same across all meshes, the relevant cell size does reach the target ratios exactly. This choice is primarily driven by the limitation that larger refinement units (boxes of 32^3 cells) give us better performance with the current AMR architecture, so we would be constrained to increase or decrease the coarsest level resolution by 32 cells or more in each direction, which wouldn't allow us to obtain the desired grid spacing ratio. The final coarse mesh is created by coarsening by a factor of 2 in all directions our fine mesh, but keeping the domain size the same, resulting in an 8 times smaller number of cells (144 M). This coarse mesh violates our best practices in that it does not have enough cells across the nozzle diameter to adequately resolve the radial gradients of the SRP motors' exhausts. And indeed, it is the only outlier in the trends of its C_A time history as shown in Figure 10. Medium, fine, and extra fine results otherwise all demonstrate clear and consistent trends in C_A : they all reach a statistically stationary state within 0.3 seconds of the SRP engines turning on, and their values are almost on top of each other. In contrast, the coarse mesh C_A results show a downward trend leading to a runaway event where diametrically opposed plumes merge and push the bow shock significantly further away from the vehicle, preventing pressure recovery on the heat shield.

An example of this is visible in Figure 11a where the cut plane through the center-left nozzles clearly shows the top and bottom nozzles' plumes merging via visualization of the Mach number. This phenomenon began around 0.25 seconds after ignition and continues dynamically with different diametrically opposed plumes merging and pushing the bow shock further away, driving larger oscillations in C_A at a lower frequency throughout the simulation. This plume merging behavior is never observed in any of the higher resolution simulations, which all look almost identical to Figure 11b. Similarly, the coarse mesh C_A history is clearly an outlier in both mean value and standard deviation



(a) Coarse mesh (144 M cells)



(b) Fine mesh (856 M cells)

Fig. 11 Snapshots from Mach 2.4 simulations on two different mesh resolution levels showing Mach number on the cut plane through the center-left nozzles. In (a), notice how diametrically opposed plumes merge and push the bow shock further away from the vehicle. This phenomenon does not occur in higher resolution simulations like (b): notice the wide range of scales captured from the fine vortices spanning only a few centimeters near the nozzles to the length of the plumes of nearly 2 vehicle diameters D or roughly 30 meters.

compared to the medium, fine, and extra fine results. We conclude from these observations that diametrically opposed plumes merging is a product of insufficient effective resolution of the plume shear layers. We specifically use the term “effective resolution” because the range of turbulent scales captured depends not only on the mesh, but on the dissipation characteristics of the method, whether the dissipation comes from a turbulence model or from the numerical scheme used for shock capturing and stabilization or a combination of both.




Fig. 12 Time-averaged $\langle C_A \rangle$ for *powered* vehicle at Mach 2.4 across all mesh refinements. Error bars indicate the unsteady standard deviation $\pm\sigma(C_A)$.

Time-averaged C_A are again computed for each simulation starting at 0.3 seconds and encompassing all available data. Figure 12 demonstrates convincing asymptotic behavior of mean aerodynamic loads. From our medium, fine and extra fine results, the order of convergence is computed according to ASME guidelines [21] and is tabulated in Table 3. It is important to note that we use the ratio of Δx directly instead of the “effective” grid spacing derived from the number of cells or the average cell volume because the domain size is not constant between all simulations. This is the first time to the authors’ knowledge that an apparent order of convergence greater than two ($p > 2$) is obtained for powered supersonic retro-propulsion simulations of this vehicle. More important is the estimated converged $\langle C_A \rangle$ and the relatively low magnitude of the estimated grid-related uncertainty (± 0.0143), which is lower than the unsteady standard deviation in C_A in both the fine ($\sigma_{\text{fine}} = 0.0203117$) and extra fine mesh simulations ($\sigma_{\text{extra fine}} = 0.0203292$) which are displayed as error bars ($\pm\sigma$) in Figure 12. It is also particularly encouraging to see that the change in time-averaged C_A between the fine mesh to the estimated converged C_A is on the order of 2%. This is a strong indication that the fine mesh resolution is sufficient to obtain an accurate engineering-level prediction of C_A . The same should hold for the other Mach numbers simulated as well because the only parameter that changes is the dynamic pressure of the oncoming flow. Given the low degree of grid-related uncertainty estimated from the mesh convergence study, one could argue that the relevant measure of uncertainty in the predicted time-averaged C_A should instead be its unsteady standard deviation as reported for each Mach number in Table 2. Either way, these results constitute the lowest levels of discretization and grid-related uncertainty reported yet for this vehicle configuration and these 3 trajectory points.

