
Oral presentation | Fluid-structure interaction

Fluid-structure interaction-II

Mon. Jul 15, 2024 2:00 PM - 4:00 PM Room A

[2-A-04] Wall-Modeled Large-Eddy Simulation of Supersonic Parachute Inflation

*Francois Cadieux¹, Michael F Barad¹ (1. NASA Ames Research Center)

Keywords: Supersonic, Parachute, Wall-Modeled Large-Eddy Simulation

Wall-Modeled Large-Eddy Simulation of Supersonic Parachute Inflation

F. Cadieux* and M. Barad*

Corresponding author: francois.cadieux@nasa.gov

* NASA Ames Research Center, Moffett Field, CA, USA.

Abstract: Supersonic parachutes have been used in nearly every robotic mission to another planetary body with an atmosphere because they are one of the most mass-efficient ways to decelerate a payload to land on the surface. Short of performing flight tests in the upper Earth atmosphere, we currently cannot reliably predict a novel parachute system's performance or potential failure modes, or even confidently explain it after the fact, as was the case with the Low Density Supersonic Decelerator flight tests and subsequent investigations. Fluid-structure interaction (FSI) simulations have the potential to bridge this gap. To this end, we present several improvements to the state-of-the-art for simulating supersonic parachutes using FSI: a higher effective resolution convective flux, an immersed boundary turbulent wall layer modeling approach to capture viscous effects, and a novel method to obtain a more realistic initial parachute shape. A recent supersonic parachute flight test is simulated and compared to measurements for the purposes of model validation.

Keywords: Fluid-Structure Interaction, Compressible Flow, Numerical Methods.

1 Introduction

Supersonic parachutes continue to be the decelerator of choice for interplanetary robotic missions for destinations with an atmosphere because of their small mass and the large drag they provide. Parachutes are a critical component of entry, descent, and landing for such missions, and represent a catastrophic single point of failure leading to loss of the mission. Supersonic parachute systems of the disk-gap-band (DGB) type are a well-proven technology thanks to a campaign of tests in the upper Earth atmosphere leading up to the two Viking spacecraft successfully landing on Mars (1968-1976) [1]. DGB chutes were subsequently used in every spacecraft NASA has sent to land on the surface of Mars. The Low-Density Supersonic Decelerator (LSD) project led to the creation and testing of a new hybrid supersonic parachute design that blended the best aspects of subsonic ring-sail parachutes with the proven DGB design. However, the novel LSD parachute ripped apart whilst inflating during its supersonic flight test. This failure demonstrated that NASA and the wider aerospace community lacked the capability to accurately predict the failure of a new supersonic parachute design without a flight test [2]. The subsequent investigation into the root cause of the LSD failure also demonstrated that we do not yet possess computational modeling tools of sufficient fidelity to reproduce the failure.

The development a modeling tool that is representative of the physics during the opening of a supersonic parachute is the overarching goal of an effort within the Entry Systems Modeling project at NASA [3]. The ultimate aim is two-fold: to aid engineers in the design of future parachute systems by assessing performance and mitigating potential failures, and to be able to assess the root cause of any past or future parachute failures. The two goals have different requirements. Any tool used in a mission design cycle is expected to have a fast time-to-solution, and the designers understand that approximations are made to the physics being modeled. In contrast, the requirements to assess the root cause of a known failure would be to reproduce all relevant physics at a level of sufficient detail to eliminate all other possible root causes. As such an analysis is retrospective, time-to-solution is less of a concern.

The Launch, Ascent and Vehicle Aerodynamics (LAVA) computational fluid dynamics (CFD) tool is being extended to enable fluid-structure interaction (FSI) simulations to fulfill the first goal of aiding the design of future parachute systems [4, 5, 6, 7]. It leverages well-established methods and faster approximate physics-based modeling approaches over more detailed and computationally expensive ones. The purpose is to simulate only enough detail to be able to reproduce quantities of interest from wind tunnel and flight test data while keeping the time-to-solution short. This has the added advantage that the FSI tool can be useful to future parachute system designers sooner. And once validated, it can be extended further to improve the level of fidelity of the underlying models and reach the second goal of assessing failure root causes.

2 Problem Statement

The Advanced Supersonic Parachute Inflation Research Experiments (ASPIRE) completed three flight tests in the upper Earth atmosphere [8, 9]. The second (SR02) and third (SR03) flight tests featured a nearly identical parachute system where the most salient difference is the canopy fabric change from IAW PIA44378D (with material properties derived from F-111 Nylon) to a proprietary material called Heathcoat for which the properties are not currently publicly available [10]. ASPIRE’s third flight test (SR03) served to certify the Mars 2020 parachute system that successfully landed the Perseverance rover on Mars and thus featured a higher dynamic pressure (1028.4 Pa) at line stretch than SR02 (744.6 Pa). The flight tests measured the parachute’s pull-force at the payload and the payload’s sensed acceleration for quantitative comparisons, and recorded high-speed video for qualitative comparisons. An approximation of the projected area of the canopy as a function of time was also extracted from the high-speed footage [11]. NASA’s Jet Propulsion Laboratory and Airborne Systems recently released the parachute system’s detailed design and material properties so that ASPIRE SR02 and SR03 could serve as benchmark cases toward validation of FSI simulations [10]. A high-level diagram of the payload and parachute system is shown in Figure 1 where each of the important components are labeled. The material properties used in these simulations are also provided for completeness in Table 1.

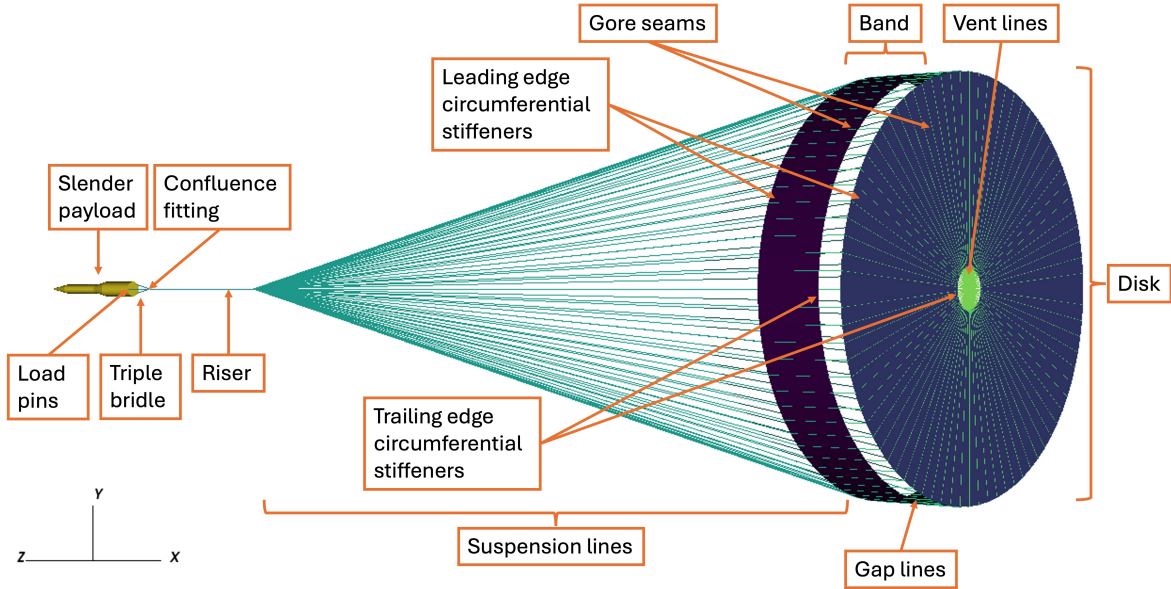


Figure 1: Diagram of ASPIRE SR03 parachute system with components labeled.

The goal of this work is to reproduce quantities of interest for supersonic parachute inflation dynamics as seen in the ASPIRE SR03 flight test from first principles. Some quantities of interest are listed below in order of increasing difficulty:

1. peak loads generated by the parachute as measured at the load pins or by accelerometers on the payload,
2. post-inflation dynamics (trough and rebound peak in loads),
3. time to inflation and inflation-rate,
4. “steady” descent phase aerodynamic coefficients as a function of M_∞ , q_∞ , α_{total} ,
5. dynamic stability of response to angle of attack changes due to wake re-contact, wind gusts, wind shear, or strong atmospheric stratification,
6. spatial distribution of stress magnitude over parachute system components with enough detail and accuracy to predict potential material failure.

Table 1: Finite element type and material properties [10] for each component of the parachute structural model.

| Component name | Element type | Young's modulus (Pa) | Poisson ratio | Density (kg/m ³) | Width (m) | Thickness (m) |
|-----------------------------|-----------------|-------------------------|---------------|------------------------------|-----------|------------------------|
| Bridles | Timoshenko beam | 1.6680×10 ¹⁰ | 0.4 | 705.13 | 0.03908 | - |
| Riser | Timoshenko beam | 1.8600×10 ¹⁰ | 0.4 | 751.99 | 0.04778 | - |
| Suspension, gap, vent lines | Timoshenko beam | 1.8202×10 ¹⁰ | 0.4 | 484.66 | 0.00506 | - |
| Gore seams | Timoshenko beam | 3.8705×10 ¹⁰ | 0.4 | 484.66 | 0.00506 | - |
| Broadcloth | MITC3 | 9.4484×10 ⁸ | 0.4 | 533.95 | - | 7.620×10 ⁻⁵ |
| Band leading edge | MITC3 | 1.3568×10 ¹⁰ | 0.4 | 454.07 | 0.0254 | 1.058×10 ⁻³ |
| Band trailing edge | MITC3 | 1.3568×10 ¹⁰ | 0.4 | 454.07 | 0.0254 | 1.058×10 ⁻³ |
| Disk leading edge | MITC3 | 1.3568×10 ¹⁰ | 0.4 | 454.07 | 0.0254 | 1.058×10 ⁻³ |
| Disk trailing edge | MITC3 | 1.0979×10 ¹⁰ | 0.4 | 525.54 | 0.0254 | 2.540×10 ⁻³ |

The first three quantities of interest are the validation targets for this work. The normalized opening force, comparable to an instantaneous maximum coefficient of drag C_D , was highly repeatable across all three ASPIRE flight tests and thus ensuring that this can be reproduced reliably via FSI simulations has been and continues to be the highest priority [7]. We thus focus on a specific segment of the ASPIRE SR03 flight test from the “line stretch” event, where the load pins that connect the bridles to the payload first record tension $t = 0.0$ s, to roughly $t = 1.1$ s where the parachute system has reached its stable deceleration phase. The initial flow conditions are described in Table 2.

