

Fluid-Structure Interaction Simulations of the ASPIRE SR03 Supersonic Parachute Flight Test

Francois Cadieux*, Jordan B. Angel*, Michael F. Barad*, and Cetin C. Kiris†
NASA Ames Research Center, Moffett Field, CA, 94035

Research into parachute performance continues to be a source of significant investment from the National Aeronautics and Space Administration to mitigate risks and to enable a variety of exploration missions, including landing on Mars as well as returning to Earth. The cost of flight tests to certify any changes to the current state-of-the-art parachute designs limits the development of next generation parachute systems. Fluid-structure interaction simulations could help accelerate this process once validated. The Launch, Ascent, and Vehicle Aerodynamics team is developing the capability to perform such fluid-structure interaction simulations by coupling a higher-order Cartesian immersed boundary computational fluid dynamics solver with adaptive mesh refinement to a finite element structural dynamics solver in space and time. We continue the effort to validate this tool with the Advanced Supersonic Parachute Inflation Research Experiments SR03 flight test featuring a strengthened parachute akin to the Mars 2020 mission that landed the Perseverance rover on Mars, and a higher freestream dynamic pressure prior to inflation. The effect of the flow conditions’ angle of attack and of the initial parachute shape are quantified. The impact of relaxing modeling assumptions with regards to radial stiffeners on the parachute canopy is also investigated. Results demonstrate improvements in agreement with the pull force recorded during the SR03 flight test as the initial conditions of the flow and parachute are brought closer to those experienced in flight, and further improved when the radial stiffener modeling assumptions are relaxed.

I. Nomenclature

M_∞	=	freestream Mach number
q_∞	=	freestream dynamic pressure (Pa)
ρ_∞	=	freestream density (kg/m^3)
α	=	freestream angle of attack (in $^\circ$)
D_0	=	parachute nominal diameter (m)
S_0	=	parachute nominal cross-sectional area (m^2)

II. Introduction

PARACHUTE systems like the one used by the Mars 2020 mission continue to be the safest and most reliable technology to decelerate robotic space exploration probes from supersonic velocities during atmospheric entry until retro-rockets can be used to land the payload safely on the surface of planets like Mars (e.g. the “sky-crane” maneuver). In order to enable larger payloads and improved landing location accuracy, continued research and development into parachute system design is necessary, from its deployment mechanism to its woven-material. For example, the low-density supersonic decelerator (LDSD) project [1] aimed to test two novel technologies: an inflatable decelerator on the circumference of the payload’s aero-shell to augment ram drag through the upper atmosphere, and a new Disksail parachute design that melded elements of traditional subsonic ringsail parachutes and more traditional disk-gap-band (DGB) supersonic parachutes [2]. Although the inflatable decelerator proved successful, the Disksail parachute showed signs of damage early on during its inflation and ultimately was ripped apart [3]. The cause of the Disksail’s failure could not be ascertained with certainty even after thorough review and investigation by NASA, JPL, and the wider aerospace engineering community. The leading hypothesis is that the Disksail design caused more severe stresses than

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

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

seen in typical DGB parachutes. Its larger shoulder region may have pulled the disk portion of the canopy flat earlier in the inflation process, causing the fabric to tear under the increased load.

Fluid-structure interaction (FSI) simulations could help bridge this gap in our ability to predict canopy stresses and thus accelerate the development of new parachute systems once validated and best practices are well established. The Launch, Ascent, and Vehicle Aerodynamics (LAVA) team is developing the capability to perform FSI simulations by coupling a higher-order Cartesian immersed boundary computational fluid dynamics (CFD) solver with adaptive mesh refinement (AMR) to a computational structural dynamics (CSD) solver based on the finite element method in space and time [4]. The tool was recently used to simulate a portion of the first flight test of the successful ASPIRE program called SR01. SR01 featured a build-to-print version of the DGB Mars Science Laboratory parachute that was inflated in supersonic conditions in earth’s upper atmosphere behind a slender payload. LAVA FSI predictions for the peak load at the payload’s load pins were within 10% of the flight test measurements and were conservative. Grid convergence with respect to the volume and structural domains were also demonstrated [5]. The subsequent ASPIRE tests, SR02

ASPIRE Flight	q_∞ at line stretch (Pa)	M_∞ at line stretch	D_0 (m)	Parachute system mass (kg)
SR01	491.68	1.79	21.35	49.8
SR02	744.57	2.00	21.45	81.6
SR03	1028.44	1.88	21.45	82.0

Table 1 Overview of some differences between ASPIRE test flights [6, 7].

and SR03 featured significantly strengthened parachutes akin to the Mars 2020 parachute system, slight differences in packing, and more importantly larger dynamic pressure just prior to inflation: see Table 1 for details and a comparison to SR01. Unfortunately, these differences largely prevent us from making a direct uncertainty assessment regarding the repeatability of the experimental measurements, and so it is difficult for us to determine whether the previous SR01 FSI simulations fall within that uncertainty.

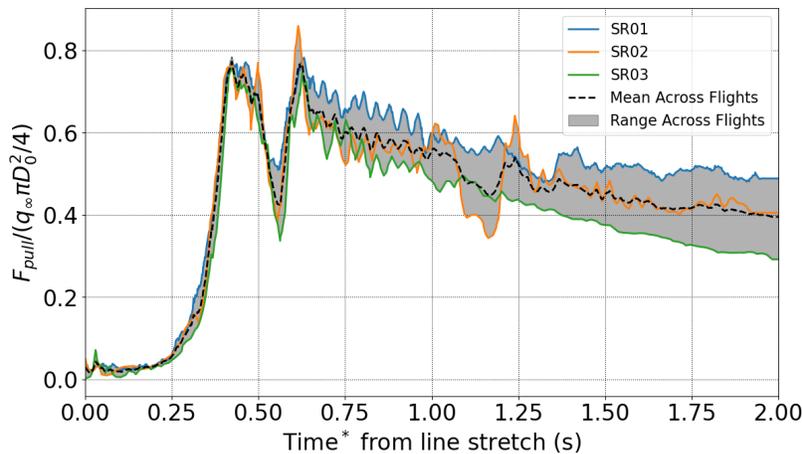


Fig. 1 ASPIRE flight to flight variation in non-dimensional pull force measured at the load pins. * Each time series has been offset so its inflation peak occurs at 0.423 s after line stretch [6, 7].

Nevertheless, if we non-dimensionalize the load pin force by the nominal parachute cross-sectional area ($S_0 = \pi D_0^2/4$, and the dynamic pressure at line stretch (q_∞), it becomes clear that the non-dimensional load pin inflation peak value is repeatable, with less than 4% variation across the three flights (ignoring slight differences in the inflation time after line stretch) – see Figure 1. Plotting the load pin force in this way also reveals that there is significant spread in the post-inflation trough and rebound peak which persists in the stable deceleration phase, even between the otherwise similarly constructed parachutes (SR02 and SR03). The average and variation across flights for the inflation peak, trough (minimum between inflation and rebound peak), and rebound peak is quantified in Table 2 demonstrating that the trough, rebound and deceleration measurements are not as repeatable as the inflation peak and carry larger flight-to-flight variability.

Flight Event	Mean Across Flights	Percent Variation Across Flights
Inflation Peak	0.773	3.42
Rebound Peak	0.798	14.93
Trough	0.409	42.10
Post-Rebound (t=2s)	0.570	27.03

Table 2 Non-dimensional pull force at important events across ASPIRE flights [6, 7].