Another form of uncertainty, model-form error, is driven by the assumptions that go into the simulations. One such source of model-form error is the assumption of “frozen chemistry” where the chemical reactions occurring in the nozzle and the plumes (e.g. after-burning) are ignored, along with solid particles like soot that are often created in carbon-based combustion processes, and the effects of radiative heat transfer. The assumption of frozen chemistry has nevertheless been used successfully time and again to predict aerodynamic loads on launch vehicles. This is because

Table 3 Convergence properties of time-averaged $\langle C_A \rangle$ for *powered* vehicle at Mach 2.4 across medium, fine and extra fine meshes computed according to ASME procedures [21].

Apparent Order (p)	Grid Convergence Index	Estimated Converged $\langle C_A \rangle$	Estimated Uncertainty in $\langle C_A \rangle$
2.336	0.039	0.368	± 0.014

the difference between such simulations and those that do take into account chemical reactions and heat radiation is generally limited to volumes inside of and in the immediate vicinity of the plumes where the temperature and peak pressure are often over-predicted when chemistry and radiative heat transfer are neglected. Unfortunately, the conceptual Mars EDL vehicle considered here is not a traditional launch vehicle, and the combustion products are being convected back toward the vehicle in this case, so we cannot eliminate the possibility that chemistry and heat radiation could affect the aerodynamic loads of the vehicle purely based on prior launch vehicle simulation experience.




Fig. 13 Snapshot from LAVA Mach 2.52 fine mesh simulation showing the logarithm of the exhaust gas mixture mass fraction where white is 100% and black is 1% or less. This highlights the plumes and the regions where strong turbulent mixing is occurring. Note the dark cylinder starting at the centerline of the nose of the heat shield and reaching 1 D upstream indicating the absence of exhaust gas in that region.

As mentioned in Section IV: the time-scale for a fluid parcel exiting the nozzle to reach the bow shock and be convected back downstream to pass by the vehicle is on the order of 0.08 to 0.09 seconds, where the majority of this time is spent in the return trip from the bow shock back toward the vehicle. The area affected by this “blow-back” effect is limited to the region outside of the outer radius of the nozzles in a band extending outward radially from each nozzle pair to the heat shield shoulder. This is particularly salient in the exhaust gas mass fraction plotted in Figure 13. The pressure in that “blow-back” region is already low with $C_p \approx -0.2$ as shown in Figure 14. The region responsible for the pressure recovery on the heat shield is indeed largely enclosed by the inner radius of the nozzles, where the surface C_p reaches positive values at least four times larger than the negative values near the shoulder. Moreover, this region in the center of the heat shield hardly sees any exhaust gas: the mass fraction of exhaust gas there at any instant in time is near zero due to the cant angle of the nozzles and the strong entrainment engendered by the high velocity plumes – refer again to Figure 13 showing a dark region (exhaust mass fraction of 1% or less) across most of the region enclosed by the nozzles ahead of the nose of the vehicle’s heat shield. All these features point to chemistry and heat radiation likely being a second order effect for aerodynamic loads like the quantity of interest C_A for these SRP cases.




Fig. 14 Time-averaged surface pressure coefficient C_p on the heat shield taken from fine mesh results, where dark blue is -0.2, green is 0.9 and red is 2 or greater.