Table 2: Flow initial conditions corresponding to line stretch in ASPIRE SR03 flight test [9].

| Name | Symbol | Value | Unit |
|-------------------------------|---------------|-------------------------|----------------------|
| Freestream Mach number | M_∞ | 1.88 | - |
| Freestream dynamic pressure | q_∞ | 1028.44 | Pa |
| Freestream density | ρ_∞ | 6.0171×10 ⁻³ | kg/m ³ |
| Wind-relative angle of attack | α | 0.0 | ° |
| Ratio of specific heats | γ | 1.4 | - |
| Gas constant | R | 287.0025 | J/(kg K) |
| Reference dynamic viscosity | μ_0 | 1.7084×10 ⁻⁵ | N s / m ² |
| Reference temperature | T_0 | 273.15 | K |
| Sutherland's temperature | C | 110.0 | K |
| Prandtl number | Pr | 0.7 | - |
| Reynolds number per meter | Re | 2.2784×10 ⁵ | 1/m |

3 Methodology

The most salient change to the approach taken here is the switch from the compressible Euler equations to the compressible large-eddy simulation (LES) equations solved in the following conservative form delineating the convective and viscous fluxes where the usual LES grid filter is implied and not denoted:

$$\frac{\partial \rho}{\partial t} = -\frac{\partial}{\partial x_j} (\rho u_j) \quad (1)$$

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

$$\frac{\partial \rho E}{\partial t} = -\frac{\partial}{\partial x_j} (H u_j) + \frac{\partial}{\partial x_j} (u_j \tau_{ij} + q_j) \quad (3)$$

where

$$\begin{aligned}\tau_{ij} &= (\mu(T) + \mu_{\text{SFS}}) \left(\frac{\partial u_j}{\partial x_i} + \frac{\partial u_i}{\partial x_j} - \frac{2}{3} \frac{\partial u_k}{\partial x_k} \delta_{ij} \right) \\ E &= e + \frac{1}{2} u_k u_k \\ e &= \frac{c_p}{\gamma} T \\ H &= h + \frac{1}{2} u_k u_k \\ h &= c_p T \\ q_j &= (\kappa + \kappa_{\text{SFS}}) \frac{\partial T}{\partial x_j}\end{aligned}$$

with the perfect gas equation of state

$$p = \rho RT, \quad (4)$$

Sutherland's Law for laminar viscosity

$$\mu(T) = \mu_0 \left(\frac{T}{T_0} \right)^{\frac{3}{2}} \left(\frac{T_0 + C}{T + C} \right) \quad (5)$$

and the constant coefficient Vreman sub-filter scale (SFS) model [12] for the turbulent eddy viscosity μ_{SFS} ($C_{\text{SFS}} = 0.07$) coupled with a constant turbulent Prandtl number ($\text{Pr}_{\text{SFS}} = 0.9$) for the turbulent heat conductivity $\kappa_{\text{SFS}} = (c_p \mu_{\text{SFS}}) / \text{Pr}_{\text{SFS}}$. The variables take their usual meaning: ρ is density, u_i is the velocity in Cartesian direction i , p is pressure, T is temperature, e is internal energy, E is total energy, h is enthalpy, H is total enthalpy, μ is dynamic viscosity, μ_0 is the reference dynamic viscosity, T_0 is the reference temperature, C is the Sutherland constant, R is the gas constant, c_p is specific heat at constant pressure, γ is the ratio of specific heats, q is the heat flux, and κ is the thermal conductivity.

The addition of the viscous fluxes is motivated by the observation that previous simulations of the rigid folded parachute did not show strong separation: parts of the wake remained attached due to the application of the adiabatic slip velocity boundary condition appropriate for inviscid flows [7]. This behavior is problematic because the lack of mixing due to the partially attached shear layer in the wake affects the pressure differential between the interior and exterior of the canopy during inflation and is thought to contribute to a faster inflation rate. Moreover, it is important to adequately resolve the turbulent wake of the payload, so as to not over-estimate the momentum deficit it causes near the canopy, which can lead to fortuitously good agreement due to compensating errors [13]. This requires some attached boundary layer turbulence be resolved on the surface of the payload to ensure the separated wake shear layer is also turbulent and that the axial length of the re-circulating flow region is not over-estimated.

Given the thin attached boundary layer developing over the slender payload, performing wall-resolved LES with a $\Delta y^+ \leq 1$ would be computationally expensive, especially with isotropic Cartesian cells. To remedy this, we leverage the wall-modeled LES (WMLES) approach where the turbulent wall-shear stress is estimated from an assumed fully developed self-similar turbulent boundary layer mean velocity profile [14] and the otherwise under-resolved viscous flux is augmented at relevant faces. This approach as deployed in the LAVA CFD software has been very successful in tackling problems like LES of aircraft in a landing configuration with flaps and slats deployed [15]. Its implementation into the LAVA Cartesian AMR module is expanded on in Section 3.3. The computational structural dynamics (CSD) solver, the contact mechanics and porosity models, as well as the nature of the coupling to the CFD solver all remain unchanged from previous work [5] and so are not described here for brevity. The methods and parameters for CFD are summarized in Table 3, and for CSD in Table 4.

3.1 Initial Parachute Shape

The ASPIRE SR03 parachute is deployed with a mortar pushing the bag containing the parachute with the suspension lines wrapped around it out of the back of the payload [8]. Inertia of the bag causes it to unroll from the suspension lines, and eventually detach from the parachute contained within. This whole segment of the flight is mostly driven by ballistic mechanics, structural dynamics, and large surfaces of

Table 3: CFD parameters for FSI simulations.

| Parameter | Value | Note/Units |
|---------------------|---|--|
| Discretization type | Conservative finite difference | In shu-Osher form |
| Convective flux | Hybrid Weighted Compact Non-linear Scheme (HWCNS) | See Section 3.2 for details |
| Flux function | AUSMPW+ [16] | With minmod limiter [17] |
| Viscous flux | Conservative 2 nd order centered | |
| Wall function | Blended log law profile | Matches Musker’s [18] |
| Time integrator | Strong Stability-Preserving Runge-Kutta (SSPRK) | Explicit 4-stage 2 nd order |
| CFL number | 2.9 | |
| Minimum mesh size | 0.01 | m |
| Maximum mesh size | 1.28 | m |
| Refinement levels | 8 | 2:1 ratio between levels |
| Mesh type | Error metric based adaption | Small-scale TKE [19] |
| Mesh size | 100-800×10 ⁶ | Cells |
| Domain size | $L_x = 307.2$ $L_y = L_z = 153.6$ | m |

Table 4: CSD parameters for FSI simulations.

| Parameter | Value | Note/Units |
|--------------------------|---|------------------------------|
| Discretization type | Finite element method | Geometrically non-linear |
| Fabric elements | MITC3 triangular shells | With additional drilling DOF |
| Cable elements | Timoshenko beams with 6 DOF | No resistance to compression |
| Mass matrix | Lumped | |
| Time integrator | Explicit 2 nd order centered | With high-freq damping |
| Critical time step ratio | 0.3 - 0.4 | |
| Minimum edge length | 0.0254 | m |
| Maximum edge length | 0.051 | m |
| Mesh size | $\approx 300 \times 10^3$ | Elements |

contact mechanics. Although it plays a role in defining the initial parachute shape, it is not unclear if it has a critical effect on the parachute loads beyond the line stretch event. It is also computationally intractable to start FSI simulations from a packed parachute state at this time. Despite their best efforts, the ASPIRE team could not reconstruct the full three-dimensional shape of the parachute from the high-speed camera video data [11], especially at early times. This would be the best starting point for FSI if it were available. The description of the parachute as-constructed is unfortunately not an acceptable starting point to predict supersonic parachute inflation dynamics because in that state the parachute already occupies close to its maximum projected area. An ideal initial parachute shape would have a cross-sectional area comparable to that seen in the test at line stretch, and would be free of large strain or stress.

In previous work, the initial parachute shape was obtained by explicitly folding the as-constructed shape into an accordion-like structure in order to reduce the initial cross-sectional area [13, 10]. This approach reduces the cross-sectional area, is computationally cheap, but induces strain and stress at the gore seams, which are known to be stiffer than the fabric. It is also clear that the parachute was never close to such a configuration during the flight test. To remedy some of the above concerns, a short structural dynamics simulation with contact and a constant outward pressure were performed with the accordion shape as a starting point, which served to smooth out the corners and relieve some of the bending stress at the gore seams. However, this did little to alleviate the unnatural near-perfect symmetry of the resulting parachute shape, or the induced slack in the suspension lines.

In order to improve the initial shape of the parachute, we draw inspiration from a technique that was

developed to obtain a more realistic initial canopy shape (without lines) for a controlled indoor test [20]. The technique starts from the as-constructed shape of a single disk gore with vent lines. A structural dynamics simulation with gravity is then performed where the nodes at the apex of the vent lines are pinned, and all the nodes along the gore seams are only allowed to move within their original plane to maintain symmetry and periodicity, until a steady state is reached. The full disk is then reconstructed via rotating and stitching copies of the drooped single gore, and the suspension lines added afterwards to complete the parachute structural model.

