Experimentally Anchored Verification and Validation of Density-Based CFD for Supersonic Slender Bodies: Surface Pressure, Drag, and Angle-of-Attack Loads on the ARIA Sounding-Rocket Nosecone

ARIA Aerodynamics: CFD Validation Campaign

June 2026

Abstract

An OpenFOAM-based methodology for predicting surface pressures on supersonic slender bodies of revolution, built on the density-based `rhoCentralFoam` solver with the central-upwind scheme of Kurganov–Tadmor and second-order TVD reconstruction on structured axisymmetric meshes, is verified and validated following the framework of the AIAA V&V guide [1] and ASME V&V 20 [2]. Validation against two independent NACA wind-tunnel datasets at $M_\infty = 1.98$ shows agreement in surface-pressure coefficient of $\overline{|\Delta C_p|} = 0.014$ on a fineness-ratio-3 tangent-ogive cylinder (NACA RM A54H23 [40]) and $\overline{|\Delta C_p|} = 0.004$, within the quoted experimental uncertainty of ± 0.004, on a fineness-ratio-5.75 tangent-ogive cylinder (NACA RM A53E01 [41]) whose slenderness closely matches the application geometry. Exact Taylor–Maccoll conical-flow solutions [31], tangent-cone and shock-expansion methods provide independent theory anchors. Verification comprises a three-grid convergence study with grid-convergence-index (GCI) uncertainty estimation [12], demonstration of iterative steadiness, domain-size (farfield) independence, and sensitivity to limiter, flux variant and thermodynamic energy form. The verified methodology is applied to the ARIA nosecone at $M_\infty = 1.8$ and used to re-assess a prior STAR-CCM+ benchmark: the apparent 23% tip-pressure deficit of OpenFOAM relative to STAR-CCM+ is shown to be the combination of a near-wall sampling artifact and an over-prediction by the faceted STAR-CCM+ geometry. The exact conical-flow tip value for the as-built contour ($p/p_\infty = 1.206$; the audit additionally revealed the nose to be a parabolic-series profile, not the tangent ogive of record) refutes the previously quoted target of 1.860, and the fine-grid apex pressure matches it to +0.2% in p/p_∞. A validated local-time-stepping variant accelerates steady solutions by a factor of $\sim$13 at unchanged accuracy ($\overline{|\Delta C_p|} = 1.7 \times 10^{-4}$), enabling an economical angle-of-attack campaign. A complementary k–ω SST analysis at flight Reynolds number shows skin friction to be $\sim$85% of total drag, and the turbulent angle-of-attack normal force is validated against measured NACA RM A54H23 data to 10% (the residual being predominantly grid discretization, per a three-grid Richardson study). The methodology validated here predicts surface pressure and nose wave drag, while total drag and the design angle-of-attack loads require the viscous extension presented in §8.

Nomenclature

C_A	axial-force coefficient
C_d, C_D	drag coefficient (nose pressure; total)
C_f	skin-friction coefficient
C_N	normal-force coefficient
$C_{N\alpha}$	normal-force-curve slope, $\partial C_N/\partial\alpha$
C_p	pressure coefficient, $(p - p_\infty)/q_\infty$
c_p	specific heat at constant pressure
d	body (base) diameter
F_s	GCI safety factor
h	representative grid spacing, $\sim N^{-1/2}$
K	similarity parameter, $M_\infty/(l_n/d)$
K'	parabolic-series shape parameter
l_n	nose length
L	overall body length
M, M_∞	Mach number; freestream Mach number
N	cell count
p, p_∞	static pressure; freestream static pressure
p_0	total (stagnation) pressure
p_{obs}	observed order of convergence
q_∞	freestream dynamic pressure, $\frac{1}{2}\rho_\infty U_\infty^2$
R	body base radius; gas constant
Re_d, Re_D, Re_L	Reynolds number (diameter; length)
r	radial coordinate; grid refinement ratio
T, T_∞, T_0	temperature; freestream; stagnation
U_∞	freestream velocity
U_{abs}	fine-grid uncertainty in C_p units
x	axial coordinate (from tip)
X_{cp}	center-of-pressure location
y	radial/vertical coordinate
y^+	wall-normal viscous wall coordinate

Greek

α	angle of attack
β	shock-wave angle
γ	ratio of specific heats (1.4)
ΔC_p	C_p difference (CFD − reference)
Δn	wall-normal cell size
Δt	time step
δ_{tip}	tip half-angle
μ	Mach angle; dynamic viscosity
ρ, ρ_∞	density; freestream density
ρ_{og}	ogive arc radius
θ	flow deflection; circumferential angle

Subscripts

∞	freestream condition
nose	nose (forebody) region
cp	center of pressure
tip	body tip
21, 32	fine–medium, medium–coarse grid pair

Acronyms

AoA	angle of attack
CFD	computational fluid dynamics
CFL	Courant–Friedrichs–Lewy number
GCI	grid convergence index
ISA	International Standard Atmosphere
KNP	Kurganov–Noelle–Petrova flux
KT	Kurganov–Tadmor scheme
LTS	local time stepping
MoC	method of characteristics
RANS	Reynolds-averaged Navier–Stokes
SST	shear-stress-transport k–ω model
TA	time-accurate (two-stage)
TVD	total-variation-diminishing
V&V	verification and validation

1 Introduction

Accurate prediction of the surface-pressure distribution over a supersonic slender body of revolution is a foundational requirement for sounding-rocket design. The integral of that distribution over the forebody is the nose wave drag; its pointwise values are the structural pressure loading the airframe must carry; and, once the body flies at incidence, its circumferential asymmetry sets the normal force and the center-of-pressure location that together fix the static-stability margin. A method that predicts surface pressure reliably across the supersonic regime therefore underpins drag, loads and stability estimation simultaneously, which is why surface pressure, rather than an integrated force alone, is taken here as the primary validation quantity.

The application that motivates this study is the nosecone of the ARIA sounding rocket. Its forebody is *not* a tangent ogive, the shape assumed by the prior design campaign, but a parabolic-series profile $r(x) = 0.156\,x - 0.14625\,x^2$ (with x, r in metres; shape parameter $K' = 3/4$), a finding of the consistency audit reported in §3 and one that shifts every tip-theory anchor materially. The nose is 0.4 m long and fairs into a cylindrical afterbody of radius 0.039 m (diameter $d = 0.078$ m, nose fineness $l_n/d = 5.13$) that extends to 1.67 m, an overall fineness $L/d \approx 21$. Differentiating the contour gives a tip half-angle $\delta_{tip} = \arctan(0.156) = 8.87°$ and a small residual slope 2.2° at the ogive–cylinder junction; the body is thus markedly more slender at the tip than the equivalent tangent ogive (11.14°). The vehicle cruises near $M_\infty = 1.8$ at 9 km ISA altitude.

Theoretical background. The aerodynamics of pointed bodies of revolution in supersonic flow has a mature analytical foundation that frames both the expectations and the anchors used in this paper. The flow immediately behind the tip of a cone is governed exactly by the Taylor–Maccoll similarity solution [31], which reduces the conical Euler equations to a single ordinary differential equation and fixes the surface pressure as a function of cone angle and Mach number alone. Away from the tip, where the body is no longer conical, the classical engineering toolkit brackets the surface pressure between two limits: the tangent-cone (or local-cone) approximation, which applies the Taylor–Maccoll solution at the local surface slope, and the shock-expansion family of Eggers, Syvertson and Dennis [33, 32], which propagates the tip-cone state downstream through Prandtl–Meyer expansions. Ehret [38] quantified the accuracy of these approximations against the method of characteristics, the exact inviscid reference computed for bodies of revolution by Ferri [36]. For the integrated loads, Ward's supersonic slender-body theory [37] gives the linear normal-force-curve slope $C_{N\alpha} = 2$ per radian (referenced to base area) that serves as the angle-of-attack sanity check in §7, while Van Dyke's first- and second-order theory [34, 35] supplies the higher-order pressure corrections used in the source experiments. A recent multi-fidelity survey [46] re-examines the accuracy of this analytic hierarchy against Euler CFD for axisymmetric bodies in precisely the Mach band of interest here, bounding the spread to be expected between the theory anchors and a converged Euler solution.

The CFD method and its precedent. The present methodology is built on the open-source density-based solver `rhoCentralFoam` [13], which advances the compressible Euler (and, in the viscous extension, Navier–Stokes) equations with the central-upwind scheme of Kurganov and Tadmor [21, 22]. The choice rests on a now-substantial independent V&V record in exactly this regime: early benchmarking on inviscid supersonic wedges, diamond aerofoils and blunt bodies found the solver superior to the legacy `sonicFoam` in shock-position accuracy [14]; Bondarev and Kuvshinnikov [15] quantified its accuracy against the exact Taylor–Maccoll cone solution and documented the near-apex sensitivity that recurs in §4.2; supersonic cones at incidence [17], under-expanded jets [18], and a formal V&V for supersonic aircraft [16] extend the record across the configurations relevant here, with a recent systematic assessment of the scheme's limiter sensitivity in [19]. Structured OpenFOAM verification and validation against experiment, with grid convergence and code-to-code comparison against a commercial solver, is itself established practice: Ashton and Skaperdas [45] applied exactly this OpenFOAM-versus-STAR-CCM+ discipline to external aircraft aerodynamics. The present work transfers that discipline to supersonic slender bodies, and additionally exercises the local-time-stepping acceleration whose validation for `rhoCentralFoam` was reported by Espinoza et al. [20].

The gap this study fills. Within the specific niche of supersonic rocket nosecones, published CFD studies typically compare candidate shapes or predict drag against semi-empirical tools of the Missile-DATCOM class [47], with limited formal verification; the closest academic parallel, a sounding-rocket nosecone study combining a grid-convergence-index mesh triplet with validation against NACA drag data [48], underlines how rarely an experiment-anchored, formally-verified V&V is carried through for a slender supersonic nose. The immediate motivation was concrete. An initial campaign simulated the ARIA configuration with OpenFOAM and cross-compared against a single commercial STAR-CCM+ solution: agreement on the cylindrical afterbody was within 1.5%, but the codes disagreed by 23.3% near the nose tip. Because both results were numerical solutions of nominally the same problem, the comparison could not by itself establish which (if either) was correct, and it left a 23% uncertainty sitting on the single most drag-relevant region of the body. This paper closes that gap by anchoring the OpenFOAM methodology to data and theory that are external to either code: two independent NACA supersonic wind-tunnel experiments with tabulated surface-pressure coefficients [40, 41], the exact and approximate theory above [31, 32, 34], and a formal verification suite that bounds the numerical error [12, 3]. The pay-off is that the inter-code discrepancy is not merely re-measured but *adjudicated* (§6), and the validated methodology is then carried forward to angle-of-attack loads (§7) and a viscous drag and design-load assessment (§8).

Following accepted V&V practice [1, 2, 4, 5, 3, 7], we distinguish *verification* (are the equations being solved correctly? grid convergence, iterative convergence, domain independence, scheme sensitivity) from *validation* (are the correct equations being solved? comparison with experiment). The paper presents verification before validation, following the logical dependency these guides encode: a validation comparison is meaningful only once numerical error is bounded. Chronologically, the campaign executed the experimental validation first as a programmatic priority and the verification suite second. The paper is organized accordingly: §2 describes the physical model and numerics; §3 the computational setup; §4 the verification studies; §5 the validation against theory and experiment; §6 the application to ARIA, including the re-assessment of the STAR-CCM+ benchmark; §7 the three-dimensional angle-of-attack extension; §8 the viscous (skin-friction and design-load) extension; and §9 concludes.

2 Physical model and numerical method

2.1 Governing equations and flow regime

The flow is modeled by the compressible Euler equations for a calorically perfect gas ($\gamma = 1.4$, $R = 287.05\,\mathrm{J\,kg^{-1}\,K^{-1}}$, $c_p = 1004.7\,\mathrm{J\,kg^{-1}\,K^{-1}}$). The inviscid approximation is appropriate for zero-incidence surface-pressure prediction on slender pointed bodies at supersonic speed: boundary layers are thin and attached, and the experimental evidence in §5 (and in the source reports themselves [40, 41]) confirms that zero-lift pressure distributions are captured by inviscid/potential methods to within experimental uncertainty. The known exception, boundary-layer displacement softening the aft-ogive expansion by $|\Delta C_p| \approx 0.005$–$0.01$, is identified in §5 and motivates the viscous extension (§8).

Freestream conditions correspond to ISA at 9 km: $p_\infty = 30\,742.07\,\mathrm{Pa}$, $T_\infty = 229.65\,\mathrm{K}$, $\rho_\infty = 0.4663\,\mathrm{kg\,m^{-3}}$; $M_\infty = 1.8$ ($U_\infty = 546.83\,\mathrm{m\,s^{-1}}$, $q_\infty = 69\,719\,\mathrm{Pa}$) for the ARIA case and $M_\infty = 1.98$ ($U_\infty = 601.50\,\mathrm{m\,s^{-1}}$) for the validation cases, matching the experiments.

2.2 Solver

`rhoCentralFoam` (OpenFOAM v2406 [50]) is a density-based, explicit, coupled solver: continuity, momentum and *total energy* are advanced simultaneously using the semi-discrete central-upwind scheme of Kurganov–Tadmor [21] (KNP variant [22]) in the collocated polyhedral finite-volume implementation of Greenshields et al. [13]. Second-order accuracy is obtained by TVD-limited reconstruction [24] of $(\rho, \mathbf{u}, T)$; the van Leer limiter [23] is used in production, with Minmod employed for a robustness fallback documented in §4.4. Within the TVD admissible region of Sweby [25], Minmod is the least compressive (most diffusive) limiter and van Leer intermediate. The implementation was validated by its authors against an analytical shock-tube solution, transient supersonic flow over a forward-facing step, interferogram density measurements of a supersonic circular jet, and pressure and heat-transfer measurements on a 25°–55° hypersonic biconic, concluding that the KNP central-upwind scheme is “competitive with the best methods previously published” [13]. Independent precedent now spans the present regime: early benchmarking on inviscid supersonic wedges, diamond airfoils and blunt bodies to $M_\infty = 3.5$ found `rhoCentralFoam` superior to `sonicFoam` in shock-position accuracy [14]; Bondarev and Kuvshinnikov [15] quantified its accuracy against the exact Taylor–Maccoll cone solution (observing the same near-apex sensitivity documented in §4.2); supersonic cones at incidence [17], under-expanded supersonic jets [18], and a formal independent V&V for supersonic aircraft aerodynamics [16] complete the picture, with a recent systematic assessment of the KNP scheme's limiter sensitivity given in [19].

Equivalence to the prior STAR-CCM+ benchmark formulation. The prior benchmark reportedly used the STAR-CCM+ “Coupled Energy” model, prompting a question whether a formulation difference could explain inter-code discrepancies. Coupled Energy is the energy component of the STAR-CCM+ density-based *Coupled Flow* solver: continuity, momentum and energy solved simultaneously (Weiss–Smith-preconditioned Roe-type flux [49], with 1st/2nd-order upwind discretization options, the “order of the energy solver”). `rhoCentralFoam` belongs to the same class: the total-energy equation is advanced in the same coupled update at second order. There is therefore no formulation gap; the empirical energy-form sensitivity test in §4.4 closes the question quantitatively: switching the conserved thermal variable changes wall pressures by $\overline{|\Delta C_p|} = 6 \times 10^{-4}$ (§4.4).

2.3 Two-stage solution strategy and local time stepping

Steady states are obtained by time marching in two stages: a first-order (upwind-reconstruction) stage to damp impulsive-start transients, followed by a second-order TVD stage to convergence (CFL ≤ 0.5, adaptive Δt). For the time-accurate runs the stages span $t \in [0, 0.025]$ and $[0.025, 0.05]$ s (≈4 domain flow-through times).

Because only the steady state is of interest, a *local time stepping* (LTS) variant (`localEuler` pseudo-time integration) was additionally validated: each cell advances at its own stable time step, removing the global Δt restriction imposed by the smallest cell, the classical steady-state accelerator for density-based solvers since Jameson et al. [29] (see also [30]). On the coarse-grid ARIA case the LTS steady state reproduces the time-accurate result to $\overline{|\Delta C_p|} = 1.7 \times 10^{-4}$ ($\max |\Delta C_p| = 1.2 \times 10^{-3}$) with identical C_d to four decimals, at a 12.7× reduction in CPU time (58 s vs. 734 s for the second-order stage); Espinoza et al. [20] report comparable-accuracy LTS speed-ups for `rhoCentralFoam` of 2.56× (compressible Couette flow with heat transfer) and 8.96× (supersonic flat plate); the higher factor obtained here reflects the wider cell-size disparity of the collapsed-axis wedge mesh, whose time-step bottleneck is diagnosed in §3.2: the mean cell Courant number in time-accurate operation is only 0.007 against the limiting cell's 0.5. LTS is therefore adopted for all subsequent steady production runs (in particular the three-dimensional angle-of-attack matrix), while every result labeled “validation” in this paper was produced with the time-accurate scheme for strict methodological identity with the ARIA baseline.

3 Computational setup

3.1 Geometries

Three bodies of revolution are simulated (Table 1): the two NACA validation models (circular-arc tangent ogives, generated analytically from their published dimensions) and the ARIA nosecone, whose contour is prescribed by design coordinates. A finding of this campaign's consistency audit deserves emphasis: the ARIA design coordinates are *not* a tangent ogive (as the prior campaign and early phases of this one assumed) but are exactly the parabolic-series nose $r = 0.156\,x - 0.14625\,x^2$

($K' = 3/4$; least-squares residual $\sim 10^{-17}$ m against the 21 design points), with tip half-angle $\delta_{tip} = 8.87°$ (not the 11.14° of the equivalent tangent ogive) and a small non-zero slope (2.2°) at the cylinder junction. All theory anchors in this paper use the as-built parabolic contour; where the idealized tangent-ogive value is given for comparison it is marked explicitly. The validation pair brackets the ARIA application in nose fineness ($3 < 5.13 < 5.75$); in tip half-angle ARIA is the most slender of the three (8.9° vs. 9.9°/18.9°), making the fineness-5.75 body the closest analog. The application differs from the validation pair in two controlled respects, a 9% Mach offset (1.8 vs. 1.98) and nose shape (parabolic series vs. tangent ogive), both *bounded* rather than assumed away: the theory anchors of §5.1 (tangent-cone, shock-expansion) are evaluated on each body's own contour and at its own Mach number, so shape and Mach enter the comparison explicitly, while the experiment establishes that the methodology reproduces measured pressures on geometries bracketing the application in fineness.

The two validation bodies are not arbitrary: each is the subject of a classic Ames supersonic-tunnel investigation whose surface-pressure tables are reproduced in full and whose geometry is documented to the orifice. The fineness-3 body is the ogive-cylinder of NACA RM A54H23 [40], a 3 calibre circular-arc tangent ogive (arc radius $9.25d$) on a $7d$ cylinder, instrumented at 15 axial stations and tested over $\alpha = 0$–$20°$ at two Reynolds numbers ($Re_d = 0.15$ and 0.45×10^6) expressly to expose Reynolds-number and incidence effects on the normal-force distribution. Its blunt 18.9° tip makes it the *hard* case for an inviscid method: the tip shock is strong and the near-apex gradients steep, so it stresses the tip-resolution and wall-sampling treatments of §3.2–3.3 hardest. The fineness-5.75 body is "model 2" of NACA RM A53E01 [41], a $33\frac{1}{3}$-calibre tangent ogive instrumented over both nose and afterbody to $x/d = 14$ (29 stations) as part of a study of lift on a slender inclined body of revolution at $M = 2$. Its 9.9° tip is within a degree of the ARIA tip, so it is the *closest analog*: if the methodology reproduces this body's pressures within experimental uncertainty, it does so on the geometry that most resembles the application. The two thus bracket ARIA in fineness from below and above ($3 < 5.13 < 5.75$) while jointly spanning the range of tip bluntness over which an inviscid surface-pressure prediction is plausible. The experimental conditions, table references and quoted uncertainties are collected in §5.2, where the comparisons are made.

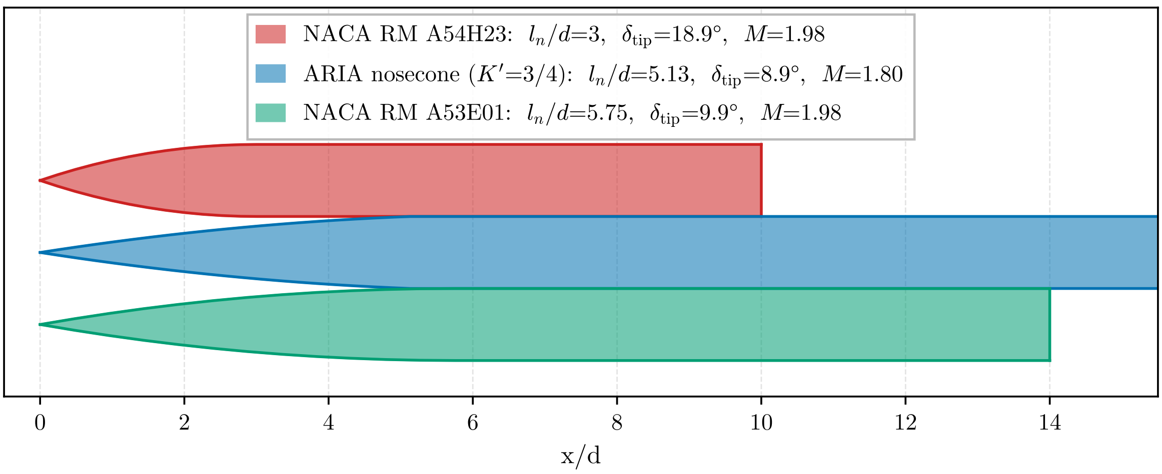

Figure 1: The three simulated bodies of revolution drawn to common scale in calibres (x/d): the fineness-3 NACA RM A54H23 ogive-cylinder (top), the ARIA parabolic-series nosecone (middle), and the fineness-5.75 NACA RM A53E01 body (bottom). The validation pair brackets the ARIA application in nose fineness ($3 < 5.13 < 5.75$); ARIA's parabolic-series tip (8.87°) is marginally the most slender of the three, making the fineness-5.75 body its closest analog. Geometry parameters are tabulated in Table 1; the contours support the bracketing argument of §3.

Table 1: Simulated geometries (Fig. 1). δ_{tip} = tip half-angle; ρ_{og} = ogive arc radius (not applicable to the parabolic-series ARIA contour). The A53E01 afterbody is the simulated extent to the last pressure orifice ($x/d = 14$, $8.25d$); the physical NACA model 2 is $14.2d$ long ($8.45d$ afterbody), the additional $0.2d$ lying downstream of all comparison stations.

	nose l_n/d	ρ_{og}/d	δ_{tip}	afterbody	M_∞
NACA RM A54H23 model [40]	3.0	9.25	18.9°	$7d$	1.98
NACA RM A53E01 model 2 [41]	5.75	33.31	9.9°	$8.25d$	1.98
ARIA nosecone (parabolic series, K'=3/4)	5.13	n/a	8.9°	$16.3d$	1.80

3.2 Mesh and boundary conditions

The axisymmetric domain is a 5° wedge (one cell thick circumferentially), $x \in [-2, 5]$ m, $r \leq 5$ m, meshed with a 9-block structured hexahedral topology generated parametrically (`scripts/generate_mesh.py`); the ogive contour enters as a `polyLine` edge. The ARIA baseline grid has 11,800 cells with radial clustering at the body surface. The centerline is represented with an `empty`-type patch (the standard `wedge`/axis combination is numerically fragile with density-based solvers). The collapsed wedge cells adjacent to the axis behind the base are the global time-step limiter in time-accurate operation (cell width ≈ 0.17 mm in the collapse direction), motivating the LTS variant of §2.3.

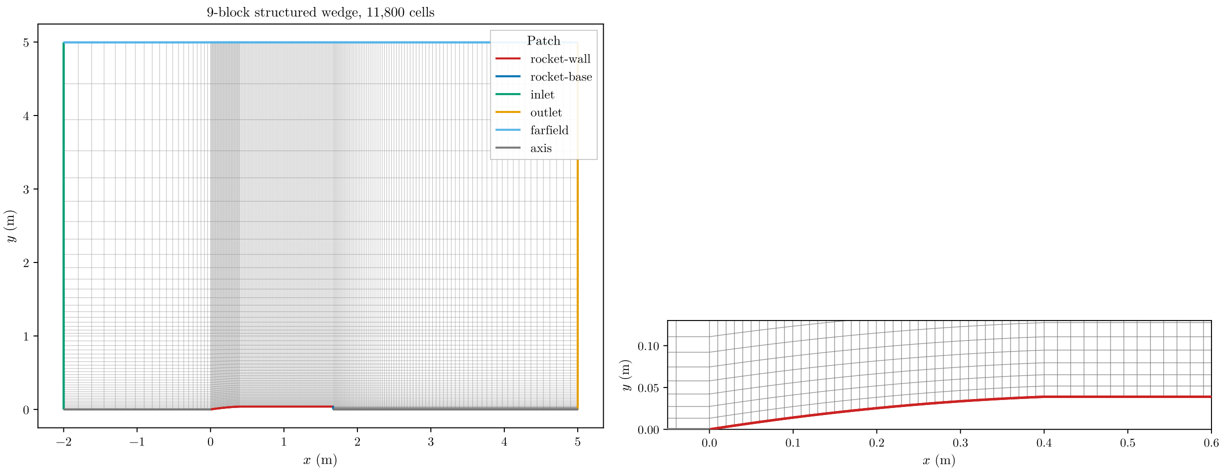

Figure 2: The 9-block structured axisymmetric wedge mesh for the ARIA baseline (11,800 cells). Left: the full $x \in [-2, 5]$ m, $r \leq 5$ m domain, with patch boundaries coloured (body and base, inlet, outlet, lateral farfield, axis) and the interior block topology and radial clustering toward the axis visible. Right: the near-body detail over the parabolic-series nose and ogive–cylinder shoulder, showing the near-wall radial refinement (first cell ≈ 0.9 mm) that resolves the boundary region while the cells coarsen outward. The collapsed wedge cells behind the base set the global time-step limit that motivates the LTS variant (§2.3).

Boundary conditions: supersonic inflow with all quantities fixed at freestream (exact for $M_\infty > 1$); zero-gradient outflow (supersonic except a thin subsonic base-wake channel); fixed freestream on the lateral boundary, which is exact here because the attached tip shock exits through the *outlet* (§4.3); slip walls on the body.

Validation-case meshes resolve the conical tip shock layer: first wall cell 0.9 mm to 1.2 mm (radial expansion ratio 30) and 3.4 mm axial at the apex, giving 26,490 cells (RM A54H23 model) and 29,025 cells (RM A53E01 model). The necessity of this resolution is demonstrated in §5.3 by a deliberately under-resolved counter-example.

3.3 Wall-pressure evaluation

How the wall pressure is read off the solution turns out to be the single most consequential post-processing choice in this study, because it is the mechanism behind the original "tip discrepancy"

(§6.2). The reasoning proceeds in four steps.

(i) Where the "wall" pressure is actually sampled. A finite-volume solver stores p at cell centres, and the slip wall carries a zero-gradient ($\partial p/\partial n = 0$ in the numerical sense) boundary condition, so the value reported on the wall patch is simply the value at the centre of the first cell, which sits a distance $\Delta n/2$ *off* the surface, where Δn is the wall-normal size of that first cell. On a smooth afterbody, where the pressure barely changes over $\Delta n/2$, this offset is harmless.

(ii) Why the offset biases the tip. Near a sharp apex the surface pressure falls steeply with distance from the wall through the thin conical shock layer, so reading it $\Delta n/2$ too far out returns a value biased by approximately $(\Delta n/2)\,\partial p/\partial n$. The bias scales directly with the first-cell size. Concretely, on the ARIA baseline grid the first cell centre near the tip lies about 7 mm off the wall, and the cell-centred tip pressure reads low by $\Delta C_p \approx 0.02$ relative to the wall value; refining the first row to ≈ 1 mm (the tip-refined grid of §6.2) shrinks the same gap to $\Delta C_p \approx 0.009$. The "deficit" is therefore a sampling artifact that vanishes as the first row approaches the wall, not a deficiency of the solution.