Another possible source of systematic or model-form error could be related to the particular use of the immersed boundary methodology (IBM) in under-resolved regions of the flow. Although the Navier-Stokes equations are solved everywhere in the domain, the slip boundary condition is selectively applied inside the nozzle walls and over most of the heat shield area. This means we effectively neglect the effects of wall shear stress over most of the heat shield area. It is unlikely that this assumption will severely affect the pressure distribution over the heat shield surface which drives the aerodynamic loads for the whole vehicle because time-averaged pressure is largely modulated by inviscid effects, even inside turbulent boundary layers. It is also important to recognize that we are necessarily under-resolving the boundary layer over the heat shield shoulder where we are applying a no-slip boundary condition despite $\Delta y^+ > 1$. In other words, we are likely under-predicting the wall shear stress there, and thus also the temperature near the shoulder. This under-prediction is expected only in the regions away from the “blow-back” phenomenon where the flow is recirculating and thus moving slowly enough for the IBM to provide an accurate no-slip boundary condition – see Figure 15 and Figure 16. This could be improved in the near future with the addition of a wall shear stress model to impose a more accurate viscous stress and heat flux contributions at cells that neighbor the vehicle surface as is done in wall-modeled LES (WMLES) approaches. But these are likely second order effects for the local values of time-averaged pressure which are again driven by the curvature of the heat shield shoulder itself, which is an inviscid effect. It is also possible that some portions of the payload mounted on the aft of the vehicle where no-slip is applied are under-resolved ($\Delta y^+ > 1$). The ghost-cell method by which these boundary conditions (adiabatic wall with slip or no-slip) are applied is locally first order accurate but globally second order (determined through the method of manufactured solutions). However, we compute loads by interpolating the near-wall pressure to triangle nodes, and average to the face. The interpolation procedure uses a cloud of CFD cells surrounding each node on the surface, and solves a least-squares problem with a quadratic basis to obtain a second order accurate pressure value. If this least-square problem becomes ill-conditioned, it drops down to a linear basis – this happens almost exclusively at concave corners, which are all located on the backside of the vehicle where the pressure is very low and has little impact on C_A . It is equally likely that the dominant error terms in the model stem from the convective discretization, which is nominally fifth order accurate in smooth flows, but often displays convergence rates of 3 to 4.5 in representative turbulent flows without shocks. Moreover, the scheme effectively drops to second order upwind at discontinuities and is thus rather dissipative. The most likely path toward reducing computational cost or increasing accuracy would be to reduce the numerical dissipation of the shock-capturing convective scheme in turbulent regions (and consider the use of a physically-derived subgrid scale model if it can be reduced enough), and to improve the near-wall treatment with a WMLES approach, given that there is little hope of designing a tractable isotropic Cartesian grid with $\Delta y^+ \leq 1$ over the entire surface of the vehicle for this problem.

VI. Computational Resources

Scale-resolving simulations are often more costly than steady and even unsteady RANS simulations for two main reasons: finer spatial and temporal resolution is needed to capture the relevant turbulent scales, and longer time intervals are needed to converge first-order statistics. The purpose of this work was not only to demonstrate that such




Fig. 15 Time-averaged velocity vectors in the x-z plane ($[< u >, 0, < w >]$) through the center left nozzles and on the vehicle surface colored by velocity magnitude from the Mach 2.4 fine mesh simulation. Notice the large recirculation regions driven by the plumes extending to the heat shield shoulder, and the smaller ones near the nose of the vehicle. Other recirculation regions exist in the wake of the heat shield, and again in the wake of the aft-mounted payload.




Fig. 16 Time-averaged velocity vectors just above the vehicle surface colored by velocity magnitude from the Mach 2.4 fine mesh simulation. Notice the peculiar flow pattern driven by the entrainment of the 4 pairs of SRP motors: flow that impacts the near nose is pulled to either the nearest plume or accelerates between plumes toward the edge of the heat shield, whereas flow in the “blow-back” region is pulled from the edge of the heat shield back toward the nozzle exits, or to the nearest higher speed in-between plume pair flow toward the heat shield shoulder.

scale-resolving simulations of SRP including a full grid sensitivity study could be completed, but also that it could be done in a potentially short-enough turnaround time to affect engineering design decisions on relatively modest HPC resources. Table 4 shows the breakdown of computational resources used for each simulation. It clearly demonstrates that at the fine mesh level, it is possible to obtain a converged mean C_A within one week’s turnaround time on a relatively modest number of x86 CPU cores (6,032 Intel Broadwell). At roughly 1 million core-hours for a full simulation at the fine mesh level, the computational cost is within reach for a limited number of cases through the EDL trajectory where the plume is over-expanded to provide added accuracy and reduce uncertainty. Thanks to good weak scaling performance, we were also able to complete one second of time integration in about one week for the extra fine mesh as well, which doubles the number of cells, and the number cores (this time on 100 AMD Rome nodes, totalling 12,800 cores). However, this is clearly pushing the envelope both in terms of the mesh size (1.6 billion cells), due to the amount of data generated, and in terms of the number of core-hours (2.3 million) for engineering level predictions potentially needed for tens of cases. Improvements in time per integration step by a factor of 2 to 4, which is our target for the next iteration of the LAVA solver, would go a long way toward making such simulations more attractive to incorporate in the design process. Similarly, research into more targeted shock capturing schemes could yield results comparable to the fine mesh on a medium mesh, again halving our turnaround time. If both goals are achieved, we believe the computational cost could be reduced to around 400,000 core-hours per second per powered SRP simulation.