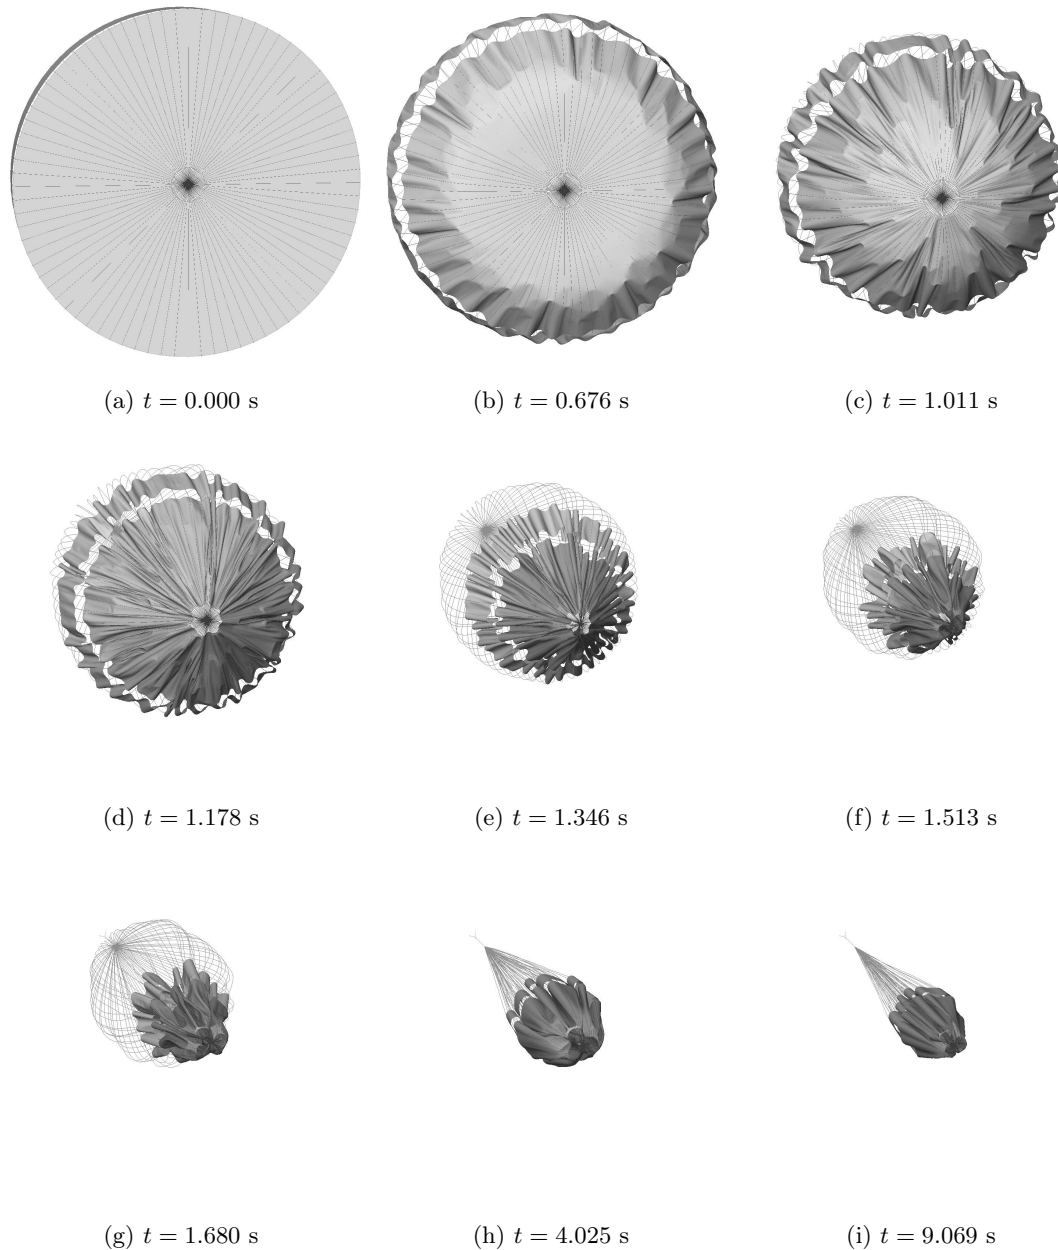


Figure 2: Snapshots from precursor structural dynamics with contact simulation where the as-constructed parachute (a) is hanging from pinned bridle attachment nodes and gravity causes it to fall under its own weight and fold in on itself. Only the components included in the structural dynamics are shown (the payload is rigid so it is not included).

We improve upon this method by considering the entire parachute system from the start including the vent, gap and suspension lines, as well as the riser and bridles. We then pin the nodes corresponding to the bridle attachment points to the rigid and stationary payload instead of the vent apex nodes. To better approximate the line stretch event with minimal strain, the parachute system effectively hangs from the payload that is being held in place, and we let the effect of gravity (in the direction of air flow)

fold the parachute canopy through a preliminary structural dynamics with contact simulation. To avoid the suspension lines' inertia during the closing process to overshoot and rebound, we apply a large drag force on the transverse motion of the suspension line elements $F_d = C_{D_0} S \frac{1}{2} \rho_\infty v_t^2$, where $C_{D_0} = 2000$ is a user input, ρ_∞ is the freestream fluid density, and $v_t = \mathbf{v} - (\mathbf{v} \cdot \mathbf{t}) \mathbf{t}$ where v_t is the transverse velocity at a structural node in the lines, \mathbf{v} is the velocity vector at that node, and \mathbf{t} is the vector from said node to the second node of the element it belongs to. This unsteady simulation is integrated forward in time until a nearly steady state is reached at $t \approx 9.1$ s. The final radius of the folded parachute's cross section is approximately 1.5 meters or 10% of the as-constructed disk diameter D . The resulting process is visualized in Figure 2. The full state of the parachute system is then saved for use as the initial conditions for FSI simulations.

3.2 Convective Flux Formulation

The proposed convective flux is loosely based on the hybrid weighted compact non-linear scheme (HWCNS) [21]. It is modified to drastically reduce artificial dissipation and improve effective resolution in turbulent regions via the use of grid quality and shock sensors. In order to maintain a formulation that can be implemented in unstructured finite volume solvers, we limit the flux to use the typical cell-centered (volume average) values neighboring the face, as well as the derivatives at those cell-centers. Following recent developments in understanding of kinetic energy and entropy preservation properties [22], we switch from interpolating the typical primitive variable vector $Q = [p, u, v, w, T]$ to the face to a modified set of primitive variables $\tilde{Q} = [\rho, u, v, w, p]$, which improves the numerical stability of the central part of the scheme – primitive variable interpolation would otherwise be more robust for the dissipative part of the numerical flux. \tilde{Q} is reconstructed (interpolated) from the left cell-center to the face using a linear approach, as well as the non-linear slope-limiter approach that avoids creating new minima or maxima:

$$\tilde{Q}_{L,j+\frac{1}{2}}^{\text{linear}} = \tilde{Q}_j + \frac{1}{2} \Delta x \frac{\partial \tilde{Q}}{\partial x} \Big|_j \quad (6)$$

$$\tilde{Q}_{L,j+\frac{1}{2}}^{\text{limiter}} = \tilde{Q}_j + \frac{1}{2} \phi(x_j) \Delta x \frac{\partial \tilde{Q}}{\partial x} \Big|_j, \quad (7)$$

where ϕ is the output of the slope limiter ranging from 0 to 1 (e.g. minmod). The linear and limited state vectors can be reconstructed from the right in a symmetric fashion. We then mix the arithmetic average of the linearly reconstructed left and right states with the slope-limited reconstruction based on the maximum of the grid quality sensor Γ and the shock sensor Θ where $\zeta = \max(\Gamma, \Theta)$ as follows:

$$\tilde{Q}_{j+\frac{1}{2}}^{\text{central}} = \frac{1}{2} (\tilde{Q}_{L,j+\frac{1}{2}}^{\text{linear}} + \tilde{Q}_{R,j+\frac{1}{2}}^{\text{linear}}) \quad (8)$$

$$\tilde{Q}_{L,j+\frac{1}{2}} = (1 - \zeta_{j+\frac{1}{2}}) \tilde{Q}_{j+\frac{1}{2}}^{\text{central}} + \zeta_{j+\frac{1}{2}} \tilde{Q}_{L,j+\frac{1}{2}}^{\text{limiter}} \quad (9)$$

The grid quality and shock sensors can all be computed directly from quantities already available in a compressible Navier-Stokes solver and are inexpensive. Their intent is to surgically introduce numerical dissipation only where necessary. All sensors vary from zero (no dissipation) to one which corresponds to a fully-upwind reconstruction of the left and right states. They are computed on a face-basis using available neighboring cell-center data. The grid quality sensor Γ goes to one if either cell-center on each side of the face is a guard cell that was filled from coarser data (representative of a 2:1 jump in cell-size). Similarly, $\Gamma = 0.5$ when any second neighbor is filled from coarser data; otherwise it is nominally zero. Additionally, the parachute features many concave crevices at the start of inflation. Without numerical dissipation in these regions at startup, spurious numerical artifacts can be generated. As such, $\Gamma = 1$ for “trapped” cells, which we define as cells for which the overall flux divergence stencil contains immersed boundary “ghost cells” on both sides of the cell center in a given direction. Cells adjacent to “trapped” cells get 50% upwind reconstructions to smoothly ramp the upwind terms in concave or thin channel-like regions. The shock sensor Θ is formulated as follows [23]:

$$\Theta_{j+\frac{1}{2}} = \left(\frac{(\nabla \cdot \mathbf{u})_{j+\frac{1}{2}}^2}{(\nabla \cdot \mathbf{u})_{j+\frac{1}{2}}^2 + \|\nabla \times \mathbf{u}\|_{j+\frac{1}{2}}^2 + \epsilon} \right)^2 \quad (10)$$

where ϵ is a small number just above numerical precision to avoid division by zero. Once the left and right primary state vector are interpolated to the face, the AUSMPW⁺ approximate Riemann solver is

used to compute the convective flux [16]. The main benefit of this approach comes from the fact that the difference between Q_L and Q_R is minimized when ζ approaches zero, because the numerical dissipation terms in AUSMPW⁺ are proportional to $(Q_L - Q_R)$. Moreover, the extension to higher-order is clear and relatively simple following the steps laid out in the HWCNS approach [21].