If we compare our previous FSI efforts to ASPIRE measurements across flights, it is clear that the SR01 FSI model inflates more quickly, over-predicts the mean inflation peak, feels slack earlier than all flights (post-inflation trough), predicts a reasonable trough value, but then under-predicts the rebound peak as shown in Figure 2. The deceleration phase is not considered since the previous simulation did not model the system’s deceleration – the payload was restrained and the flow conditions kept constant in time. Putting the previous LAVA FSI results in the broader

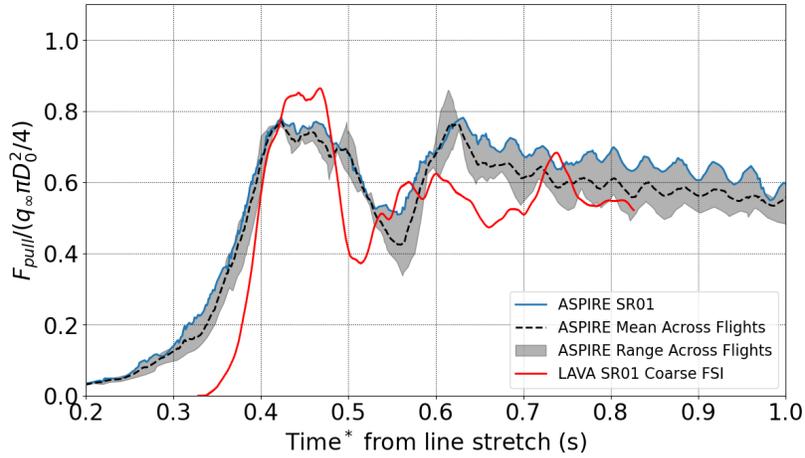


Fig. 2 Comparison of previous LAVA SR01 FSI simulation results [5] compared to ASPIRE measurements [6, 7]. * Each time series has been offset so its inflation peak occurs at 0.423 s after line stretch.

context of the repeatability of the ASPIRE flights helps to reinforce one of the conclusions drawn from our previous grid convergence study: having demonstrated less than 1% variability in the inflation peak across consistent refinement both in the CFD and CSD domains, the remaining error must lie with the assumptions made in the model. One likely source of modeling error stems from the assumptions made about the construction of the parachute and its material properties in the absence of detailed information available in literature. For example, the presence and properties of the band and disk leading and trailing edge circumferential stiffeners were unknown and thus completely ignored. The lack of data regarding the radial stiffeners led us to make the assumption that they could be modeled by having the suspension lines connect through the band, gap, disk, and vent lines as 40 single threads starting and ending at the riser. The properties of these suspension lines and radial stiffeners were also unknown and approximated as a type of Kevlar with a large Young’s modulus (stiffness) [5]. Another source of modeling error that could affect the inflation time and peak load is the initialization procedure for the FSI simulations. This can be broken down into the following aspects:

- Initial conditions of the flow: $M_\infty, q_\infty, \rho_\infty, \alpha,$
- Initial shape of the parachute (whether and how it is structurally “pre-inflated”) which determines its initial cross-sectional area S_i and its level of asymmetry,
- Length of time the parachute is held static and rigid to allow the flow to react to its presence and produce a physically-representative pressure distribution over its surface before starting the FSI simulation.

Stanford/JPL recently completed a thorough review of the as-built design, construction, and material properties of the ASPIRE SR02 and SR03 parachutes along with their modeling simplifications and initialization procedure for their own FSI simulations to enable external FSI validation efforts [8]. In the spirit of collaboration and code-to-code verification, the focus of this work is to perform and analyze FSI simulations of the SR03 flight test using these newly

released data about the parachute system construction and material properties. The effect of the initial conditions of the flow, the initial shape of the parachute, and the modeling of radial stiffeners are investigated. We first detail the choices made in the construction of a new parachute structural mesh to better represent the as-built design of the SR02 and SR03 parachutes. Changes in our methodology from that described in previous works [4, 5] are then presented and justified. We then present results for a sequence of four FSI simulations of ASPIRE SR03 exploring the effects of initial flow conditions (angle of attack α), initial parachute shape (introducing asymmetry in the inflation process), and modeling of the radial stiffeners.

III. Methodology

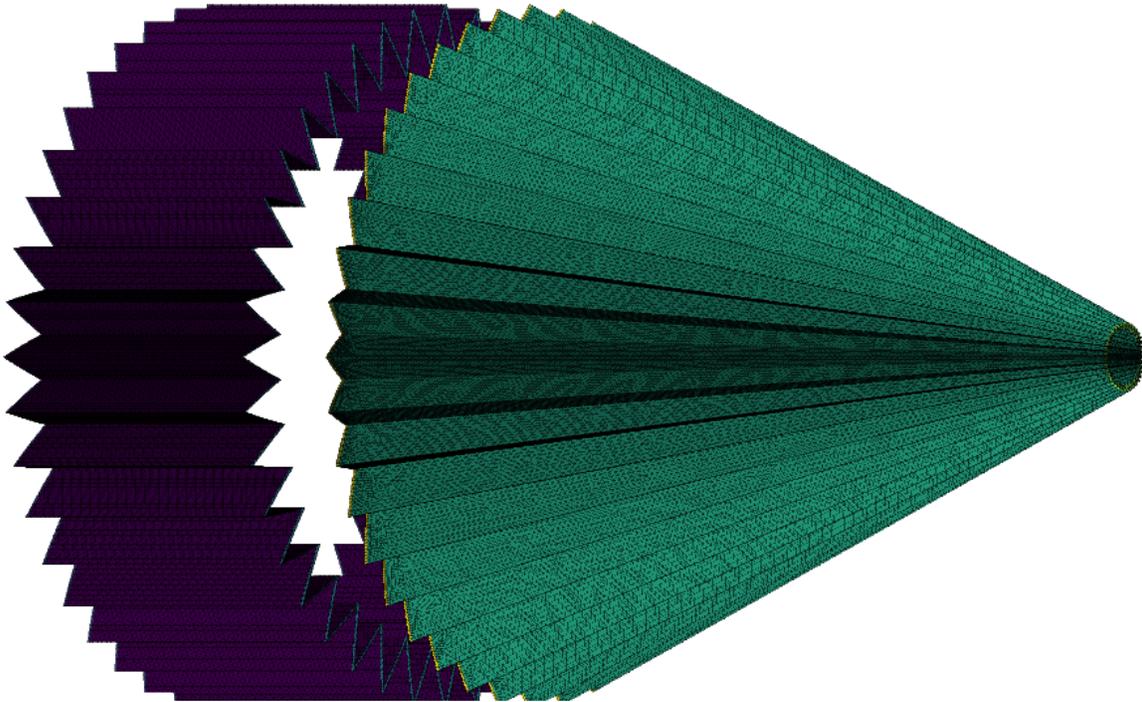


Fig. 3 Iso view of SR03 parachute structural mesh colored by component showing the leading and trailing edge regions used for circumferential stiffeners.

Following the description of the ASPIRE SR02/SR03 parachute system [8], we used STAR-CCM+ to design a triangulation with clear feature lines through the gores, and different components for each of the leading and trailing edge stiffeners of the band and disk as shown in Figure 3. The mesh is constructed in a way that the folded diameter is 40% of the nominal diameter of the disk D_0 by folding each gore into peaks and valleys, and rotating the gores of the disk vertically into an accordion like shape. The mean triangle edge length is 0.05 m. The minimum triangle edge length of 0.025 m is used at the leading and trailing edges of the band and disk. The maximum edge length growth rate is 1.2. This yields 267,668 triangular MITC3 shell elements on the canopy structural mesh. The parachute mesh is then inflated to a thickness of 0.001 m representative of the thickness of the circumferential stiffeners to yield a water-tight CFD mesh with 541,416 triangles with a minimum edge length of 0.001 m where the two inflated surfaces connect, but otherwise identical node distributions to the original surface definition. The parachute mesh is augmented with 76,954 line (beam) elements to model the suspension lines, band radial stiffeners (gore seams), gap lines, disk radial stiffeners, and vent lines as shown in Figure 4. Based on our previous best practice from SR01 and from previous Stanford/JPL publications [9], we initially connect the 80 suspension lines from the riser through the canopy to vent back to riser into 40 threads to model the radial stiffeners on the band and disk. To test the influence of this assumption, we also produced a structural mesh where the radials are not modeled. This second structural mesh is identical in every way except for that the suspension lines do not go through the band and disk, and thus has fewer line elements (61,193). In both structural meshes, the line (beam) element length is kept constant at roughly 0.05 m. The payload is assumed to be

