

RPDS-X: A Reduced-Order, Inverse-First Framework for Conceptual Design of Two-Ramp External-Compression Scramjet Inlets

Anush Badii^{*1}

¹Institution, Department, Address, Country

October 22, 2026

Abstract

We present RPDS-X, a self-contained, open-source Python framework for the conceptual design of two-ramp external-compression scramjet inlets. The framework is built on an *inverse-first* architecture: given a mission defined by freestream Mach number, geometric altitude, and required net thrust, the framework solves for the isolator height that closes the thrust requirement exactly, using a Brent root-find on a fully coupled forward cycle. The forward model integrates an exact perfect-gas oblique-shock solver, a closed-form laminar Blasius skin-friction correction for the isolator, a calibrated constant-area energy-balance combustor, and an isentropic perfectly expanded nozzle. We validate the theta-beta-M shock solver against the NACA 1135 reference tables at $M_\infty = 2.0$; the maximum angular error is 0.112° and the maximum pressure-ratio error is 0.16 %. We demonstrate the framework on a 30 kN mission at $M_\infty = 5.00$ and $h = 3000$ m, producing a thrust-closed design with $\pi_{\text{total}} = 0.9006$, $H_c = 5.908$ cm, $\dot{m}_{\text{air}} = 29.7470$ kg s⁻¹, $\dot{m}_{\text{H}_2} = 737.1885$ g s⁻¹, and total (axial) engine length $L_{\text{total}} = 128.8$ cm. A Mach sweep at fixed mission ($M_\infty = 2\text{--}6$) locates the compression floor of the two-shock inlet at $M_\infty \approx 2.19$ and finds that the declared manufacturing limit is not reached. A Cantera ignition-delay check shows that hydrogen does not autoignite at the isolator-exit state of any of three demonstration missions, so the heat release assumed by the calibrated combustor is conditional on an ignition source that the model does not contain. Inviscid two-dimensional Eilmer 4 cross-checks agree with the inlet model at the capture plane (Mach number within 1 %, pressure within 2 %) but show a 9 % lower Mach number at the isolator exit, which the model's friction-only isolator cannot reproduce; a reacting run with a numerical ignition source, whose combustion efficiency varies by about ± 0.1 with the grid, is reported as a feasibility demonstration, not as a performance prediction. The framework explicitly enumerates unclaimed physics and reports the Barraza et al. second-mode critical N-factor as a transition threshold. A one-parameter calibrated N-factor model is used to show how an illustrative wedge estimate depends on that threshold and on an unknown amplification coefficient, and a compressible linear-stability solver, verified against published eigenvalues, is then used to compute the second-mode N-factor, first on that illustrative flat plate and then on the engine's own two ramps. The resulting transition locations are model results under stated assumptions (assumed wall temperature, two-dimensional waves, flat-plate similarity flow, thresholds applied outside the range of their fit); they are not predictions for the vehicle.

^{*}Corresponding author: anushbadii@gmail.com

1 Introduction

The design of hypersonic air-breathing engines for a specified mission remains a challenging multi-disciplinary problem. Unlike rocket propulsion, where the flow path is largely decoupled from the freestream, scramjet engines are tightly coupled to their flight condition: the inlet must capture and compress freestream air across a wide Mach range, the isolator must accommodate the pressure rise of combustion without unstating, and the nozzle must expand the exhaust against a varying ambient pressure. This coupling makes forward-only design tools expensive to use in conceptual studies, since every candidate geometry requires a full forward evaluation before its performance can be assessed.

The classical approach to scramjet inlet design is the external compression ramp, in which the freestream is compressed through a series of oblique shocks before entering the isolator. The two-ramp configuration is the canonical building block for higher-order inlet geometries. The Oswatitsch optimization criterion [13] provides a classical benchmark: for a given number of ramps and a given freestream Mach number, an equal-strength condition on the successive shock strengths maximizes the total-pressure recovery. However, the classical equal-strength criterion is not always feasible for two attached shocks at low Mach numbers (section 3.4).