3.3 Viscous Flux Correction With Wall Model

The viscous fluxes are computed using standard second-order central difference derivative approximations along with interpolations to the face centers. The viscous stress computed at each face that uses immersed boundary ghost-cell information in its stencil is first rotated to the local coordinate frame [24] composed from the wall normal vector of the nearest point on the surface and the relative tangential velocity direction from the matching point. The components of the rotated stress tensor that contain the derivative of tangential velocity are then corrected using an algebraic wall model that approximates Musker’s zero-pressure-gradient turbulent boundary layer mean velocity profile [18], along with the distance from the face center to the surface d to evaluate the energy equation contribution $u_j \tau_{ij}(d)$. The face center can land inside the solid, and where this occurs, the distance input to the wall model to evaluate the relative tangential velocity is set to zero. The corrected stress tensor and relative tangential velocity at the face center are then rotated back to Cartesian coordinates (with their other components untouched). The relative tangential velocity and density input to the algebraic wall model comes from a point a fixed distance away (matching point) from the surface and in the fluid set to $h_{wim} = \sqrt{3}\Delta x$ or larger to ensure that the input does not include information from inside the solid or from boundary condition cells. Since this matching point may not coincide with a cell center, tri-linear interpolation of the primary variables is leveraged to obtain its state vector. This has the added advantage to provide a small amount of spatial filtering (or an equivalent time filter in the LES sense) to the wall model inputs which has been shown to be beneficial to reduce log-layer mismatch [25]. The location of the faces where the viscous flux is corrected and an example of the matching point interpolation are shown in Figure 3 for a relatively simple two-dimensional case.

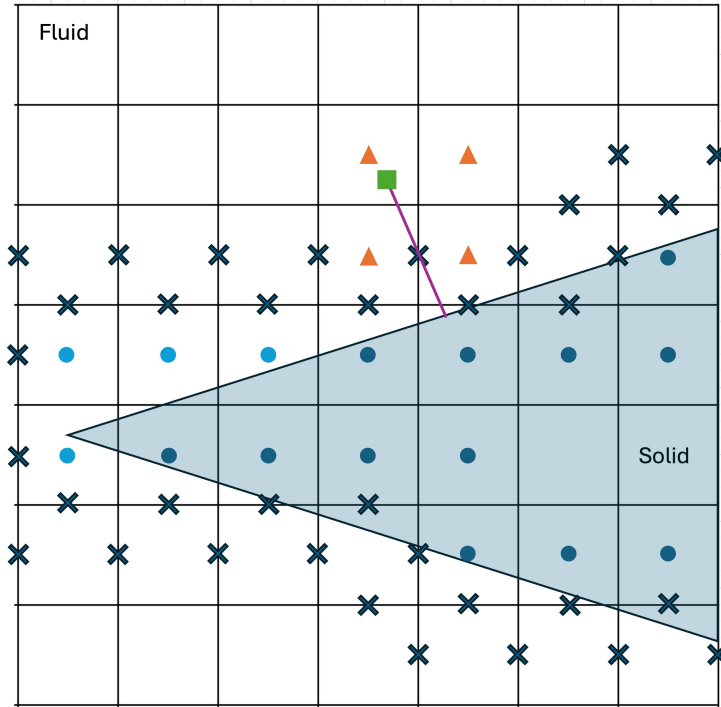


Figure 3: Diagram of the immersed boundary viscous flux correction scheme using an algebraic wall model. Immersed boundary “ghost” cells: dark blue circles; immersed boundary “ghost-in-fluid” cells: light blue circles; faces where viscous flux correction is performed: dark X symbols; algebraic wall model matching point input: green square; vector from closest point on surface from face center to matching point: purple line; cell values used for interpolation to matching point: orange triangles.

3.4 Eulerian Free Flight Motion

In order to model the deceleration of the system, a free flight module has been implemented in LAVA in the fixed Eulerian frame. The free flight module solves the rigid body motion equations for the payload using the same time integration scheme as the CFD, using the loads sampled at each Runge-Kutta (RK) sub-step to update the payload's velocity and position. The payload thus moves with respect to the Cartesian background grid in the same way the deforming parachute does. The payload's motion is coupled to that of the parachute through the bridle attachments which share nodes with the surface grid of the rigid payload. The CSD solver provides point (reaction) forces at these nodes for the free flight module. The free flight module provides velocity boundary conditions at the same shared nodes for the CSD solver. The pure Eulerian approach is simple, effective and can seamlessly handle multiple objects in relative motion. However, it entails that the immersed boundary conditions must be recomputed at each RK sub-step, which adds significant computational cost. As rigid objects are advected by the flow, there is also additional cost due to grid adaption to follow the bodies, such as the payload and parachute. Thankfully, this cost is mitigated by performing the simulation in a steadily moving Eulerian frame at the freestream velocity U_∞ . We start the simulation with the rigid payload by itself to establish the flow and let the payload's wake develop beyond the parachute's location in this typical CFD frame of reference where the payload is not moving with respect to the Cartesian grid. Then we restart with the parachute with FSI and with free flight turned on. The payload then quickly starts to move in the direction of the flow due to drag, decreasing its relative velocity. Some of this additional cost can be alleviated by performing the simulation in the non-inertial frame of reference of the rigid body instead of the Eulerian frame and this will be explored in future work.

4 Results: WMLES of Payload Prior To FSI

Prior to starting the FSI simulation, we perform WMLES where the payload is fixed and rigid in the typical steadily moving Eulerian CFD frame of reference to mimic the state prior to the deployment of the chute. This CFD-only simulation is integrated for 0.2 seconds, which is more than sufficient for the aerodynamic loads on the payload to reach a statistically stationary state as shown in Figure 4.

The loads off the body axis oscillate around zero as expected for this zero angle of attack case, whereas the axial force quickly reaches a stationary state 0.05 seconds into the simulation. The time-averaged drag coefficient value also aligns well with previous unstructured body-fitted CFD simulations [26, 27] demonstrating that the present Cartesian WMLES approach is able to capture the viscous drag accurately. This is also evidenced in Figure 5 where the instantaneous Δy^+ is visualized on the surface. The maximum Δy^+ values do not exceed 300 and so are well within the range where WMLES is expected to be accurate, except at the nose of the payload where the flow is laminar and the boundary layer extremely thin. The mesh spacing is constant over the surface of the payload at $\Delta x = 0.01$ m. Δy^+ is proportional to the wall shear stress returned by the wall model, and thus indicates through its spatial structure that the flow transitions from laminar to turbulent over the conical expansion the payload, likely triggered by oblique shock boundary layer interaction and instabilities therein. The normalized axial velocity on the slide also shows resolved turbulence starting just downstream of the conical expansion. The drag decomposition across inviscid and viscous contributions is also in good agreement with previous CFD studies of the payload where the viscous contribution was found to be approximately 20% of the total drag [26]. The time interval is chosen for the payload wake to reach beyond the location of the parachute as evidenced in Figure 6a.




Figure 4: Time history of aerodynamic coefficients on the payload during WMLES before FSI.

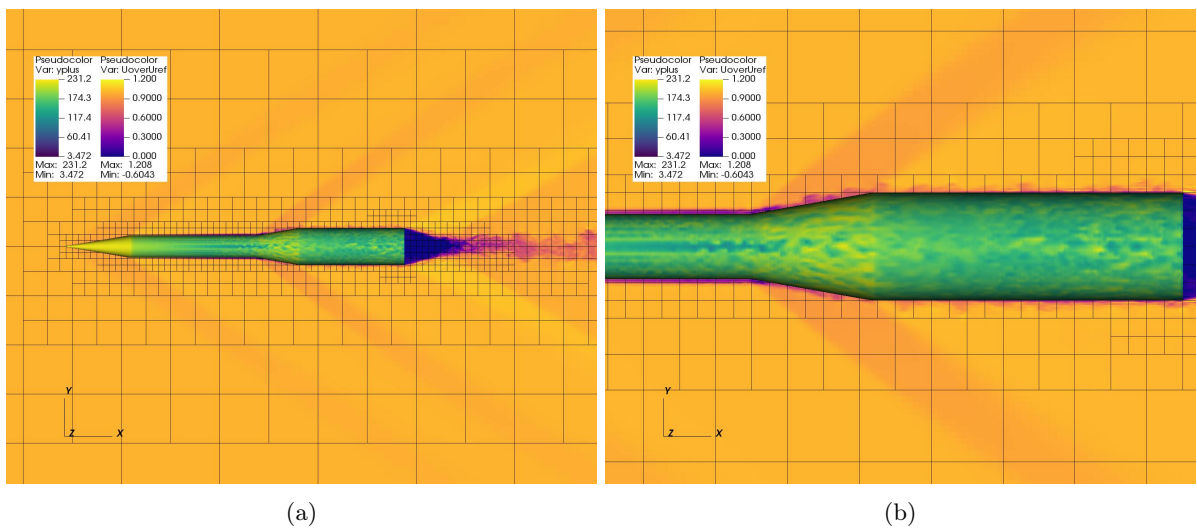


Figure 5: (a) Snapshot of the Δy^+ sampled onto the surface of the payload, and normalized axial velocity on a slice. Black boxes each contain 12^3 cells. (b) is zoomed in further to highlight the onset of resolved turbulent fluctuations in the boundary layer.