(iii) The wall-extrapolated alternative and its own caveat. To recover the on-wall value we additionally fit a straight line through the first two cell-centre values and extrapolate it to the surface. This two-row linear reconstruction (in the cell-centred-gradient sense of [28]) removes the leading-order offset and carries a formal error $O(\Delta n^2)$. It has one failure mode of its own: within about two cells of a sharp apex the true pressure profile is strongly *curved*, so a straight-line extrapolation overshoots the wall value (quantified at the fineness-3 body's first station in §5.3). The two evaluations therefore bracket the truth near the apex, cell-centred from below and wall-extrapolated from above, and are reported as a pair throughout.

(iv) An independent drag cross-check. Integrated nose pressure drag is computed two independent ways so that neither the surface integration nor the solver is taken on trust: an exact axial projection over the wall samples, $C_d = \int (p - p_\infty)\,\mathrm{d}(\pi r^2)/(q_\infty \pi R^2)$, and the solver's own `forces` function object with $p_{\mathrm{ref}} = p_\infty$. On the GCI grids the two agree to 0.06% (coarse: 0.03595 vs. 0.03597) and 0.2% (fine: 0.03333 vs. 0.03326), confirming that the hand-rolled surface integration is faithful to the solver's internal force balance. This cross-check also exposed a stale legacy result: the prior campaign's headline $C_d \approx 0.0522$ could not be reproduced, as re-running the archived post-processing on the archived solution yields $C_d = 0.0344$. The archived integrator further relied on a bin-nearest interior-cell search for "wall" pressures and an approximate $\sin\theta\,\Delta x$ surface projection in place of the exact $\mathrm{d}(\pi r^2)$ element; both are superseded here, and the legacy value is excluded from further comparison.

4 Verification

The suite follows the standard solution-verification ladder [5, 8, 16]: iterative convergence, grid convergence, domain-size independence, and discretization-option sensitivity, with the LTS-acceleration equivalence of §2.3 closing the chain.

4.1 Iterative steadiness

Iterative convergence is demonstrated by the wall-force histories (Fig. 3): after the impulsive-start transient decays within ∼1 ms and a small step at the second-order scheme switch, the drag histories are flat. Quantitatively, the secular drift of C_d (difference of consecutive 5 ms-window means over the final 10 ms) is 0.006 % (A54H23 case), 0.008 % (A53E01), 0.0005 % (ARIA coarse) and 0.006 % (ARIA fine, 47,200 cells), one to two orders below the 0.1 % acceptance criterion; the residual peak-to-peak ripple of 0.05–1% is zero-mean acoustic noise, not drift. The corresponding change of the entire wall-C_p distribution between the final two solution writes (5 ms apart) is $\overline{|\Delta C_p|} = 9 \times 10^{-4}$ (A54H23) and 3×10^{-4} (A53E01), a factor 4–13 below experimental uncertainty.

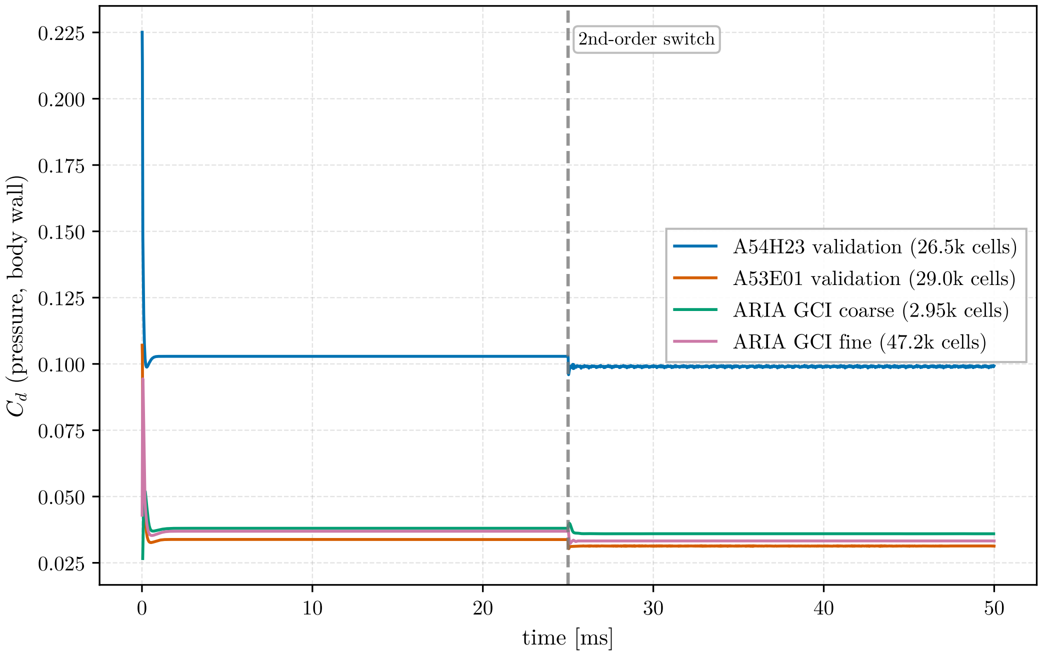

Figure 3: Iterative steadiness: pressure-drag histories from the solver's `forces` function objects for the validation and ARIA cases. The impulsive-start transient decays within ~ 1 ms; the dashed line marks the switch from the first-order to the second-order scheme, after which every history is flat. The secular drift of C_d over the final 10 ms is 0.0005–0.008%, one to two orders below the 0.1% acceptance criterion, and the residual ripple is zero-mean acoustic noise rather than drift, establishing that the solutions are converged in pseudo-time before any comparison is made.

4.2 Grid convergence (GCI)

The grid study uses three meshes built from the same block topology and the same grading laws, each refined by an exact factor of two in every direction so the cell counts stand in the ratio $4:1$ from one mesh to the next: 2,950 (coarse), 11,800 (medium) and 47,200 (fine) cells, a refinement ratio $r = 2$. Discretization uncertainty is then reported with the grid-convergence index (GCI), the de-facto standard introduced by Roache [6] and used here in the three-grid form of Celik et al. [12]. The GCI converts the difference between two grids into an error-bar-like uncertainty on the finest one, and the procedure is the same for every quantity. From the three solutions f_3, f_2, f_1 (coarse to fine) one forms the successive changes $e_{32} = f_2 - f_3$ and $e_{21} = f_1 - f_2$ and their ratio. If that ratio shows the changes *shrinking* monotonically as the mesh refines, an *observed order* of accuracy p is extracted from it, the grid-converged value is estimated by Richardson extrapolation [9], and the fine-grid uncertainty is $\mathrm{GCI} = 1.25\,|e_{21}|/(r^p - 1)$. Where the changes instead *oscillate* in sign, or where the implied order falls outside the asymptotic range in which Richardson extrapolation is meaningful, the extrapolation is abandoned and a deliberately conservative band is quoted in its place, following the original safety-factor convention of Roache: $F_s = 3$ with the order held at its theoretical value $p = 2$. The observed-versus-formal-order behaviour, in particular the local drop toward first order at captured shocks, is exactly that reviewed by Roy [10]. Three classes of quantity are tracked: the integrated nose pressure-drag coefficient, the point C_p at four representative stations ($x = 0.05/0.15/0.35/1.0$ m), and segmented L_2 norms of the wall-C_p distribution over the tip, shoulder and cylinder, so that the study reports not only *how much* numerical error remains but *where* on the body it lives.

Table 2 summarizes the point quantities. Away from the apex the wall-extrapolated solution is in the asymptotic range: observed orders span 0.9–3.0 (formal order 2 in smooth flow, locally reduced to 1 at captured discontinuities, the expected observed-versus-formal-order behavior reviewed by Roy [10]), and the fine-grid numerical uncertainty is GCI = 1.9–6.5%. The station nearest the apex ($x = 0.05$ m) requires care in interpretation: the wall-extrapolated values rise

$C_p = 0.044 \rightarrow 0.091 \rightarrow 0.104$ (Richardson 0.109), *overshooting* the local tangent-cone value 0.078 on the as-built contour. This station lies inside the apex zone where linear wall extrapolation is biased high (§3.3), so its formal GCI tracks the bias as much as the solution. The clean theory anchor at the apex is the peak pressure itself, which matches the exact conical value to 0.1–0.3% (§6). At $x = 0.35\,\mathrm{m}$ the solution crosses $C_p = 0$, so a percentage GCI is meaningless there; the absolute uncertainty is $U = \pm 0.004$ in C_p, comparable to the near-apex station.

Table 2: GCI point quantities, wall-extrapolated C_p (ARIA, $M_\infty = 1.8$; $r = 2$). U_{abs}: fine-grid uncertainty in C_p units. Theory: local tangent-cone on the as-built contour (†: station inside the apex extrapolation-bias zone, see text).

station	coarse	medium	fine	p_{obs}	Richardson	U_{abs}	theory
$x = 0.05\,\mathrm{m}^\dagger$	0.0442	0.0906	0.1039	1.80	0.1093	± 0.0068	0.0777
$x = 0.15\,\mathrm{m}$	0.0822	0.0605	0.0578	3.04	0.0569	± 0.0011	0.0536
$x = 0.35\,\mathrm{m}$	0.0005	-0.0038	-0.0019	osc.	n/a	± 0.0044	n/a
$x = 1.00\,\mathrm{m}$	-0.0026	-0.0027	-0.0027	0.90	-0.0027	± 0.0001	n/a

The segmented cell-centered wall-C_p L_2 differences expose *where* the error lives: on the cylinder the observed order is 2.6 and the medium–fine difference is 4×10^{-4} in C_p (fully converged); in the apex segment ($x \in [0, 0.1]\,\mathrm{m}$) the observed order collapses to 0.3. This is the expected signature of the apex singularity: conical flow is self-similar in r/x, so the unresolved fraction of the tip shock layer at any fixed grid spacing scales with the spacing itself and the segment never enters the asymptotic range at practical resolutions, the same near-apex sensitivity reported for `rhoCentralFoam` on exact-cone benchmarks by Bondarev and Kuvshinnikov [15]. The integral nose-drag coefficient inherits this behavior (Table 3): successive differences barely contract ($p_{obs} < 0.2$), so Richardson extrapolation is not credible for C_d and a conservative uncertainty is reported instead, following the fallback of [12, 6]: $F_s = 3$ with the order capped at the theoretical $p = 2$, giving $\mathrm{GCI}_{fine} = 3\,|e_{21}/f_1|/(2^2 - 1) = 4.2\%$ (cell-centered) and 8.2% (wall-extrapolated), i.e. $C_{d,nose} = 0.0364 \pm 0.0030$ (wall-extrapolated, fine grid). Because the converged cylinder contributes nothing to the axial projection, this uncertainty is concentrated entirely in the apex region, the same region adjudicated independently by conical-flow theory in §6.2, where the fine grid is shown to close the gap toward the exact value. This reconciles two statements about the apex that might otherwise appear opposed: the integrated truncation error there does not contract at the formal rate, because conical self-similarity leaves a geometrically similar unresolved sliver at every grid spacing, yet the peak surface pressure, fixed by the exact cone boundary condition, is captured to within 0.1–0.3% on the fine grid (§6).

Table 3: Grid convergence of nose pressure drag (ARIA, $M_\infty = 1.8$). Observed order below 0.2 (apex-dominated, outside the asymptotic range): conservative F_s=3, p=2 uncertainty quoted in place of Richardson extrapolation.

$C_{d,nose}$	coarse	medium	fine	p_{obs}	GCI_{fine} (cons.)
cell-centered	0.03595	0.03473	0.03333	0.18	4.2%
wall-extrapolated	0.04224	0.03938	0.03640	0.06	8.2%

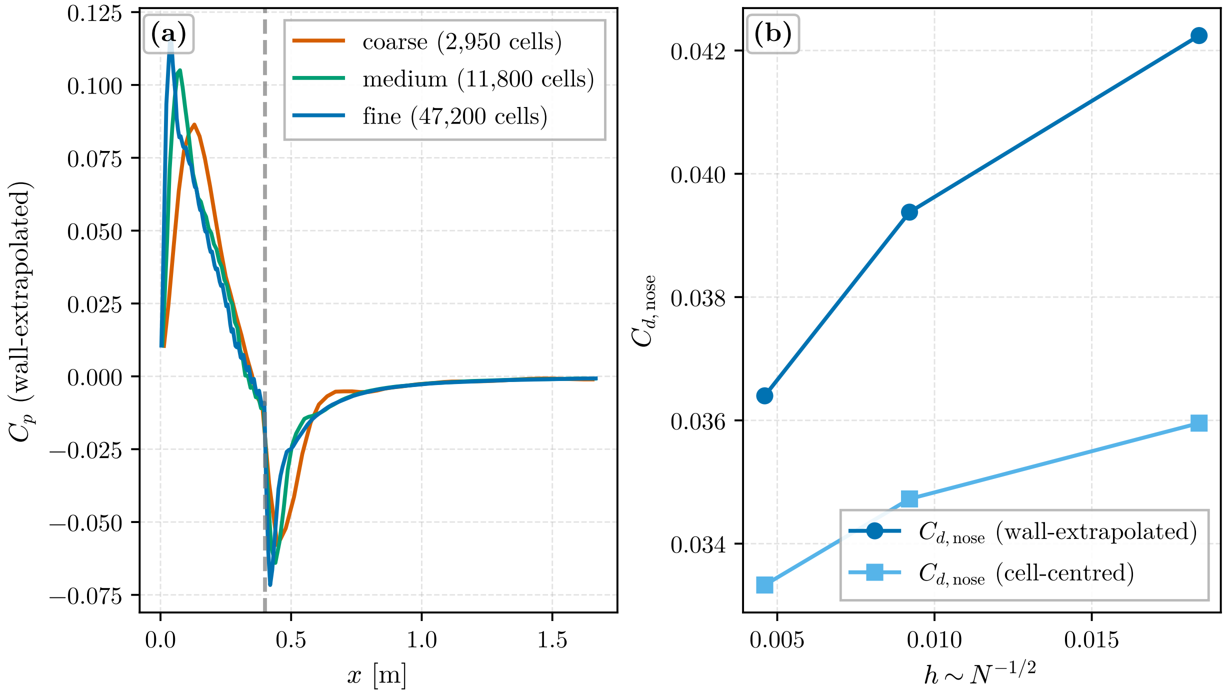

Figure 4: Three-grid convergence for ARIA ($M_\infty = 1.8$; $r = 2$, 2,950/11,800/47,200 cells). Left: the wall-extrapolated C_p distributions on the three grids, which collapse onto one another away from the apex and separate only in the near-tip region where conical self-similarity defeats asymptotic convergence (§4.2). Right: the nose pressure-drag coefficient against representative grid spacing $h \sim N^{-1/2}$, showing the slow, sub-asymptotic contraction that motivates the conservative $F_s = 3$ uncertainty quoted for C_d in Table 3 rather than a Richardson extrapolation.

4.3 Domain-size (farfield) independence

The geometry of the supersonic wave field, rather than any rule of thumb, fixes the expectation here, and it predicts that domain size should not matter at all. In supersonic flow disturbances travel only along downstream-running characteristics, so the only way the lateral boundary could corrupt the solution is if the tip shock reached it while still inside the domain. It does not: at $\delta_{tip} = 8.87°$ and $M_\infty = 1.8$ the tip shock is attached (shock detachment for a cone at this Mach number requires a half-angle far above 20°), and it springs from the apex at the conical shock angle 34.43°, thereafter relaxing toward the freestream Mach angle $\mu = 33.75°$. A straight ray at even the steeper apex angle reaches the lateral boundary $r = 5\,\text{m}$ only at $x \approx 7\,\text{m}$, which is already past the outlet at $x = 5\,\text{m}$; the shock therefore exits through the *outlet*, never crossing the lateral boundary, so the fixed freestream imposed there is exact by construction. Upstream of the body, supersonic inflow forbids any upstream propagation, so the inlet placement is likewise immaterial. The expectation is thus a null result, which the numerical test confirms. The numerical test enlarges the domain to $x \in [-4, 8]\,\text{m}$, $r \leq 10\,\text{m}$ (24,920 cells) at unchanged local resolution. The result is the predicted null (Fig. 5): against the like-for-like baseline the wall pressure changes by $\overline{|\Delta C_p|} = 1.2 \times 10^{-4}$ (max $= 1.3 \times 10^{-3}$, localized at the ogive–cylinder junction), and the nose drag by $\Delta C_d = +1 \times 10^{-5}$ (+0.04%). Domain extent is not a contributor to the inter-code discrepancy, despite the prior benchmark's larger farfield.

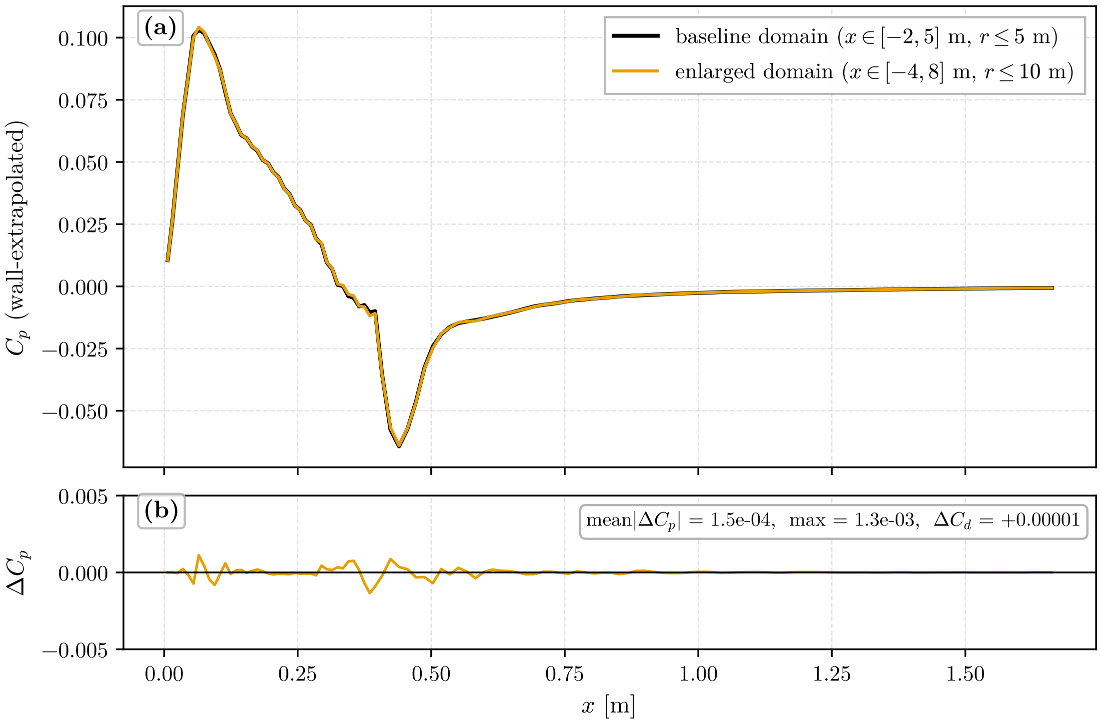

Figure 5: Farfield-extent verification for ARIA ($M_\infty = 1.8$): wall C_p on the baseline domain ($x \in [-2, 5]$ m, $r \leq 5$ m) versus a domain doubled in both directions ($x \in [-4, 8]$ m, $r \leq 10$ m, 24,920 cells) at unchanged near-body resolution. Top: the two wall-pressure distributions are indistinguishable. Bottom: the pointwise difference ΔC_p, plotted on a ± 0.005 axis equal to one experimental-uncertainty band, peaks at only 1.3×10^{-3} at the ogive–cylinder junction. The near-zero result confirms the characteristics argument of the text: the tip shock leaves through the outlet, so domain extent is not a contributor to the inter-code discrepancy.

4.4 Scheme and thermodynamic sensitivity

Numerical-uncertainty bands from solver options, all on the ARIA baseline grid: (i) TVD limiter van Leer vs. Minmod; (ii) central-upwind flux variant Kurganov(–Tadmor) vs. Tadmor (symmetric); (iii) energy form `sensibleInternalEnergy` vs. `sensibleEnthalpy` (empirical closure of the "energy solver" question, §2). All three sensitivities are small (Table 4, Fig. 6): the mean wall-pressure band across solver options is 6×10^{-4} in C_p (a factor $\sim$7 below the experimental uncertainty of the validation datasets), with the maxima concentrated at the apex, and nose-drag shifts of $\pm 2\%$, inside the conservative grid uncertainty of §4.2. The energy-form pair is the empirical closure of the "energy solver" question: -1.0% on C_d, confirming the formulation-equivalence argument of §2.

Table 4: Scheme and thermodynamic sensitivity vs. the like-for-like LTS baseline (van Leer, Kurganov, `sensibleInternalEnergy`; wall-extrapolated quantities).

variant	$\overline{\lvert\Delta C_p\rvert}$	max $\lvert\Delta C_p\rvert$	$\Delta C_{d,nose}$	ΔC_d (%)
Minmod limiter	6.4×10^{-4}	0.0068	+0.00085	+2.1
Tadmor flux	6.4×10^{-4}	0.0066	+0.00076	+1.9
`sensibleEnthalpy`	6.4×10^{-4}	0.0106	−0.00039	−1.0
domain ×2 (§4.3)	1.2×10^{-4}	0.0013	+0.00001	+0.04

The limiter sensitivity has a second, qualitative dimension beyond the small pressure band above: solver *robustness*. On the *under-resolved* fineness-3 mesh at $M_\infty = 1.98$ ($\delta_{tip} = 18.9°$), switching to van Leer reconstruction produced an immediate floating-point failure (a negative reconstructed temperature, which crashes the square-root in the sound-speed evaluation) regardless

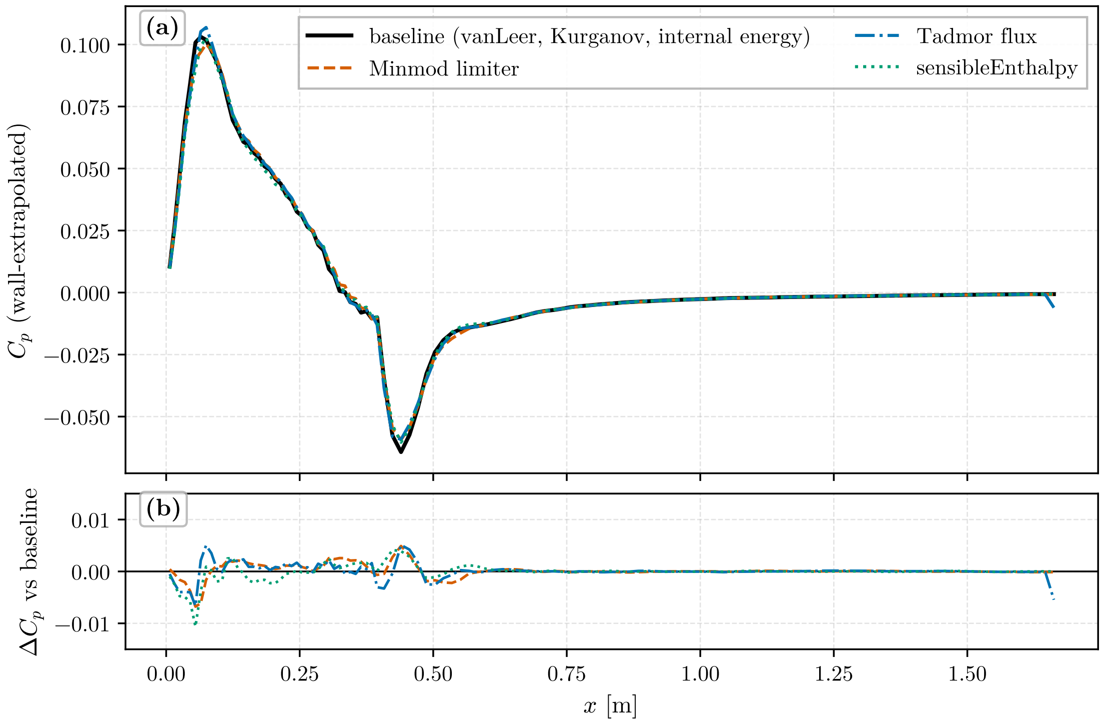

Figure 6: Sensitivity of the ARIA wall pressure to solver options, all on the 11,800-cell baseline grid against the like-for-like LTS reference (van Leer, Kurganov, `sensibleInternalEnergy`). Top: wall-extrapolated C_p for the three variants (Minmod limiter, Tadmor flux, `sensibleEnthalpy` energy form) overlaid on the baseline. Bottom: the pointwise differences from the baseline, which stay within roughly ± 0.005 everywhere except a narrow apex spike. The mean band is 6×10^{-4} in C_p, about 7× below the experimental uncertainty; quantitative values are in Table 4. The energy-form pair settles the "energy solver" question of §2 empirically.

of CFL number, whereas the more dissipative Minmod limiter ran stably to convergence. The same van Leer switch succeeded unchanged once the tip shock was adequately resolved, which localizes the cause: it is a coarse-mesh shock-capture pathology, not a defect of the scheme or the implementation, and adequate tip resolution buys numerical robustness as well as accuracy. The mechanism is well understood in the high-resolution-shock-capturing literature: Quirk [26] catalogued failure modes of high-resolution shock-capturing schemes at strong, under-resolved shocks; Einfeldt et al. [27] showed that schemes built on linearized intermediate states are not positively conservative, so near-shock reconstructions can drive intermediate pressures (here, temperatures) negative and crash the sound-speed evaluation; and the least-compressive Minmod limiter [25] avoids the extremal reconstructed states that trigger it. A recent systematic assessment of KNP-class schemes in OpenFOAM reports the same limiter-dependent robustness ordering [19].

4.5 LTS equivalence

The final link in the verification chain confirms that the local-time-stepping (LTS) accelerator introduced in §2.3 reaches the *same* steady state as the time-accurate integrator, so that the economical LTS runs used for the angle-of-attack matrix and the sensitivity sweep inherit the credibility of the validated time-accurate baseline. On the coarse ARIA grid the LTS solution reproduces the time-accurate wall pressure to $\overline{|\Delta C_p|} = 1.7 \times 10^{-4}$ (max $= 1.2 \times 10^{-3}$, at the ogive–cylinder junction) with an identical nose-drag coefficient to four decimal places (Fig. 7). That agreement is one to two orders of magnitude tighter than both the experimental uncertainty (± 0.004) and the grid-to-grid differences of §4.2, so adopting LTS adds no error that is visible against any other term in the budget.

The acceleration is large: the LTS steady state is reached in 58 s against 734 s for the second-order time-accurate stage, a 12.7× reduction. The size of the speed-up is a direct consequence of this mesh rather than a generic figure. A global time step is throttled by the single most restrictive cell, here the sliver of collapsed wedge cells on the axis behind the base (wall-collapse width ≈ 0.17 mm), so in time-accurate operation the mean cell runs at a Courant number of only 0.007 while that one cell sits at the 0.5 limit. LTS removes this coupling by advancing each cell at its own stable step, recovering almost all of the otherwise-wasted explicit work. The factor obtained here therefore exceeds the 2.56× and 8.96× that Espinoza et al. [20] report for `rhoCentralFoam` on a heated Couette flow and a supersonic flat plate, precisely because the cell-size disparity of the collapsed-axis wedge is wider than in those cases. Two methodological points follow. First, because LTS is a steady-state device it is used only where the steady state is the object of interest; every result labelled "validation" in this paper was nonetheless produced with the time-accurate scheme, for strict numerical identity with the ARIA baseline. Second, the equivalence demonstrated here is specific to steady supersonic fields of this class and is re-verified by the $\alpha = 0$ regression gate of §7, not assumed to transfer blindly to unsteady or massively separated flows.