Prior reduced-order frameworks for scramjet inlet design include the analytical methods of Billig [3], the engineering correlations of Heiser and Pratt [9], and the more recent physics-based tools of Curran and Murthy [5]. What has been less explored is the systematic construction of an *inverse-first* design framework, in which the required mission output is imposed as a constraint and the geometric design variables are solved.

The present paper makes three contributions: (1) a complete, self-contained, open-source Python framework for two-ramp external-compression scramjet inlet design; (2) an inverse-first design architecture in which the isolator height H_c is the primary inverse variable and the required net thrust is imposed as a scalar constraint; and (3) a constrained max-recovery two-ramp search that maximizes total-pressure recovery over the feasible ramp-angle domain subject to explicit compression, recovery and manufacturing constraints. The search is verified against the classical Oswatitsch equal-strength solution, with which it coincides wherever the latter exists (section 2.3), and it additionally reports feasibility.

This paper treats the engine flowpath only. It does not size a complete vehicle, and it makes no claim about vehicle-level thrust–drag balance, lift, mass or range. Two checks beyond the reduced-order model are reported: a homogeneous ignition-delay analysis (section 4) and inviscid Eilmer 4 cross-checks (section 5).

2 Methodology: The RPDS-X Framework

RPDS-X is organised as a sequence of coupled sub-models, each of which is validated independently. The design variables, mission inputs, and fixed design choices are summarised in table 1.

2.1 Mission, Freestream, and Standard Atmosphere

The mission is defined by exactly three numbers: M_∞ , h , and F_{req} . All other mission quantities are derived. The freestream static state is computed from the ISA 1976 standard atmosphere. The freestream velocity is $V_\infty = M_\infty a_\infty$ with $\gamma_\infty = 1.4$ and $R_\infty = 287.05 \text{ J kg}^{-1} \text{ K}^{-1}$.

Table 1: Summary of mission inputs, design variables, design choices, and derived outputs.

Category	Quantity	Symbol	Status
Mission input	Freestream Mach	M_∞	Given
Mission input	Geometric altitude	h	Given
Mission input	Required net thrust	F_{req}	Given
Design variable	Ramp-1 angle	θ_1	Optimized
Design variable	Ramp-2 angle	θ_2	Optimized
Design variable	Isolator height	H_c	Inverse solved
Design choice	Duct width / isolator height	W/H_c	Declared
Derived output	Ramp lengths, hinge-2 position	L_1, L_2, x_{h2}	Shock-on-lip
Design choice	Isolator length ratio	L_{iso}/D_h	Declared
Derived output	Net thrust	F_{net}	Constraint
Derived output	Total-pressure recovery	π_{total}	Objective
Derived output	Capture height	H_{cap}	Continuity
Derived output	Hydrogen mass flow	$\dot{m}_{\text{H2}}$	Fuel burn

2.2 Two-Ramp External-Compression Inlet

The inlet is a two-ramp external-compression configuration. Ramp 1 turns the freestream by θ_1 ; ramp 2 then turns the already-shocked flow by the additional local angle $\Delta\theta_2 = \theta_2 - \theta_1$, where θ_2 is the absolute second-ramp angle measured from the freestream direction. Each ramp initiates a weak oblique shock. The theta-beta-M relation for each shock is

$$\tan \theta = \frac{2 \cot \beta (M^2 \sin^2 \beta - 1)}{M^2(\gamma + \cos 2\beta) + 2}, \quad (1)$$

solved on the weak-shock branch by bracketed Brent root-find. The jump conditions are the exact Rankine-Hugoniot relations. The total-pressure ratio follows from the standard shock-loss formula, and the inlet π_{total} is the product of the two successive shock recoveries.

Ramp-angle selection. The ramp angles are selected as described in section 2.3. In the present run the ramp mode is `maxrecovery` and the solver returns $\theta_1 = 6.966^\circ$ and $\theta_2 = 15.029^\circ$ (absolute), i.e. a second-ramp local turn of $\Delta\theta_2 = 8.063^\circ$, giving $p_2/p_\infty = 5.000$. The isolator is fed by the ramp-2 exit state (station 2), so both shocks enter the cycle. Hinge 2 is placed so that both shocks meet at the cowl lip (shock-on-lip); the lip height H_{cap} follows from continuity.