5 Results: FSI-WMLES of Payload and Parachute System

At $t = 0.0$ s we restart the simulation in full FSI-WMLES mode with the parachute system present. The parachute immediately starts to generate some drag and the disk starts to inflate as shown in Figure 6b. Oblique shocks form at the band and disk leading edges due to the abrupt change in cross-sectional area felt by the air flow. An expansion fan develops at the shoulder of the disk leading to a re-compression shock that limits the width of the wake. At this point, the flow through the parachute vent hole is not choked. This quickly changes as the disk inflates further as seen in Figures 6c to 6e, where the flow through the vent hole is choked and forms a strong terminating shock reaching a maximum Mach number greater than 3. This barrel shock increases in intensity as the parachute reaches its maximum cross-sectional area, and then decreases as the whole system decelerates and the effective freestream dynamic pressure is reduced. As evidenced in Figures 6e and 6f, the disk is pushed nearly flat by the oncoming flow which pulls the band to its full diameter. The peak loads occur just after $t = 0.3$ s as the pressure has built up to a maximum inside the canopy. The bow shock then moves further upstream in Figure 6g and relieves some of that pressure, allowing some small portions of the band, and to a lesser extent the disk, to depress from their maximum inflated radius. This corresponds to the trough in the pull force. The bow shock then travels back and settles a bit closer to the leading edge of the band as observed in Figure 6h, leading to higher loads. From there on, the parachute oscillates about a stationary state and enters the stable deceleration phase.

Figure 7 displays the comparison between the pull force predicted by the FSI simulation and the ASPIRE flight test measurements. The flight test pull force is calculated as the sum of the load measured at each of the three load pins located where the bridles attach to the payload. The FSI derived values are computed using the total force acting on the payload as follows:

$$F_{\text{pull}} = F_{\text{p,total}} - F_{\text{p,aero}} = -m_{\text{p}}a_{\text{p}} - F_{\text{p,aero}}, \quad (11)$$

where m_{p} is the mass of the payload, a_{p} is the acceleration of the payload and $F_{\text{p,aero}}$ is the aerodynamic drag on the payload – all quantities that are easily computed from the simulation output. FSI results have been shifted forward in time by 0.1 s to approximately align the peak pull force to the flight test data in Figure 7. Excellent agreement of the FSI predictions with the flight measurements of pull force is observed in the entire inflation process up to just prior to the inflation peak. So although the total inflation time is predicted to be shorter by 0.1 s, the new initial parachute shape allows the FSI to match the ASPIRE inflation rate accurately. The flight test features sustained a maximum pull force from $t = 0.4$ s to $t = 0.45$ s of around 0.75 whereas the FSI predictions reach a peak of 0.9 and then drop back down to values in line with the ASPIRE flight measurements by $t = 0.45$ s. Although the FSI inflation peak is 17% larger than what was measured in the ASPIRE flights, the SR02 rebound peak at $t = 0.6$ s with a value of 0.86 shows that the FSI peak pull force prediction is not necessarily an outlier. The FSI does not capture the pronounced post-inflation trough seen in the ASPIRE flights. Similarly, the flight test pull force increases dramatically to a value comparable to the inflation peak at $t = 0.61$ s in what is termed the rebound peak, whereas the FSI predicts a much more muted rebound. There is excellent correlation between flight test and FSI predictions for the phase, amplitude and frequency of the oscillations in the pull force curve after the inflation peak that represent the so-called suspension line mode of the parachute system. The downward slope due to the system's rapid deceleration after inflation from $t = 0.65$ s to $t = 1.1$ s seen the ASPIRE SR03 data is also well predicted, despite a small constant deficit.

Figure 8 summarizes the performance of the current FSI-WMLES simulation compared to the flight test for the quantities of interest derived from the pull force that are outlined in Section 2. The most salient differences are the lack of a significant rebound peak and the faster inflation in the FSI.

The flight test's payload carried an accelerometer onboard which recorded the sensed deceleration. The FSI simulation recorded the velocity of the payload as a function of time, from which we can compute a second-order accurate acceleration $a_{j+\frac{1}{2}} = du/dt|_{j+\frac{1}{2}} = (u_{j+1} - u_j)/(t_{j+1} - t_j)$. The comparison of the two is shown in Figure 9. Note again excellent agreement from $t = 0.1$ s to just before $t = 0.4$ s at which point we see a 17% over-prediction of the deceleration in the FSI. The agreement improves again after the rebound peak at 0.613 s but shows a small constant deficit compared to the flight test measurements. Indeed, the payload velocity at the inflation peak ($t = 0.41$ s) in the FSI is 566.1 m/s compared to 573.2 m/s as measured in the flight test [28].

Given the drag of the payload is negligible compared to that of the parachute near peak inflation, it appears that the parachute as modeled in the FSI provides significantly more drag from the onset, causing an additional $\Delta v \approx -7$ m/s of velocity deficit by the time peak deceleration is reached. This is

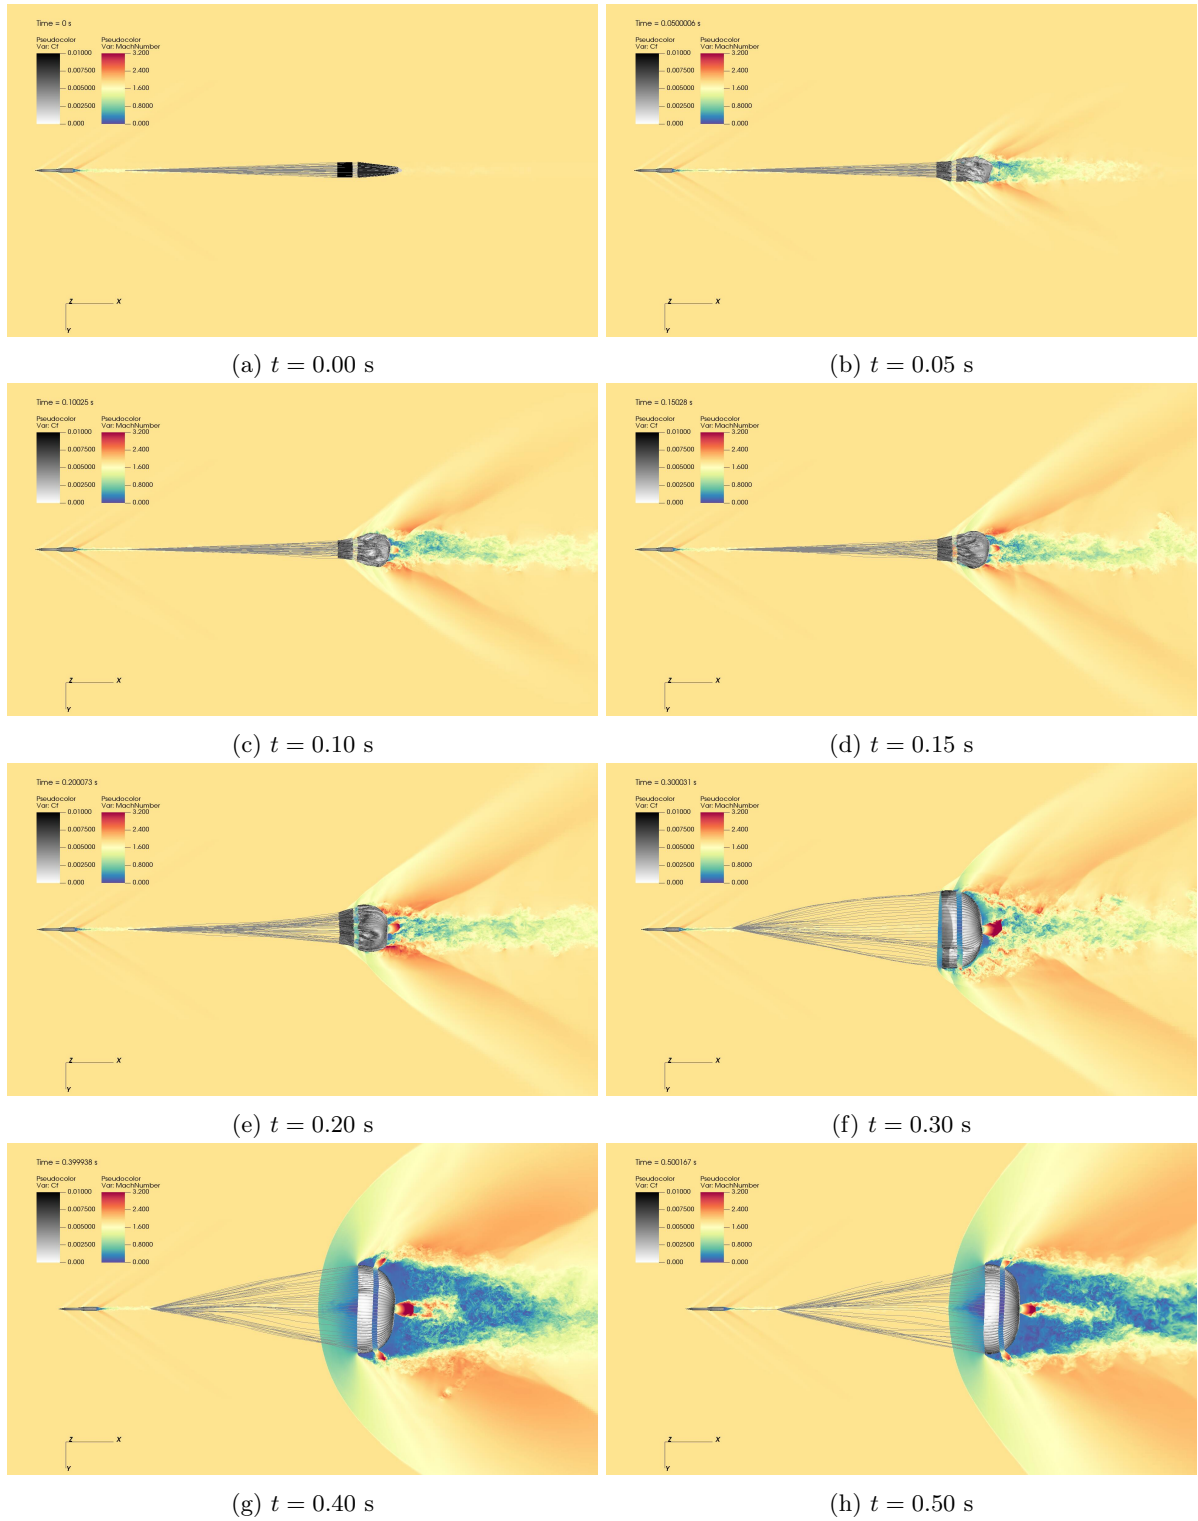


Figure 6: Snapshots from FSI-WMLES simulation of ASPIRE SR03 showing Mach number on the cut plane, as well as the coefficient of skin friction on the payload and the parachute surface. Suspension lines are shown in gray.