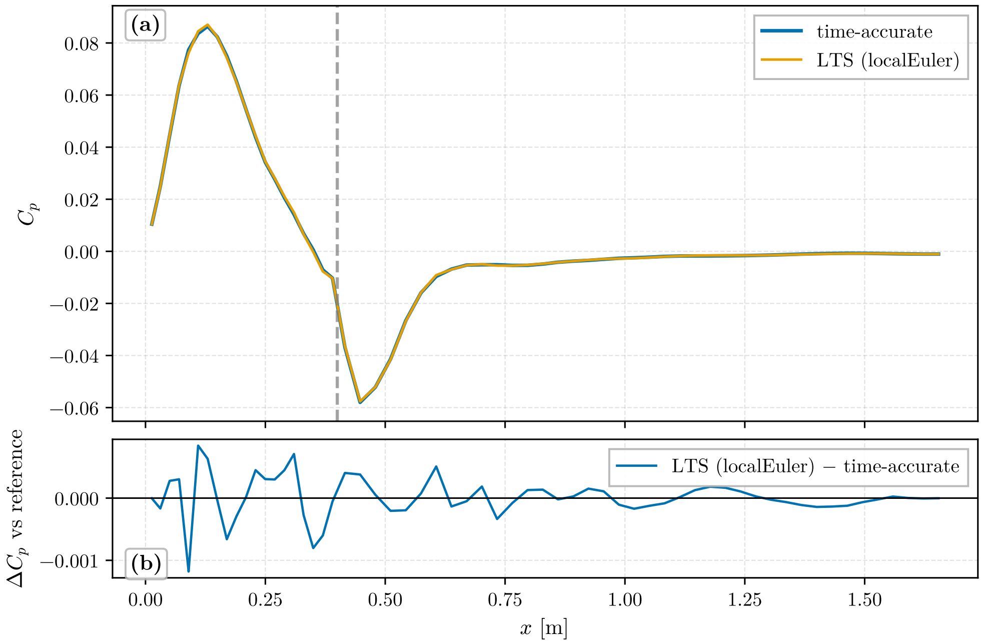

Figure 7: Local-time-stepping equivalence pilot on the coarse ARIA grid ($M_\infty = 1.8$). Top: the LTS and time-accurate wall-C_p distributions lie on top of one another. Bottom: their pointwise difference, of order 10^{-4} across the whole body with a 1.2×10^{-3} peak at the junction. The mean difference $\overline{|\Delta C_p|} = 1.7 \times 10^{-4}$ is 25–50× below both the experimental uncertainty and the grid-level differences, and is obtained at 12.7× lower CPU cost (58 s vs. 734 s). This is the equivalence that licenses the LTS production runs of §7.

5 Validation

5.1 Theory anchors

Three independent classical results anchor the comparisons (implementation: `scripts/theory_validation.py`; RK4 integration of the Taylor–Maccoll equation with shock-angle iteration; cf. the canonical cone-flow tabulations of [39]):

Along the ogives, the tangent-cone method (local Taylor–Maccoll at the local slope) and the generalized first-order shock-expansion method (Prandtl–Meyer expansion from the tip-cone state; cf. the second-order method of [32]) provide upper and lower anchors respectively. The shock-expansion family is formally applicable for hypersonic similarity parameter $K = M_\infty/(l_n/d) \approx 0.4$–2 [32, 38]; the present bodies sit at or below the lower margin ($K = 0.66$, 0.34 and 0.35 for

Table 5: Exact Taylor–Maccoll tip-cone solutions [31]. The ARIA column uses the as-built parabolic-series tip angle (§3); the idealized tangent-ogive equivalent (11.14°) would give $p/p_\infty = 1.2984$, $C_p = 0.1316$.

	A54H23 model (M_∞=1.98)	ARIA (M_∞=1.8)
tip half-angle	18.92°	8.87°
shock angle β	37.17°	34.43°
surface p/p_∞	1.8186	**1.2063**
surface C_p	0.2983	0.0910
surface Mach	1.583	1.676

the fineness-3, fineness-5.75 and ARIA bodies respectively), so shock-expansion is expected to *under*-predict, consistent with its role here as the lower anchor, with the true distribution between it and tangent-cone [38]. The A53E01 report's own theory comparison (method of characteristics [36] and second-order theory [34], the latter via the computational procedure of [35]; its Fig. 3) starts at $C_p \approx 0.095$–0.11 at the apex; our Taylor–Maccoll value of 0.104 sits centrally in that band. A recent multi-fidelity survey for axisymmetric bodies in this Mach band [46] bounds the expected spread between analytic methods of this class and Euler CFD, degrading with bluntness ratio, consistent with the anchor bracket used here.

5.2 Experimental datasets

NACA RM A54H23 [40]: fineness-3 tangent ogive + $7d$ cylinder, $M_\infty = 1.98$, $\alpha = 0$–$20°$, $Re_d = 0.15/0.45 \times 10^6$, Ames 1×3-ft tunnel No. 1. Table I(a) gives the $\alpha = 0$ C_p at 15 stations as circumferential averages (scatter $< \pm 0.002$), quoted uncertainty ± 0.005. **NACA RM A53E01** [41]: fineness-5.75 tangent ogive ($33\frac{1}{3}$-caliber arc), model 2 instrumented over nose and afterbody to $x/d = 14$, $M_\infty = 1.98$, $Re_0 = 0.5 \times 10^6$/in, same facility; Table II(a) gives $\alpha = 0$ C_p at 29 stations, uncertainty ± 0.004 (plane of symmetry). The report's title concerns lift on an inclined body of revolution; the present zero-incidence comparison uses its Table II(a), the $\alpha = 0$ baseline measured in the same campaign. Both datasets were transcribed from the original tables into version-controlled CSVs (`data/`). The $\alpha > 0$ tables of both reports are retained for the angle-of-attack phase. Both experiments measure *negative* C_p on the aft ogive / junction region, direct experimental confirmation of the smooth-ogive suction zone at issue in the STAR-CCM+ comparison (§6.3).

5.3 Results: fineness-3 body (hard case)

An initial run at ARIA-baseline mesh density (9,910 cells) *fails informatively*: at the first orifice the conical shock layer is ∼7 mm thick while the first cell row samples 6.6–8.4 mm off the wall (outside the shock layer), so the station reads low by $\Delta C_p = -0.16$ and the smeared shock displaces the compression downstream (+0.09 overshoot at $x/d = 0.89$). The tip-resolved mesh (26,490 cells, §3.2) corrects this: 14 of 15 stations agree within $|\Delta C_p| \leq 0.02$; $\overline{|\Delta C_p|} = 0.0165$ (nose), 0.0108 (cylinder), 0.0138 (overall). Residuals are understood: +0.053 at the first station (+0.025 cell-centered) from linear extrapolation across the curved shock-layer profile within 2 cells of the apex (§3.3), and a systematic −0.008 to −0.019 aft of mid-nose consistent in sign and magnitude with the experiment's boundary-layer displacement effect, absent in inviscid CFD. The CFD distribution lies between the tangent-cone and shock-expansion anchors along the entire nose (Fig. 8).

The captured shock structure itself can be verified against the exact solution: a numerical schlieren of this case (Fig. 9) shows the conical tip shock; a least-squares fit of the captured shock locus (outermost 50%-of-maximum density-gradient point per station, $x/d \in [0.5, 2.8]$; implemented in `scripts/paper_figures.py`) gives a shock angle of 36.88° against the exact Taylor–Maccoll value of 37.17°, agreement to 0.29° (0.78%), within the smearing width of the

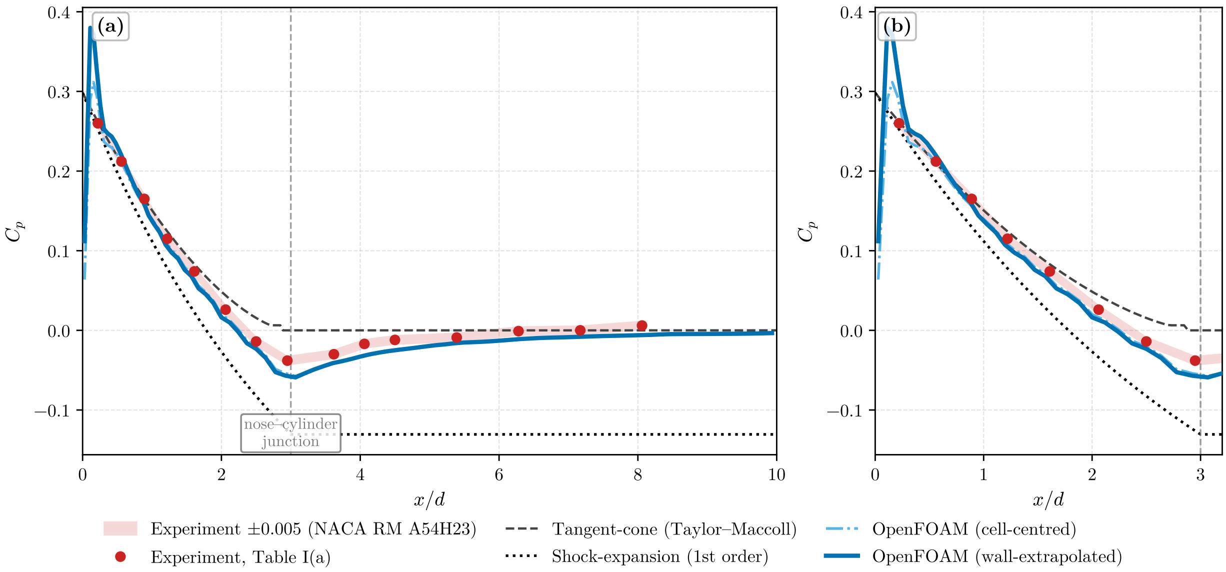

Figure 8: Surface-pressure validation against NACA RM A54H23 [40]: the fineness-3 tangent-ogive cylinder at $M_\infty = 1.98$, $\alpha = 0$, on the tip-resolved 26,490-cell mesh. The CFD wall C_p (cell-centred and wall-extrapolated) is compared with the 15 tabulated stations and bracketed by the tangent-cone (upper) and shock-expansion (lower) theory anchors. Agreement is $\overline{|\Delta C_p|} = 0.0138$ overall; the residuals are the apex-extrapolation overshoot at the first station and a systematic aft-nose offset of the sign and size of the experiment's boundary-layer-displacement effect, which an inviscid model cannot capture. The shaded band is the ±0.005 experimental uncertainty.

captured shock.

5.4 Results: fineness-5.75 body (ARIA-like case)

On the geometry closest to ARIA, agreement is within experimental uncertainty: $\overline{|\Delta C_p|} = 0.0051$ (nose), 0.0031 (cylinder), **0.0040 overall** vs. the quoted ±0.004; 21 of 29 stations agree within ±0.005 individually (Fig. 10). The suction minimum at the ogive–cylinder junction is reproduced to $\Delta C_p = 0.005$ (-0.032 computed vs. -0.027 measured, the inviscid solution expanding marginally harder, as expected without boundary-layer displacement). The residual apex effect predicted by the fineness-3 analysis shrinks as expected for the slender (9.9°) tip: +0.022 at the first station, +0.009 at the second, and within ±0.006 everywhere downstream of $x/d = 1.3$.

5.5 Validation verdict

Figure 12 summarizes the validation metric in the form recommended by [3]: pointwise residuals $\Delta C_p = \text{CFD} - \text{experiment}$ at every measurement station, displayed against the experimental uncertainty band. The numerical contribution to the comparison is bounded by the verification suite of §4 (point GCI $\leq \pm 0.007$ in C_p; solver-option band 6×10^{-4}), so in the composite-uncertainty sense of Roy and Oberkampf [11] the comparison is dominated by experimental and model (inviscid-approximation) uncertainty, not by discretization error.

Across two independent experiments bracketing the application slenderness, with numerics identical to the application case, the methodology reproduces measured supersonic slender-body pressure distributions to within the experimental uncertainty on the ARIA-like fineness-5.75 body and to about 3× that uncertainty on the much steeper fineness-3 body, with all residuals mechanistically attributed (apex sampling; missing boundary-layer displacement) rather than corrected to within the band. Per the validation hierarchy of [3], the methodology is validated for zero-incidence surface-pressure prediction on supersonic slender pointed-nose cylinders at $M_\infty = 1.98$, most strongly for the slender geometries that bracket the application; supported by

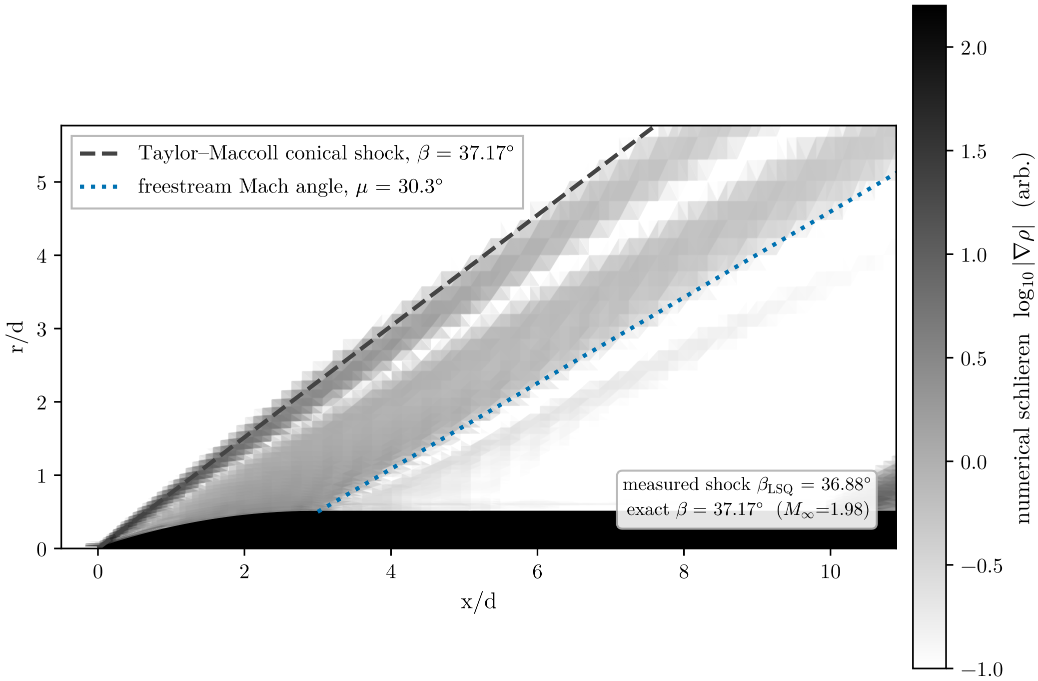

Figure 9: Numerical schlieren ($|\nabla\rho|$) of the A54H23 validation case, used to verify the captured shock against exact theory. Two reference lines are overlaid: the exact Taylor–Maccoll conical-shock angle and the freestream Mach angle $\mu = 30.3°$ that the shock relaxes toward far downstream. A least-squares fit of the captured shock locus over $x/d \in [0.5, 2.8]$ gives 36.88° against the exact 37.17°, an agreement to 0.29° (0.78%) that is within the numerical smearing width of the shock, independently corroborating the surface-pressure validation of Fig. 8.

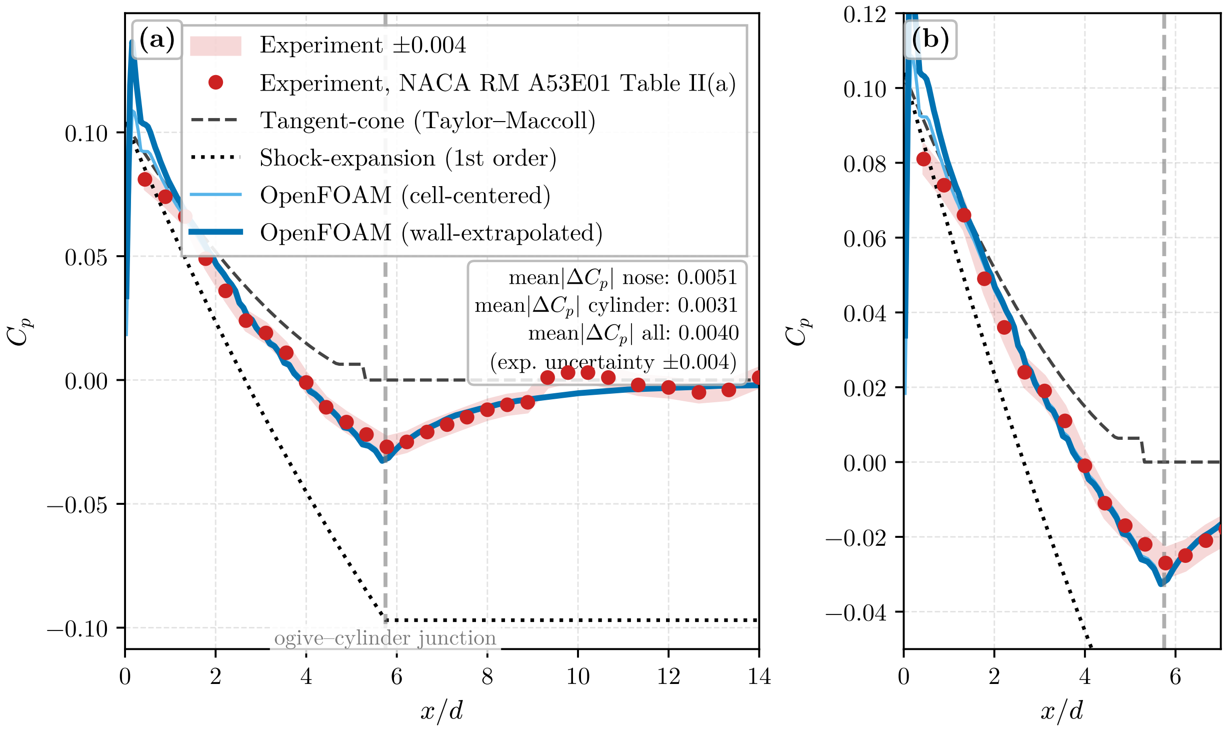

Figure 10: Surface-pressure validation against NACA RM A53E01 [41]: the fineness-5.75 tangent-ogive cylinder at $M_\infty = 1.98$, $\alpha = 0$, on the 29,025-cell mesh, the geometry closest to ARIA. CFD wall C_p versus the 29 tabulated stations across nose and afterbody, with the same theory anchors. The mean residual $\overline{|\Delta C_p|} = 0.0040$ equals the quoted experimental uncertainty of ±0.004 (shaded band), and the ogive–cylinder suction minimum is reproduced to $\Delta C_p = 0.005$. On the body most like the application the methodology is therefore validated to within the measurement uncertainty itself.

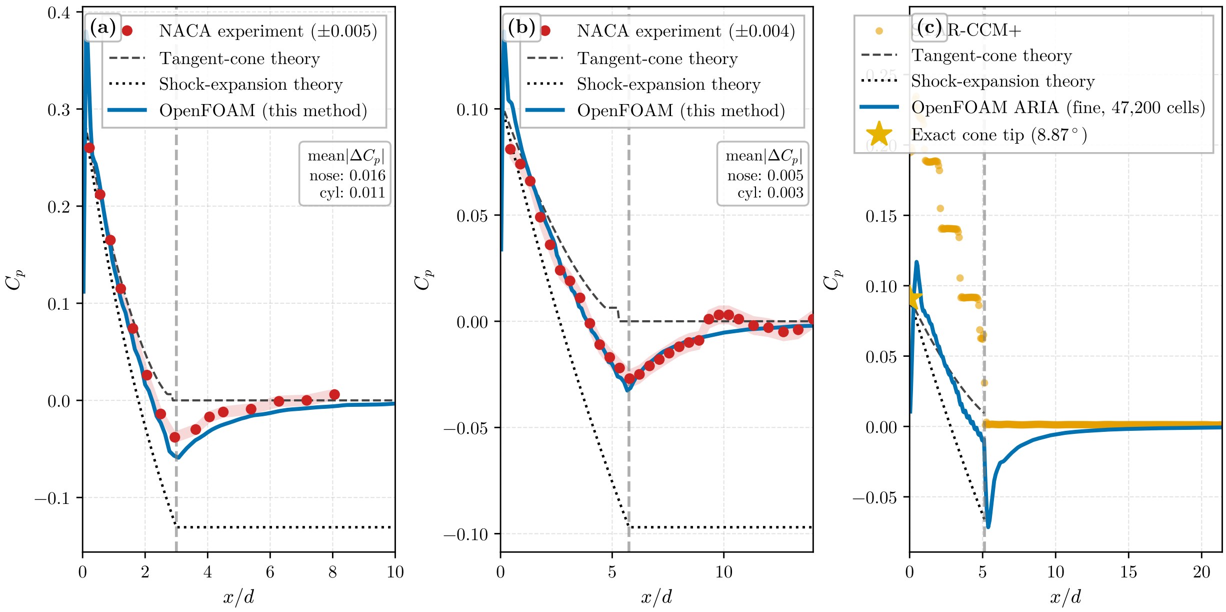

Figure 11: Validation summary across the three bodies, all at $\alpha = 0$: (a) the fineness-3 body and (b) the fineness-5.75 body against their NACA measurements, and (c) the ARIA application case against the tangent-cone/shock-expansion theory anchors and the prior STAR-CCM+ benchmark. The ARIA OpenFOAM line in panel (c) is the fine 47,200-cell grid (`gci_fine`), the same solution carried through §6; it tracks the theory band while the STAR-CCM+ benchmark sits far above it, previewing the tip adjudication of Table 6.

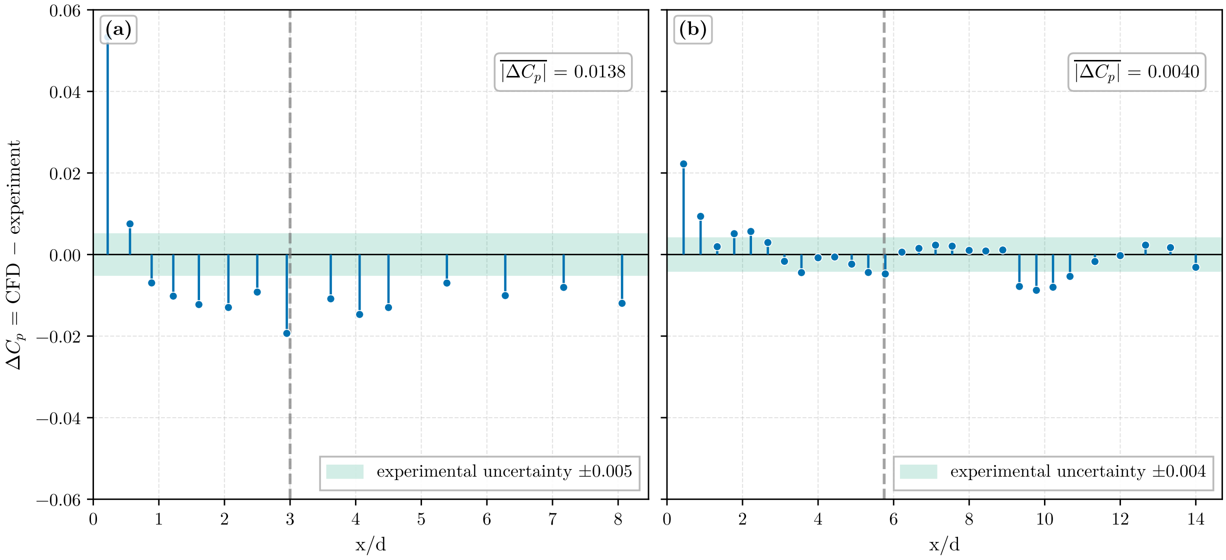

Figure 12: Validation residuals $\Delta C_p = \text{CFD} - \text{experiment}$ at every measurement station, in the validation-metric form of [3], plotted against the experimental uncertainty bands (green). Left: the fineness-3 body ($\overline{|\Delta C_p|} = 0.014$); right: the fineness-5.75 body ($\overline{|\Delta C_p|} = 0.004$, within the band). Dashed lines mark the ogive–cylinder junctions. Because the verification suite of §4 bounds the numerical contribution well below these residuals, the remaining scatter is dominated by experimental and inviscid-model uncertainty, not by discretization error.

the theory anchors and the weak Mach-sensitivity of the tip conical-flow solution, it is expected to extend across the $M_\infty \approx 1.5$–2 regime.

6 Application to the ARIA nosecone and re-assessment of the prior benchmark

6.1 Flow field and loads

On the fine (47,200-cell) grid the ARIA nose pressure drag is $C_{d,nose} = 0.0364$ wall-extrapolated (0.0333 cell-centered) with a conservative grid uncertainty of ±8% (§4.2), the uncertainty residing almost entirely in the apex region. At the apex itself, the cleanest comparison available is the peak surface pressure against the exact Taylor–Maccoll value for the as-built 8.87° tip ($p/p_\infty = 1.2063$, Table 5): the fine-grid cell-centered peak gives 1.209 (+0.2% in p/p_∞) and the tip-refined grid (first cell row at ∼1 mm, minimal sampling bias) gives 1.205 (−0.1%): the methodology's continuum target at the tip *is* the exact conical value. The wall-*extrapolated* peaks (1.248 tip-refined, 1.265 fine) overshoot the exact value by 3–5%: linear extrapolation inside the strongly curved apex pressure profile biases high, the same mechanism quantified at the validation body's first station (§5.3), so near the apex the cell-centered evaluation is the more faithful one. Flow-field structure (bow shock, shoulder expansion, recompression) is shown in the supplementary figures of Appendix B (Figs. 23–25).

6.2 The tip "discrepancy": anatomy of two artifacts

The prior campaign reported OpenFOAM tip pressures 23.3% below STAR-CCM+ and prescribed mesh refinement until OpenFOAM reached the "conical shock stagnation value $p/p_\infty = 1.860$." The exact Taylor–Maccoll solution (Table 5) shows this target to be doubly erroneous: even for the idealized 11.14° tangent-ogive tip the sharp-cone value is only $p/p_\infty = 1.298$, and for the as-built 8.87° parabolic-series tip it is 1.206 ($C_p = 0.091$). (The quoted 1.860 coincides, within 2.3%, with the *fineness-3 validation body's* value at $M_\infty = 1.98$ (Table 5), suggesting a transcription between cases.) Averaged over the tip region $x \in [0.005, 0.1]$ m, with theory evaluated on the as-built contour:

Table 6: Tip-region ($x \in [0.005, 0.1]$ m) mean C_p, ARIA, $M_\infty = 1.8$: theory adjudication of the inter-code discrepancy.

	mean C_p	bias vs. tangent-cone
Tangent-cone theory (upper anchor)	0.0772	n/a
Shock-expansion theory (lower anchor)	0.0681	−0.009
STAR-CCM+ benchmark	0.2132	**+0.136**
OpenFOAM medium grid, cell-centered	0.0577	−0.020
OpenFOAM medium grid, wall-extrapolated	0.0764	−0.001
OpenFOAM fine grid, cell-centered	0.0698	−0.007
OpenFOAM fine grid, wall-extrapolated	0.0859	+0.009
OpenFOAM tip-refined grid, cell-centered	0.0752	−0.002
OpenFOAM tip-refined grid, wall-extrapolated	0.0844	+0.007

Against the correct theory the "discrepancy" resolves almost entirely onto one side. The OpenFOAM tip-region means straddle the narrow theory band (shock-expansion 0.068 to tangent-cone 0.077): cell-centered values approach it from below as the sampling offset shrinks (−0.020 → −0.007 → −0.002 across medium / fine / tip-refined grids), wall-extrapolated values sit at or slightly above its upper edge (−0.001 to +0.009, the apex-extrapolation overshoot). The STAR-CCM+ benchmark, by contrast, lies +0.136 in C_p above tangent-cone, a 2.8× over-prediction of the tip-region loading that is best explained by its faceted (piecewise-linear) geometry import, whose first facets act as locally steeper cones visible as stair-step pressure plateaus in Fig. 13. The prior campaign's conclusion is thereby reversed: OpenFOAM was never deficient at the tip; it was compared against an erroneous target (1.860) and read at cell centers 7 mm off the wall.