Validation. The solver is validated against the NACA 1135 tables [1] at $M_\infty = 2.0$, $\gamma = 1.4$ (table 2).

Table 2: Validation against NACA 1135 at $M_\infty = 2.0$, $\gamma = 1.4$.

θ [°]	$\beta_{\text{RPDS-X}}$ [°]	β_{ref} [°]	$\Delta\beta$ [°]	$\Delta p_2/p_1$ (rel.)
5.0	34.3016	34.19	0.1116	5.4×10^{-4}
10.0	39.3139	39.31	0.0039	1.3×10^{-5}
15.0	45.3436	45.34	0.0036	2.1×10^{-5}
20.0	53.4229	53.42	0.0029	1.6×10^{-3}

The maximum angular error is set by the $\theta = 5^\circ$ row alone; the remaining rows agree with the tabulated angles to within 0.0039° . The reference values are the tabulated (rounded) entries as transcribed in the reproducibility notebook.

2.3 Max-Recovery Two-Ramp Optimization

The ramp angles are chosen by a direct maximum-recovery search,

$$\max_{(\theta_1, \theta_2)} \pi_{\text{total}} = \pi_1 \pi_2, \quad (2)$$

subject to attached weak shocks, $M_3 > 1$, and the engineering constraints $\pi_{\text{total}} \geq 0.55$, $p_2/p_\infty \geq 5$ and $H_c \geq 1.5$ cm. The last constraint depends on the thrust-closed sizing and is checked on the solved design.

For fixed θ_1 , π_{total} decreases and p_2/p_∞ increases monotonically with $\Delta\theta_2$, so the compression constraint is active at the optimum. The search therefore scans θ_1 on a grid, solves $\Delta\theta_2$ from $p_2/p_\infty = 5$ by Brent root-find, and refines the best bracket with a fine grid and a bounded scalar minimiser. A brute-force $(\theta_1, \Delta\theta_2)$ grid check is part of the validation summary (table 21).

Oblique-shock pressure and total-pressure ratios depend only on the shock-normal Mach number, so this problem is the classical Oswatitsch problem: its optimum is the equal-strength solution $M_{n,1} = M_{n,2}$ [13] wherever that solution exists. The search is therefore *not* a different optimum; it is a numerical realisation that also enforces the additional constraints and reports feasibility. We verify the equivalence against a closed-form equal-strength benchmark at the same compression requirement (table 4).

2.4 Isolator: Closed-Form Laminar Blasius Skin Friction

The laminar Blasius drag closure in closed form is

$$D_f = 0.664 \rho_e U_e^2 P_w \sqrt{\frac{\mu_e L_{\text{iso}}}{\rho_e U_e}}. \quad (3)$$

It is evaluated at the ramp-2 exit state and checked against adaptive quadrature of the local skin friction (relative error below 10^{-15}). The reported D_f is a documented lower bound on internal drag.

2.5 Combustor: Calibrated Energy Balance

The energy balance is $(1 + F)h_{t,4} = h_{t,3} + F \cdot \text{LHV}$, solved for T_4 by Brent root-find. The calibration constants are collected in table 3.

The balance fixes the exit static pressure by α and does not enforce axial momentum. In a frictionless constant-area duct, with fuel that carries no axial momentum, heat addition leaves the impulse function $(p + \rho U^2)A$ unchanged. Across this combustor the model's impulse function rises by 54 %. The factor α therefore stands in for a pressure force that the constant-area geometry does not supply, and the exit state is not one that a momentum-conserving duct of the same area reaches. The thrust values that rest on this exit state should be read as optimistic until the combustor is replaced by a momentum-consistent model. This is also why the Eilmer cases of section 5, which conserve momentum, do not reproduce the model's combustor exit.