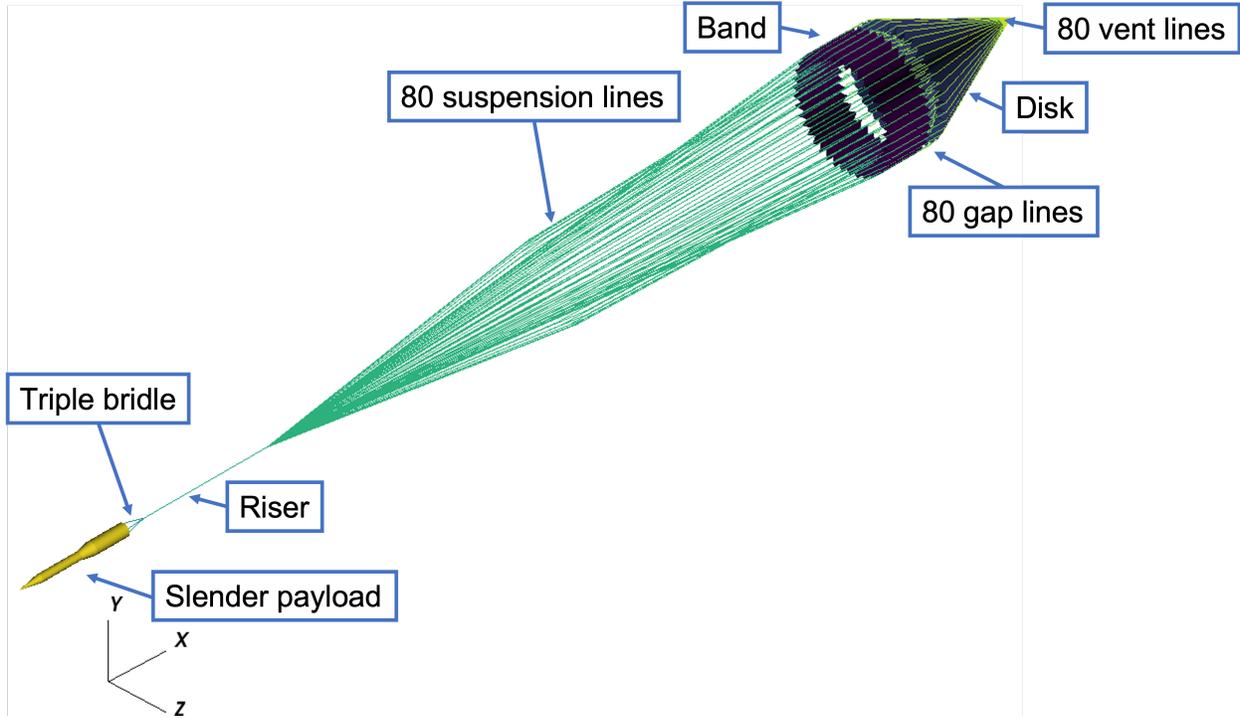


Fig. 4 Complete structural dynamics mesh of ASPIRE system colored by component, showing the added bridles, riser, suspension lines, gap and vent lines.

rigid and static, and its water-tight CFD representation was re-meshed from its previous iteration [4] to roughly match the minimum edge length of the parachute 0.025 m away from its blunt nose. This reduced the number of triangles from 186,320 to 43,346. Figure 5 displays the ASPIRE payload mesh used in CFD-only and FSI simulations.

The CSD solver used in LAVA treats the bridles, riser, suspension lines, radial stiffeners, gap and vent lines as Timoshenko beams that do not resist compressive forces. The disk and band broadcloth, along with their circumferential leading and trailing edge stiffeners are modeled with MITC3 triangular elements augmented with a drilling degree of freedom to alleviate shear-locking. Porosity of the canopy fabric is modeled by allowing a certain amount of mass flow through the canopy based on the pressure difference across its thickness by analytically solving the Darcy-Forcheimer equations. An effective porosity of 4.5% is imposed based on the PIA-C-7020D fabric specification and experiments done by Cruz et al. [10]. However, due to an oversight, all simulations presented in this work have an effective porosity of 3.92% instead. We expect this reduced porosity will slightly increase the canopy wake deficit and thus generate more shear-driven turbulence, making the parachute slightly more unsteady and prone to larger oscillations in effective cross-sectional area. The inputs to the porosity model will be corrected to reach the target 4.5% in future work. Contact mechanics are handled by an efficient ray-tracing algorithm that identifies if any triangular structural elements will come into contact and imposes an equal and opposite momentum (coefficient of restitution of 1) to ensure the elements do not cross. For more details on the structural dynamics model and solver, refer to Boustani et al. [5]. The material properties of each component of the structural dynamic mesh are presented in Table 3.

Having demonstrated less than 1% sensitivity for the CFD volume mesh in prior publications, we limit the scope of this work to the coarsest resolution previously used $\Delta x = 0.033$ m. The time integration algorithm for the CFD was improved to the three-stage third order Strong Stability Preserving Runge-Kutta (SSPRK) to enable a higher stable CFL of 0.9 during FSI simulations. The convective scheme was also improved from the incremental-stencil fifth-order weighted essentially non-oscillatory (IS-WENO5) to the less dissipative Z-WENO5 [11]. The CFD mesh is static in time and is defined by a few refinement regions in the shape of cylinders, and of frustums – see Figure 6a for an overview of the mesh. Solution-based adaptive mesh refinement (AMR) is not used to ensure greater consistency between runs, and will be studied in future work. We use the finest mesh resolution to cover the payload, its wake all the way to the parachute, and a generous cylinder around the parachute as evident in Figure 6a. The cylinder refinement region around the parachute is designed to encompass the parachute at its widest, and capture the supersonic flow coming out of the

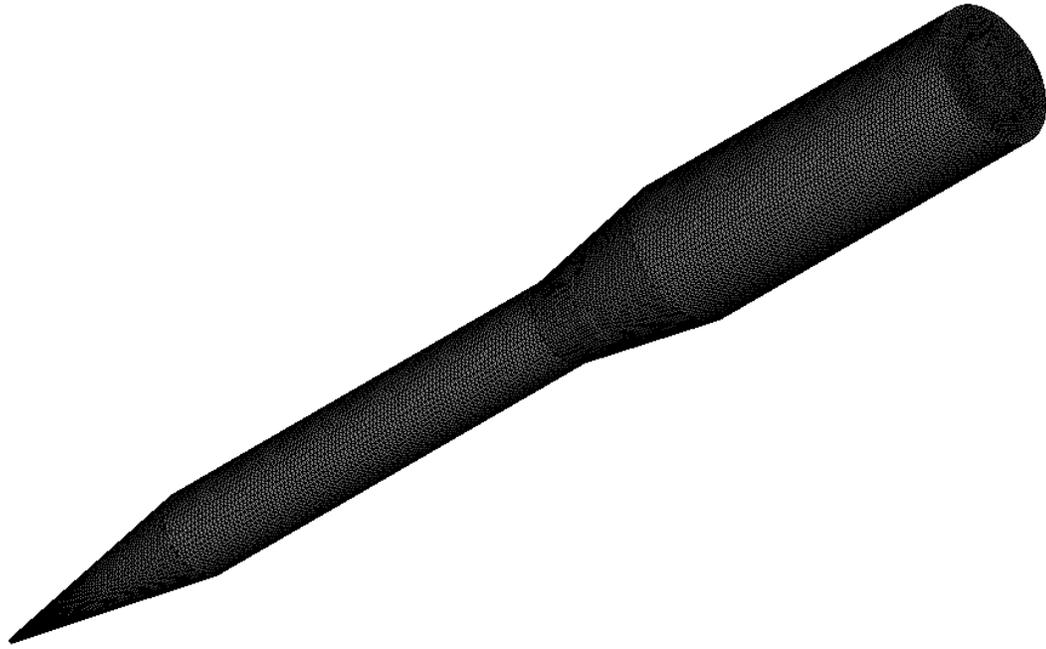


Fig. 5 Iso view of the ASPIRE payload CFD surface mesh.

Component name	Element type	Young's modulus (Pa)	Poisson ratio	Density (kg/m ³)	Width/ Radius (m)	Thickness (m)
Bridles	beam	1.6680×10^{10}	0.4	705.13	0.01954	-
Riser	beam	1.8600×10^{10}	0.4	751.99	0.02389	-
Suspension lines	beam	1.8202×10^{10}	0.4	484.66	0.00253	-
Radial stiffeners	beam	1.8202×10^{10}	0.4	484.66	0.00253	-
Broadcloth	MITC3	9.4484×10^8	0.4	533.95	-	7.6200×10^{-5}
Disk leading edge	MITC3	1.3568×10^{10}	0.4	454.07	0.0254	1.0583×10^{-3}
Band leading edge	MITC3	1.3568×10^{10}	0.4	454.07	0.0254	1.0583×10^{-3}
Band trailing edge	MITC3	1.3568×10^{10}	0.4	454.07	0.0254	1.0583×10^{-3}
Disk trailing edge	MITC3	1.0979×10^{10}	0.4	525.54	0.0254	2.5400×10^{-3}

Table 3 Material properties used in finite element model are consistent with Rabinovitch et al. [8].

gap with a radius of roughly 9 m. It extends downstream to ensure the shock of the supersonic flow going through the vent hole is also fully encompassed on the finest mesh resolution. The upstream extent is buffered to ensure the bow shock in front of the canopy is inside the cylindrical refinement region even at its furthest during the post-inflation trough, as shown in Figure 6b. The total cell count for the CFD mesh is 305 million.