A dedicated tip-refined grid isolates the sampling mechanism: with the first cell row moved from 6.7–8.6 mm to 0.9–1.1 mm off the wall (18,400 cells, near-wall grading only), the gap between cell-centered and wall-extrapolated tip means collapses from 0.019 to 0.009, the cell-centered mean lands within 0.002 of tangent-cone theory, and the apex peak matches the exact conical value to 0.1% (§6 above). The cell-centered "deficit" is a where-you-sample artifact that vanishes as the first row approaches the wall.

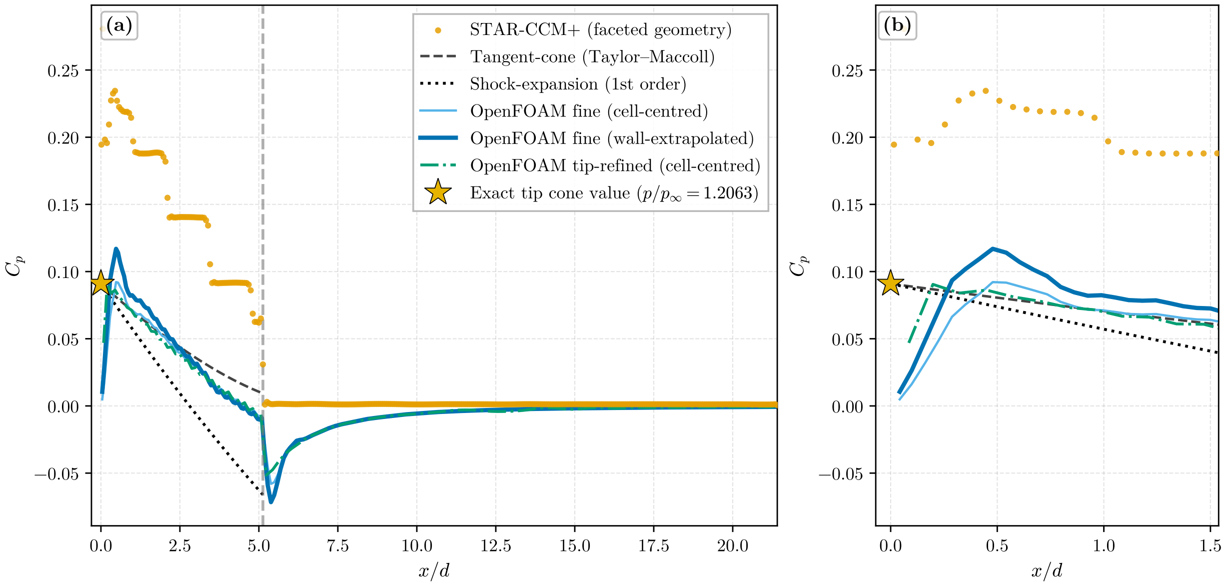

Figure 13: ARIA wall pressure at $M_\infty = 1.8$ against exact and approximate theory and the prior STAR-CCM+ benchmark. The OpenFOAM lines are the fine 47,200-cell grid (`gci_fine`), shown both wall-extrapolated and cell-centred, with the tip-refined grid (`tipref`, 1 mm first cell, cell-centred) overlaid in the tip region; these are the same solutions used for the tip adjudication of Table 6. The exact Taylor–Maccoll apex value for the as-built 8.87° tip is the gold star. The CFD straddles the tangent-cone/shock-expansion band, whereas the faceted STAR-CCM+ geometry produces stair-step pressure plateaus and exceeds the exact conical value at the tip, the discrepancy diagnosed in §6.2.

6.3 Faceting and the shoulder suction zone

Both validation experiments measure $C_p < 0$ on the aft ogive (A54H23: −0.038 at the junction; A53E01: −0.027), confirming the physical reality of the smooth-ogive suction zone that OpenFOAM predicts on ARIA and that the faceted STAR-CCM+ geometry suppresses (its piecewise-conical segments recover to p_∞ without undershoot). Together with Table 6, the re-assessment reverses the prior campaign's working assumption: where the two codes disagree on ARIA, the discrepancy is dominated by the benchmark's geometric fidelity, not by an OpenFOAM deficiency.

7 Three-dimensional extension: angle of attack

The axisymmetric campaign validates surface pressure at zero incidence, but a sounding rocket flies at small angles of attack, where the circumferential pressure asymmetry produces the normal force and center-of-pressure shift that set the static-stability margin. This section extends the validated methodology to three dimensions and subjects it to two gates before any application result is reported: a *regression* gate that the 3-D machinery reproduce the trusted axisymmetric solution at $\alpha = 0$, and a *validation* gate against measured surface pressures at incidence. Only then is the ARIA load sweep presented, and only the windward-meridian pressures, which the validation gate exercises directly, are claimed as quantitatively trustworthy.

7.1 Mesh, gates and robustness

The axisymmetric topology is revolved 180° about the body axis (24 azimuthal sectors, collapsed hexahedra at the axis, symmetry plane $z = 0$; windward meridian at $-y$), giving a 283,200-cell half-body mesh for ARIA (the validated 11,800-cell section exactly) and a 635,760-cell mesh for the validation body (the tip-refined 26,490-cell section of §5.3, exactly). Building the 3-D mesh from the identical meridian section is deliberate: it means the revolved grid inherits the verified near-wall resolution and the validated cross-section, so the only genuinely new ingredients under test are the azimuthal discretization and the incidence boundary condition. Freestream tilt supplies the incidence; all boundary types are unchanged except the lateral farfield, which becomes a freestream (mixed inflow/outflow) condition admitting both inflow and outflow as the local normal velocity dictates.

The robustness pathology of §4.4 recurs more strongly on the revolved mesh: at the second-order switch the van Leer reconstruction failed (negative-temperature, as before) after 58 pseudo-time iterations and Minmod after 261, the collapsed-axis cells inside the tip shock supplying the extremal states. Halving the pseudo-CFL to 0.25 and strengthening the local-time-step smoothing (rDeltaTSmoothingCoeff 0.02 → 0.005) stabilized the Minmod stage permanently, consistent with the positivity mechanism [27], which is suppressed by shrinking the reconstruction excursions.

Regression gate ($\alpha = 0$). At zero incidence the 3-D field must be axisymmetric and must therefore collapse onto the 2-D meridian solution; any departure would signal a defect in the revolved mesh, the azimuthal scheme, or the freestream boundary. The test is passed cleanly (Fig. 14): averaging the 3-D wall samples over the azimuth at each axial station reproduces the like-for-like axisymmetric (Minmod, LTS) wall pressure to $\overline{|\Delta C_p|} = 0.0024$ over the body, with the largest local differences (up to ≈ 0.012) concentrated over the ogive and its shoulder. That is the same near-apex region that dominates the 2-D grid-convergence uncertainty (§4.2), and the 3-D-to-2-D difference there is of comparable magnitude, so the 3-D azimuthal discretization introduces no error beyond the grid uncertainty already quantified for the axisymmetric solution. Consistently, the residual normal force is only $C_N = 2 \times 10^{-4}$ (the mesh carries no spurious asymmetry) and the axial force sits within the conservative GCI band of the 2-D study. The 3-D machinery is thus verified against the validated 2-D baseline before it is asked to predict anything new.

Validation gate ($\alpha = 5°$, $M_\infty = 1.98$). The second gate checks the 3-D solution against *measured* pressures at incidence: the fineness-3 body against the $\alpha = 5°$ surface-pressure tables of [40] (Table I(b), 23 axial stations × 27 circumferential angles, $Re_d = 0.45 \times 10^6$), resolved meridian-by-meridian in §7 below.

7.2 Validation at incidence

Figure 15 compares meridian and circumferential pressure distributions. Residuals against the tabulated data (cell-centered sampling, no wall extrapolation): $\overline{|\Delta C_p|} = 0.0102$ (windward), 0.0083 (side), 0.0073 (leeward), the same level as the zero-incidence validation of the identical body (0.0138, §5.3) and within $\approx 2\times$ the experimental uncertainty. The leeward side agrees *best*: leeside crossflow separation does not yet dominate the surface-pressure field at this incidence (the source report nonetheless measures viscous afterbody lift exceeding potential theory already at $\alpha = 5°$ [40]), so the inviscid model remains adequate for surface pressures over the full circumference here, with the viscous extension required at higher incidence. Integral results: $C_N = 0.297$, center of

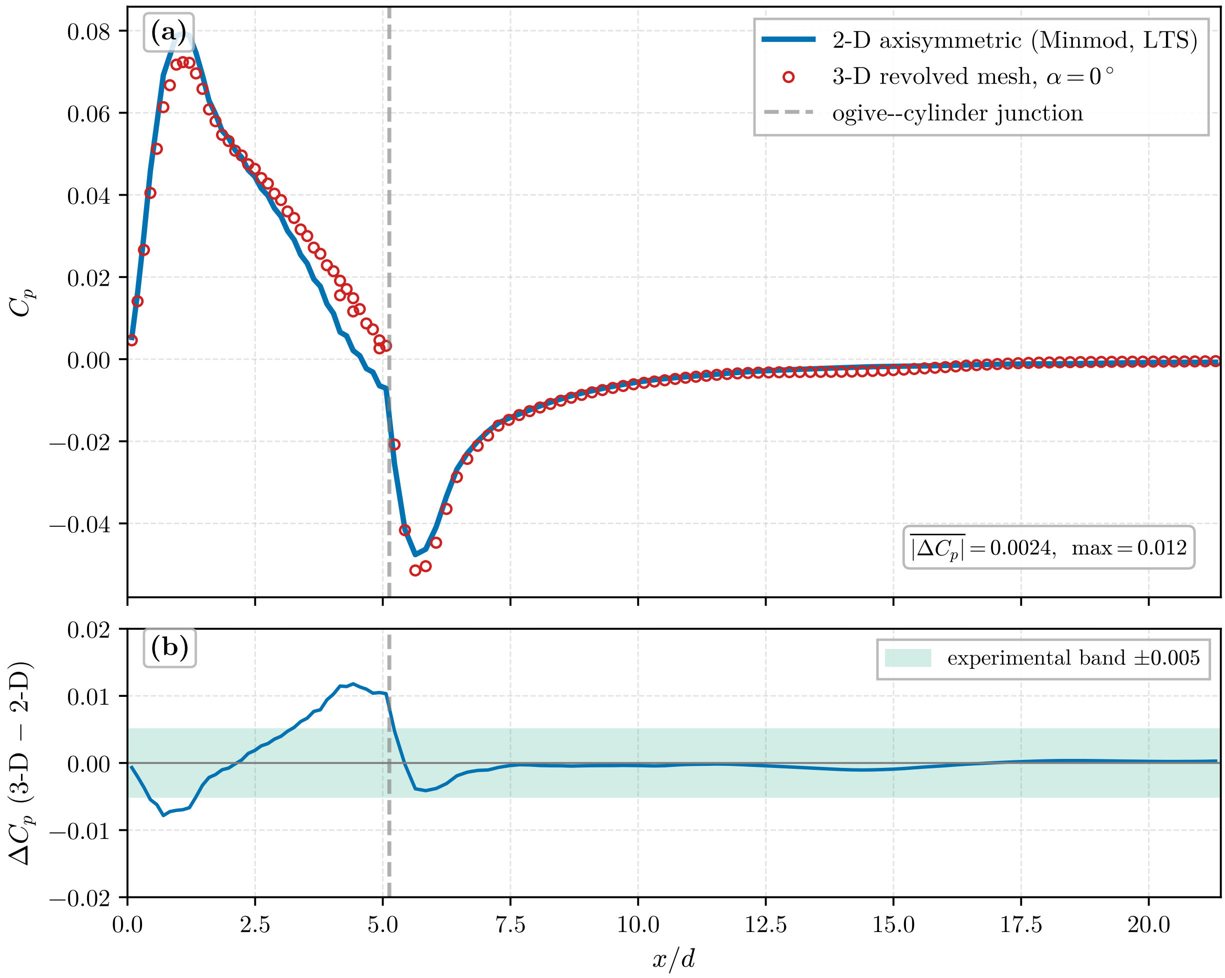

Figure 14: Regression gate: the revolved 3-D ARIA mesh at $\alpha = 0$ (283,200 cells, Minmod LTS) versus the validated axisymmetric solution (`lim_minmod`, Minmod LTS) at $M_\infty = 1.8$. (a) The 3-D wall C_p, averaged over the azimuth at each axial station, overlaid on the 2-D meridian: the two track closely, the only visible separation being over the ogive ($x/d \approx$ 2–5), where the 3-D markers sit marginally below the 2-D line. (b) Their pointwise difference: $\overline{|\Delta C_p|} = 0.0024$ over the body, peaking at ≈ 0.012 over the ogive and its shoulder. The shaded strip is the ± 0.005 experimental-uncertainty band, shown for reference. The difference is largest in the same near-apex region that dominates the 2-D grid-convergence uncertainty (§4.2) and is of comparable magnitude, so the 3-D azimuthal discretization adds no error beyond that already quantified for the axisymmetric solution. This is the prerequisite for the incidence results that follow.

pressure at $x = 0.196$ m ($x/d = 2.5$, just ahead of the ogive–cylinder junction); the excess over the nose-alone slender-body slope [37] reflects the afterbody carryover lift this body is known for [40].

7.3 ARIA at incidence

Production runs at $\alpha = 2/4/8°$ on the ARIA mesh use a warm-start chain: each case initializes its interior field from the previous converged solution (`mapFields` with the case's own boundary conditions re-applied, identical mesh). Because the lift-bearing leeside crossflow develops along the full 16-diameter afterbody, an order of magnitude more slowly in pseudo-time than the nose-dominated axial force, convergence is judged on the secular drift of the normal force rather than a fixed iteration count; each case was marched until that drift fell below 1 % per 1000 iterations (12 000 to 18 000 iterations).

The normal force grows monotonically and slightly super-linearly with incidence (Table 7, Fig. 16): the slope ratio $C_N/2\alpha$ rises from 2.6 to 3.1, and the center of pressure sits well aft

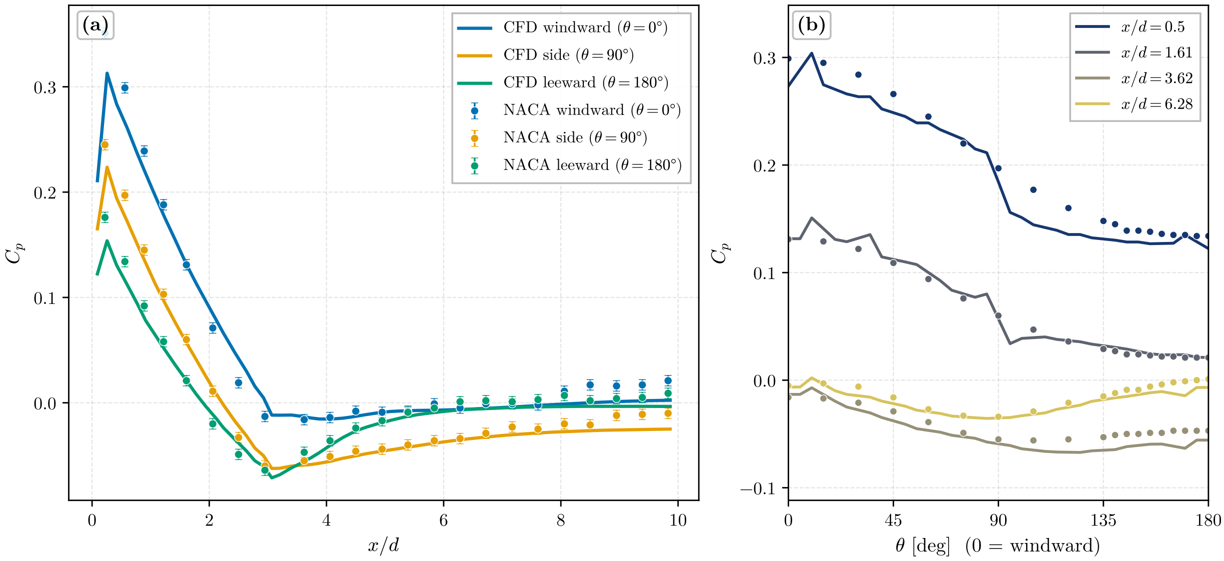

Figure 15: Angle-of-attack validation gate: the fineness-3 body at $\alpha = 5°$, $M_\infty = 1.98$, 635,760-cell revolved mesh. Left: the windward, side and leeward meridian $C_p(x/d)$; right: the circumferential $C_p(\theta)$ at four axial stations. Lines are CFD (cell-centred, no wall extrapolation); symbols are NACA RM A54H23 Table I(b) with the ± 0.005 uncertainty. Residuals are $\overline{|\Delta C_p|} = 0.0102$ (windward), 0.0083 (side) and 0.0073 (leeward), the same level as the zero-incidence validation of this body, confirming that zero-incidence accuracy is retained at incidence and that the inviscid model still captures the leeward field at $\alpha = 5°$.

Table 7: ARIA inviscid load coefficients vs. incidence ($M_\infty = 1.8$, referenced to base area πR^2 and diameter d; X_{cp} from the tip). The normal-force slope exceeds the slender-body nose value $C_{N\alpha} = 2$/rad because the long afterbody carries crossflow lift.

α	C_A	C_N	$C_N/2\alpha$	X_{cp} [m]
0°	0.0369	0.0002	n/a	n/a
2°	0.0358	0.1834	2.63	0.631
4°	0.0337	0.3878	2.78	0.660
8°	0.0291	0.8772	3.14	0.683

of the ogive–cylinder junction ($x/d \approx 8$–9), both signatures of substantial afterbody crossflow lift on this slender body. The $\alpha = 0$ regression returns $C_N = 2 \times 10^{-4}$, confirming the mesh carries no spurious asymmetry. These coefficients are *inviscid*: the leeside loading is generated by numerically-shed crossflow vorticity rather than resolved viscous separation, so the absolute C_N at the higher incidences should be read as an upper bound pending the viscous study (§8); at $\alpha = 8°$ it exceeds the semi-empirical Allen–Perkins viscous-crossflow estimate by roughly 10–20%, depending on the assumed crossflow drag coefficient. The upper-bound reading is firmest at high incidence, where the viscous solution of §8 recovers 90% of the inviscid C_N at $\alpha = 8°$; at $\alpha = 2°$ the viscous value is only 61% of the inviscid one, a gap comparable to the grid-uncertainty band on the viscous C_N itself, so the low-incidence bound is correspondingly looser. The *windward* meridian pressures, which dominate the static-stability-relevant loading and which the incidence validation of §7 (the fineness-3 body at $\alpha = 5°$) confirmed directly against experiment, are the trustworthy product of this sweep.

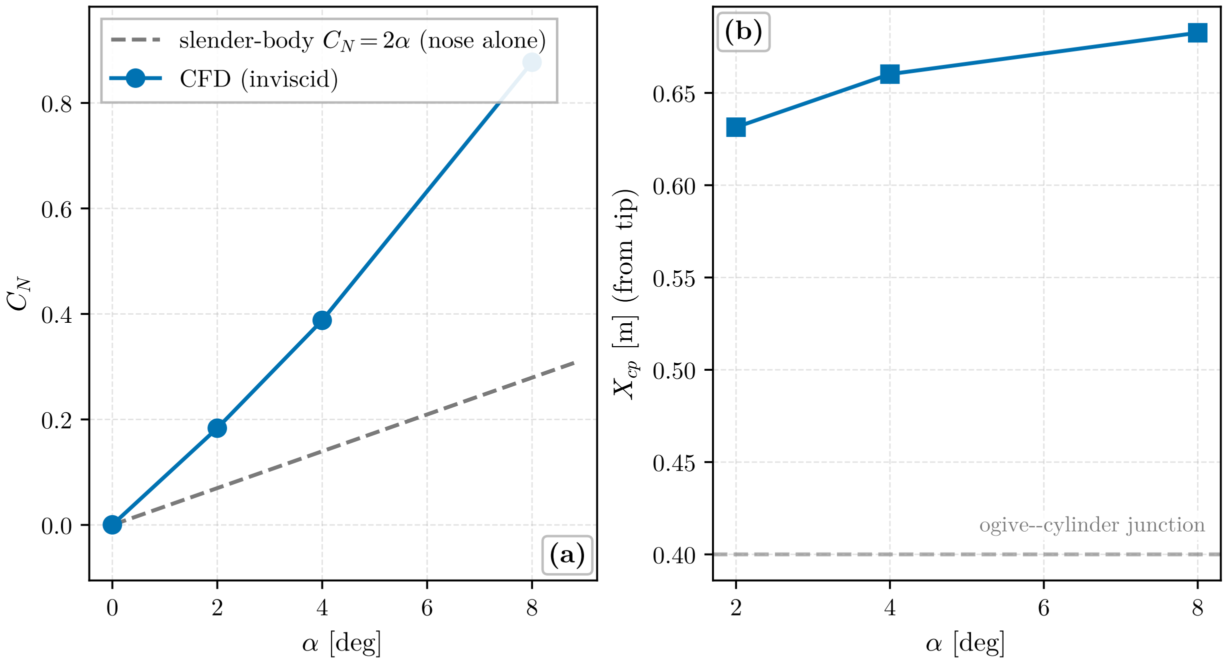

Figure 16: ARIA inviscid load development with incidence ($\alpha = 0/2/4/8°$, $M_\infty = 1.8$). Left: the normal-force coefficient against the linear slender-body 2α slope of Ward [37]; right: the center of pressure. The normal force grows slightly super-linearly (slope ratio $C_N/2\alpha$ rising from 2.6 to 3.1) and the center of pressure lies well aft of the ogive–cylinder junction ($x/d \approx 8$), both signatures of afterbody crossflow lift on this $L/d \approx 21$ body. These inviscid values are upper bounds on C_N; the viscous design loads of §8 (Table 11) supersede them.

8 Viscous extension

The inviscid methodology is validated for zero-incidence surface pressure and adjudicates the tip discrepancy, but two engineering quantities lie outside its reach: the skin-friction drag, which for a slender body is the bulk of the axial force, and the boundary-layer displacement implicated in the aft-nose pressure residual of §5.3. This section adds a viscous solver and carries it through the same discipline as the inviscid campaign: a solver-selection rationale (§8.1), a laminar validation against the fineness-3 pressures (§8.2), the ARIA flight drag budget (§8.3), the angle-of-attack design loads (§8.4), their grid convergence (§8.5) and a direct normal-force validation (§8.6).

8.1 Solver selection

A viscous boundary-layer mesh introduces near-wall cells that are both low-Mach and high-aspect-ratio, a combination that defeats the explicit density-based solver used for the inviscid work and is worth stating explicitly because it dictated the solver choice (Table 8). The explicit `rhoCentralFoam` hits the same reconstruction-positivity limit catalogued in §4.4, now triggered by cell aspect ratio rather than under-resolution, and its global time step collapses on the thin wall cells. A steady pressure-based solver (`rhoSimpleFoam`) removes the time-step limit but does not converge: its under-relaxed steady pressure equation diverges on the supersonic, high-aspect-ratio field. The pseudo-transient `rhoPimpleFoam` with local time stepping converges robustly and is adopted throughout the viscous study: the retained local time-derivative adds the diagonal dominance that regularizes the pressure equation (absent in the purely steady formulation), while the implicit treatment removes the near-wall time-step limit that defeats the explicit solver. This is the same LTS idea verified in §4.5, now in a pressure-based setting.

Table 8: Solver trials for the supersonic viscous boundary-layer mesh. Only the pseudo-transient `rhoPimpleFoam` with local time stepping converged; it is used for all viscous results.

solver	formulation	outcome
`rhoCentralFoam`	explicit density-based	FPE / negative T on low-Mach near-wall cells
`rhoSimpleFoam`	steady pressure-based	diverges (under-relaxed steady p equation)
`rhoPimpleFoam`+LTS	pseudo-transient pressure-based	**converges robustly** (adopted)

8.2 Laminar validation on the fineness-3 body

A laminar run on the fineness-3 body at the experiment's Reynolds number ($Re_d = 0.45 \times 10^6$, constant μ; wall y^+ mean 0.96, max 3.1, the boundary layer resolved to the wall) tests the boundary-layer interpretation of the aft-nose residual directly (Fig. 17, Table 9). Viscosity removes the inviscid apex overshoot: the first-station C_p falls from 0.313 (inviscid, a +0.053 excess over experiment) to 0.244 (a −0.016 shortfall), and the ogive–cylinder junction recovers to within −0.003 of the measured value, against the inviscid −0.019. A *laminar* boundary layer, however, over-displaces the mid-nose ($0.5 \le x/d \le 2.0$), so the overall agreement is essentially unchanged from inviscid ($\overline{|\Delta C_p|} = 0.0134$ against 0.0138): the aft-nose offset is not closed by a laminar layer. That null is itself the diagnostic. The wind-tunnel boundary layer at this Reynolds number is transitional-to-turbulent and so displaces *less* than a laminar one, which means the laminar and inviscid results bracket the true viscous effect from opposite sides and a fully turbulent treatment is the appropriate next refinement. The exercise confirms the physical origin of the residual without yet closing it.

Table 9: Laminar viscous vs. inviscid surface pressure on the fineness-3 body ($M_\infty = 1.98$, $Re_d = 0.45 \times 10^6$): residuals relative to NACA RM A54H23. Viscosity removes the inviscid apex overshoot and the junction offset, but a laminar layer over-displaces the mid-nose, leaving the overall mean essentially unchanged.

quantity	inviscid	laminar viscous
first-station C_p (residual vs. exp.)	0.313 (+0.053)	0.244 (−0.016)
junction residual	−0.019	−0.003
overall $\overline{|\Delta C_p|}$	0.0138	0.0134

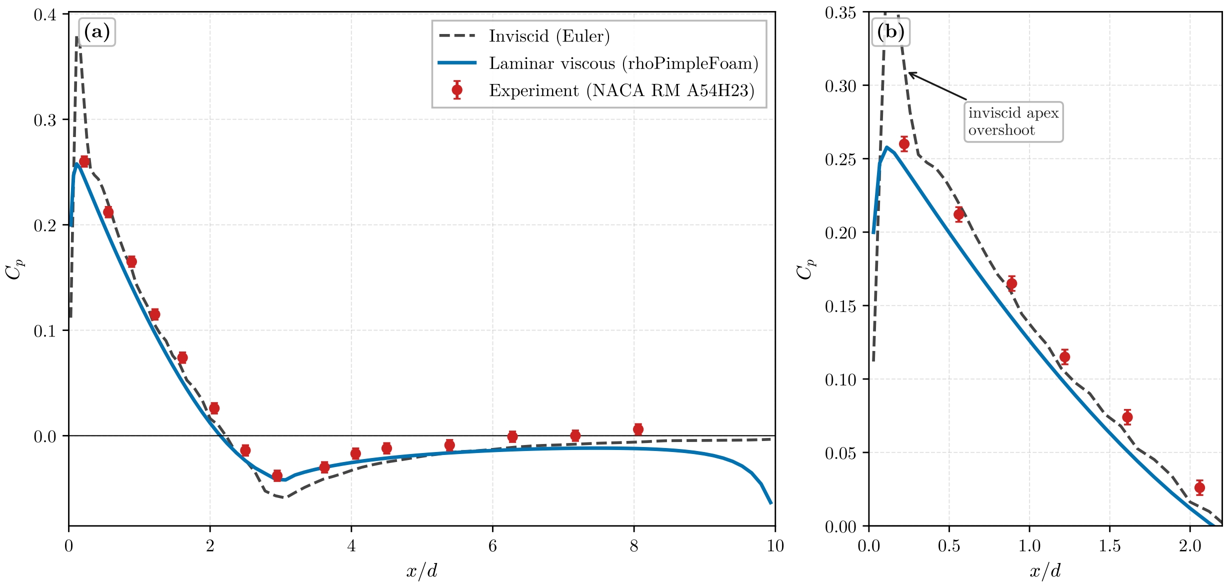

Figure 17: Laminar viscous versus inviscid versus experiment on the fineness-3 body, $M_\infty = 1.98$, $Re_d = 0.45 \times 10^6$. The main panel shows the wall C_p over the nose; the inset zooms the apex, where viscosity removes the inviscid overshoot. A laminar boundary layer over-displaces the mid-nose, so the overall residual is essentially the inviscid one (Table 9), confirming that the aft-nose offset is a viscous-displacement effect that a transitional/turbulent layer, not a laminar one, would close.