VII. Conclusion

The Launch, Ascent, and Vehicle Aerodynamics (LAVA) framework’s Cartesian AMR module was leveraged to simulate a conceptual Mars EDL vehicle at three Mach numbers in the over-expanded plume regime: 2.4, 2.52, and 2.78. Results confirm the expectation that for a fixed SRP motor thrust, the higher Mach number of the incoming flow moves the mean bow shock location closer to the vehicle, increases the pressure recovery on the heat shield, and thus

Table 4 Computational resources used for SRP simulations for 1 second of time integration.

Case	Cell Count ($\times 10^6$)	Core Count	Wall Time (Days)	Core-Hours
Unpowered Fine	188	2,800	0.500	33,600
Powered Medium	575	6,032	4.375	633,400
Powered Fine	856	6,032	7.368	1,066,600
Powered Extra Fine	1,615	12,800	7.610	2,337,800

increases the time-averaged integrated coefficient of axial force $\langle C_A \rangle$. The specific values of $\langle C_A \rangle$ obtained for different Mach numbers fall mostly within the range of previous CFD predictions. The time-averaged surface pressure distribution is largely symmetric, resulting in near-zero time-averaged side forces $\langle C_y \rangle$, and moments $\langle C_{M_y} \rangle$. Excursions from the mean in these quantities, like their reported unsteady standard deviation, are small enough that they are unlikely to negatively impact the trajectory or attitude of the vehicle. In order to provide a stronger metric of the uncertainty with these simulations, a grid convergence study was completed for the Mach 2.4 case. Time-averaged C_A was found to converge at greater than second order from the medium to the extra fine mesh, and the associated uncertainty with the estimated converged $\langle C_A \rangle$ was found to be smaller than the unsteady standard deviation $\sigma(C_A)$. Moreover, the difference between the fine mesh C_A and the estimated converged C_A was also within the estimated uncertainty, thus justifying our use of the fine mesh for other Mach numbers and our choice of the unsteady standard deviation as a conservative estimate of the grid-related uncertainty in those predictions. The computational cost to complete one second of time integration was documented and demonstrates that performing such simulations for a subset of trajectory points is within reach. Furthermore, with code optimizations and numerical method improvements, the turnaround time and overall computational cost could be reduced by a factor of more than 2, making this methodology potentially usable in the later stages of an engineering design cycle. This work thus demonstrates an alternate path to reliable CFD predictions for SRP in the over-expanded plume regime using high-order scale-resolving simulations on Cartesian AMR grids.

Acknowledgments

Thank you to Ashley Korzun for providing the vehicle geometry and motor exhaust gas thermodynamic properties. Computing resources supporting this work were provided by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center. Special thanks to James Jensen and Elisha Makarevich for preparing the surface mesh and the cut planes. The authors are grateful to Jeff Housman and Emre Sozer for their insights and guidance throughout this project.

References

- [1] Edquist, K. T., Korzun, A. M., Dyakonov, A. A., Studak, J. W., Kipp, D. M., and Dupzyk, I. C., "Development of supersonic retropropulsion for future mars entry, descent, and landing systems," *Journal of Spacecraft and Rockets*, Vol. 51, No. 3, 2014, pp. 650–663.
- [2] Love, E. S., and Grigsby, C. E., "Some Studies of Axisymmetric Free Jets Exhausting From Sonic and Supersonic Nozzles Into Still Air and Into Supersonic Streams," *NACA Research Memorandum NACA-RM-L54L31*, 1955.
- [3] Peterson, V. L., and McKenzie, R. L., *Effects of Simulated Retrorockets on the Aerodynamic Characteristics of a Body of Revolution at Mach Numbers from 0.25 to 1.90*, NASA Technical Note D-1300, Ames Research Center, Moffett Field, CA, 1962.
- [4] Keyes, J. W., and Hefner, J. N., "Effect of forward-facing jets on aerodynamic characteristics of blunt configurations at Mach 6," *Journal of Spacecraft and Rockets*, Vol. 4, No. 4, 1967, pp. 533–534.
- [5] Adams, R., *The effects of retrorockets on the aerodynamic characteristics of conical aeroshell planetary entry vehicles*, 8th Aerospace Sciences Meeting, AIAA, 1970.
- [6] Daso, E., Pritchett, V., Wang, T.-S., Ota, D., Blankson, I., and Auslender, A., *The Dynamics of Shock Dispersion and Interactions in Supersonic Freestreams with Counterflowing Jets*, 45th AIAA Aerospace Sciences Meeting and Exhibit, AIAA, 2007, p. 1423.