2.6 Hydrogen Consumption

The hydrogen mass flow follows from the declared equivalence ratio and the stoichiometric fuel/air ratio:

$$\dot{m}_{\text{H}_2} = \phi F_{\text{ST}} \dot{m}_{\text{air}} = F \dot{m}_{\text{air}}, \quad (4)$$

Table 3: Declared calibration constants.

Constant	Symbol	Value	Units
Static-pressure ratio	α	2.5	—
Equivalence ratio	ϕ	0.85	—
Hot-side γ	γ_p	1.28	—
Hot-side R	R_p	380.0	J/(kg·K)
Hot-side c_p	$c_{p,p}$	1700.0	J/(kg·K)
Cold-side c_p	$c_{p,a}$	1004.5	J/(kg·K)
LHV (H ₂)	—	120×10^6	J/kg
Stoichiometric F	F_{ST}	0.0291552	—

where $F = \phi F_{ST}$ is the fuel/air mass ratio. The gas mass flow through the combustor and nozzle is $\dot{m}_{\text{gas}} = \dot{m}_{\text{air}} + \dot{m}_{\text{H}_2}$. The hydrogen mass flow is conserved through the nozzle, so the total engine-level hydrogen consumption equals the value computed at the combustor inlet.

2.7 Nozzle: Isentropic Perfect Expansion

Isentropic expansion to $p_e = p_{\text{amb}}$. The gross thrust is the momentum-flux thrust net of the freestream ram momentum and before internal friction,

$$F_{\text{gross}} = \dot{m}_{\text{gas}} U_e - \dot{m}_{\text{air}} V_{\infty} + (p_e - p_{\infty}) A_e, \quad (5)$$

whose pressure term vanishes for perfect expansion. Reported F_{gross} is an optimistic upper bound; a physical single-expansion-ramp nozzle and its thrust cost are treated in section 3.6.

2.8 Inverse-Thrust Closure

Solve find H_c such that $F_{\text{gross}}(H_c) - D_f(H_c) = F_{\text{req}}$, by bracketed Brent root-find, closing to 10^{-7} relative.

2.9 Design Choices

Fixed design choices: $W/H_c = 2.0$ ($W = 11.816$ cm at the solved design), $L_{\text{iso}}/D_h = 4.0$ for the baseline and $L_{\text{iso}}/D_h = 6.0$ for the promoted reference design. The combustor length is $2H_c$ and the nozzle half-angle is 12° ; L_{total} is the axial extent from the leading edge to the nozzle exit. All are stated explicitly and overridable.

3 Results

3.1 Baseline Design

Baseline mission: $M_{\infty} = 5.00$, $h = 3000$ m, $F_{\text{req}} = 30$ kN, $W = 11.816$ cm. The two-ramp inlet is configured in `maxrecovery` mode; the solver returns $\theta_1 = 6.966^\circ$ and $\theta_2 = 15.029^\circ$ (absolute), with $H_c = 5.908$ cm.




Figure 1: Baseline two-ramp engine profile. Ramp 1 is at $\theta_1 = 6.966^\circ$ and ramp 2 at $\theta_2 = 15.029^\circ$ absolute ($\Delta\theta_2 = 8.063^\circ$ local turn). The red and purple dashed lines are the two oblique-shock rays; both land on the cowl lip. The yellow inset reports the hydrogen mass flow, air and gas mass flows, and fuel-performance metrics. Total engine length is $L_{\text{total}} = 128.8$ cm.

Table 4 compares the selected ramp schedule with the classical equal-strength benchmark at the same compression requirement and with a declared $8^\circ/16^\circ$ reference schedule.

Table 4: Ramp-angle schedules at the baseline Mach number.

Schedule	θ_1 [$^\circ$]	θ_2 [$^\circ$]	π_{total}	p_2/p_∞
Selected (maxrecovery)	6.966	15.029	0.9006	5.000
Oswatitsch equal-strength	6.966	15.029	0.9006	5.000
Declared reference (fixed)	8.000	16.000	0.8832	5.447




Figure 2: Station physics for the baseline design. Each panel shows the property at the canonical stations; connecting lines are visual guides.