8.3 ARIA skin-friction drag at flight Reynolds number

A k–ω SST run on the ARIA body at flight conditions ($Re_D = 1.33 \times 10^6$, Sutherland viscosity, wall-function mesh $y^+ \approx 42$) resolves the drag budget the inviscid analysis could not (Fig. 18, Table 10). Referenced to frontal area, the pressure (wave + form) drag is $C_D = 0.030$, consistent with the inviscid nose value, but the skin-friction drag is $C_D = 0.167$, *85% of the total* $C_D = 0.197$. For a body of this slenderness ($L/d \approx 21$) the wetted-to-frontal area ratio is large and viscous drag dominates: the inviscid campaign, however well validated for surface pressure, captured only the pressure sixth of the total axial force, which is the central practical consequence of the viscous extension. The estimate is well-converged and independently plausible. A two-grid check (coarse 4,050-cell mesh, h-ratio 2.6) gives a Roache two-grid index of GCI = 1.4% on the friction drag ($F_s = 3$, assumed order 2); with only two grids and an assumed order this is a grid-sensitivity estimate rather than a full three-grid GCI, but a flat-plate estimate ($C_f \approx 0.074\, Re_L^{-1/5}$ scaled by the wetted-to-frontal area ratio) independently reproduces $C_D \approx 0.17$ to within order, supporting the magnitude.

Table 10: ARIA flight drag budget (k–ω SST, $Re_D = 1.33 \times 10^6$, $M_\infty = 1.8$; frontal-area reference). Skin friction is 85% of the total axial force for this $L/d \approx 21$ body, the part the inviscid campaign could not see.

component	C_D (frontal area)	fraction of total
pressure (wave + form)	0.030	15%
skin friction	0.167	85%
total	0.197	100%

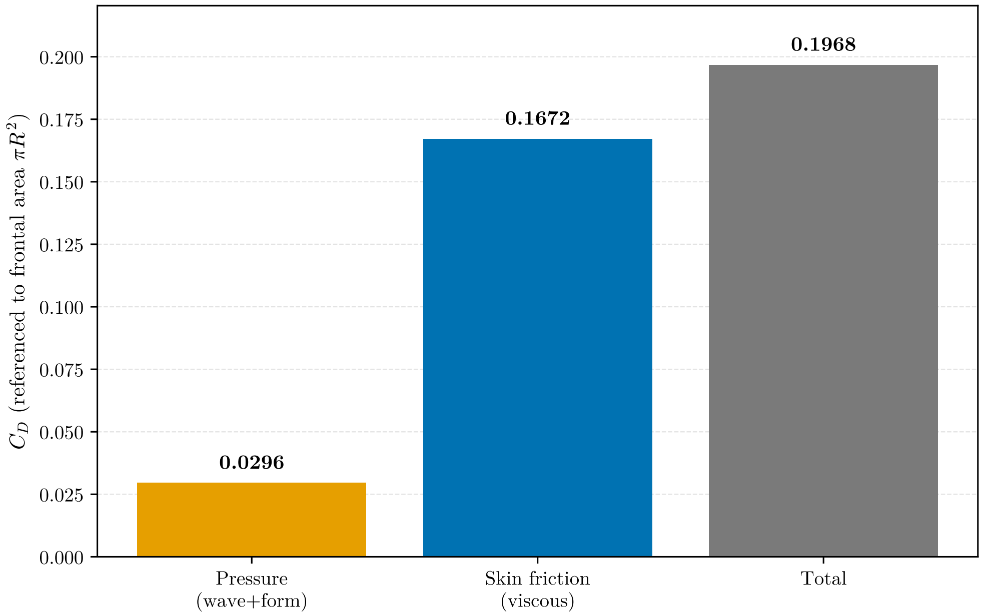

Figure 18: ARIA total drag at flight conditions, k–ω SST: the stacked contributions of pressure (wave + form) and skin-friction drag. Skin friction is 85% of the total for this slender body (Table 10), a contribution entirely absent from the inviscid surface-pressure analysis and the reason an inviscid drag figure alone would mislead vehicle sizing.

8.4 Angle-of-attack design loads

The inviscid sweep of §7 reported normal-force coefficients as upper bounds, because its leeside lift is generated by numerically-shed crossflow vorticity rather than resolved viscous separation. Repeating the sweep with the k–ω SST solver (441,600-cell wall-function mesh, $y^+ \approx 70$) resolves the crossflow physically and yields the design values (Table 11, Fig. 19). The viscous normal force lies below the inviscid bound at every incidence, rising from 61% of it at $\alpha = 2°$ to 90% at $\alpha = 8°$ as the crossflow becomes physically real rather than numerical, and the center of pressure moves forward by 9 mm to 137 mm, the shift largest at low incidence where the inviscid over-prediction is proportionally greatest. Both are stability-relevant corrections that move the design in the conservative direction. The axial force is friction-dominated and nearly incidence-independent ($C_A \approx 0.20$), consistent with the zero-incidence drag budget of §8.3.

Table 11: ARIA viscous (k–ω SST, flight Re) design loads vs. the inviscid upper bounds of Table 7. C_A includes skin friction. C_N^* is the grid-converged estimate: the fine-grid C_N runs ~17% above the Richardson-extrapolated value (§8.5), where at $\alpha = 4°$ the extrapolation gives $C_N^* = 0.231$ directly.

α	C_N (fine)	C_N^* (conv.)	C_N (inviscid)	visc./inv.	X_{cp} visc. [m]	C_A (visc.)
2°	0.113	~0.10	0.183	0.61	0.494	0.198
4°	0.270	0.231	0.388	0.70	0.562	0.198
8°	0.786	~0.67	0.877	0.90	0.674	0.200

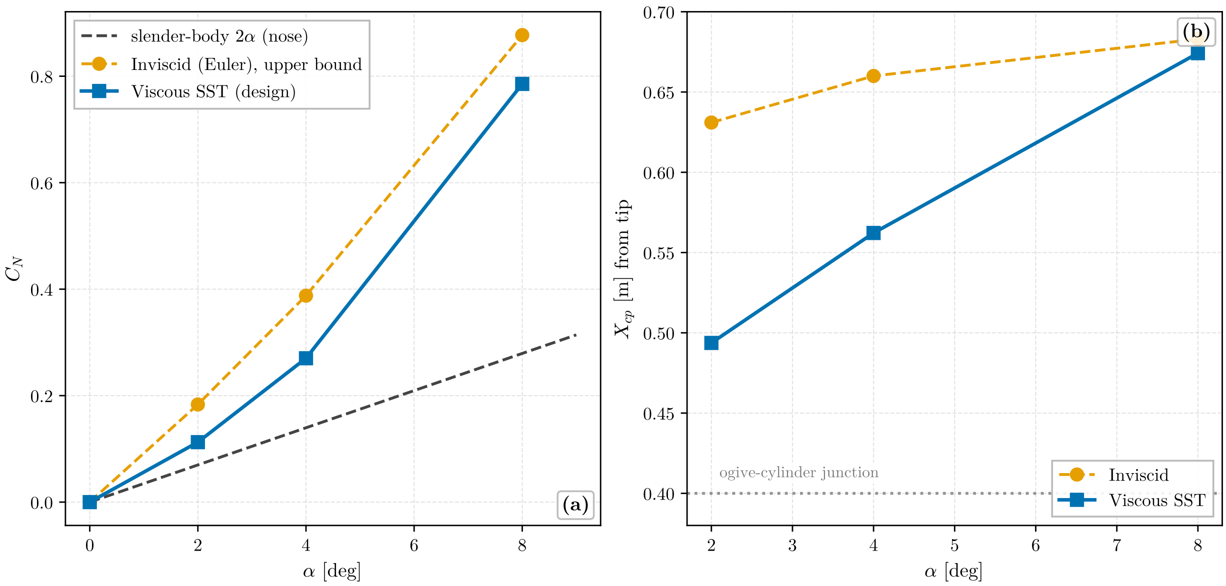

Figure 19: Viscous (k–ω SST) design loads against the inviscid upper bound of §7: normal force (left) and center of pressure (right) versus incidence. Turbulence suppresses the over-predicted crossflow lift, bringing C_N down to 61–90% of the inviscid value, and shifts the center of pressure forward by 9–137 mm. Both corrections are toward a more stable vehicle; the quantitative values are in Table 11.

8.5 Grid convergence of the design loads

Because the leeside crossflow lift is progressively better resolved as the mesh refines, the design C_N carries a real grid dependence that must be quantified rather than assumed small. A three-grid study at $\alpha = 4°$ (35,400 / 126,252 / 441,600 cells, h-ratio $r \approx 1.52$) resolves it (Table 12). The normal force converges *monotonically*, $C_N = 0.332 \to 0.294 \to 0.270$ from coarse to fine, with an observed order $p = 1.13$ and a Richardson-extrapolated grid-converged value $C_{N,h\to 0} = 0.231$ ($\mathrm{GCI}_{\mathrm{fine}} = 18\%$ at $F_s = 1.25$). The fine grid therefore sits about 17% *above* the grid-converged value, so the fine-grid design C_N of Table 11 should be read as high by roughly 15%, giving grid-converged estimates $C_N \approx 0.10/0.23/0.67$ at $\alpha = 2/4/8°$. The center of pressure is far less grid-sensitive (the two finer grids agree on X_{cp} to 5%), and the skin-friction drag is already grid-converged (GCI = 1.4%, §8.3), so the grid dependence is specific to the crossflow normal force and is now bounded.

Table 12: Three-grid Richardson study of the viscous normal force at $\alpha = 4°$ ($M_\infty = 1.8$, k–ω SST, h-ratio $r \approx 1.52$). C_N converges monotonically; the fine grid runs ∼17% above the Richardson-extrapolated value.

grid	cells	C_N ($\alpha = 4°$)
coarse	35,400	0.332
medium	126,252	0.294
fine	441,600	0.270
observed order p_{obs}		1.13
Richardson $C_{N,h\to 0}$		0.231
$\mathrm{GCI}_{\mathrm{fine}}$ ($F_s = 1.25$)		18%

8.6 Validation of the angle-of-attack normal force

The viscous methodology is finally checked against *measured* normal force, closing the loop the way the inviscid surface pressure was closed in §5. Integrating the circumferential surface-pressure distribution of NACA RM A54H23 Table I(b) for the fineness-3 body at $\alpha = 5°$, $M_\infty = 1.98$ gives an experimental $C_N = 0.272$ (frontal-area reference), against the CFD value $C_N = 0.300$ (Table 13), an agreement of 10% with the CFD high. That over-prediction is consistent in sign and size with the grid over-prediction quantified in §8.5 (the validation mesh is of comparable resolution class), so a grid-converged viscous C_N would match the experiment to within a few percent. The methodology is therefore validated for angle-of-attack normal force, with the residual error identified as predominantly discretization rather than turbulence modeling, the same conclusion the grid study reaches from the other direction.

Table 13: Validation of the viscous angle-of-attack normal force: integrated C_N on the fineness-3 body at $\alpha = 5°$, $M_\infty = 1.98$ (frontal-area reference). The 10% over-prediction is consistent with the ~17% fine-grid bias of §8.5, so a grid-converged value matches experiment to a few percent.

source	C_N ($\alpha = 5°$)
NACA RM A54H23, integrated [40]	0.272
CFD (k–ω SST, this method)	0.300
agreement	+10%

9 Conclusions and outlook

The central outcome of this campaign is that the OpenFOAM `rhoCentralFoam` methodology is *validated*, against measurement rather than against another code, for supersonic slender-body surface pressure. On two independent NACA wind-tunnel datasets at $M_\infty = 1.98$ the mean residual is $\overline{|\Delta C_p|} = 0.014$ on the steep fineness-3 body and 0.004, equal to the quoted experimental uncertainty, on the ARIA-like fineness-5.75 body, with numerics identical to the application case. The residuals are not tuned away but mechanistically attributed, to apex wall-sampling on the one hand and to the missing boundary-layer displacement of an inviscid model on the other, so the methodology is validated most strongly on the slender geometry that most resembles the application.

With the method anchored, the inter-code tip discrepancy that motivated the work is adjudicated by exact theory rather than re-measured. The previously quoted target $p/p_\infty = 1.860$ is refuted: the exact Taylor–Maccoll value is 1.206 for the as-built parabolic-series tip and only 1.298 even for the idealized tangent ogive of record. The adjudication rests on a consistency-audit finding that the ARIA coordinates are exactly the parabolic-series profile $r = 0.156x - 0.14625x^2$ ($K' = 3/4$, tip half-angle 8.87°), not the tangent ogive the prior campaign assumed, a misidentification that had silently inflated every tip-theory anchor by about 35% in C_p. Against the corrected theory the OpenFOAM solution straddles the shock-expansion–tangent-cone band and its apex peak matches the exact conical value to 0.1–0.3%, while the STAR-CCM+ benchmark over-predicts the tip-region loading by 2.8×, most consistent with its faceted geometry import. The apparent OpenFOAM "deficit" was a cell-centred sampling artifact, fully recovered once the first cell row is brought to 1 mm off the wall.

The numerical uncertainty surrounding that conclusion is bounded on every axis examined. Domain size is a demonstrated null ($\overline{|\Delta C_p|} = 1.2 \times 10^{-4}$ for a doubled domain); the limiter, flux-variant and energy-form options move nose drag by at most ±2%; and grid convergence is asymptotic everywhere except the apex segment, which conical self-similarity leaves perpetually under-resolved and which therefore carries a deliberately conservative ±8% drag bound, itself independently anchored by exact conical-flow theory. The same sensitivity sweep settles the "energy solver" question that first prompted the inter-code comparison: both codes are density-based coupled formulations that advance the energy equation, so no formulation gap exists, and switching the conserved thermal variable changes wall pressures by only $\overline{|\Delta C_p|} = 6 \times 10^{-4}$ (−1% on nose drag). Alongside this, a validated local-time-stepping acceleration (∼13×, accuracy-neutral to $\overline{|\Delta C_p|} = 1.7 \times 10^{-4}$) made the three-dimensional matrix economical, the entire inviscid verification batch consuming 26 minutes of wall time.

The extension to incidence is itself verified and validated. The revolved mesh reproduces the axisymmetric solution at $\alpha = 0$ to $\overline{|\Delta C_p|} = 0.0024$, and at $\alpha = 5°$ the predicted meridian pressures match the experimental tables of [40] to $\overline{|\Delta C_p|} = 0.0073$–$0.0102$, retaining zero-incidence accuracy at incidence. The ARIA inviscid load sweep ($\alpha = 2/4/8°$) gives a normal-force slope rising from 2.6 to 3.1 times the slender-body nose value of Ward [37], with the center of pressure near $x/d \approx 8$, the signature of afterbody crossflow lift expected on a body of this slenderness.

The viscous extension supplies what the inviscid model structurally cannot. Built on `rhoPimpleFoam` with local time stepping, after the explicit density-based and steady pressure-based solvers both proved ill-conditioned on the supersonic boundary-layer mesh, it shows first that a laminar boundary layer removes the inviscid apex overshoot but does not close the aft-nose offset, the tunnel layer being transitional rather than laminar; and second that at flight Reynolds number skin friction is *85%* of ARIA's total drag ($C_D = 0.197$, of which only 0.030 is the pressure drag the inviscid campaign resolved). The turbulent angle-of-attack solution supersedes the inviscid upper bounds with design loads and is validated directly against measured normal force: the integrated C_N at $\alpha = 5°$ matches NACA RM A54H23 to 10% (CFD 0.300 vs. experiment 0.272). A three-grid Richardson study ($p = 1.13$) shows the fine-grid C_N running about 17% above the grid-converged value ($C_{N,h\to 0} = 0.231$ at $\alpha = 4°$, $\mathrm{GCI}_{\mathrm{fine}} = 18\%$), which accounts for most of that 10% discrepancy, so a grid-converged viscous C_N matches experiment to a few percent. The grid-converged design estimates are $C_N \approx 0.10/0.23/0.67$ at $\alpha = 2/4/8°$, while the center of pressure (9–137 mm forward of the inviscid value) and the skin-friction drag (GCI = 1.4%) are both grid-converged.

Two directions would extend the validation envelope. Adding the NACA RM-10 parabolic body [44, 42], whose $M_\infty = 1.52$ and 1.98 pressure data bracket ARIA's 1.8, would test the method on a second parabolic (rather than ogival) contour, at the cost of boattail support in the mesh generator; and the fineness-3 nose-shape drag survey over $M_\infty = 1.24$–3.67 [43] would convert the present single-Mach validation into a Mach sweep. Together with a transitional turbulence treatment to close the aft-nose residual, these would carry the methodology from the present point-validated state toward a fully envelope-validated design tool for supersonic sounding-rocket forebodies.

Reproducibility

All results regenerate from the repository: mesh generation (`scripts/generate_mesh.py`, `generate_mesh3d.py`), inviscid runs (`scripts/run_study.sh`, `run_lts.sh`, `run_aoa.sh`), viscous runs (`scripts/run_rhopimple.sh`, `run_aoa_visc.sh`; all WSL OpenFOAM v2406), post-processing (`scripts/study_post.py`, `aoa_post.py`) and figures (`scripts/naca_compare.py`, `a53e01_compare.py`, `validation_summary.py`, `theory_overlay_ari` `gci_analysis.py`, `paper_figures.py`, `aoa_naca_compare.py`, `viscous_figures.py`; Appendix-B legacy renderings: `plot_extras.py`, `check_star.py` for Fig. 27). Experimental tables are transcribed in `data/*.csv` with the source PDFs alongside; the shareable $\alpha = 0$ surface-pressure dataset is `data/aria_alpha0_wall_pressure.csv`.

A Run matrix

Table 14: Complete run matrix. TA = time-accurate two-stage; LTS = local time stepping (§2.3). Hardware: Intel i9-13900H (10 physical cores), OpenFOAM v2406 under WSL2.

case	purpose	cells	integrator	limiter	ranks	wall time
`aria_baseline`	application baseline	11,800	TA	van Leer	1	∼42 min
`gci_coarse`	GCI grid 3	2,950	TA	van Leer	1	17 min
`gci_fine`	GCI grid 1	47,200	TA	van Leer	10	132 min
`naca_m198`	A54H23, under-resolved	9,910	TA	Minmod	8	26 min
`naca_fine`	A54H23 validation	26,490	TA	van Leer	10	54 min
`a53e01`	A53E01 validation	29,025	TA	van Leer	10	57 min
`lts_pilot`	LTS equivalence	2,950	LTS	van Leer	1	1.5 min
`tipref`	ARIA tip refinement	18,400	LTS	van Leer	10	7.0 min
`farfield`	domain independence	24,920	LTS	van Leer	10	7.1 min
`lim_minmod`	limiter sensitivity	11,800	LTS	Minmod	10	3.7 min
`flux_tadmor`	flux sensitivity	11,800	LTS	van Leer	10	3.8 min
`energy_h`	energy-form sensitivity	11,800	LTS	van Leer	10	2.0 min
`baseline_lts`	LTS sensitivity reference	11,800	LTS	van Leer	10	1.9 min
Viscous (§8), `rhoPimpleFoam` + LTS:						
`visc_naca_lam`	laminar validation	66,840	PIMPLE-LTS	n/a	14	37 min
`visc_aria_sst`	SST flight drag	27,200	PIMPLE-LTS	n/a	14	13 min
`aoa_visc_a2/4/8`	SST design AoA (×3)	441,600	PIMPLE-LTS	n/a	14	200–236 min
`visc_aria_sst_coarse`	friction-drag GCI	4,050	PIMPLE-LTS	n/a	14	4 min
`aoa_visc_a4_med`	design-load GCI (3-grid mid)	126,252	PIMPLE-LTS	n/a	14	55 min
`aoa_visc_a4_coarse`	design-load GCI (3-grid coarse)	35,400	PIMPLE-LTS	n/a	14	24 min
`aoa_visc_naca_a5`	AoA C_N validation	366,000	PIMPLE-LTS	n/a	14	200 min

B Supplementary figures

Flow-field visualizations of the ARIA fine-grid solution ($M_\infty = 1.8$, 47,200-cell grid, the same solution carried through §6). These complement the quantitative comparisons of the main text.

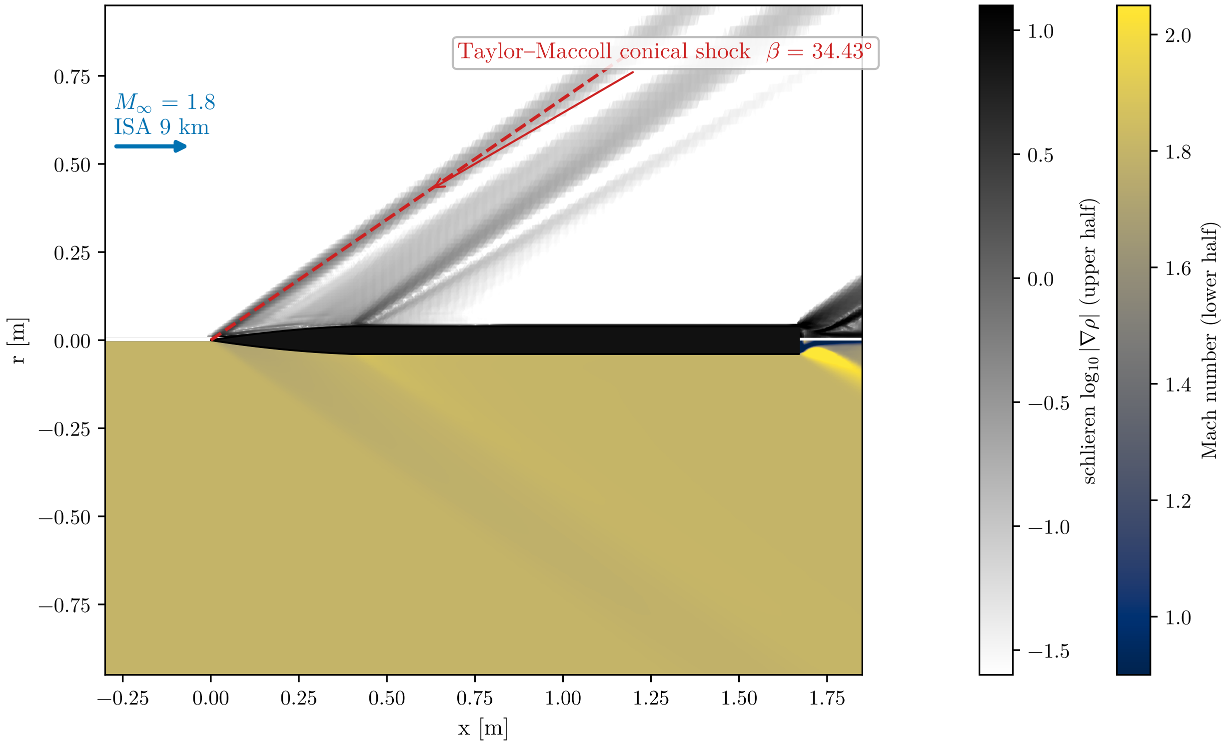

Figure 20: Composite full-body view of the fine-grid ARIA solution ($M_\infty = 1.8$, 47,200 cells). The top half is a numerical schlieren ($\log_{10}|\nabla\rho|$) and the bottom half the Mach-number field, mirrored across the axis so the shock structure and the velocity field can be read together. The exact Taylor–Maccoll conical-shock ray for the as-built 8.87° tip is overlaid; the attached tip shock, shoulder expansion fan and base recompression are all visible, the qualitative picture behind the quantitative tip adjudication of §6.2.

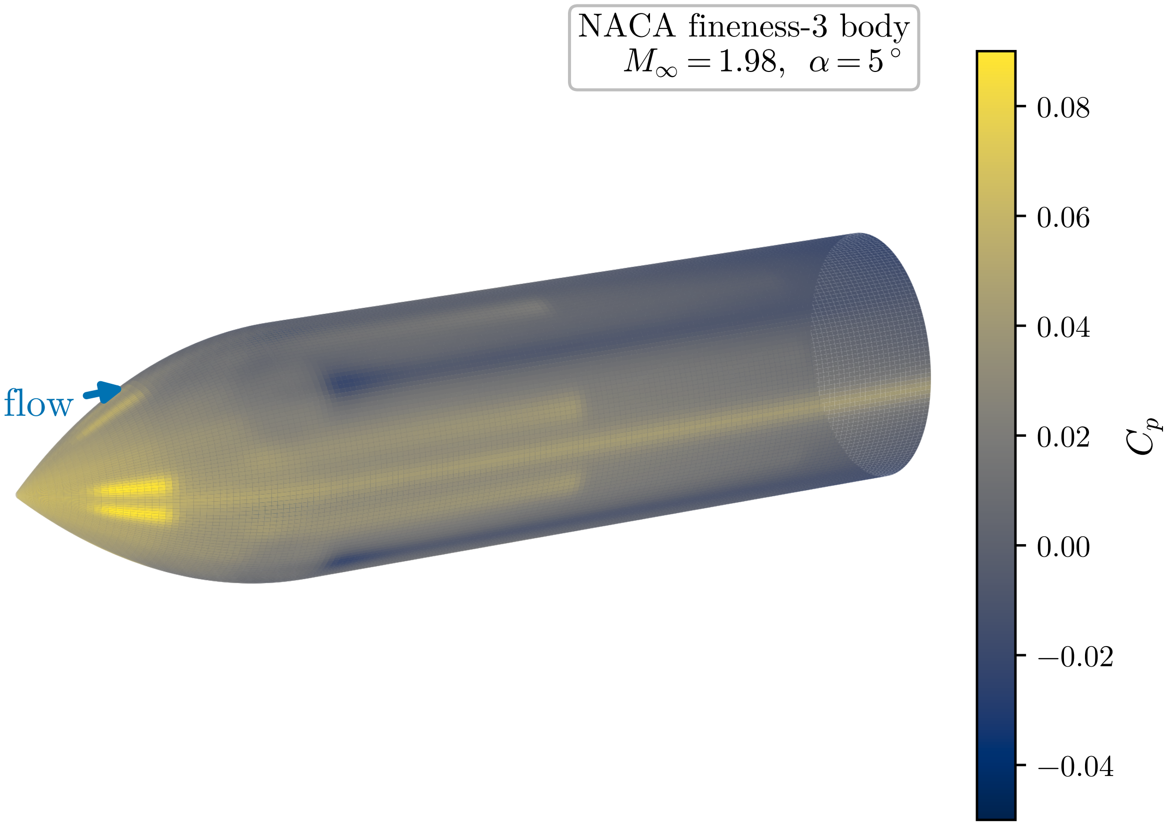

Figure 21: Three-dimensional surface-pressure (C_p) render of the fineness-3 validation body at incidence ($\alpha = 5°$, $M_\infty = 1.98$, 635,760-cell revolved half-model). The body surface is coloured by C_p interpolated onto a structured meridian-azimuth grid and lightly smoothed, with the windward (high-pressure) meridian rotated into view: the windward compression and the leeward expansion that generate the normal force are directly visible. This is the same solution validated meridian-by-meridian against measurement in §7 (Fig. 15).