illustrated in Figure 10 where the parachute force is computed from flight measurements as

$$F_{\text{chute}} = F_{\text{pull}} - m_{\text{chute}}a_p, \quad (12)$$

where m_{chute} is the mass of the parachute system. Within the context of FSI, we can also reconstruct

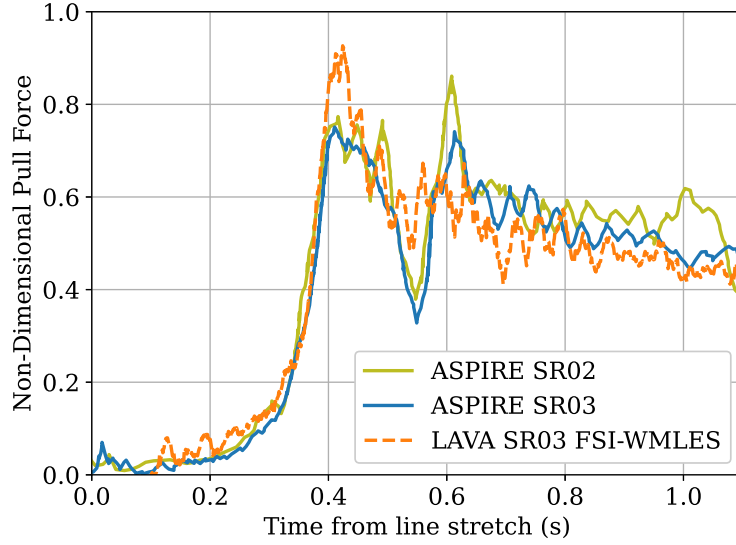


Figure 7: Comparison of non-dimensional pull force as a function of time from line stretch between the ASPIRE SR02, SR03 flight test measurements, and the LAVA SR03 FSI-WMLES predictions. SR02 and LAVA curves shifted by -0.0375 and 0.1 s respectively to remove inflation time phase difference with SR03.

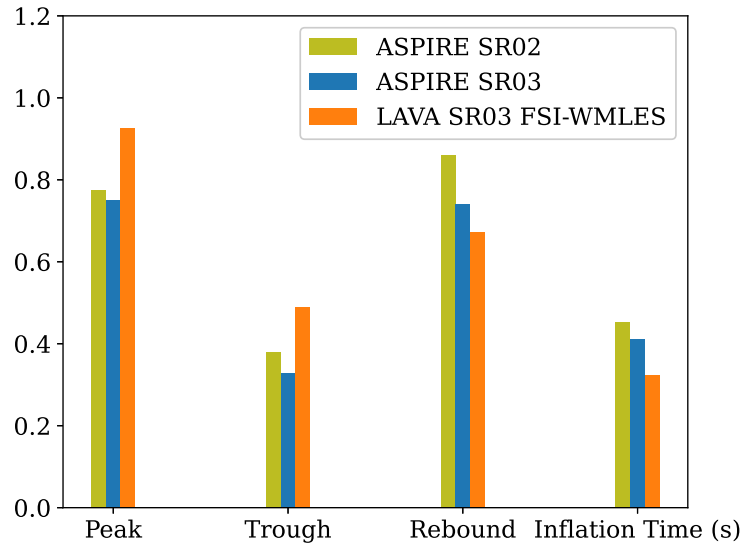


Figure 8: Comparison of quantities of interest derived from the non-dimensional pull force as a function of time from line stretch between the ASPIRE SR02 and SR03 flight test measurements, and the LAVA SR03 FSI-WMLES prediction. The y-axis represents non-dimensional pull force and time.

F_{chute} as above, but we also have direct access to the aerodynamic drag of the parachute as the integral of pressure and viscous forces over its triangulated surface F_{aero} . Note how the predicted aerodynamic drag of the parachute does not display the oscillations associated with the suspension line mode. This is expected as it is only dependent on the changes in the cross-sectional area of the parachute and the dynamic pressure. The downward slope after $t = 0.6$ s is driven by the reduction in dynamic pressure, whereas the lower frequency oscillations in parachute force after the peak are driven by the projected area changes. The correlation with projected area is made clear in Figure 11. By $t = 0.613$ s, the wind-relative velocity of the vehicle has reduced to 524.4 m/s compared to the 537.6 m/s recorded on the flight test. This deficit in wind relative velocity in the FSI directly translates to a lower dynamic pressure q_{∞} than experienced in flight. Given that the coefficient of drag of the parachute can be considered close to

constant once fully open, this q_∞ deficit explains why the predicted F_{chute} and F_{pull} have a small constant deficit in the stable deceleration phase despite projected areas on par or larger than those measured in the ASPIRE flight tests.

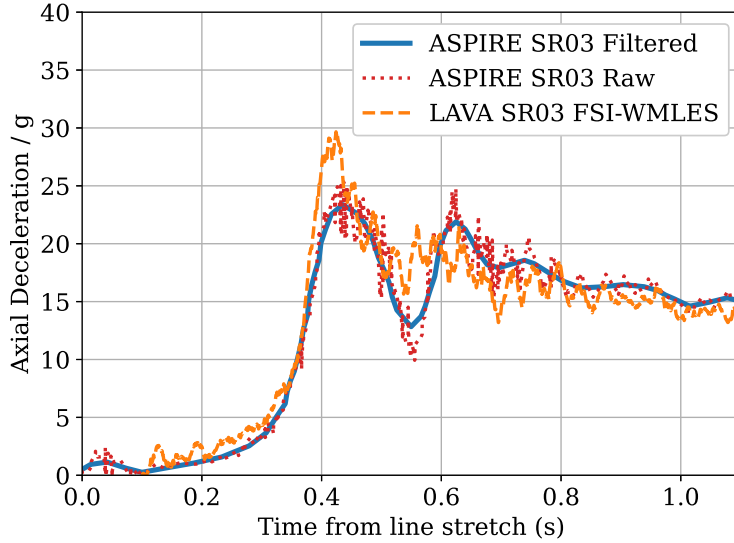


Figure 9: Comparison of non-dimensional axial deceleration as a function of time from line stretch of the ASPIRE SR03 flight test measurements with the LAVA SR03 FSI-WMLES prediction. FSI results shifted by 0.1s to align inflation peak.




Figure 10: Comparison of non-dimensional parachute force F_{chute} as a function of time from line stretch between the ASPIRE SR02, SR03 flight test measurements, and the LAVA SR03 FSI-WMLES predictions. SR02 and LAVA curves shifted by -0.0375 and 0.1 s respectively to remove inflation time phase difference with SR03.

What is the reason for the approximately 17% over-prediction in parachute drag? The simplest explanation is that the disk – the primary drag surface – was pulled to its maximum area sooner due to its different initial conditions and time evolution. This is supported by the non-dimensional projected area time history shown in Figure 11. The projected area of the flight tests and FSI were calculated using the convex hull procedure. The FSI projected area matches the flight tests well until $t = 0.4$ s after which its growth rate increases and reaches a peak just over 100% of the disk area 0.02 s earlier




Figure 11: Comparison of projected area non-dimensionalized by disk area as a function of time from line stretch between the ASPIRE SR02, SR03 flight test measurements in black and blue [11], and the LAVA SR03 FSI-WMLES prediction in orange. The ASPIRE SR02 time is shifted by -0.0325 s to align it to the SR03 peak, and the FSI is similarly shifted by 0.1 s.

than the peaks in the flight tests, which reach 0.96 and 0.88 for SR02 and SR03 respectively. The peak area reached by the FSI is more in line with the SR02 flight than SR03, which could be due to the change in canopy fabric material between the two flights, or more likely due to flight-to-flight variation. Similarly, the subsequent peaks in FSI and in the SR02 data reach 100% whereas SR03 does not reach the full disk area until 0.9 s. The post-inflation dip in projected area for FSI is however more consistent with the SR03 flight which saw a fairly small change in post-inflation area compared to SR02 whose area reduced by nearly 25% and had a slower recovery to the rebound peak.

Another possible contributing factor is that this simulation ignores the foreshortening of the vent lines: in the flight test, the vent lines are 10% shorter than the diameter of the vent hole in the disk [10]. Vent line foreshortening nominally reduces the effective cross-sectional area the disk can occupy by 1.75%. Similarly, the stitching in gore seams typically reduces the fabric length by 1 to 2% due to material “take-up”, which could lead up to a 4% drop in effective disk area in the flight test article. The tape stitched at the leading and trailing edges of the disk also leads to material take-up and changes the inflated shape of the disk gores near the leading edge by forcing additional circumferential and radial curvature, further reducing its effective cross-sectional area. Taken all together, the foreshortening of the vent lines and of all stitching could be responsible for up to a 7% reduction in effective disk area.

The FSI shows a trend consistent with the two flights where the peak pull force and peak parachute force occurs before the peak in projected area – see Figures 7, 10 and 11. It was conjectured that this trend is driven by the fact that the disk part of the canopy is the primary drag contributor and it is pulled taut first, and that the subsequent area increase is due to the band inflation but contributes much less to the overall drag due to its angle to the flow [11]. Although this is exactly how it occurs in FSI and SR02, where the disk is fully inflated by the time of the inflation peak, it was not the case for SR03 where the disk was not fully inflated until the rebound peak, as can be gleaned from Figure 11 and Figure 12.