Table 5: Complete cycle state at the baseline design point ($L_{\text{iso}}/D_h = 4.0$).

Station	Quantity	Value	Units
Freestream	T_∞	268.65	K
	p_∞	70108.3	Pa
	ρ_∞	0.909128	kg/m ³
	V_∞	1642.88	m/s
Ramp 1 exit	M_1	4.29931	—
	p_1	156767	Pa
	T_1	343.19	K
	U_1	1596.65	m/s
Ramp 2 exit	M_2	3.65833	—
	p_2	350541	Pa
	T_2	438.412	K
	U_2	1535.56	m/s
	π_{total}	0.900606	—
Isolator exit	p_2/p_∞	5	—
	L_{iso}	0.315085	m
	D_h	0.0787713	m
	D_f	65.5212	N
	$\dot{m}_{\text{air}}$	29.747	kg/s
Combustor exit	M_4	2.95216	—
	p_4	876353	Pa
	T_4	1173.47	K
	U_4	2230.35	m/s
Nozzle exit	M_e	4.51798	—
	U_e	2589.42	m/s
	p_e	70108.3	Pa
	A_e	0.0430932	m ²
Summary	$\dot{m}_{\text{gas}}$	30.4842	kg/s
	F_{gross}	30065.5	N
	F_{net}	30000	N
	residual	2.78669e-08	N

3.1.1 Hydrogen Consumption and Fuel-Performance Metrics

Applying section 2.6 at the baseline design point gives a hydrogen mass flow of $\dot{m}_{\text{H}_2} = 737.1885 \text{ g s}^{-1}$ ($2653.8787 \text{ kg h}^{-1}$), against an air mass flow of $\dot{m}_{\text{air}} = 29.7470 \text{ kg s}^{-1}$. The specific fuel consumption and specific impulse are

$$SFC = \frac{\dot{m}_{\text{H}_2}}{F_{\text{net}}} = 8.846\,262 \times 10^{-2} \text{ kg N}^{-1} \text{ h}^{-1}, \quad I_{\text{sp}} = \frac{F_{\text{net}}}{\dot{m}_{\text{H}_2} g_0} = 4149.75 \text{ s}. \quad (6)$$

Table 6 collects the fuel-performance summary. The gas mass flow $\dot{m}_{\text{gas}} = \dot{m}_{\text{air}} + \dot{m}_{\text{H}_2}$ is the mass flow used by the nozzle.

Table 6: Hydrogen-consumption and fuel-performance summary at the baseline design point.

Quantity	Symbol	Value	Units
Air mass flow	$\dot{m}_{\text{air}}$	29.747	kg/s
Hydrogen mass flow	$\dot{m}_{\text{H}_2}$	0.737189	kg/s
Hydrogen mass flow	$\dot{m}_{\text{H}_2}$	737.189	g/s
Hydrogen mass flow	$\dot{m}_{\text{H}_2}$	2653.88	kg/h
Gas mass flow (air + fuel)	$\dot{m}_{\text{gas}}$	30.4842	kg/s
Equivalence ratio	ϕ	0.85	—
Fuel/air mass ratio	F	0.0247819	kg/kg
Stoichiometric ratio	F_{ST}	0.0291552	—
Specific fuel consumption	SFC	0.0884626	kg/(N·h)
Specific fuel consumption	SFC	2.4573e-05	kg/(N·s)
Specific impulse	I_{sp}	4149.75	s

3.2 Ramp-Angle Design Landscape




Figure 3: Ramp-angle design landscape at the baseline mission. Every plotted point is thrust-closed. The red dashed line is the compression constraint $p_2/p_\infty = 5$ and the hatched region violates it. The star marks the selected design. Blank cells have no attached-shock solution, $\theta_2 \leq \theta_1$, or no thrust closure.

3.3 Mach Sweep: The Design Envelope




Figure 4: Mach sweep with the max-recovery ramps. Red bars mark Mach points at which no design exists under the declared constraints.