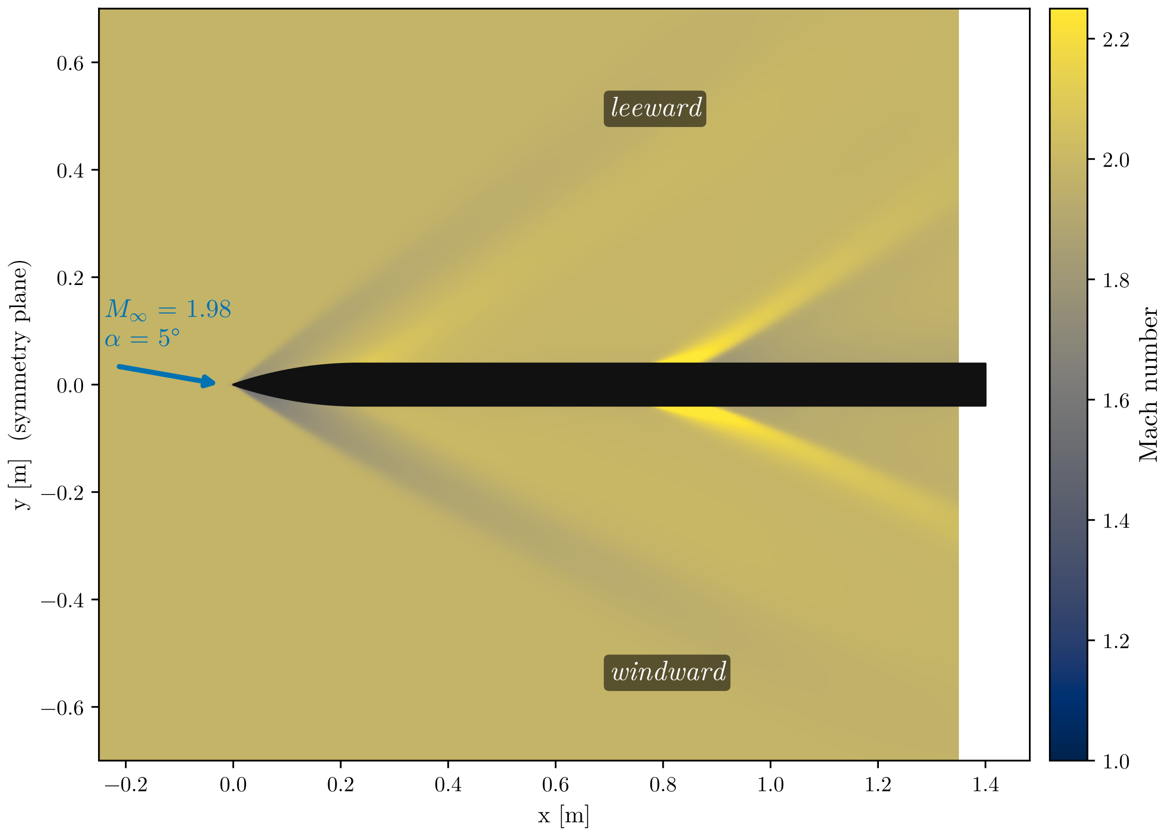

Figure 22: Symmetry-plane ($z = 0$) Mach field of the fineness-3 body at $\alpha = 5°$, $M_\infty = 1.98$. The windward tip shock (lower meridian) is stronger and lies at a steeper angle than the leeward one, and the leeward expansion is faster, the flow asymmetry that produces the circumferential pressure difference and hence the normal force. This field underlies the meridian pressure distributions of Fig. 15.

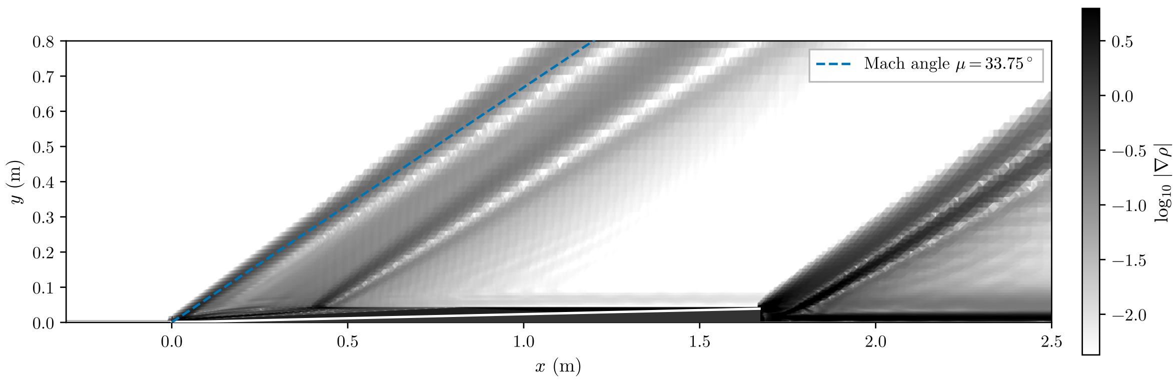

Figure 23: Numerical schlieren of the ARIA fine-grid solution ($M_\infty = 1.8$): the shaded field is $\log_{10}|\nabla\rho|$, with the display window clamped from the field median (freestream, white) to the strong-shock tail so that weak features remain visible. The attached conical tip shock, the ogive-shoulder expansion fan and the base recompression all read as crisp dark lines; the freestream Mach-angle ray $\mu = 33.75°$ is overlaid for reference, and the captured tip shock sits just inside it as expected for a finite-strength shock relaxing toward the Mach wave downstream.

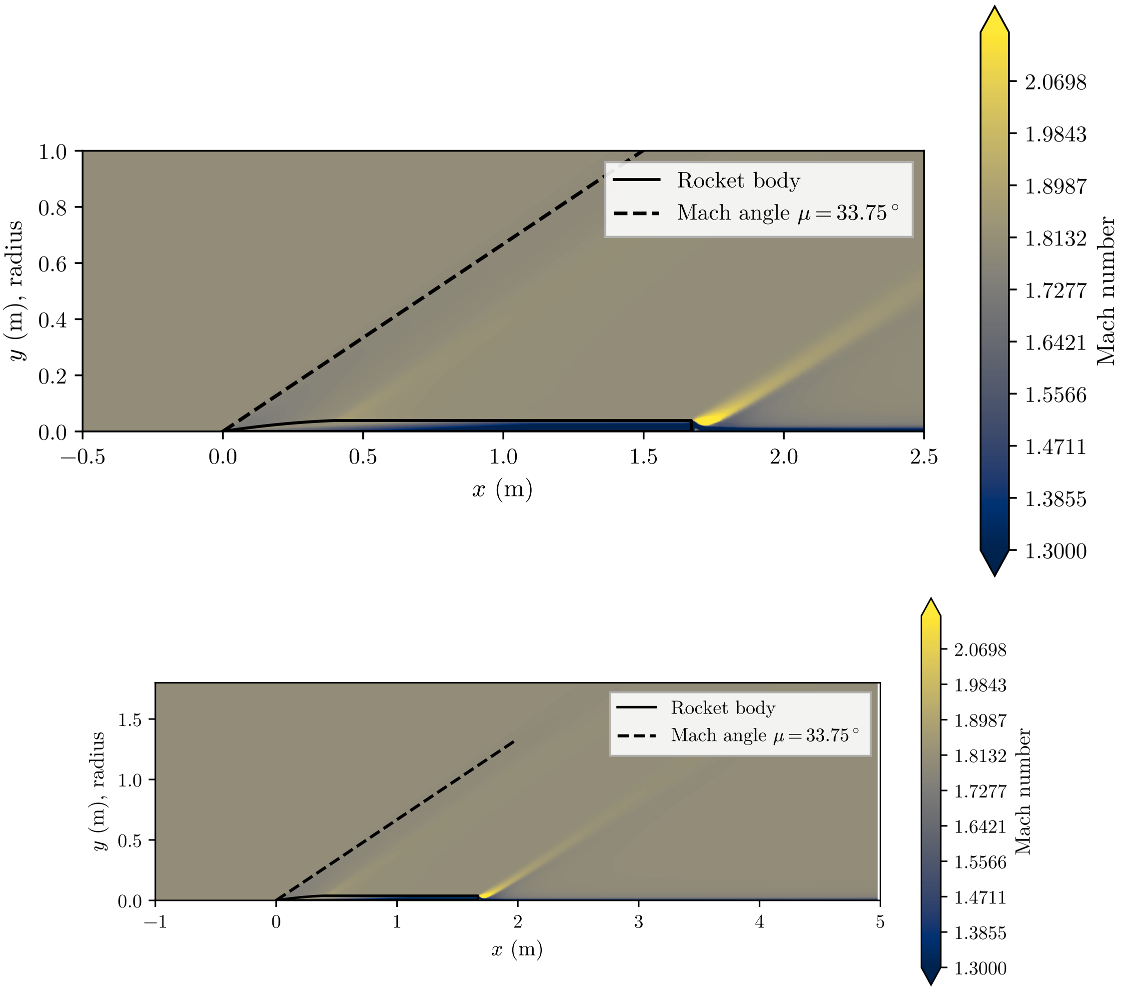

Figure 24: Mach-number contours of the ARIA fine-grid solution ($M_\infty = 1.8$, 47,200 cells): a near-body zoom (top) and the full domain (bottom). Both use a Mach window tightened around the supersonic field (1.3–2.15) so the conical tip shock (behind which M drops), the shoulder expansion ($M > 1.9$) and the base wake are distinguishable rather than washed to a single tone. The dashed line is the analytical freestream Mach angle $\mu = \arcsin(1/1.8) = 33.75°$, which the captured wave system tracks, an independent check on the wave geometry.

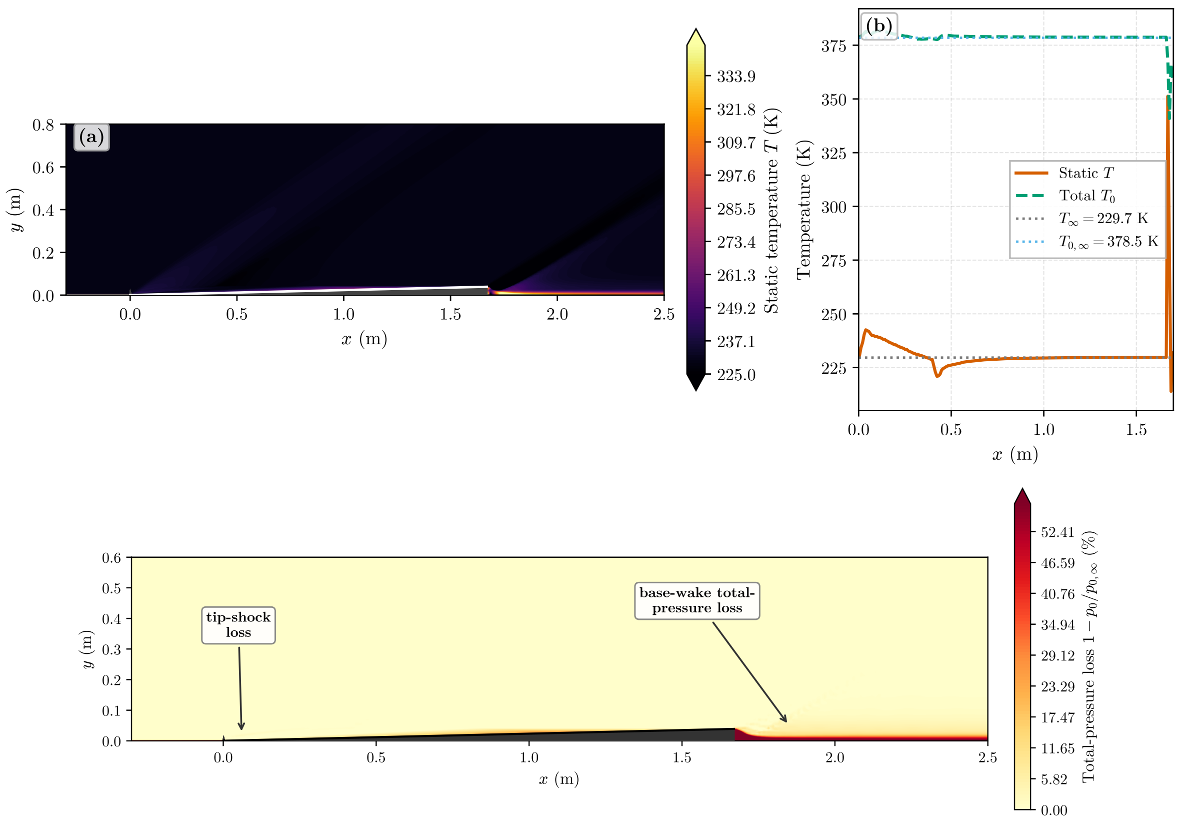

Figure 25: Thermodynamic fields of the ARIA fine-grid solution ($M_\infty = 1.8$). Top (Fig. a–b): the static-temperature field and the wall temperature distribution, with the static and total (T_0) temperatures and their freestream values ($T_\infty = 229.7\,\mathrm{K}$, $T_{0,\infty} = 378.5\,\mathrm{K}$) marked; the wall heats toward T_0 at the stagnation tip and recovers downstream. Bottom: the total-pressure *loss* $1 - p_0/p_{0,\infty}$ in percent, plotted on a sequential white-to-red scale clamped to the actual loss magnitude so the entropy-generating regions stand out, the tip shock and the separated base wake (both annotated) being the only places total pressure is lost.

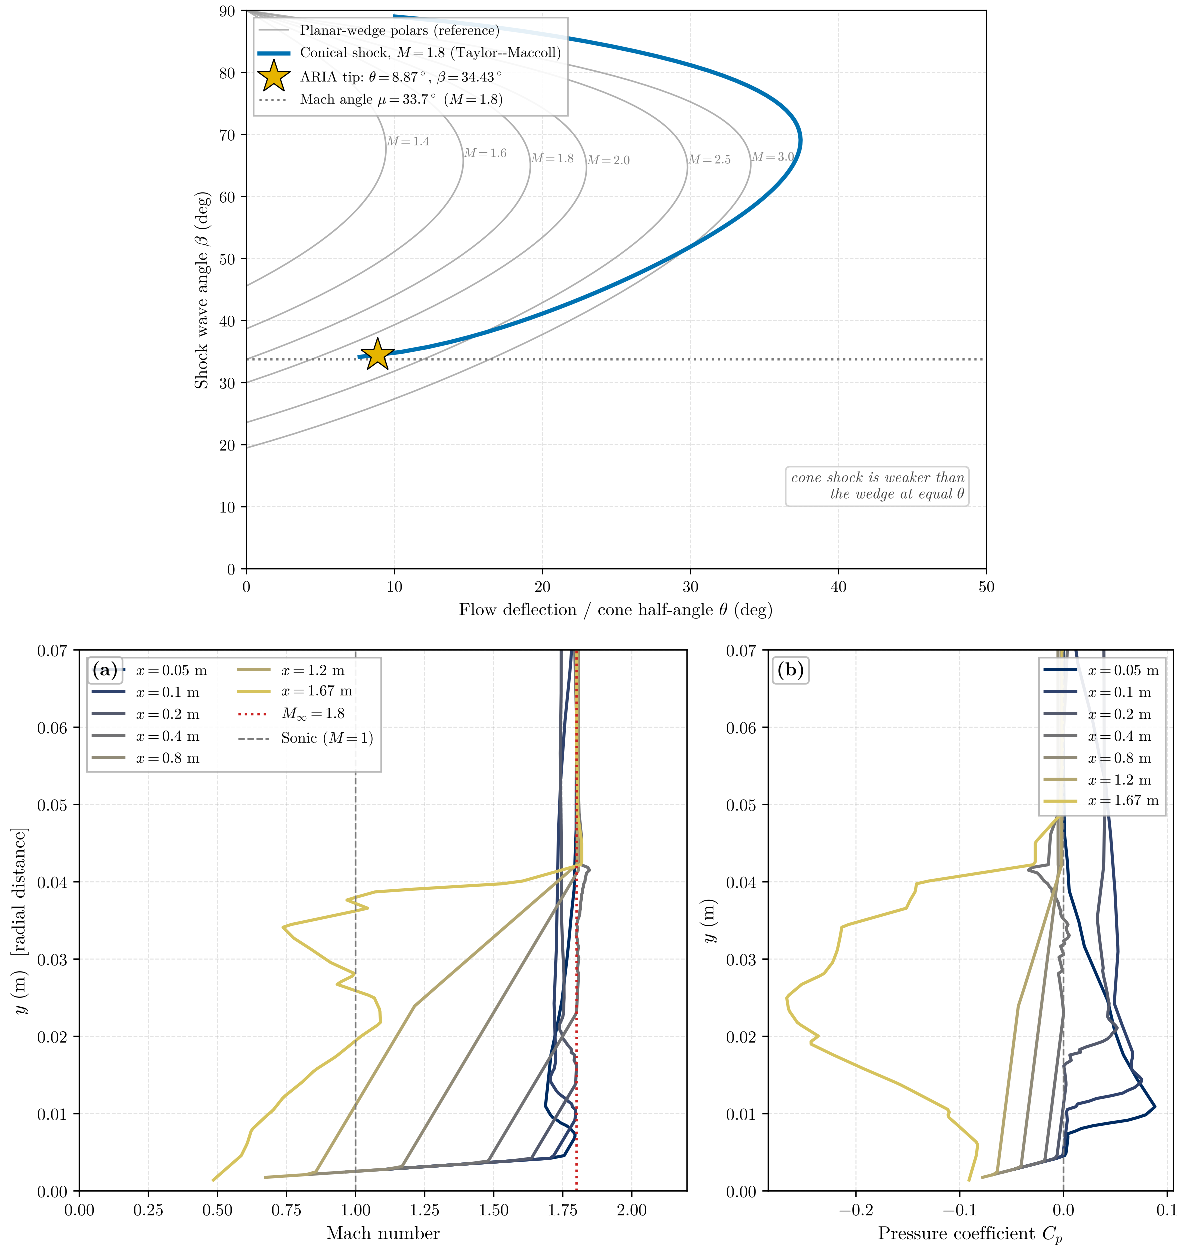

Figure 26: Top: shock-angle diagram for the ARIA tip at $M_\infty = 1.8$. The grey curves are the planar-wedge θ–β–M polars for a range of Mach numbers (reference only); the heavy blue curve is the *conical* θ_c–β locus at $M = 1.8$ computed by Taylor–Maccoll integration, which is the curve a slender axisymmetric nose actually follows. The gold star is the ARIA tip ($\theta = 8.87°$, $\beta = 34.43°$): it sits squarely on the conical curve, and below the wedge curves because a cone shock is weaker than a wedge shock at equal deflection. Bottom: radial profiles of Mach number (Fig. a) and C_p (Fig. b) at successive axial stations, with the radial axis zoomed to the near-wall shock-layer region ($y \leq 0.07$ m) where the profiles differ; above it every station has relaxed to the uniform freestream.

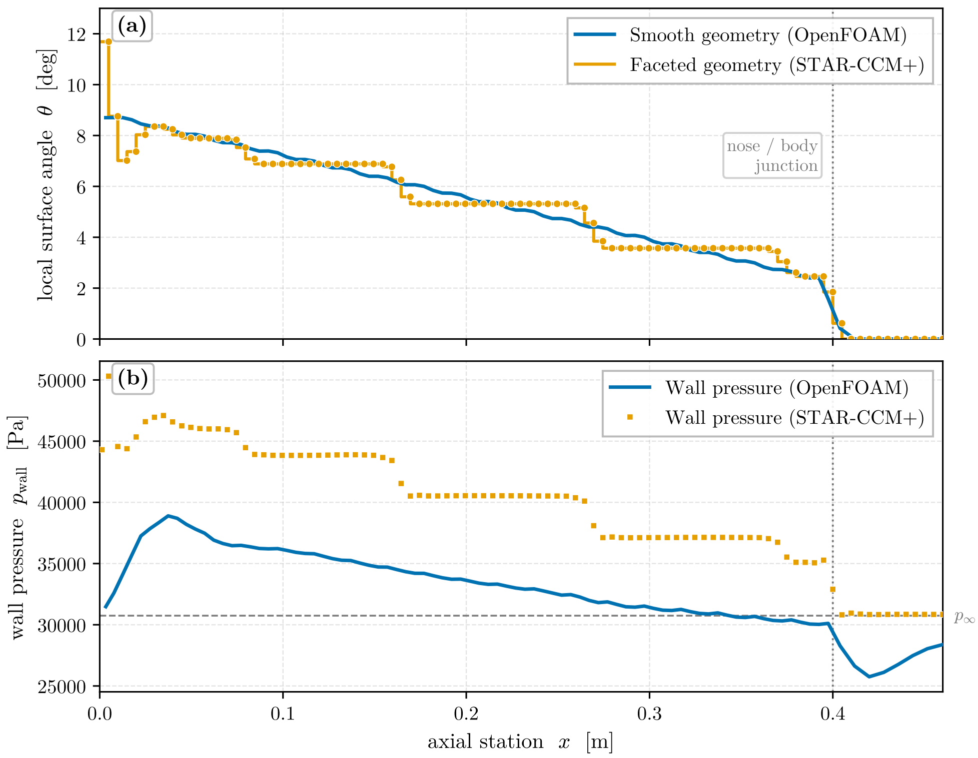

Figure 27: Evidence that the STAR-CCM+ benchmark geometry was faceted. The surface slope $\mathrm{d}r/\mathrm{d}x$ derived from the benchmark contour is piecewise-constant (stair-stepped) rather than smooth, and each slope step aligns with a pressure plateau in the benchmark wall C_p. The first facets act as locally steeper cones, which raises the tip-region pressure and explains the 2.8× over-prediction diagnosed in §6.2 and §6.3 without invoking any OpenFOAM deficiency.

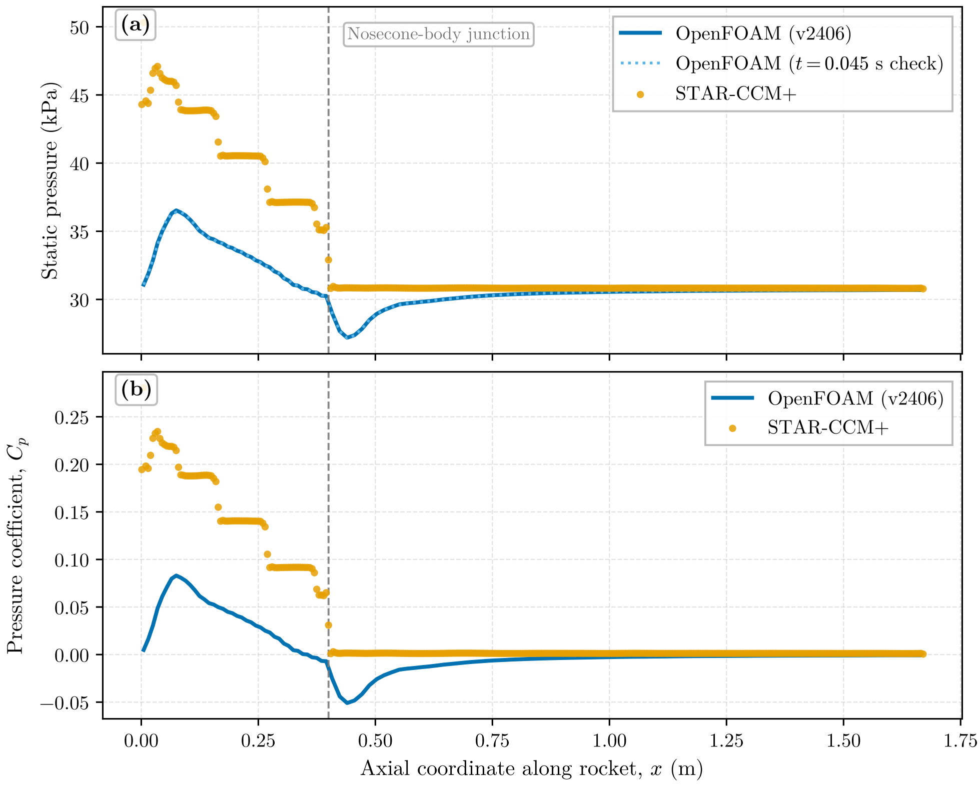

Figure 28: The original inter-code comparison that motivated this campaign: the OpenFOAM baseline wall pressure against the STAR-CCM+ benchmark, showing the 23.3% tip disagreement and the within-1.5% afterbody agreement described in §1. With neither code anchored to an external reference, this comparison could not say which was correct; it is superseded by the theory-anchored assessment of Fig. 13 and Table 6, which shows the OpenFOAM solution to be the faithful one.

C Per-run visual catalogue

This appendix presents, run by run, the mesh and flow-field visualizations for every case in the run matrix (Table 14), so that each quantitative result in the main text can be traced to the solution it came from. Two render types are used: a near-body Mach-number field (the structured wedge or its revolved meridian) for the steady inviscid and viscous fields, and a three-dimensional surface-C_p render for the angle-of-attack runs. The seven grids for which the volume mesh was retained additionally show the mesh itself alongside the field. All images are regenerated by `scripts/appendix_gallery.py` from the same archived solutions used throughout. The inviscid three-dimensional angle-of-attack runs of §7 are catalogued here for completeness although they post-date the tabulated inviscid matrix.

C.1 ARIA baseline and grid-convergence triplet

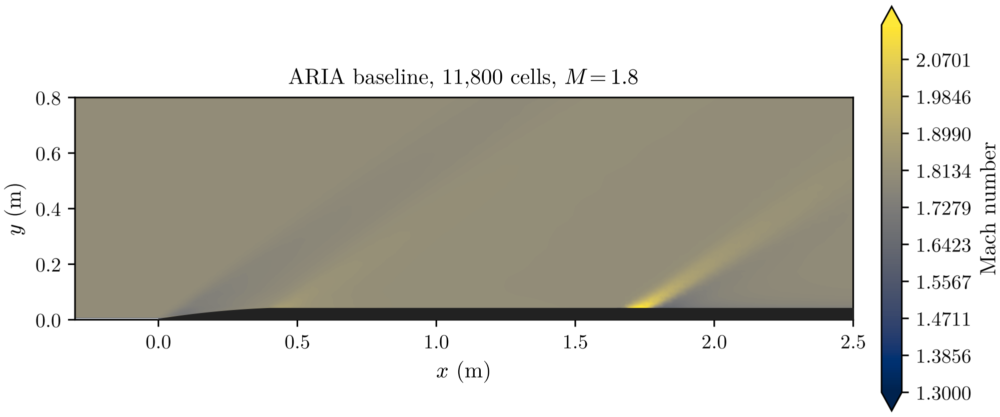

Figure 29: `aria_baseline` (11,800 cells, time-accurate, van Leer): the application baseline at $M_\infty = 1.8$. Near-body Mach field showing the attached tip shock, the shoulder expansion and the base wake. This is the grid on which the scheme- and energy-sensitivity studies of §4.4 are run.

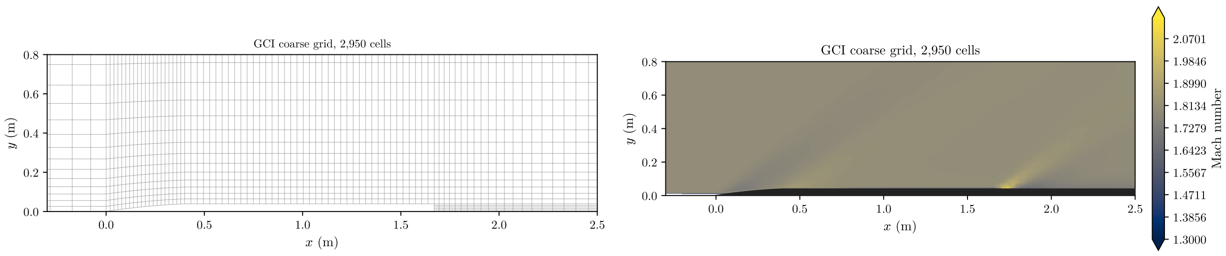

Figure 30: `gci_coarse` (2,950 cells), the coarse member of the GCI triplet (§4.2). Left: the structured wedge mesh; right: the Mach field, visibly more smeared at the tip shock than the finer grids below.

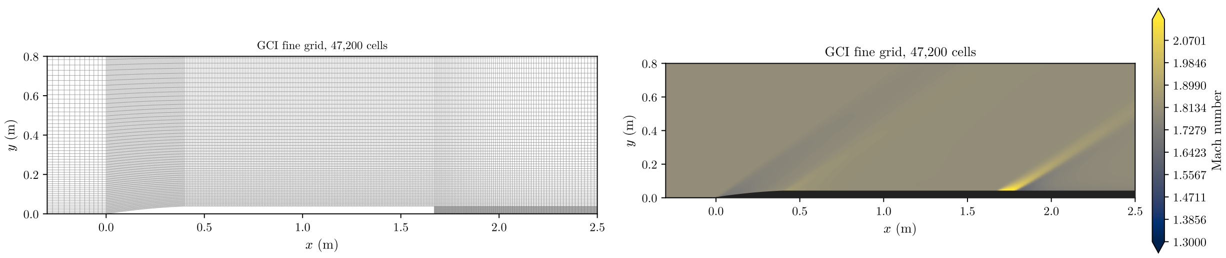

Figure 31: `gci_fine` (47,200 cells), the fine member of the GCI triplet and the solution carried through §6. Left: mesh; right: Mach field, with the sharpest captured tip shock of the triplet.