- [7] Daso, E. O., Pritchett, V. E., Wang, T.-S., Ota, D. K., Blankson, I. M., and Auslender, A. H., "Dynamics of shock dispersion and interactions in supersonic freestreams with counterflowing jets," *AIAA journal*, Vol. 47, No. 6, 2009, pp. 1313–1326.
- [8] Codoni, J., and Berry, S., "Analysis of Dynamic Data from Supersonic Retropropulsion Experiments in NASA Langley's Unitary Plan Wind Tunnel," *42nd AIAA Fluid Dynamics Conference and Exhibit*, 2012, p. 2706.
- [9] Berry, S., Rhode, M., Edquist, K., and Player, C., "Supersonic Retropropulsion Experimental Results from the NASA Langley Unitary Plan Wind Tunnel," *42nd AIAA Thermophysics Conference*, American Institute of Aeronautics and Astronautics, 2011, p. 3489.
- [10] Berry, S. A., Rhode, M. N., and Edquist, K., "Supersonic Retropropulsion Validation Experiment in the NASA Langley Unitary Plan Wind Tunnel," *Journal of Spacecraft and Rockets*, Vol. 51, No. 3, 2014, pp. 664–679.
- [11] Berry, S. A., Rhode, M. N., and Edquist, K. T., "Supersonic Retropropulsion Experimental Results from NASA Ames 9×7 Foot Supersonic Wind Tunnel," *Journal of Spacecraft and Rockets*, Vol. 51, No. 3, 2014, pp. 724–734.
- [12] Bakhtian, N., and Aftosmis, M., "Analysis of Inviscid Simulations for the Study of Supersonic Retropropulsion," *29th AIAA Applied Aerodynamics Conference*, 2011, p. 3194.
- [13] Kleb, W., Schauerhamer, D., Trumble, K., Sozer, E., Barnhardt, M., Carlson, J.-r., and Edquist, K., "Toward Supersonic Retropropulsion CFD Validation," *42nd AIAA Thermophysics Conference*, 2011, p. 3490.
- [14] Schauerhamer, D. G., Zarchi, K. A., Kleb, W. L., Carlson, J.-R., and Edquist, K. T., "Supersonic Retropropulsion Computational Fluid Dynamics Validation with Langley 4×4 Foot Test Data," *Journal of Spacecraft and Rockets*, Vol. 51, No. 3, 2014, pp. 693–714.
- [15] Schauerhamer, D. G., Zarchi, K. A., Kleb, W. L., and Edquist, K. T., "Supersonic Retropropulsion Computational-Fluid-Dynamics Validation with Ames 9×7 Foot Test Data," *Journal of Spacecraft and Rockets*, Vol. 51, No. 3, 2014, pp. 735–749.
- [16] Korzun, A. M., Tang, C., Rizk, Y., Canabal, F., Childs, R., Van Norman, J. W., Tynis, J. A., and Bibb, K., "Powered Descent Aerodynamics for Low and Mid Lift-to-Drag Human Mars Entry, Descent and Landing Vehicles," *AIAA Scitech 2020 Forum*, 2020, p. 1510.
- [17] Korzun, A. M., Nielsen, E. J., Walden, A. C., Jones, W. T., Carlson, J.-R., Moran, P. J., Henze, C. E., and Sandstrom, T. A., "Effects of Spatial Resolution on Retropropulsion Aerodynamics in an Atmospheric Environment," *AIAA Scitech 2020 Forum*, 2020, p. 1749.
- [18] Kiris, C. C., Housman, J. A., Barad, M. F., Brehm, C., Sozer, E., and Moini-Yekta, S., "Computational framework for launch, ascent, and vehicle aerodynamics (LAVA)," *Aerospace Science and Technology*, Vol. 55, 2016, pp. 189–219.
- [19] Brehm, C., Barad, M. F., Housman, J. A., and Kiris, C. C., "A comparison of higher-order finite-difference shock capturing schemes," *Computers & Fluids*, Vol. 122, 2015, pp. 184–208.
- [20] Castiglioni, G., Domaradzki, J., Pasquariello, V., Hickel, S., and Grilli, M., "Numerical simulations of separated flows at moderate Reynolds numbers appropriate for turbine blades and unmanned aero vehicles," *International Journal of Heat and Fluid Flow*, Vol. 49, 2014, pp. 91–99. <https://doi.org/https://doi.org/10.1016/j.ijheatfluidflow.2014.02.003>, URL <https://www.sciencedirect.com/science/article/pii/S0142727X14000216>.
- [21] Celik, I. B., Ghia, U., Roache, P. J., and Freitas, C. J., "Procedure for estimation and reporting of uncertainty due to discretization in CFD applications," *Journal of fluids Engineering-Transactions of the ASME*, Vol. 130, No. 7, 2008.