Figure 12 shows the ASPIRE SR03 flight test images recorded by the black and white high-speed camera mounted on the back of the payload. Just after line stretch in Figure 12(a), the parachute is so tightly folded that it is barely larger than the bundle of suspension lines from this vantage point. This is a clear difference from our starting point at 10% of the disk diameter D . The folded canopy also appears to be at an angle with respect to the axis of the payload. Both of these facts could explain why the inflation time is 0.1 s longer in the flight test than in the FSI.

Figure 13 shows a similar view at comparable times from the FSI simulation. Good qualitative agreement in the shape and structure that the parachute takes on at various times is observed given that no two parachute inflation sequences are ever the same, even for the same parachute system deployed at the same freestream conditions due to the chaotic nature of the turbulent flow and of the parachute’s

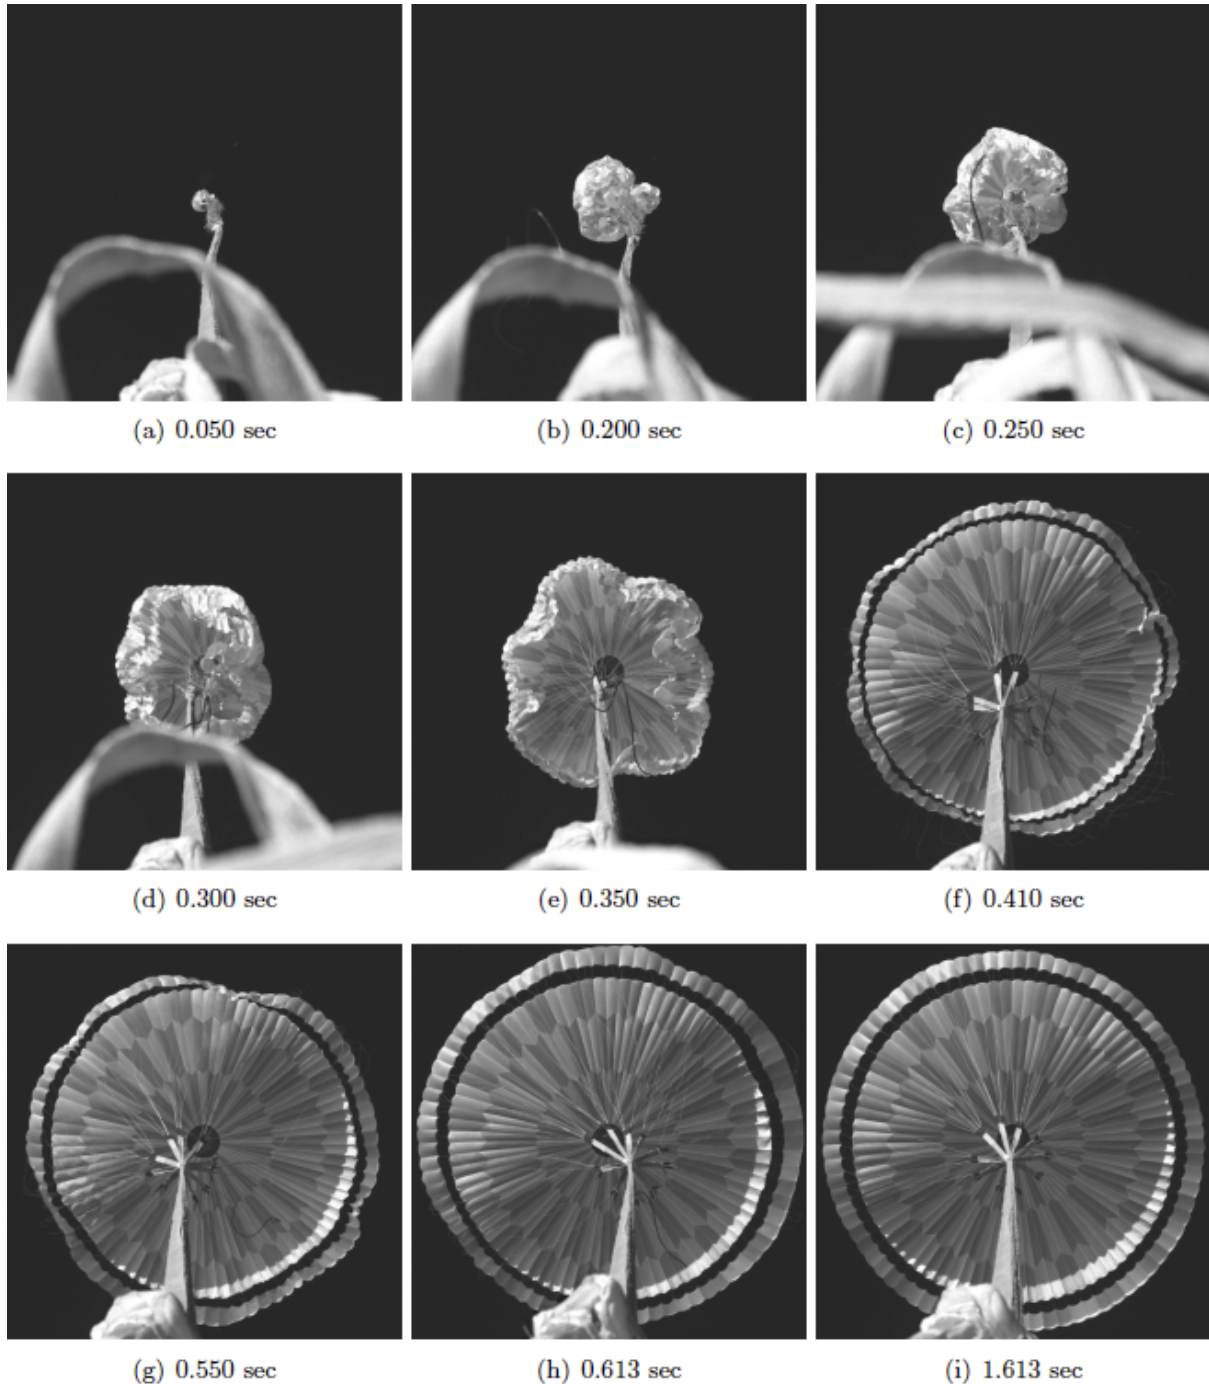


Figure 12: Frames from the black and white high-speed camera mounted on the back of the ASPIRE SR03 flight test payload [28]. Times are relative to line stretch.

exact shape as it exits the bag. Notice the similarities in the multi-lobed structure of the parachute at $t = 0.3$ s across the FSI and the flight test. Observe also that the peak load recorded at the payload occurs before the parachute reaches its maximum cross-sectional area, as several depressions appear in the band across FSI and flight test at $t = 0.410$ s. At the post-inflation trough of 0.550 s, the flight test shows several depressions in the band, whereas the FSI only displays three, likely due to slightly different timing of the trough for the FSI – see Figures 10 and 11. The band is then fully inflated at the flight test's rebound peak at $t = 0.613$ s onward and that is well mirrored in the FSI.

The computational cost of each portion of the simulation is described in Table 5. As massively parallel architectures come into production, the question of computational resources is becoming less about absolute cost or resource usage (CPU-hours), and more about time-to-solution. This emphasizes




Figure 13: Snapshots of parachute structural model in FSI viewed from the payload axis. Times have been shifted by $\Delta t = 0.1$ s to line up inflation peak to SR03 flight test.

strong-scaling and load balancing which is particularly challenging in coupled multi-physics problems like supersonic parachute inflation. LAVA made use of dynamic load balancing every time the mesh was adapted during the simulation. Unfortunately, mesh adaption and dynamic load balancing require collective operations for which the communication cost cannot easily be overlapped and hidden. The time-to-solution is still too long to use simulation results in a serial process to inform future parachute designs (e.g. based on simulation 1 results, change parameter ξ for simulation 2, and so on). However, this simulation only required moderate computing resources (5120 cores) so it is conceivable that many such simulations could soon be run in parallel exploring various parameters and still provide useful knowledge to a mission design team within a reasonable turnaround time.

Table 5: Computational resources for each segment of the simulation.

| Segment | Integrated Time (s) | Nodes | Node Type | Cores /Node | Wall Clock Time (Hours) | Core-Hours ($\times 10^3$) |
|---------------|---------------------|-------|-----------|-------------|-------------------------|------------------------------|
| CSD + Contact | 9.0 | 6 | AMD Rome | 128 | 192.0 | 148.5 |
| CFD | 0.2 | 40 | AMD Rome | 128 | 1.3 | 6.8 |
| FSI | 1.0 | 40 | AMD Rome | 128 | 145.8 | 746.5 |

6 Conclusion and Future Work

We present a novel method to obtain an initial parachute shape that better approximates line stretch flight observations than previous approaches. This new approach further reduces the initial cross-sectional area of the canopy that is open to the flow. The use of this new initial parachute shape directly improves FSI predictions for time to inflation, inflation rate, and indirectly improves predictions of peak pull force when compared to flight test measurements. An improved numerical method to simulate supersonic parachute inflation that reduces numerical dissipation in the convective flux and incorporates viscous effects via the use of immersed boundary adapted WMLES was also presented. The result of these changes to our CFD method primarily improve predictions of viscous drag and flow separation both in the wake of the payload and the parachute during inflation. Excellent agreement was observed between FSI and flight test measurements for the inflation rate, and the stable deceleration stage. The pull force and parachute force predicted by the FSI also generally agrees well with the flight test measurements, except for the peak and the post-inflation trough which are both over-predicted by 17%. This is most likely due to the parachute reaching a larger projected area sooner than what was observed in flight tests due to its different initial conditions and time evolution. The FSI projected area otherwise correlated well with high-speed imagery computed from the flight tests, both showing that the peak pull force precedes the peak projected area. CFD and CSD grid sensitivity studies are planned to help bound the uncertainty associated with discretization error in the simulation’s results.