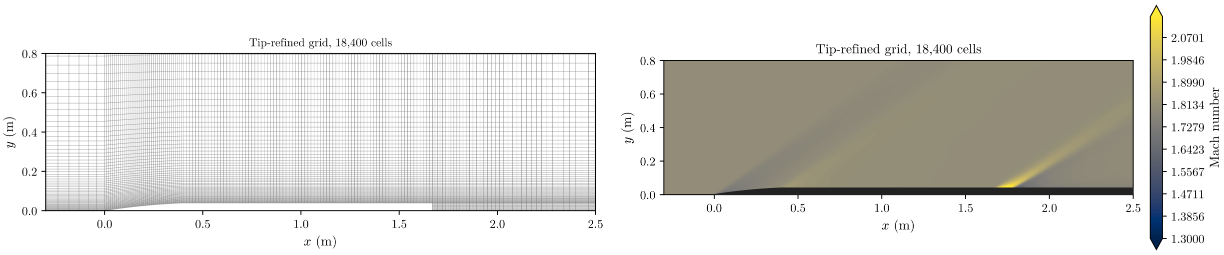

Figure 32: `tipref` (18,400 cells, near-wall grading only): the tip-refined grid that moves the first cell row to $\approx 1\,\mathrm{mm}$ off the wall and collapses the cell-centred/wall-extrapolated tip gap (§6.2). Left: mesh, showing the near-wall clustering; right: Mach field.

C.2 Verification and sensitivity runs

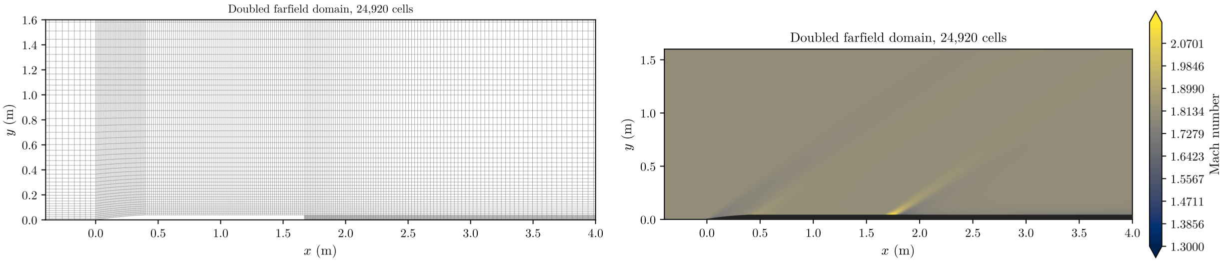

Figure 33: `farfield` (24,920 cells, LTS): the doubled-domain grid of the domain-independence test (§4.3). Left: the enlarged mesh ($x \in [-4, 8]\,\mathrm{m}$, $r \leq 10\,\mathrm{m}$); right: Mach field, which reproduces the baseline to $\overline{|\Delta C_p|} = 1.2 \times 10^{-4}$.

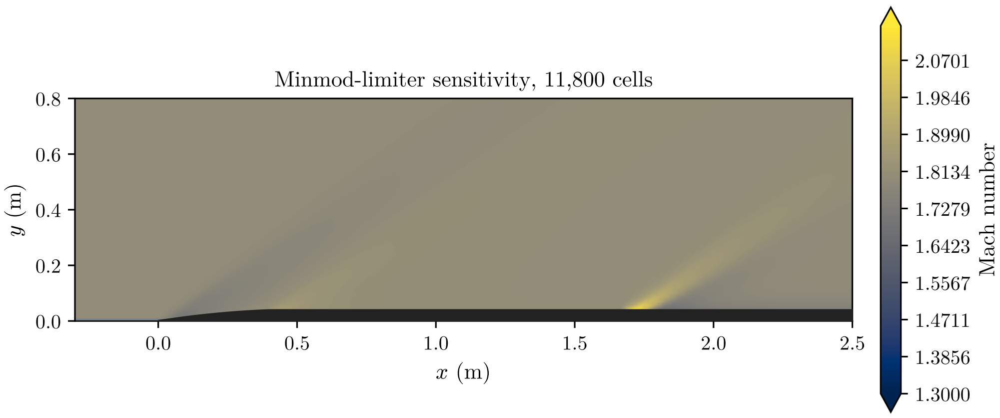

Figure 34: `lim_minmod` (11,800 cells, LTS, Minmod limiter): the limiter-sensitivity run of §4.4. The Mach field is visually indistinguishable from the van Leer baseline; the quantitative difference is $\overline{|\Delta C_p|} = 6.4 \times 10^{-4}$.

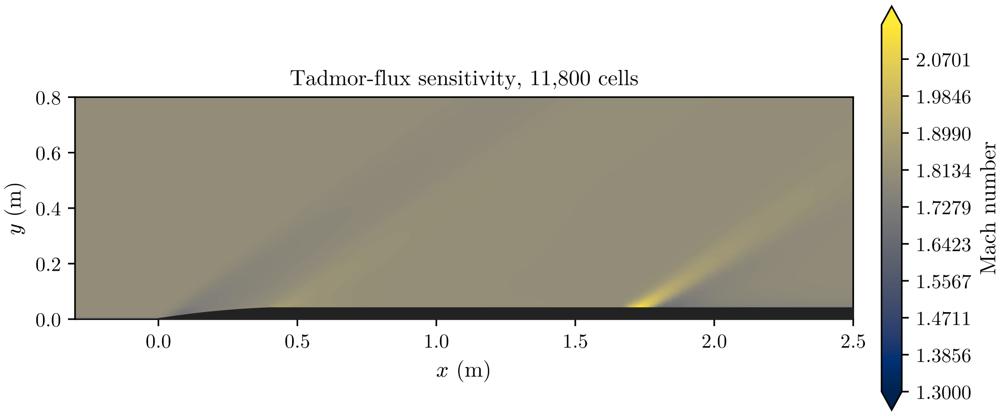

Figure 35: `flux_tadmor` (11,800 cells, LTS, Tadmor symmetric flux): the flux-variant sensitivity run. The field matches the Kurganov baseline to $\overline{|\Delta C_p|} = 6.4 \times 10^{-4}$.

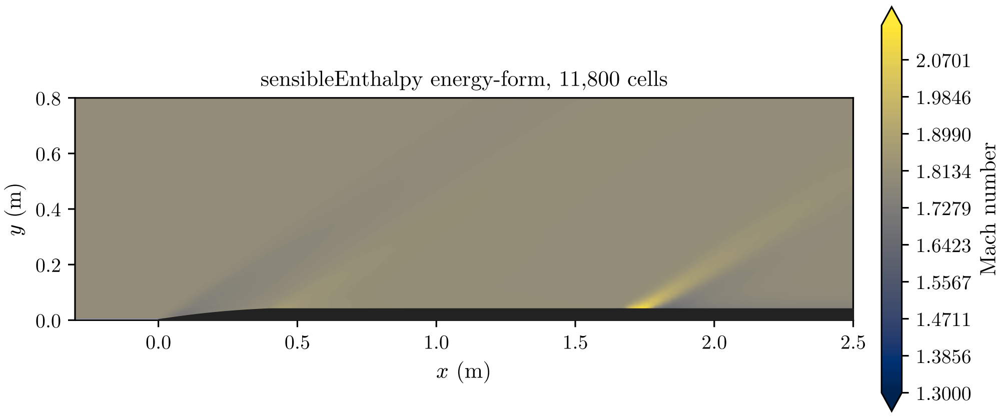

Figure 36: `energy_h` (11,800 cells, LTS, `sensibleEnthalpy`): the energy-form sensitivity run that closes the "energy solver" question empirically (−1% on nose drag, §4.4).

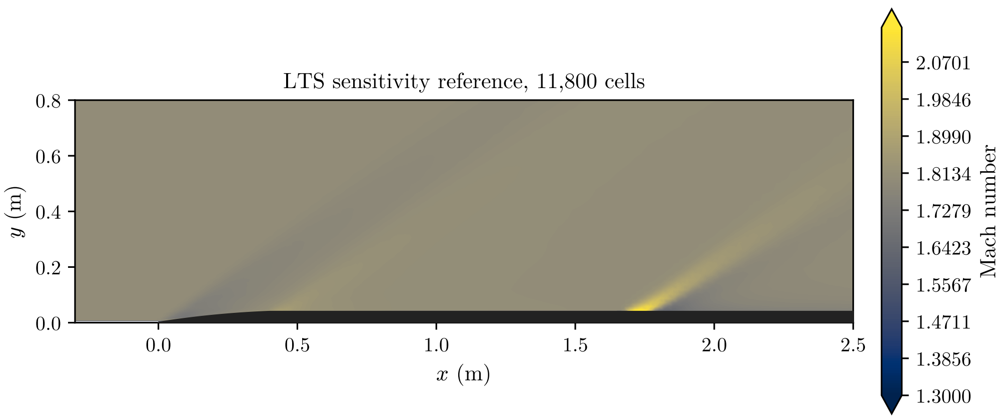

Figure 37: `baseline_lts` (11,800 cells, LTS, van Leer): the like-for-like LTS reference against which the three sensitivity runs above are differenced.

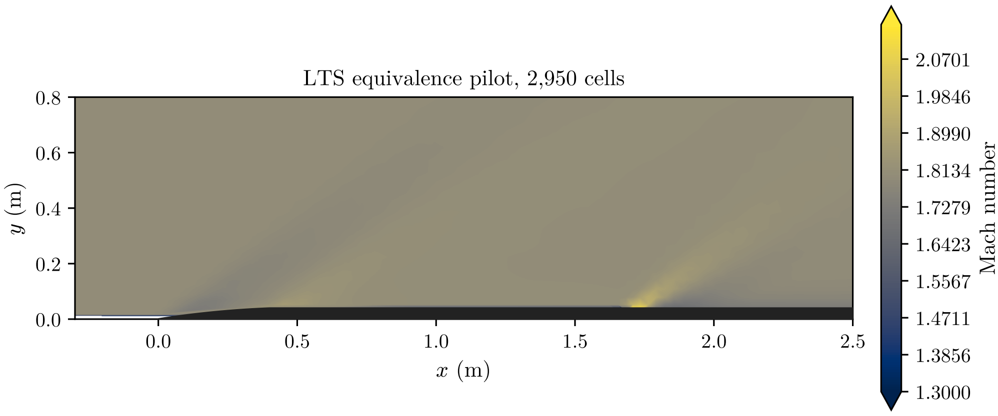

Figure 38: `lts_pilot` (2,950 cells, LTS): the coarse-grid pilot that established LTS equivalence to the time-accurate solution ($\overline{|\Delta C_p|} = 1.7 \times 10^{-4}$ at 12.7× lower cost, §4.5).

C.3 NACA validation bodies

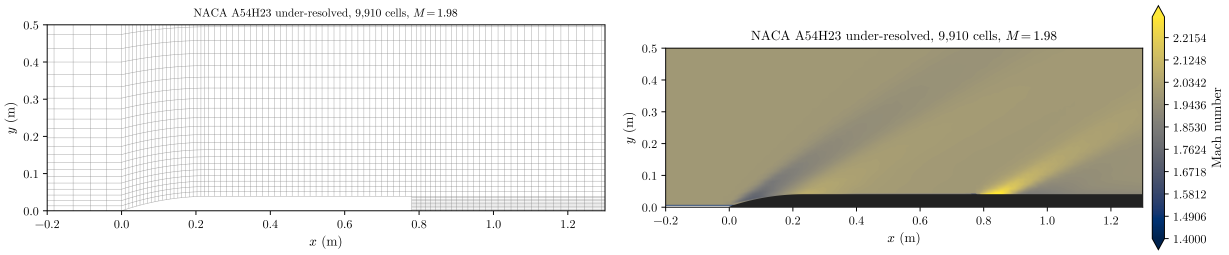

Figure 39: `naca_m198` (9,910 cells, time-accurate, Minmod): the deliberately under-resolved fineness-3 case (§5.3). Left: mesh; right: Mach field, in which the first cell row lies outside the thin conical shock layer, the informative failure that motivates the tip-resolved grid.

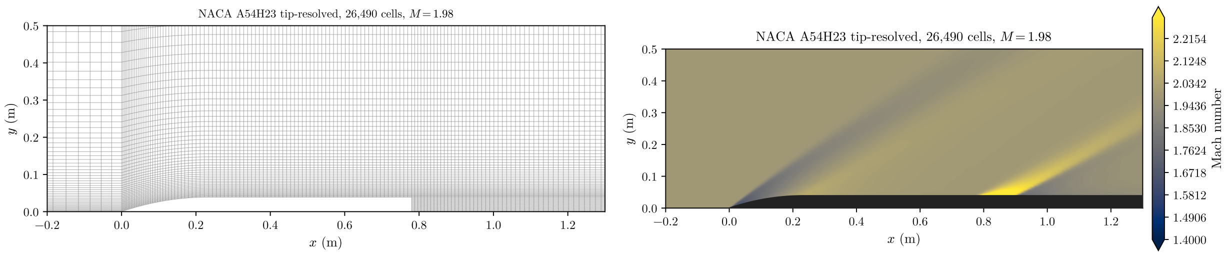

Figure 40: `naca_fine` (26,490 cells, time-accurate, van Leer): the tip-resolved fineness-3 validation grid. Left: mesh, with the refined tip-shock layer; right: Mach field. This run yields $\overline{|\Delta C_p|} = 0.0138$ against NACA RM A54H23.

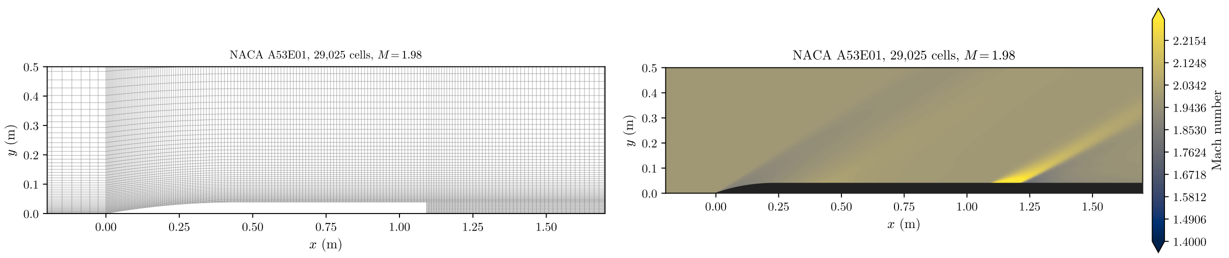

Figure 41: `a53e01` (29,025 cells, time-accurate, van Leer): the fineness-5.75 validation grid, the body closest to ARIA. Left: mesh; right: Mach field. This run agrees with NACA RM A53E01 to $\overline{|\Delta C_p|} = 0.0040$, within the experimental uncertainty.

C.4 Inviscid angle of attack (§7)

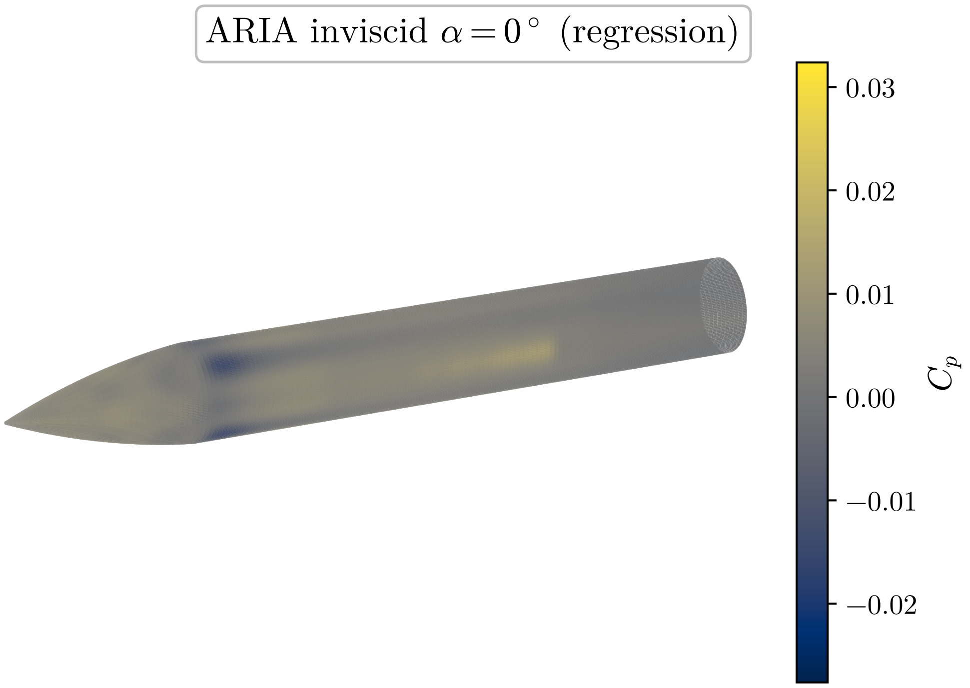

Figure 42: `aoa_a0` (283,200 cells, revolved 3-D, LTS): the $\alpha = 0$ regression gate. The surface-C_p render is a smooth, near-uniform field (axisymmetric, as it must be at $\alpha = 0$); quantitatively it reproduces the 2-D meridian to $\overline{|\Delta C_p|} = 0.0024$ (Fig. 14).

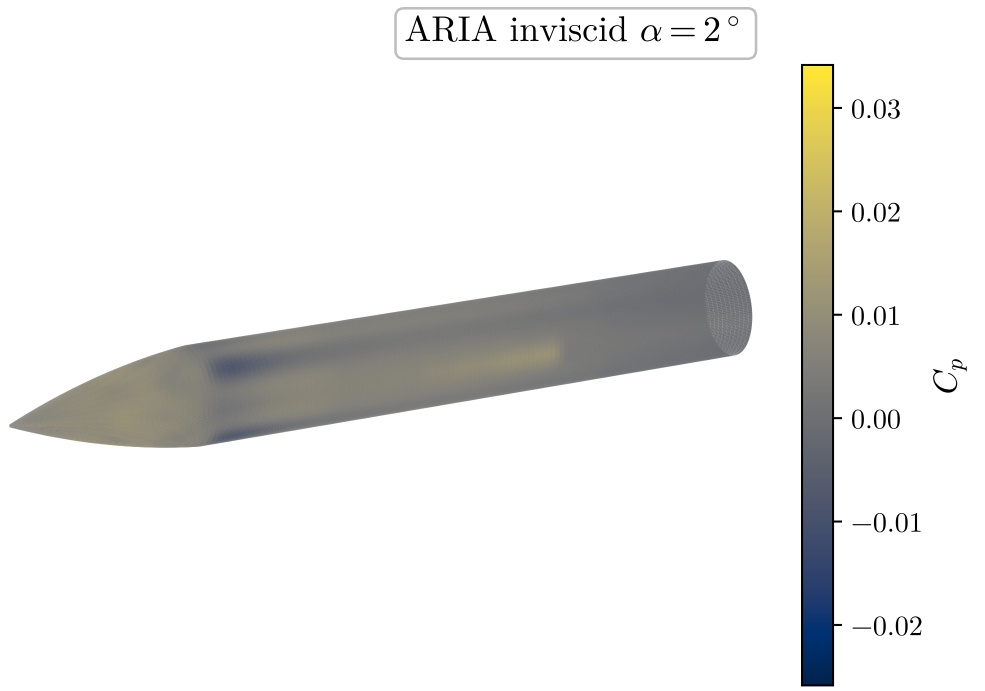

Figure 43: `aoa_a2` (283,200 cells): ARIA inviscid surface C_p at $\alpha = 2°$, the mild windward/leeward asymmetry just becoming visible.

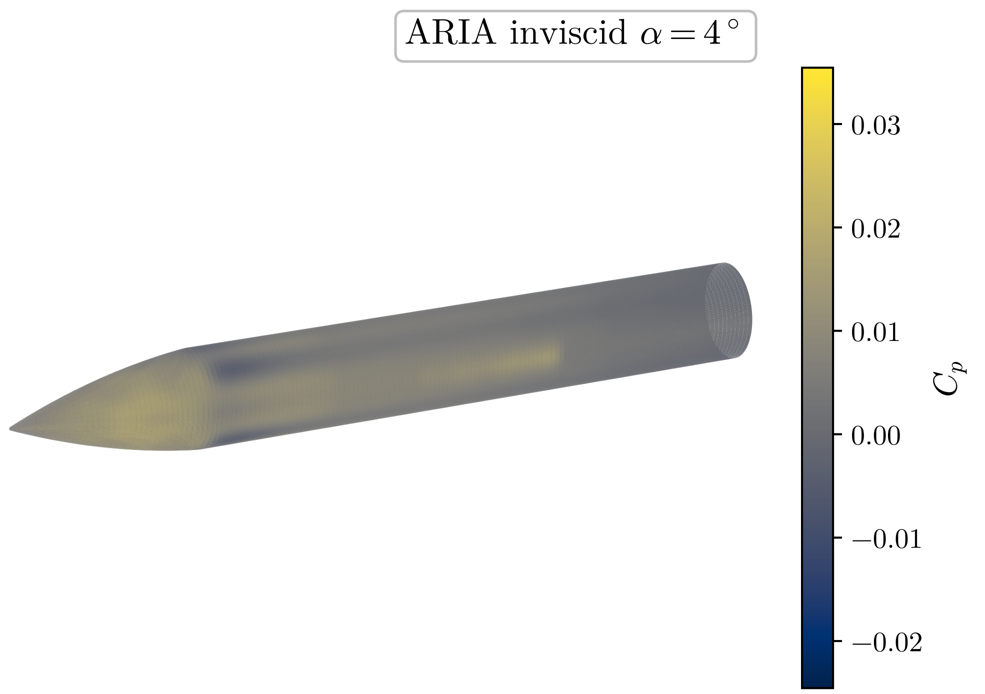

Figure 44: `aoa_a4` (283,200 cells): ARIA inviscid surface C_p at $\alpha = 4°$, windward compression and leeward expansion now clearly developed.

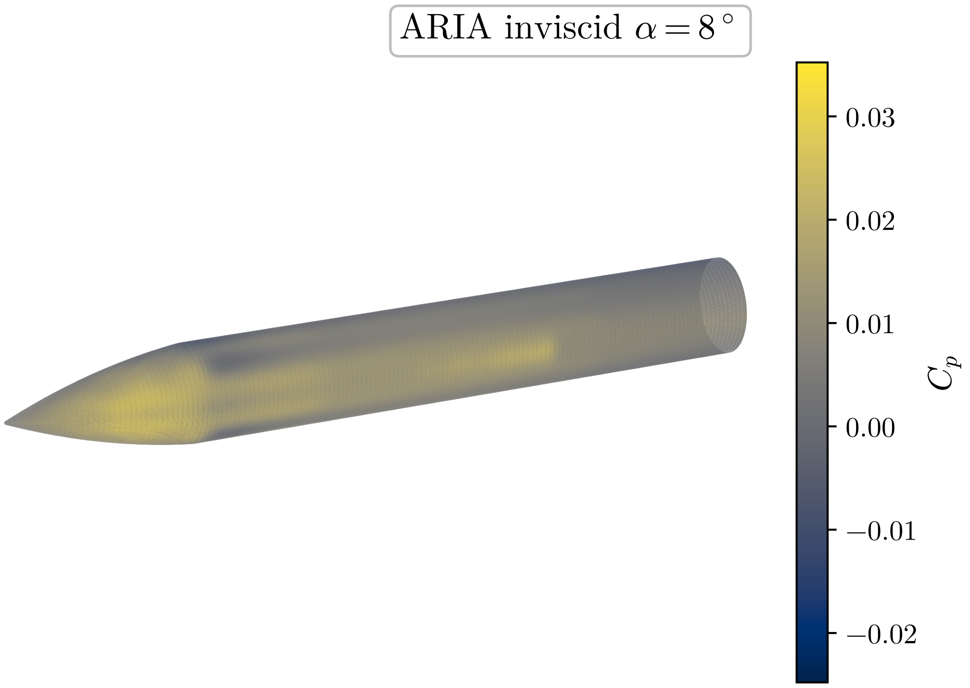

Figure 45: `aoa_a8` (283,200 cells): ARIA inviscid surface C_p at $\alpha = 8°$, the strongest incidence in the inviscid sweep (Table 7).

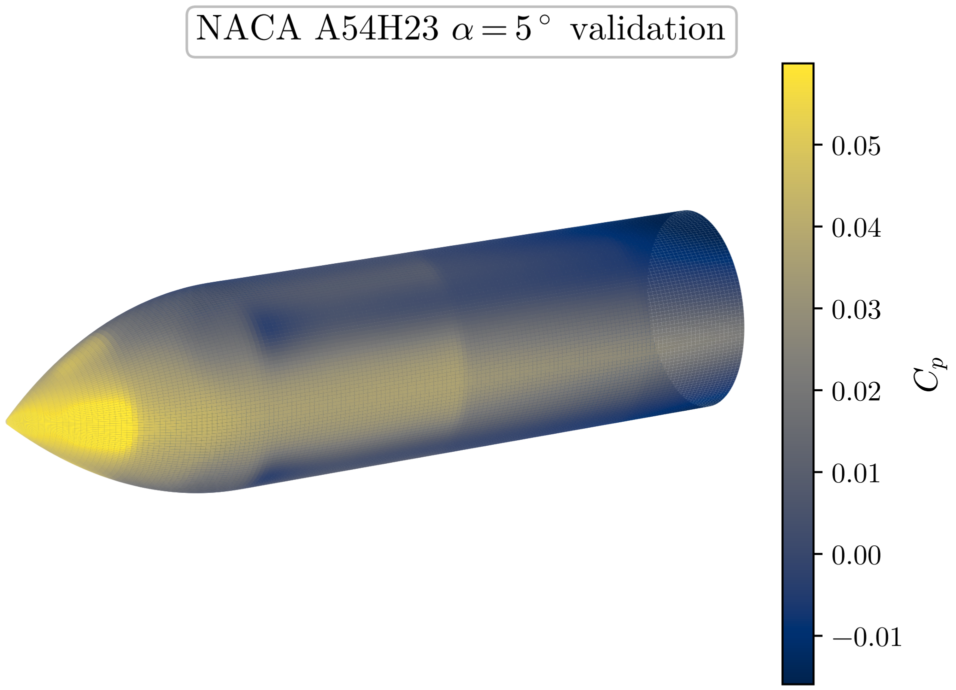

Figure 46: `aoa_naca_a5` (635,760 cells): the fineness-3 body at $\alpha = 5°$, the inviscid angle-of-attack validation case whose meridian and circumferential pressures are compared with NACA RM A54H23 in Fig. 15.

C.5 Viscous runs (§8)

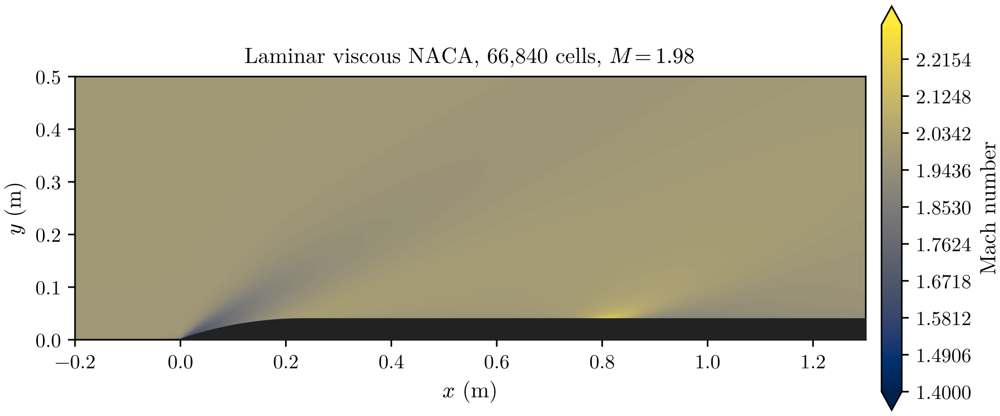

Figure 47: `visc_naca_lam` (66,840 cells, `rhoPimpleFoam`+LTS, laminar): the laminar fineness-3 validation (§8.2), boundary layer resolved to the wall ($y^+ < 3$). Near-body Mach field.