3.4 Feasibility Boundaries

The sweep covers $M_\infty = 2$ –6 (13 points) at $h = 3000$ m and $F_{\text{req}} = 30$ kN, with the max-recovery ramp search and every constraint of section 2.3 applied at each point; 12 of 13 points close.

Compression floor. Below $M_\infty \approx 2.19$ no pair of attached weak shocks reaches $p_2/p_\infty \geq 5$ (infeasible at $M_\infty = 2$). The classical equal-strength benchmark is infeasible below $M_\infty \approx 2.19$; above that the two methods return the same ramp angles.

Manufacturing floor. The declared minimum $H_c = 1.5$ cm is *not* reached within the swept range: H_c decreases monotonically from 10.37 cm at $M_\infty = 2.2$ to 5.56 cm at $M_\infty = 6$.

Recovery. $\pi_{\text{total}} = 0.9006$ at every feasible point: with the compression constraint active, the optimal two-shock recovery depends only on the required pressure ratio, not on M_∞ .

Combustor-exit Mach. M_4 rises from 1.003 to 3.49; it stays above unity at every feasible point, so no ram-mode (subsonic-combustion) point appears in this sweep. The lowest-Mach feasible point is close to the $M_4 = 1$ boundary.

3.5 Isolator-Length Family




Figure 5: Isolator-length family at fixed ramp angles $\theta_1 = 6.966^\circ$, $\theta_2 = 15.029^\circ$ (absolute). Each point is thrust-closed; the inverse solver adjusts H_c to compensate for the added friction of a longer isolator. Total length is axial. Table 7 lists the Billig criterion $L_{iso}/H_c \geq 8$ for each member.

Table 7: Isolator-length family at fixed ramp angle $\theta_1 = 6.966^\circ$.

L_{iso}/D_h	H_c (cm)	L_{iso} (cm)	L_{iso}/H_c	$\dot{m}$ (kg/s)	D_f (N)	L_{total} (cm)	Billig
4.0	5.908	31.51	5.33	29.7470	65.52	128.8	SHORT
6.0	5.912	47.29	8.00	29.7617	80.33	144.6	OK
8.0	5.915	63.09	10.67	29.7741	92.83	160.4	OK
10.0	5.918	78.91	13.33	29.7850	103.87	176.2	OK
12.0	5.921	94.73	16.00	29.7949	113.85	192.0	OK

Three observations follow.

First, the isolator height is essentially independent of the isolator length. Across the family, H_c varies by only 0.22%. The physical reason is that the added laminar friction of the longer isolator is a small fraction of gross thrust, so the inverse solver needs only a negligible increase in H_c to compensate. This is a reassuring robustness result.

Second, the total engine length grows approximately linearly with L_{iso}/D_h .

Third, the empirical Billig criterion $L_{iso}/H_c \geq 8$ for avoiding thermal choking is *not* satisfied by the baseline $L_{iso}/D_h = 4.0$ ($L_{iso}/H_c = 5.33$). It is first satisfied at $L_{iso}/D_h = 6.0$. We therefore promote $L_{iso}/D_h = 6.0$ as the paper's reference design (axial length 144.6 cm); the baseline is retained as the shortest member of the family for comparison.

The framework does not itself select a preferred member of the family. The choice of L_{iso}/D_h is a declared design decision.

3.6 SERN Nozzle and Vehicle Integration

The ideal perfectly-expanded nozzle of section 2.7 is not buildable, so the framework also designs a planar single-expansion-ramp nozzle (SERN) on the combustor-exit state. The cowl is horizontal; the flow expands internally to $M_{\text{lip}} = 3.19$ over a straight ramp inclined at 8° , then through a centred Prandtl–Meyer fan from the cowl lip onto an external ramp whose contour is the wave-cancelling (ideal) shape, truncated at a chosen length. Thrust is obtained from a control-volume momentum balance, with mass closure along the last Mach line (relative error $6.7\text{e-}10$) and agreement of the full-length ramp with the one-dimensional limit (relative error $1.0\text{e-}08$) as checks. The design rule is the shortest nozzle whose net-thrust loss against the ideal nozzle is within a declared budget of 10 %, subject to a length cap of 64 cm and a thrust-vector cap of 25° (meets the thrust-loss budget).