Continued research and development in several subdomains may be necessary to accurately predict a novel parachute design’s potential failure. For example, modeling vent lines and fabric seams foreshortening could help predict the canopy’s inflated shape, effective area, and stress distribution more accurately. Improving woven fabric material models to represent the orthotropic nature of the bias construction of the canopy (with fibers at 45°) could also improve stress predictions. It may also be beneficial to consider developing material models beyond Saint Venant–Kirchhoff capable of reflecting the non-linear stress-strain behavior, hysteresis, and energy dissipation in woven fabrics and braided chords recently documented [29, 30, 31]. Similarly, the porosity of woven fabrics has recently been shown to vary with strain [32] which could motivate the development of porosity models that vary not only with the flow states on either side of the fabric, but also with the local strain experienced in the structural model.

Acknowledgments

This work was supported by the NASA Entry systems Modeling (ESM) project. Computational resources were provided by the High-End Computing Capability (HECC) program on NASA Advanced supercomputing (NAS) systems at NASA Ames Research Center. This work would not have been possible without the guidance and support of Jared Duensing. Thank you to Aaron Brandis, Justin Haskins, Ben Tutt, and Giovanni Guecha-Ahumada for insightful discussions and for helping guide the focus of our efforts. Special thanks to Jason Rabinovitch for organizing the FSI workshop which spurred great cross-discipline conversations and ideas for future collaboration. We are also grateful to Brandon Lowe and Aaron Brandis for carefully reviewing this manuscript, as well as all current and former LAVA team members whose contributions made this work possible.

References

- [1] Ian Clark and Christopher Tanner. A historical summary of the design, development, and analysis of the disk-gap-band parachute. In *2017 IEEE Aerospace Conference*, pages 1–17, 2017.
- [2] Ian G. Clark, John C. Gallon, and Allen Witkowski. Parachute decelerator system performance during the low density supersonic decelerator program’s first supersonic flight dynamics test. In *23rd AIAA Aerodynamic Decelerator Systems Technology Conference*, 2015.

- [3] Aaron M. Brandis, Michael Barnhardt, Thomas K. West, and Monica Hughes. New developments in NASA’s entry systems modeling project. In *AIAA SCITECH 2023 Forum*. AIAA, 2023.
- [4] Cetin C Kiris, Jeffrey A Housman, Michael F Barad, Christoph Brehm, Emre Sozer, and Shayan Moini-Yekta. Computational framework for launch, ascent, and vehicle aerodynamics (LAVA). *Aerospace Science and Technology*, 55:189–219, 2016.
- [5] Jonathan Boustani, Francois Cadieux, Gaetan K. Kenway, Michael F. Barad, Cetin C. Kiris, and Christoph Brehm. Fluid-structure interaction simulations of the ASPIRE SR01 supersonic parachute flight test. *Aerospace Science and Technology*, 126:107596, 2022.
- [6] Jonathan Boustani, Gaetan Kenway, Francois Cadieux, Michael F. Barad, Christoph Brehm, and Cetin C. Kiris. Fluid-structure interaction simulations of the ASPIRE SR01 supersonic parachute. In *AIAA SCITECH 2022 Forum*. AIAA, 2022.
- [7] Francois Cadieux, Jordan B. Angel, Michael F. Barad, and Cetin C. Kiris. Fluid-structure interaction simulations of the ASPIRE SR03 supersonic parachute flight test. In *AIAA SCITECH 2023 Forum*. AIAA, 2023.
- [8] Clara O’Farrell, Chris Karlgaard, Jake A Tynis, and Ian G Clark. Overview and reconstruction of the ASPIRE project’s sr01 supersonic parachute test. In *2018 IEEE Aerospace Conference*, pages 1–18. IEEE, 2018.
- [9] Clara O’Farrell, Bryan S. Sonneveldt, Chris Karlgaard, Jake A. Tynis, and Ian G. Clark. Overview of the ASPIRE project’s supersonic flight tests of a strengthened DGB parachute. In *2019 IEEE Aerospace Conference*, pages 1–18. IEEE, 2019.
- [10] Jason Rabinovitch, Faisal As’ad, Philip Avery, Charbel Farhat, Navid Ataei, and Marcus Lobbia. Update: Modeling supersonic parachute inflations for mars spacecraft. In *26th AIAA Aerodynamic Decelerator Systems Technology Conference*, 2022.
- [11] Jason Rabinovitch, Gregory S. Griffin, William Seto, Clara O’Farrell, Christopher L. Tanner, and Ian G. Clark. ASPIRE supersonic parachute shape reconstruction. In *AIAA Scitech 2019 Forum*. AIAA, 2019.
- [12] AW Vreman. An eddy-viscosity subgrid-scale model for turbulent shear flow: Algebraic theory and applications. *Physics of fluids*, 16(10):3670–3681, 2004.
- [13] Faisal As’ad, Philip Avery, Charbel Farhat, Jason Rabinovitch, and Marcus Lobbia. Validation of a high-fidelity supersonic parachute inflation dynamics model and best practice. In *AIAA SCITECH 2022 Forum*, 2022.
- [14] Johan Larsson, Soshi Kawai, Julien Bodart, and Ivan Bermejo-Moreno. Large eddy simulation with modeled wall-stress: recent progress and future directions. *Mechanical Engineering Reviews*, 3(1):15–00418, 2016.
- [15] Oliver MF Browne, Jeffrey A Housman, Gaetan K Kenway, Aditya S Ghate, and Cetin C Kiris. Numerical investigation of cl, max prediction on the NASA high-lift common research model. *AIAA Journal*, 61(4):1639–1658, 2023.
- [16] Kyu Hong Kim, Chongam Kim, and Oh-Hyun Rho. Methods for the accurate computations of hypersonic flows: I. AUSMPW+ scheme. *Journal of computational physics*, 174(1):38–80, 2001.
- [17] Marsha Berger, Michael Aftosmis, and Scott Murman. Analysis of slope limiters on irregular grids. In *43rd AIAA Aerospace Sciences Meeting and Exhibit*, page 490, 2005.
- [18] Antony J Musker. Explicit expression for the smooth wall velocity distribution in a turbulent boundary layer. *AIAA Journal*, 17(6):655–657, 1979.
- [19] Siavash Toosi and Johan Larsson. Anisotropic grid-adaptation in large eddy simulations. *Computers & Fluids*, 156:146–161, 2017. Ninth International Conference on Computational Fluid Dynamics (ICCFD9).
- [20] Benjamin Tutt, Scott Roland, Greg Noetscher, and Richard Charles. Finite mass simulation techniques in LS-DYNA. In *21st AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, 2011.
- [21] X. Deng, M. Mao, Y. Jiang, and H. Liu. New High-Order Hybrid Cell-Edge and Cell-Node Weighted Compact Nonlinear Schemes. In *20th AIAA Computational Fluid Dynamics Conference, Honolulu, Hawaii*, June 2011. AIAA-2011-3857.
- [22] Yoshiharu Tamaki, Yuichi Kuya, and Soshi Kawai. Comprehensive analysis of entropy conservation property of non-dissipative schemes for compressible flows: KEEP scheme redefined. *Journal of Computational Physics*, 468:111494, 2022.
- [23] F Ducros, F Laporte, Th Soulères, Vincent Guinot, Ph Moinat, and B Caruelle. High-order fluxes for conservative skew-symmetric-like schemes in structured meshes: application to compressible flows. *Journal of Computational Physics*, 161(1):114–139, 2000.

- [24] Douglas L Sondak and Richard H Pletcher. Application of wall functions to generalized nonorthogonal curvilinear coordinate systems. *AIAA journal*, 33(1):33–41, 1995.
- [25] Xiang IA Yang, George Ilhwan Park, and Parviz Moin. Log-layer mismatch and modeling of the fluctuating wall stress in wall-modeled large-eddy simulations. *Physical review fluids*, 2(10):104601, 2017.
- [26] Suman Muppidi, Clara O’Farrell, Christopher Tanner, John W Van Norman, and Ian G Clark. Modeling and flight performance of supersonic disk-gap-band parachutes in slender body wakes. In *2018 Atmospheric Flight Mechanics Conference*, page 3623, 2018.
- [27] Suman Muppidi, Clara O’Farrell, John W Van Norman, and Ian G Clark. ASPIRE aerodynamic models and flight performance. In *AIAA Aviation 2019 Forum*, page 3376, 2019.
- [28] Bryan S Sonneveldt, Ian G Clark, and Clara O’Farrell. Summary of the advanced supersonic parachute inflation research experiments (ASPIRE) sounding rocket tests with a disk-gap-band parachute. In *AIAA Aviation 2019 Forum*, page 3482, 2019.
- [29] Alireza Amirkhizi, Marinos Blanas, Fernando Campos, Kayla Krantz, Patrick Drane, James Sherwood, and Allen Witkowski. Material characteristics of F-111 (PIA-C-44378) canopy fabric in uniaxial and biaxial conditions. In *26th AIAA Aerodynamic Decelerator Systems Technology Conference*. AIAA, 2022.
- [30] Catherine P. Barry, Bradford G. Olson, Keith Bergeron, David J. Willis, and James A. Sherwood. Modeling tensile tests of a braided parachute suspension line using a mesomechanical finite element model. In *AIAA Scitech 2020 Forum*. AIAA, 2020.
- [31] Bradford G. Olson, James Sherwood, David J. Willis, and Keith Bergeron. Investigation of tensile properties of braided parachute suspension line. In *24th AIAA Aerodynamic Decelerator Systems Technology Conference*. AIAA, 2017.
- [32] Cutler Phillippe, Marco Mattei, Francesco Panerai, and Laura Villafañe Roca. In-situ micro-scale characterization of parachute textiles with micro-tomography and machine learning. In *AIAA SCITECH 2023 Forum*. AIAA, 2023.