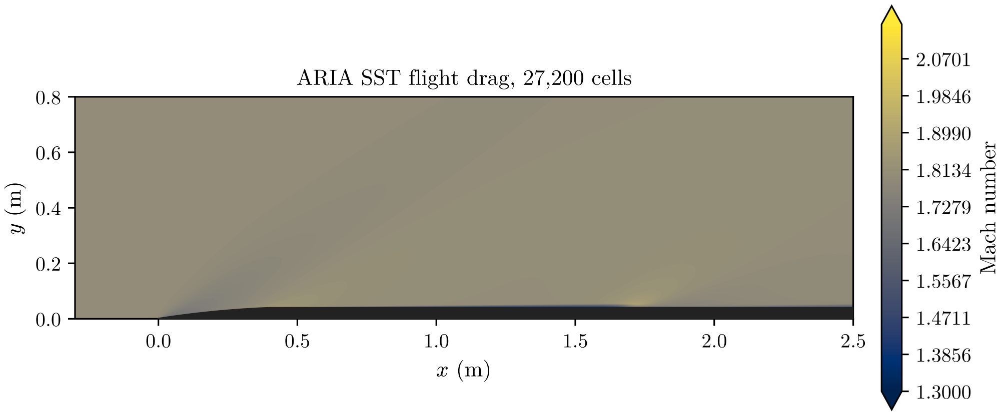

Figure 48: `visc_aria_sst` (27,200 cells, SST, $y^+ \approx 42$): the ARIA flight-drag run that resolves the 85%-skin-friction budget (§8.3). Near-body Mach field.

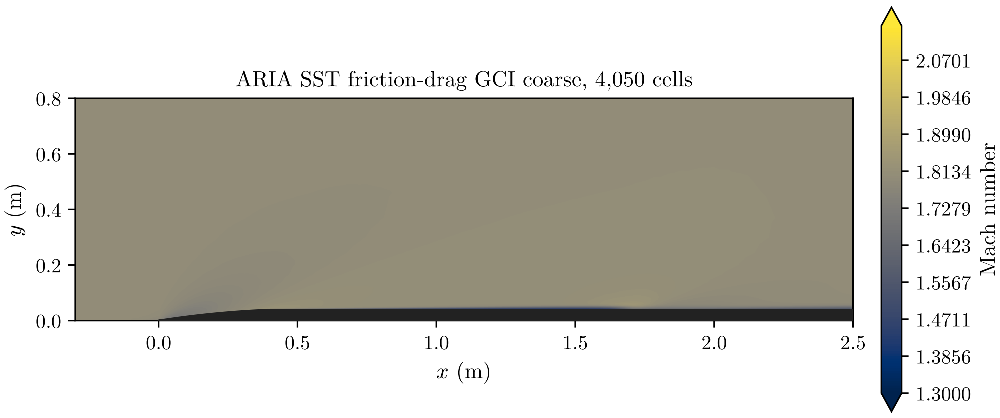

Figure 49: `visc_aria_sst_coarse` (4,050 cells, SST): the coarse mesh of the two-grid friction-drag check (GCI = 1.4%, §8.3).

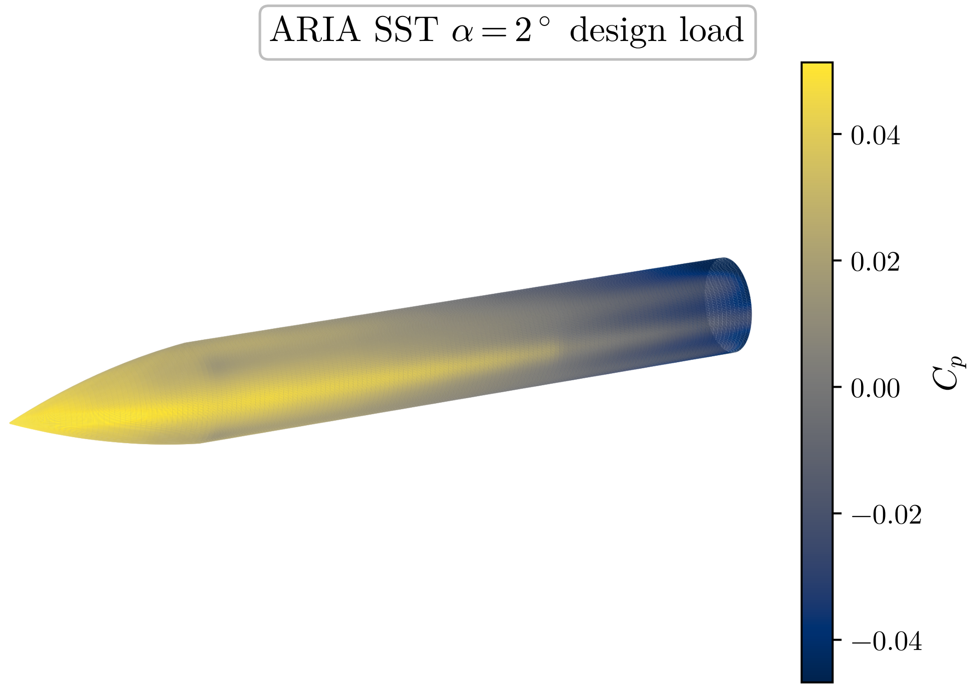

Figure 50: `aoa_visc_a2` (441,600 cells, SST): viscous design-load surface C_p at $\alpha = 2°$ (§8.4).

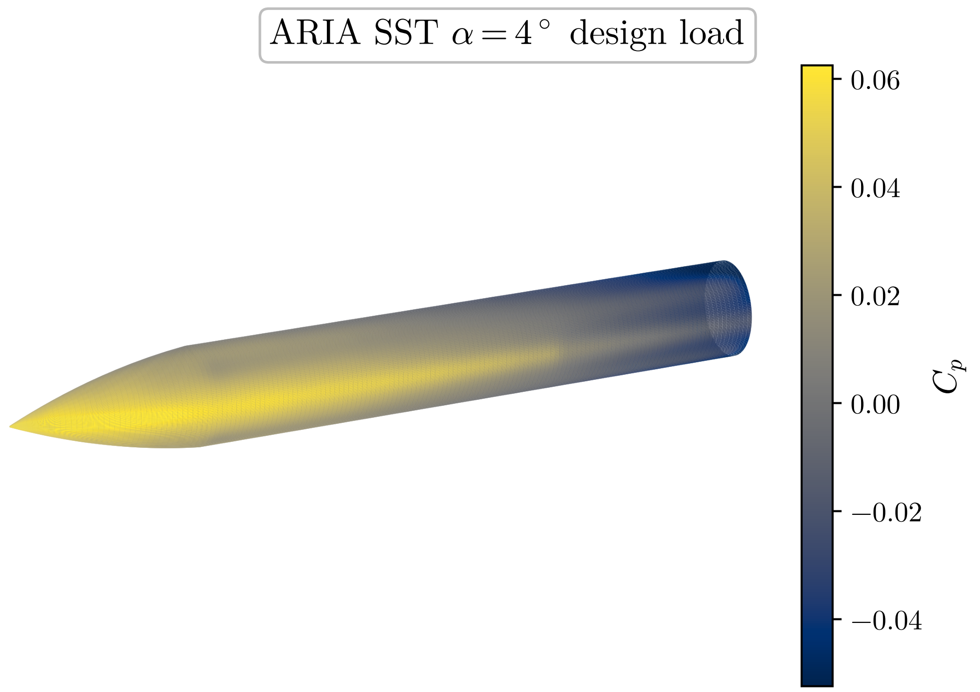

Figure 51: `aoa_visc_a4` (441,600 cells, SST): viscous design-load surface C_p at $\alpha = 4°$, the fine grid of the three-grid design-load study (§8.5).

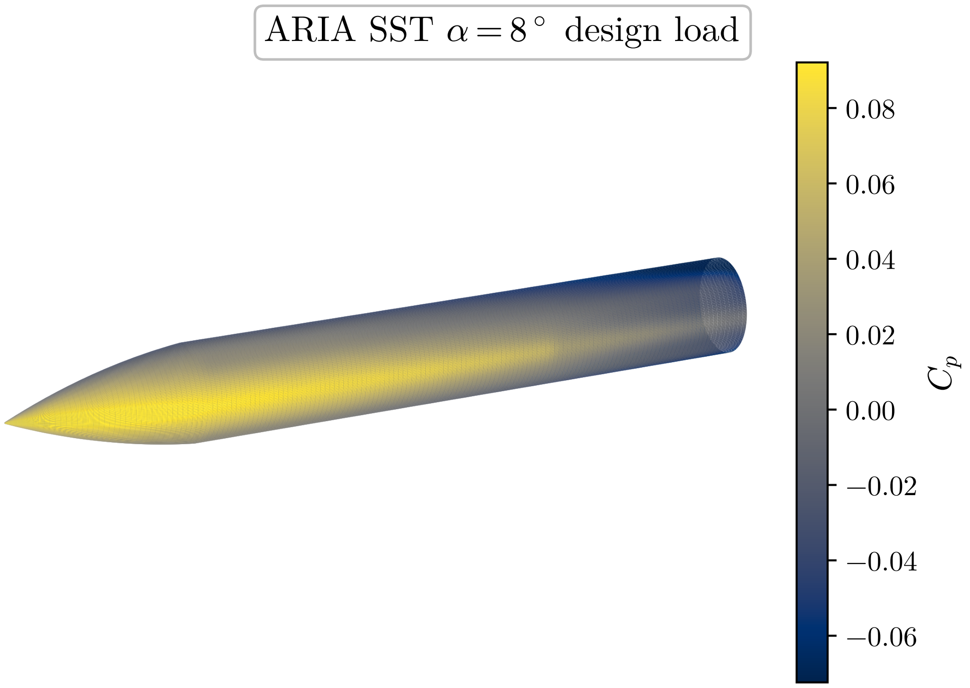

Figure 52: `aoa_visc_a8` (441,600 cells, SST): viscous design-load surface C_p at $\alpha = 8°$, where the viscous C_N reaches 90% of the inviscid bound.

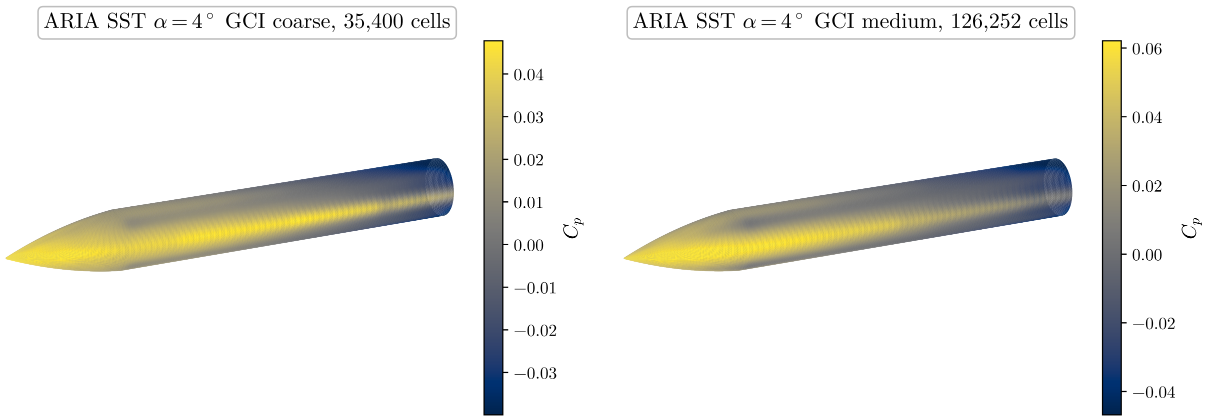

Figure 53: The two coarser grids of the $\alpha = 4°$ design-load GCI study (§8.5): `aoa_visc_a4_coarse` (35,400 cells, left) and `aoa_visc_a4_med` (126,252 cells, right), surface C_p. With the fine grid (Fig. 51) these give the monotone Richardson sequence $C_N = 0.332 \to 0.294 \to 0.270$.

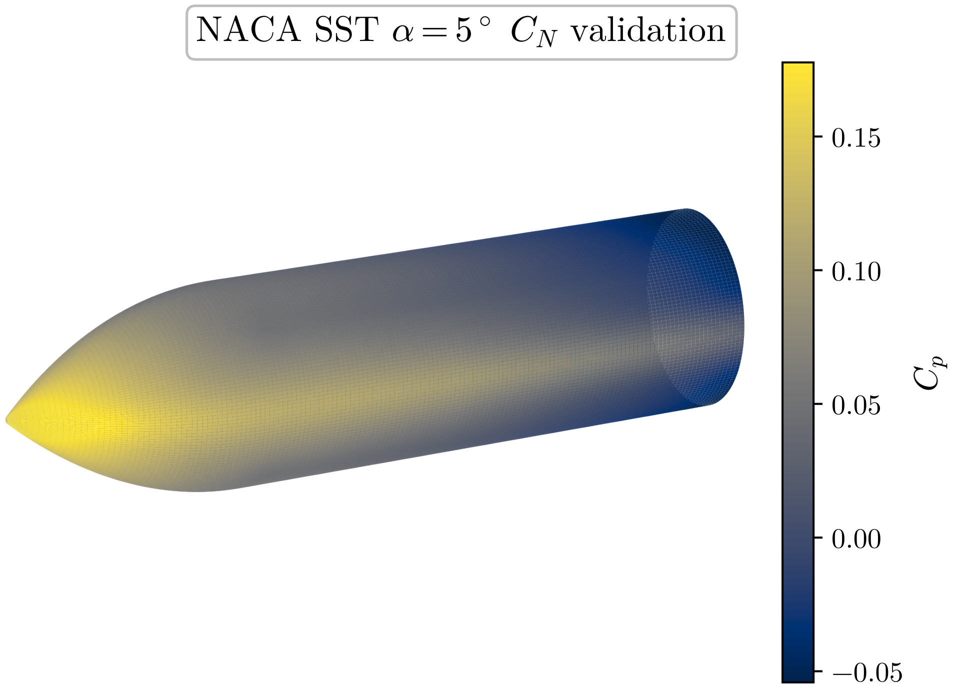

Figure 54: `aoa_visc_naca_a5` (366,000 cells, SST): the viscous fineness-3 body at $\alpha = 5^\circ$ used for the normal-force validation ($C_N = 0.300$ vs. measured 0.272, §8.6).

References

[1] AIAA, "Guide for the Verification and Validation of Computational Fluid Dynamics Simulations," AIAA G-077-1998(2002), 1998. doi:10.2514/4.472855.

[2] ASME, "Standard for Verification and Validation in Computational Fluid Dynamics and Heat Transfer," ASME V&V 20-2009 (reaffirmed 2016), 2009.

[3] Oberkampf, W. L., and Trucano, T. G., "Verification and validation in computational fluid dynamics," *Progress in Aerospace Sciences*, Vol. 38, No. 3, 2002, pp. 209–272. doi:10.1016/S0376-0421(02)00005-2.

[4] Roache, P. J., *Verification and Validation in Computational Science and Engineering*, Hermosa Publishers, Albuquerque, 1998.

[5] Roache, P. J., "Quantification of Uncertainty in Computational Fluid Dynamics," *Annual Review of Fluid Mechanics*, Vol. 29, 1997, pp. 123–160. doi:10.1146/annurev.fluid.29.1.123.

[6] Roache, P. J., "Perspective: A Method for Uniform Reporting of Grid Refinement Studies," *ASME Journal of Fluids Engineering*, Vol. 116, No. 3, 1994, pp. 405–413. doi:10.1115/1.2910291.

[7] Oberkampf, W. L., and Roy, C. J., *Verification and Validation in Scientific Computing*, Cambridge University Press, Cambridge, UK, 2010. doi:10.1017/CBO9780511760396.

[8] Stern, F., Wilson, R. V., Coleman, H. W., and Paterson, E. G., "Comprehensive Approach to Verification and Validation of CFD Simulations: Part 1, Methodology and Procedures," *ASME Journal of Fluids Engineering*, Vol. 123, No. 4, 2001, pp. 793–802. doi:10.1115/1.1412235.

[9] Richardson, L. F., "The Approximate Arithmetical Solution by Finite Differences of Physical Problems Involving Differential Equations, with an Application to the Stresses in a Masonry Dam," *Philosophical Transactions of the Royal Society of London A*, Vol. 210, 1911, pp. 307–357. doi:10.1098/rsta.1911.0009.

[10] Roy, C. J., "Review of Code and Solution Verification Procedures for Computational Simulation," *Journal of Computational Physics*, Vol. 205, No. 1, 2005, pp. 131–156. doi:10.1016/j.jcp.2004.10.036.

[11] Roy, C. J., and Oberkampf, W. L., "A Comprehensive Framework for Verification, Validation, and Uncertainty Quantification in Scientific Computing," *Computer Methods in Applied Mechanics and Engineering*, Vol. 200, No. 25–28, 2011, pp. 2131–2144. doi:10.1016/j.cma.2011.03.016.

[12] Celik, I. B., Ghia, U., Roache, P. J., Freitas, C. J., Coleman, H., and Raad, P. E., "Procedure for Estimation and Reporting of Uncertainty Due to Discretization in CFD Applications," *ASME Journal of Fluids Engineering*, Vol. 130, No. 7, 2008, 078001. doi:10.1115/1.2960953.

[13] Greenshields, C. J., Weller, H. G., Gasparini, L., and Reese, J. M., "Implementation of semi-discrete, non-staggered central schemes in a colocated, polyhedral, finite volume framework, for high-speed viscous flows," *International Journal for Numerical Methods in Fluids*, Vol. 63, No. 1, 2010, pp. 1–21. doi:10.1002/fld.2069.

[14] Gutiérrez Marcantoni, L. F., Tamagno, J. P., and Elaskar, S. A., "High Speed Flow Simulation Using OpenFOAM," *Mecánica Computacional*, Vol. XXXI, No. 16, 2012, pp. 2939–2959. (CIMEC).

[15] Bondarev, A. E., and Kuvshinnikov, A. E., "Analysis of the Accuracy of OpenFOAM Solvers for the Problem of Supersonic Flow Around a Cone," *Computational Science – ICCS 2018*, Lecture Notes in Computer Science, Vol. 10862, Springer, Cham, 2018, pp. 221–230. doi:10.1007/978-3-319-93713-7_18.

[16] Stoldt, H. H., Korobenko, A., Ziadé, P., and Johansen, C. T., "Verification and Validation of High-Fidelity Open-Source Simulation Tools for Supersonic Aircraft Aerodynamic Analysis," *ASME Journal of Verification, Validation and Uncertainty Quantification*, Vol. 6, No. 4, 2021, 041005. doi:10.1115/1.4052150.

[17] Lorenzon, D., and Elaskar, S., "OpenFOAM Simulations of the Supersonic Flow Around Cones at Angles of Attack," *IOSR Journal of Mechanical and Civil Engineering*, Vol. 16, No. 5, 2019, pp. 66–80. doi:10.9790/1684-1605036680.

[18] Zang, B., Vevek, U. S., Lim, H. D., Wei, X., and New, T. H., "An Assessment of OpenFOAM Solver on RANS Simulations of Round Supersonic Free Jets," *Journal of Computational Science*, Vol. 28, 2018, pp. 18–31. doi:10.1016/j.jocs.2018.07.002.

[19] Ray, A. A., Bhaduri, S., Das, S., and De, A., "Assessment of Modern Shock Capturing Schemes for All-Speed Flows in the OpenFOAM Framework," arXiv preprint 2510.24146, 2025. doi:10.48550/arXiv.2510.24146.

[20] Espinoza, D., Scanlon, T., and Brown, R., "Validation of Tools to Accelerate High-Speed CFD Simulations Using OpenFOAM," AIAA Paper 2015-3566, 20th AIAA International Space Planes and Hypersonic Systems and Technologies Conference, Glasgow, July 2015. doi:10.2514/6.2015-3566.

[21] Kurganov, A., and Tadmor, E., "New High-Resolution Central Schemes for Nonlinear Conservation Laws and Convection–Diffusion Equations," *Journal of Computational Physics*, Vol. 160, No. 1, 2000, pp. 241–282. doi:10.1006/jcph.2000.6459.

[22] Kurganov, A., Noelle, S., and Petrova, G., "Semidiscrete Central-Upwind Schemes for Hyperbolic Conservation Laws and Hamilton–Jacobi Equations," *SIAM Journal on Scientific Computing*, Vol. 23, No. 3, 2001, pp. 707–740. doi:10.1137/S1064827500373413.

[23] van Leer, B., "Towards the ultimate conservative difference scheme. II. Monotonicity and conservation combined in a second-order scheme," *Journal of Computational Physics*, Vol. 14, No. 4, 1974, pp. 361–370. doi:10.1016/0021-9991(74)90019-9.

[24] Harten, A., "High Resolution Schemes for Hyperbolic Conservation Laws," *Journal of Computational Physics*, Vol. 49, No. 2, 1983, pp. 357–393. doi:10.1016/0021-9991(83)90136-5.

[25] Sweby, P. K., "High Resolution Schemes Using Flux Limiters for Hyperbolic Conservation Laws," *SIAM Journal on Numerical Analysis*, Vol. 21, No. 5, 1984, pp. 995–1011. doi:10.1137/0721062.

[26] Quirk, J. J., "A Contribution to the Great Riemann Solver Debate," *International Journal for Numerical Methods in Fluids*, Vol. 18, No. 6, 1994, pp. 555–574. doi:10.1002/fld.1650180603.

[27] Einfeldt, B., Munz, C. D., Roe, P. L., and Sjögreen, B., "On Godunov-Type Methods Near Low Densities," *Journal of Computational Physics*, Vol. 92, No. 2, 1991, pp. 273–295. doi:10.1016/0021-9991(91)90211-3.

[28] Barth, T. J., and Jespersen, D. C., "The Design and Application of Upwind Schemes on Unstructured Meshes," AIAA Paper 89-0366, 27th AIAA Aerospace Sciences Meeting, Reno, NV, Jan. 1989. doi:10.2514/6.1989-366.

[29] Jameson, A., Schmidt, W., and Turkel, E., "Numerical Solution of the Euler Equations by Finite Volume Methods Using Runge–Kutta Time-Stepping Schemes," AIAA Paper 81-1259, 14th AIAA Fluid and Plasma Dynamics Conference, Palo Alto, CA, June 1981. doi:10.2514/6.1981-1259.

[30] Blazek, J., *Computational Fluid Dynamics: Principles and Applications*, 3rd ed., Butterworth-Heinemann (Elsevier), 2015, Chap. 9. ISBN 978-0-08-099995-1.

[31] Taylor, G. I., and Maccoll, J. W., "The Air Pressure on a Cone Moving at High Speeds. I," *Proceedings of the Royal Society of London A*, Vol. 139, No. 838, 1933, pp. 278–297. doi:10.1098/rspa.1933.0017.

[32] Syvertson, C. A., and Dennis, D. H., "A Second-Order Shock-Expansion Method Applicable to Bodies of Revolution Near Zero Lift," NACA Report 1328, 1957. NTRS 19930091011.

[33] Eggers, A. J., Jr., and Savin, R. C., "Approximate Methods for Calculating the Flow About Nonlifting Bodies of Revolution at High Supersonic Airspeeds," NACA TN 2579, 1951. NTRS 19930083341.

[34] Van Dyke, M. D., "First- and Second-Order Theory of Supersonic Flow Past Bodies of Revolution," *Journal of the Aeronautical Sciences*, Vol. 18, No. 3, 1951, pp. 161–178. doi:10.2514/8.1896.

[35] Van Dyke, M. D., "Practical Calculation of Second-Order Supersonic Flow Past Nonlifting Bodies of Revolution," NACA TN 2744, 1952. NTRS 19930083480.

[36] Ferri, A., "The Method of Characteristics for the Determination of Supersonic Flow Over Bodies of Revolution at Small Angles of Attack," NACA Report 1044, 1951. NTRS 19930090959.

[37] Ward, G. N., "Supersonic Flow Past Slender Pointed Bodies," *Quarterly Journal of Mechanics and Applied Mathematics*, Vol. 2, No. 1, 1949, pp. 75–97. doi:10.1093/qjmam/2.1.75.

[38] Ehret, D. M., "Accuracy of Approximate Methods for Predicting Pressures on Pointed Nonlifting Bodies of Revolution in Supersonic Flow," NACA TN 2764, 1952. NTRS 19930083543.

[39] Ames Research Staff, "Equations, Tables, and Charts for Compressible Flow," NACA Report 1135, 1953. NTRS 19930091059.

[40] Perkins, E. W., and Jorgensen, L. H., "Comparison of Experimental and Theoretical Normal-Force Distributions (Including Reynolds Number Effects) on an Ogive-Cylinder Body at Mach Number 1.98," NACA RM A54H23, 1954. NTRS 19930088493.

[41] Perkins, E. W., and Kuehn, D. M., "Comparison of the Experimental and Theoretical Distributions of Lift on a Slender Inclined Body of Revolution at $M = 2$," NACA RM A53E01, 1953. NTRS 19930087544.

[42] Perkins, E. W., Gowen, F. E., and Jorgensen, L. H., "Aerodynamic Characteristics of the NACA RM-10 Research Missile in the Ames 1- by 3-Foot Supersonic Wind Tunnel No. 2: Pressure and Force Measurements at Mach Numbers of 1.52 and 1.98," NACA RM A51G13, 1951. NTRS 19930086835.

[43] Perkins, E. W., and Jorgensen, L. H., "Investigation of the Drag of Various Axially Symmetric Nose Shapes of Fineness Ratio 3 for Mach Numbers from 1.24 to 3.67," NACA RM A52H28, 1952. (UNT/NTRS).

[44] Luidens, R. W., and Simon, P. C., "Aerodynamic Characteristics of NACA RM-10 Missile in 8- by 6-Foot Supersonic Wind Tunnel at Mach Numbers from 1.49 to 1.98. I: Presentation and Analysis of Pressure Measurements (Stabilizing Fins Removed)," NACA RM E50D10, 1950. NTRS 19930086239.

[45] Ashton, N., and Skaperdas, V., "Verification and Validation of OpenFOAM for High-Lift Aircraft Flows," *Journal of Aircraft*, Vol. 56, No. 4, 2019, pp. 1641–1657. doi:10.2514/1.C034918.

[46] Abraham, T. A., Lazzara, D. S., and Hunsaker, D. F., "Multifidelity Comparison of Supersonic Wave Drag Prediction Methods Using Axisymmetric Bodies," *Aerospace*, Vol. 11, No. 5, 2024, 359. doi:10.3390/aerospace11050359.

[47] Sooy, T. J., and Schmidt, R. Z., "Aerodynamic Predictions, Comparisons, and Validations Using Missile DATCOM (97) and Aeroprediction 98 (AP98)," *Journal of Spacecraft and Rockets*, Vol. 42, No. 2, 2005, pp. 257–265. doi:10.2514/1.7814.

[48] Qaumi, T. S., *Aerodynamic Optimization of the Von Karman Nose Cone for a Supersonic Sounding Rocket*, B.Eng. thesis, Dept. of Aerospace Engineering, Ryerson University (now Toronto Metropolitan University), Toronto, 2021. doi:10.32920/RYERSON.14635638.V1.

[49] Weiss, J. M., and Smith, W. A., "Preconditioning applied to variable and constant density flows," *AIAA Journal*, Vol. 33, No. 11, 1995, pp. 2050–2057. doi:10.2514/3.12946.

[50] OpenCFD Ltd., *OpenFOAM v2406*, ESI Group, 2024, `https://www.openfoam.com`.