The startup procedure for the FSI simulations is as follows :

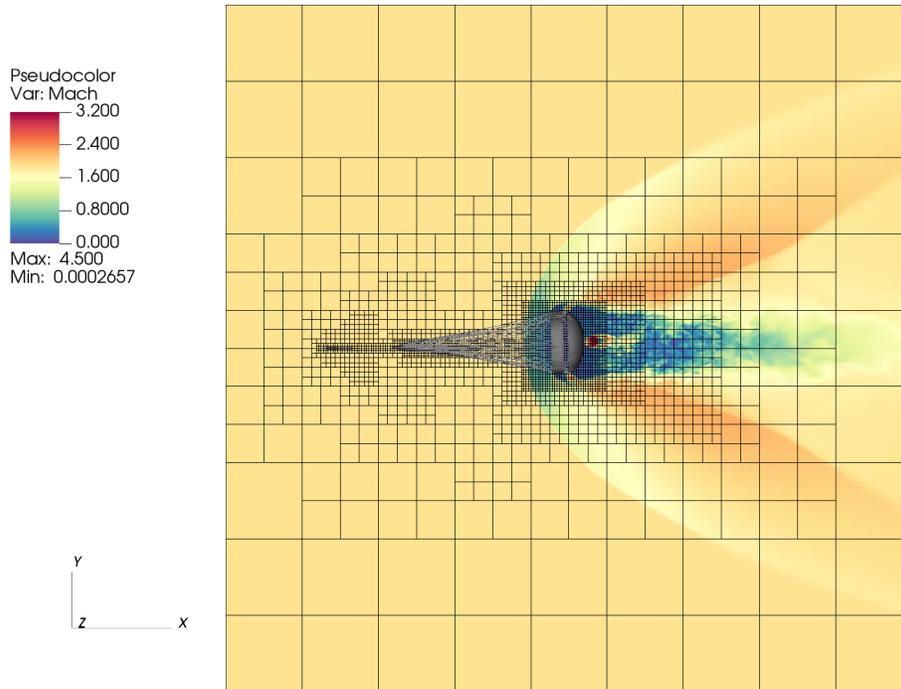
- 1) Perform a CFD simulation of the rigid and static ASPIRE payload until its aerodynamic loads have reached a statistically stationary state and its wake has reached beyond the location of the parachute (55 m downstream), roughly 0.25 seconds.
- 2) Pre-inflate the parachute using structural dynamics solver coupled with contact mechanics module using prescribed force normal to the canopy surface, and in certain cases augmented with a prescribed force in the radial direction with asymmetric circumferential variation for 0.016 seconds.

- 3) Restart CFD simulation of the rigid and static ASPIRE payload with the rigid and static pre-inflated parachute mesh and advance it in time until aerodynamic loads on the parachute canopy have reached a statistically stationary state (typically another 0.25 seconds).
- 4) Perform coupled FSI simulation.

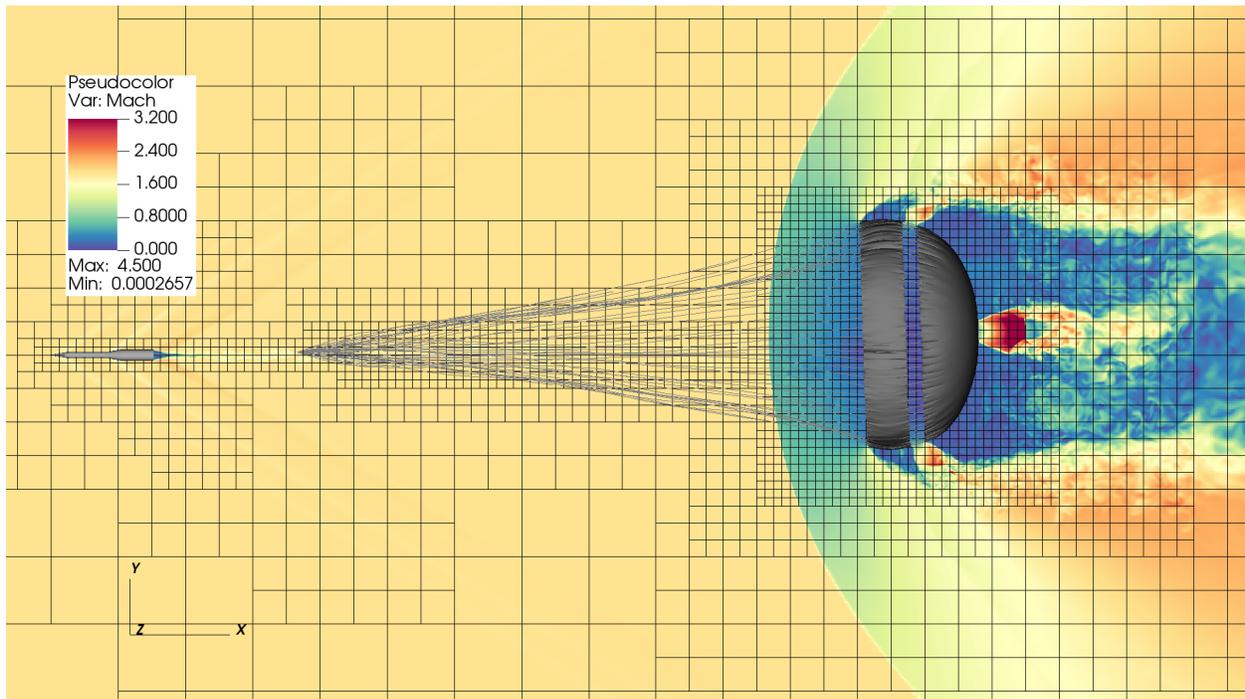
The flow conditions for the simulations shown in Table 4 are taken from the first moment at which the load pins detected loads above measurement uncertainty, otherwise called the line stretch event, which occurred 1.027 seconds after the mortar fire that freed the parachute system from the aft of the payload. These conditions are chosen despite our inability to match the small parachute cross-sectional area and complex asymmetric shape at that time because the flight recorded fairly little difference in conditions between line stretch and inflation peak load. However, the deceleration after peak load is so high, nearly 25 times the Earth's gravity, that we suspect it will be very difficult to match the timing and loads experienced by the flight test during the post-inflation trough and rebound without modeling the system's deceleration. So the focus is limited primarily to predicting the inflation peak load knowing that the remainder of the simulation will correspond to a different scenario than what transpired during the flight test given the restrained motion of the payload.

Quantity	Value	Unit
M_∞	1.88	-
q_∞	1028.44	Pa
ρ_∞	6.0171×10^{-3}	kg/m ³
α	0.9	°
Reynolds number per meter	2.2784×10^5	1/m
Prandtl number (Pr)	0.7	-
Ratio of specific heat γ	1.4	-
Gas constant (R)	287.0025	J/(kg K)
Sutherland's Law μ_{ref}	1.7084×10^{-5}	N s / m ²
Sutherland's Law T_{ref}	273.15	K
Sutherland's Law T_S	110.0	K

Table 4 Flow conditions selected for FSI simulations corresponding to line stretch in ASPIRE SR03 flight test [7].



(a) View of entire domain



(b) Zoomed-in view of payload and parachute system

Fig. 6 Slice through CFD mesh boxes (dark lines) showing location of static refinement zones. Each box contains 16^3 cells. Mach number is also shown on the slice to justify the placement of the refinement zones. The payload and parachute are displayed in gray. The image is taken from an FSI simulation with $\alpha = 0.9^\circ$, asymmetric pre-inflation, and without radials at a time corresponding to the post-inflation trough where the canopy bow shock moves furthest upstream.

IV. Results

FSI simulations are performed to investigate the effect of the initial flow conditions (taking into account α at line stretch or not), the initial parachute shape (symmetric vs asymmetric), and the presence of radial stiffeners (with the same material properties as the suspension lines). To reduce the number of simulations, we make each choice binary, start from our previous best practices, and only continue to use a given parameter value if it improves the agreement between predicted pull force and measure pull force during ASPIRE SR03. This strategy reduced the overall number of simulations from 8 down to 4. Table 5 displays the breakdown of the parameters chosen for each simulation.

Name	α ($^\circ$)	Pre-inflation	Radial stiffeners
Radials	0.0	Symmetric	Yes
Radials + AoA	0.9	Symmetric	Yes
Radials + AoA + Asym	0.9	Asymmetric	Yes
AoA + Asym	0.9	Asymmetric	No

Table 5 FSI simulation parameters.

The choice of angle of attack α is driven by the measurements of the ASPIRE payload at the instance of line stretch [7]. The radial stiffeners are assumed to have the same properties as the suspension lines in such a way that they connect from the riser, go through the band and disk, to the center of the vent hole and back through the canopy down to the riser in 40 unique threads. This was a reasonable starting point in the absence of details regarding the material properties of the gore seams which make up the radial stiffeners.

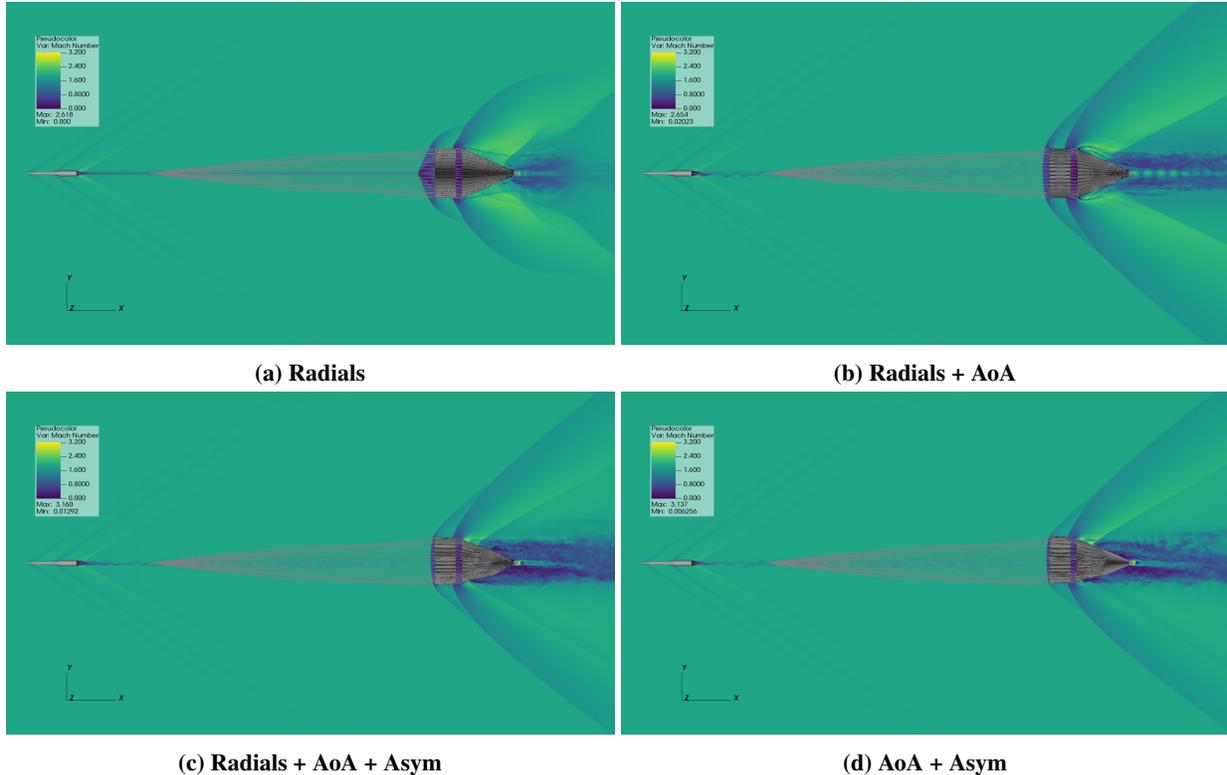


Fig. 7 Slice showing Mach number and parachute system (gray) at the start of the four FSI simulations.

A. Effect of Angle of Attack

The simulation at $\alpha = 0^\circ$ produces a stronger wake deficit at the parachute plane with virtually no turbulence. This is a striking departure from SR01 simulations, which prominently featured a turbulent wake but relatively little mean

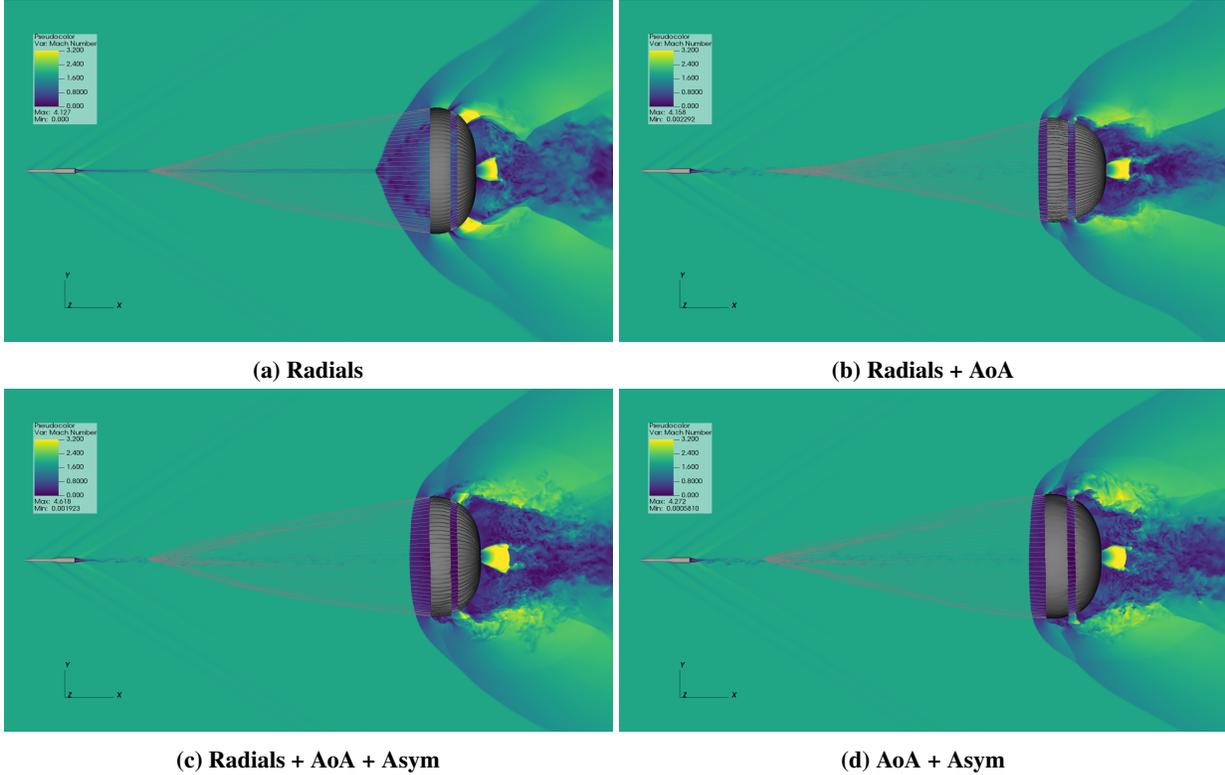


Fig. 8 Slice showing Mach number and parachute system (gray) at the inflation peak of each of the four FSI simulations.

velocity deficit. We expect this is due to the higher dynamic pressure, the relatively coarse mesh, the absence of any freestream turbulence or payload boundary layer turbulence, and, to a lesser extent, the coarser but more symmetric payload mesh. However, we do not believe this numerical solution to be physically representative: wakes at such high Reynolds numbers are unstable due to the inflection point in the mean velocity profile following the aft end of the payload and should quickly breakdown to turbulence. This stronger coherent (and nearly laminar) wake has a much stronger impact on the parachute bow shock, even at the starting point of the FSI simulation – see Figure 7a. Note that this simulation followed more closely the initialization procedure of As’ad et al. [9] where step 3 in III (rigid and static parachute added to CFD-only pre-FSI simulation) was performed for only 0.05 seconds, which explains why the shocks have not reached their conical shape all the way to the edge of the domain. Despite this difference in CFD-only time integration over the rigid and static parachute, the pressure distribution on the parachute is nearly identical and provides aerodynamic loads within less than 5% of the statistically stationary loads, leaving inflation dynamics mostly unaffected. The wake deficit interacts with the bow shock in Figure 7a, causes a dip in the shock front, and creates a well-defined quasi two-dimensional vortex structure that sheds into the canopy and can prevent flow from going supersonic through the disk-band gap in Figure 8a. This interaction also causes the bow shock to move further upstream from the parachute, reducing the pressurization effect on the interior of the canopy shown in Figure 9a. This is reflected in the pull force experienced at the load pins and is especially salient in the deep post-inflation trough as shown in Figure 10. Also note the presence of high frequency oscillations in the pull force that were not observed in previous SR01 simulations.

The simple and rather minor change from $\alpha = 0^\circ$ to $\alpha = 0.9^\circ$ causes the wake to break down to turbulence much faster and reduces the wake deficit at the plane of the parachute to levels comparable to that of previous simulations of SR01. Figures 7b, 8b, and 9b demonstrate the impact of the change across various important instances in time: the bow shock stays rounded, the quasi two-dimensional vortex structure is replaced by axially-flattened turbulence, and the bow shock stand-off distance is reduced during the post-inflation trough and rebound. With the added angle of attack, the FSI predicts a qualitative trend in pull force in better agreement with the ASPIRE measurements if we smooth over the large relatively high frequency oscillations. These oscillations are particularly concerning in the initial rise to the inflation peak. The suspension lines must be experiencing synchronized tension waves during inflation, which causes

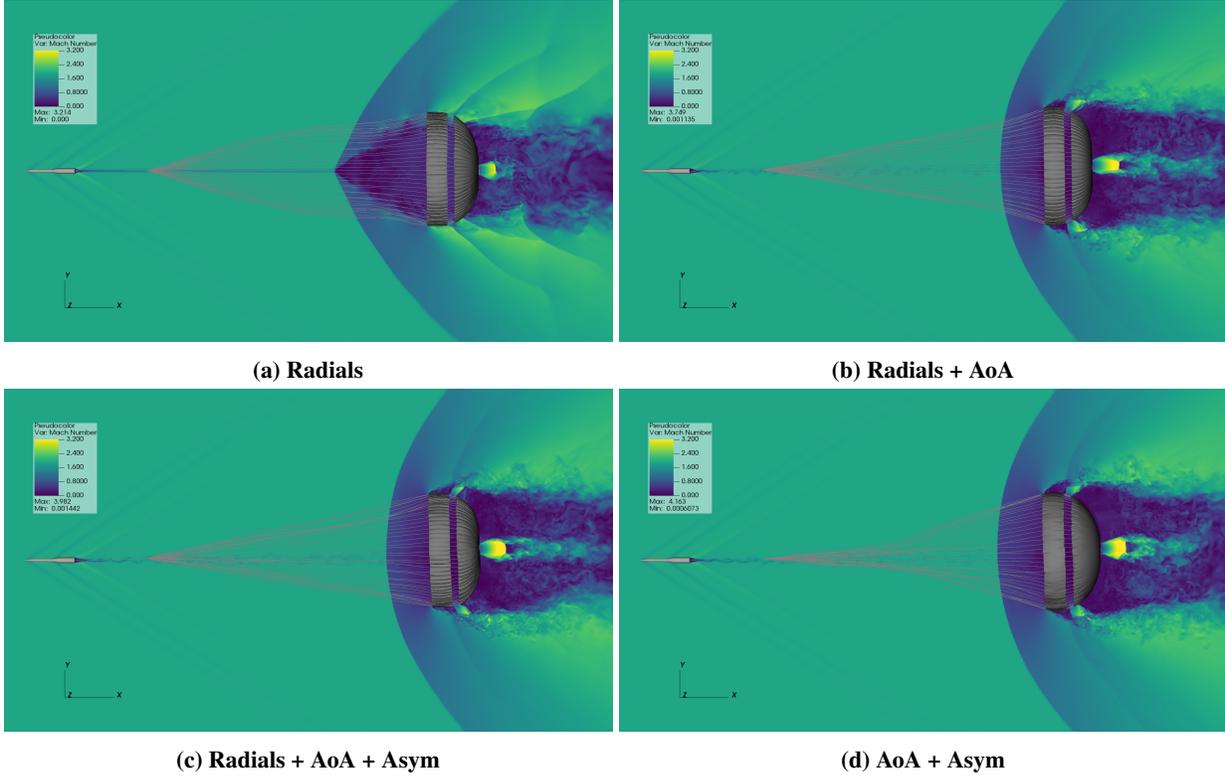


Fig. 9 Slice showing Mach number and parachute system (gray) at the post-inflation trough of each of the four FSI simulations.

constructive interference at the riser and is then transmitted to the load pins. Having seen the impact that breaking the flow symmetry had on the pull force in Figure 10, it seems probable that breaking the near-perfect symmetry of the pre-inflated parachute might have an outside effect on the pull force at the load pins by reducing the synchronization between suspension lines.

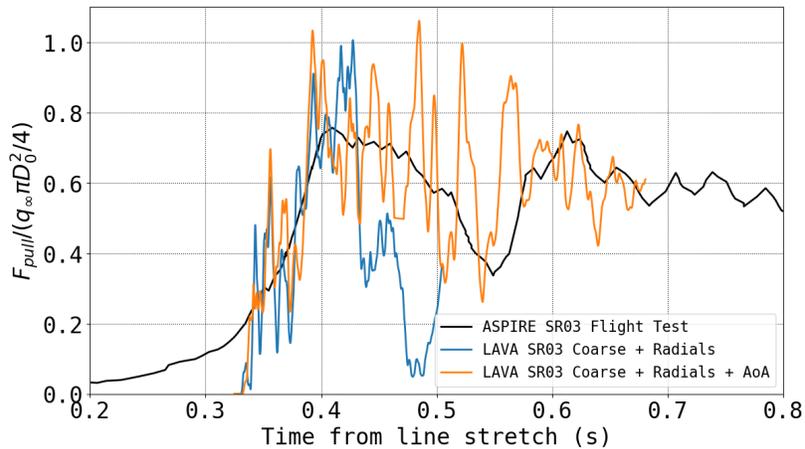


Fig. 10 Non-dimensional pull force comparison between ASPIRE measurements and LAVA FSI simulations with $\alpha = 0^\circ$ and $\alpha = 0.9^\circ$.

B. Effect of Asymmetric Pre-inflation

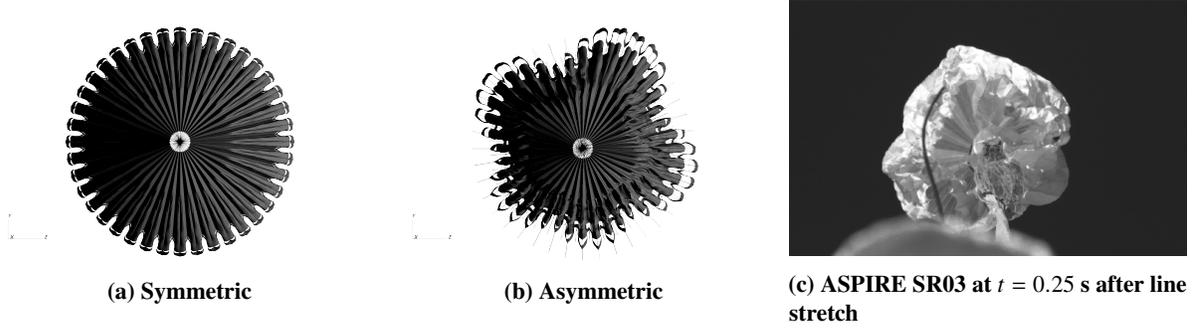


Fig. 11 View from below of the pre-inflated parachute canopy compared to ASPIRE SR03 shortly after line stretch [7].

During the pre-inflation process, instead of applying a constant force normal to the parachute surface, we instead apply a combination of the afore-mentioned normal force with a circumferentially and axially varying force. The intent is to break the symmetry of the initial parachute shape and move toward a lobed shape closer to what is seen in photographs from ASPIRE tests near the line stretch event. The new force F_{asym} is defined as follows:

$$F_{asym}(x, y, z) = (|\sin(2.5\theta + \pi/6) + 0.5 \sin(1.5\theta + \pi/4)| \cos(2\pi(x + 3.5)/40) - 1) \begin{bmatrix} 0, y/\sqrt{y^2 + z^2}, z/\sqrt{y^2 + z^2} \end{bmatrix}, \quad (1)$$

where $\theta = \arctan(z/y)$. And it is combined with the normal force F_{normal} to yield a total force:

$$F_{total} = F_{mag}(0.1F_{normal} + 10F_{asym}) \quad (2)$$

where F_{mag} is a user-specified parameter set to 1000 (Newtons), and the origin is at the centerline of the payload in y and z , and the nose of the payload is located at $x = -55.612$ m, which makes the $x = 0$ plane land at the leading edge of the folded disk of the canopy. This yields three distinct lobes and a depressed region as shown in Figure 11. It also does not increase the overall cross-sectional area of the parachute because the force is negative in many regions and pushed the parachute inward. This change in pre-inflation procedure thus has two distinct effects: it slightly reduces the cross-sectional area of the initial parachute shape, and it introduces significant asymmetry in the band and disk which in turn produces a circumferentially-varying amount of slack in the suspension lines.

Figures 7c, 8c, and 9c demonstrate how this asymmetric canopy shape affects the FSI simulation. At line stretch in Figure 7c, the differences are subtle compared to Figure 7b. The bow shock stand-off distance is smaller, the wake narrower, and the re-compression shocks occur further upstream. The peak tension event is delayed slightly and thus occurs when the parachute has opened more fully and pushed the bow shock further upstream – see Figure 8c. The effect of asymmetric pre-inflation continues to be prominent even at the post-inflation trough (as in Figure 9c) where the parachute has moved laterally and angled itself further to the left – likely due to the larger area open to the flow in the upper y half-plane than the lower y half-plane at the onset. The effect of these changes on the pull force are to effectively reduce the number and amplitude of high frequency oscillations. Specifically, the first peak near $t = 0.36$ s nearly disappears, and the large inflation double-peak at $t = 0.38$ s and $t = 0.4$ s are severely damped. The slew of peaks and troughs between 0.45 and 0.6 seconds are also damped. This effect confirms the previous explanation that synchronized tension waves were causing these high-amplitude, high-frequency oscillations because introducing asymmetry in the canopy shape (and amount of slack in the lines) was enough to significantly reduce the recorded amplitudes. It is also clear from the comparison to the ASPIRE SR03 recorded pull force that this new simulation better follows the qualitative trends and is in better quantitative agreement with the flight test data. This result is expected because we know the ASPIRE SR03 parachute shape was not perfectly symmetric shortly after the line stretch event as seen in Figure 11.

C. Effect of Removing Radial Stiffeners

We investigate the sensitivity of pull force to the assumption that we can model the radials as uninterrupted suspension lines. A new structural mesh is generated, identical in every aspect to the previous one with the exception

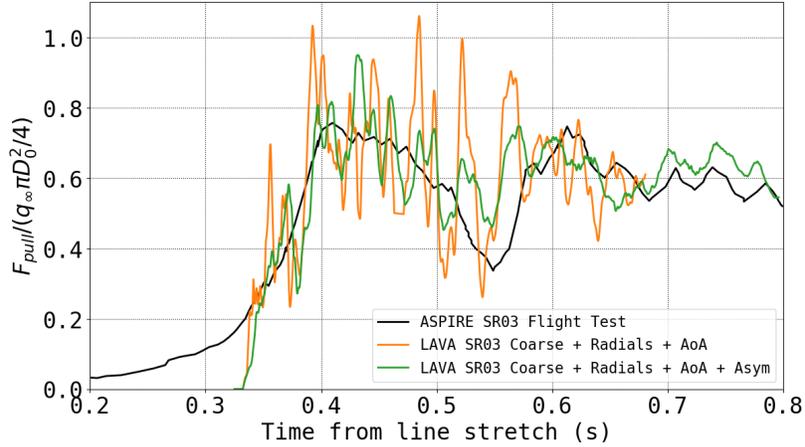


Fig. 12 Non-dimensional pull force comparison between ASPIRE measurements and LAVA FSI simulations with and without asymmetric pre-inflation.

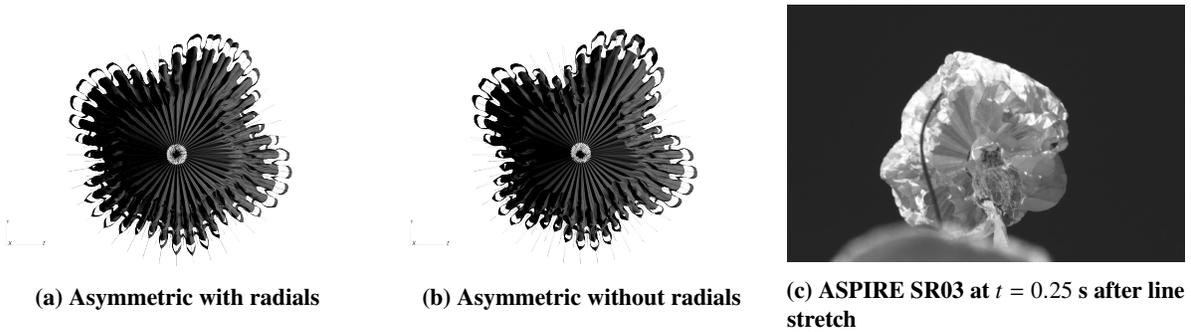


Fig. 13 View from below of the pre-inflated parachute canopy compared to ASPIRE SR03 shortly after line stretch [7].

that the beam elements spanning through the band and disk are removed. The new mesh is then pre-inflated using the same asymmetric force described in the previous section, but this time without radial stiffeners, yielding a slightly different shape. The two pre-inflated canopies are shown side-by-side for comparison in Figure 13. As expected, having reduced the overall stiffness of the canopy but kept the pre-inflation force and time interval constant, we see slightly more displacements without radials: more pronounced lobe structures and a deeper depression in the lower left quadrant. Although the pre-inflated shapes show slight differences, the overall structure and cross-sectional area open to the flow are nearly identical. It is thus of little surprise that the flow at line stretch is very similar. Due to the stronger depression between lobes near the $z = 0$ cut on the $+y$ side, we see earlier and stronger flow acceleration on the left side of Figure 7d at the start of the FSI (line stretch) compared to previous cases. Figure 8c clearly displays the effect of the presence of the radials on the band and disk shape when compared to the simulation in which they are removed in Figure 8d. The radials effectively restrain the parachute fabric from fully extending radially on the band and axially on the disk by providing added stiffness. Absent this added stiffness, the broadcloth can stretch more. The suspension lines show the most evidence of slack out of all simulations at the post-inflation trough stage – see Figure 9d. These differences translate into the non-dimensional pull-force felt at the load pins. Figure 14 displays the effect of removing the radial stiffeners (modeled as suspension lines) by comparing two simulations (with and without) to the ASPIRE measurements. Removing the radial stiffeners further damps the amplitude of previously observed high frequency oscillations. This is explained by the fact that the suspension lines are now interrupted by different materials with different properties (20 times less stiff). Due to the large difference in stiffness between the suspension lines and the canopy broadcloth, a portion of the energy in the tension waves would be transmitted (albeit at a different propagation velocity), and the rest reflected back. Despite this, the inflation peak is nearly identical between the simulations with and without radials, both over-predicting the ASPIRE SR03 pull force measurement by approximately 20%. The post-inflation trough value is

over-predicted by 13% and occurs earlier in time compared to the flight test. The drop from peak to trough measured in the SR03 flight was approximately 52%, and the simulation closely matches this drop at 55%. The rebound peak occurs earlier in time without radials (at $t = 0.55$ s) compared to with radials (at $t = 0.575$ s), and to the flight test data (at $t = 0.6$ s). However, the rebound pull force value is in excellent agreement with that measured in flight. Similarly, values in the post-rebound phase from $t = 0.6$ to $t = 0.8$ s also show good agreement with flight data, and largely fall within the flight-to-flight variability observed in Figure 1. We are not considering the deceleration phase at this time, since the current simulations do not model the deceleration of the system, and as such, the non-dimensional pull force will eventually reach a statistically stationary state (oscillations about a mean value of about 0.6) instead of the characteristic downward slope seen in flight tests.

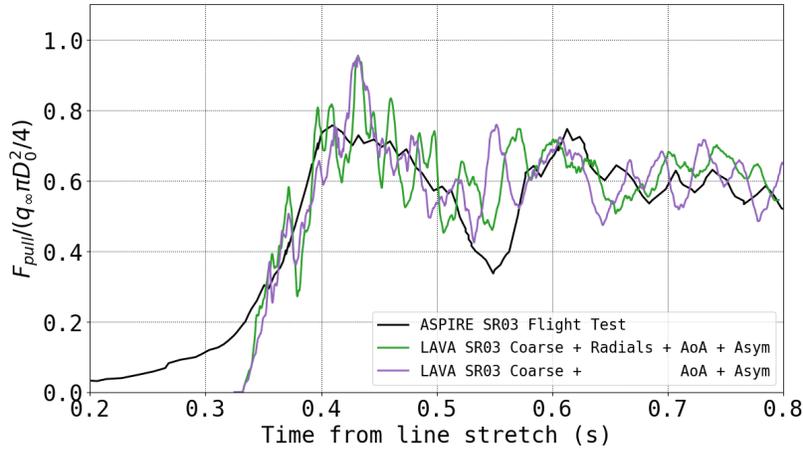


Fig. 14 Non-dimensional pull force comparison between ASPIRE measurements and LAVA FSI simulations with and without radial stiffeners (modeled as suspension lines).

Putting this latest FSI result in the greater context of all the ASPIRE flights, Figure 15 demonstrates that we can obtain improved agreement as we refine the initial conditions and material models to be more representative of the flight tests. The inflation rate quickly matches that of tests after an initial overshoot due to the impulsive start of the FSI simulation (where the fluid forces on the parachute are far from in equilibrium with those in the structural model). The SR03 FSI predicts a higher instantaneous peak than the SR01 simulation, but the LAVA SR03 curve otherwise lies closer to the measured flight test values almost across the board. The two share in common a shorter peak inflation plateau, and an earlier trough and rebound compared to the flight tests. Nevertheless, the results presented here show closer agreement in the peak to trough drop, rebound peak, and post-rebound values with the ASPIRE flight recorded data.

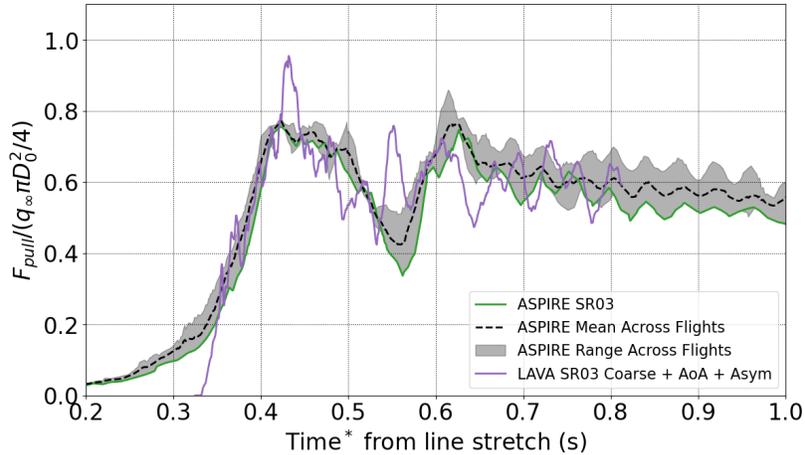


Fig. 15 Comparison of best LAVA SR03 FSI simulation results compared to ASPIRE measurements [6, 7]. * Each time series has been offset so its inflation peak occurs at 0.423 s after line stretch.

V. Conclusions

Fluid-structure interaction simulations of the ASPIRE SR03 parachute flight test were performed to better understand the effects of modeling choices. We demonstrated the importance of breaking symmetry both in the flow initial conditions, and in the initial parachute shape in order to obtain pull-force predictions in better agreement with the flight tests – because the flight tests never experienced the perfect symmetry that can be achieved in a simulation environment, and because these asymmetries were documented in terms of total angle of attack and in pictures of the canopy at line stretch. We also learned that assuming the suspension lines go through the band and disk to mimic the radial stiffeners (gore seams) yields spurious standing waves that produce high frequency and high amplitude oscillations in the pull-force. Choosing more representative initial flow conditions, parachute shape, and removing a simplifying assumption has led to better qualitative and quantitative agreement between FSI predictions and the measured ASPIRE non-dimensional pull force. We plan to refine our assumptions, best practices, and material models to further validate this new FSI capability. Effort is also focused on making these simulations as “turn-key” as possible with a turnaround time of 24 hours or less via algorithmic and end-user-focused enhancements like solution-based adaptive mesh refinement.

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. Special thanks to Marcus A. Lobbia for providing the design, construction and material properties of the SR02/SR03 parachutes and their corresponding inputs for FSI simulations. Thank you to Michael Barnhardt and Aaron Brandis for insightful discussions on the sources of uncertainty in the ASPIRE flight data and for helping guide the focus of our development efforts. The authors also express their gratitude to Jonathan Boustani who developed the geometrically non-linear structural dynamics capability within LAVA and to Gaetan Kenway for his contributions to improving its parallel efficiency.

References

- [1] Gallon, J., Witkowski, A., Clark, I. G., Rivellini, T., and Adams, D. S., “Low density supersonic decelerator parachute decelerator system,” *AIAA Aerodynamic Decelerator Systems (ADS) Conference*, 2013, p. 1329.
- [2] Clark, I. G., Gallon, J. C., and Witkowski, A., “Parachute Decelerator System Performance During the Low Density Supersonic Decelerator Program’s First Supersonic Flight Dynamics Test,” *23rd AIAA Aerodynamic Decelerator Systems Technology Conference*, 2015, p. 2130.
- [3] Clark, I. G., Manning, R., and Adler, M., “Summary of the First High-Altitude, Supersonic Flight Dynamics Test for the Low-Density Supersonic Decelerator Project,” *23rd AIAA Aerodynamic Decelerator Systems Technology Conference*, 2015, p. 2100.

- [4] Boustani, J., Kenway, G., Cadieux, F., Barad, M. F., Brehm, C., and Kiris, C. C., “Fluid-Structure Interaction Simulations of the ASPIRE SR01 Supersonic Parachute,” *AIAA SCITECH 2022 Forum*, 2022, p. 0904.
- [5] Boustani, J., Cadieux, F., Kenway, G. K., Barad, M. F., Kiris, C. C., and Brehm, C., “Fluid-structure interaction simulations of the ASPIRE SR01 supersonic parachute flight test,” *Aerospace Science and Technology*, 2022, p. 107596.
- [6] Dutta, S., Karlgaard, C. D., Tynis, J. A., O’Farrell, C., Sonneveldt, B. S., Queen, E. M., Bowes, A. L., Lylek, E. A., and Ivanov, M. C., “Advanced supersonic parachute inflation research experiment preflight trajectory modeling and postflight reconstruction,” *Journal of Spacecraft and Rockets*, Vol. 57, No. 6, 2020, pp. 1387–1407.
- [7] O’Farrell, C., Sonneveldt, B. S., Karhgaard, C., Tynis, J. A., and Clark, I. G., “Overview of the ASPIRE Project’s Supersonic Flight Tests of a Strengthened DGB Parachute,” *2019 IEEE Aerospace Conference*, IEEE, 2019, pp. 1–18.
- [8] Rabinovitch, J., As’ad, F., Avery, P., Farhat, C., Ataei, N., and Lobbia, M., “Update: Modeling Supersonic Parachute Inflations for Mars Spacecraft,” *26th AIAA Aerodynamic Decelerator Systems Technology Conference*, 2022, p. 2746.
- [9] As’ad, F., Avery, P., Farhat, C., Rabinovitch, J., and Lobbia, M., “Validation of a High-Fidelity Supersonic Parachute Inflation Dynamics Model and Best Practice,” *AIAA SCITECH 2022 Forum*, 2022, p. 0351.
- [10] Cruz, J. R., O’Farrell, C., Hennings, E., and Runnells, P., “Permeability of Two Parachute Fabrics-Measurements, Modeling, and Application,” *24th AIAA Aerodynamic Decelerator Systems Technology Conference*, 2017, p. 3725.
- [11] 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.