Figure 6: SERN design: selected geometry (left) and the length–thrust-loss trade (right), each curve being a ramp truncation at fixed lip Mach number.

The selected nozzle has a total length of 36.3 cm (14.0 cm internal, 22.2 cm external ramp) against 29.2 cm for the ideal nozzle, and ends at $M_W = 3.19$ of $M_e = 4.52$ ($p_W/p_\infty = 8.38$). Net thrust falls from 30 000 N to 27 381 N, a loss of 8.7 %. The nozzle also exerts a normal force of 3.82 kN on the vehicle (thrust vector 7.9° off the engine axis); the effect on trim has not been assessed. To first order (thrust scaling with size squared; friction is not exactly self-similar), recovering F_{req} with the SERN requires all dimensions scaled by 1.047.

RPDS-X vehicle integration: $M_\infty=5$, $h=3000$ m, $F_{\text{req}}=30$ kN; SERN costs 8.7% of net thrust (concept drawing, not a validated design)
all dimensions in mm; values tagged [assumed] are layout assumptions, all others come from the model




Figure 7: Dimensioned vehicle-integration concept (flight orientation, millimetres). Engine and SERN dimensions come from the model; the body outline and plan-form widening are assumed layout parameters and are tagged as such.

In this concept the forebody underside is the two compression ramps, so the vehicle nose is the ramp leading edge and the cowl hangs below the body; the SERN ramp is the vehicle aft underside. Side fences of the engine width ($W = 118.2$ mm) run from nose to tail, because the inlet and nozzle are calculated as planar and would otherwise lose pressure to lateral spillage; the plan form is engine-wide up to the cowl lip (563 mm) and widens behind it with an assumed half-angle of 6° . The SERN model is inviscid and planar, ignores jet-boundary and freestream interaction, and takes uniform flow at the lip plane; fence-corner flow, fence boundary layers and three-dimensional relief are not modelled.

3.7 Transition Threshold: Barraza et al. Second-Mode N_{crit}

The framework reports the Barraza et al. second-mode critical N-factor as a *threshold*, not as a computed transition location. The expression below is Eq. (21) of Barraza et al. [2], and was checked against that paper. The authors determined it empirically from their own test cases (Mach 6 flat plate, cone and flared cone), so applying it at other conditions is an extrapolation of

their calibration. The critical N-factor is

$$N_{\text{crit,SM}} = [1.1 Tu_{\infty}(\%) + 0.18]^{-3/2} + 1.5, \quad (7)$$

where $Tu_{\infty} = u'/U_{\infty} \times 100$. When only the wind-tunnel noise level $NL = \text{RMS}(p')/P$ is known, the turbulence level is taken from the conversion of Eq. (22) of the same paper, $Tu_{\infty}(\%) = NL(\%)/(\gamma_{\infty} M_{\infty})$, which they attribute to Wang et al.




Figure 8: Barraza et al. transition-threshold authority. Left: critical second-mode N-factor as a function of freestream turbulence intensity. Right: equivalent noise level at each turbulence level, evaluated across the Mach envelope.

Table 8: Barraza et al. critical N-factor threshold at selected turbulence intensities.

Tu_{∞} (%)	$N_{\text{crit,SM}}$	NL at $M_{\infty} = 3.4$ (%)
0.10	7.903	0.476
0.20	5.453	0.952
0.40	3.548	1.904
0.80	2.416	3.808
1.60	1.870	7.616

3.7.1 Comparison with the classical Mack correlation

The classical e^N criterion of Mack [11] relates the critical N-factor to the freestream turbulence level,

$$N_{\text{crit,Mack}} = -8.43 - 2.4 \ln\left(\frac{Tu_{\infty}(\%)}{100}\right). \quad (8)$$

It is a low-speed calibration against first-mode (Tollmien–Schlichting) transition, whereas the expression above is a second-mode correlation for hypersonic boundary layers. Both are shown in fig. 9 and table 9.




Figure 9: Critical N-factor against freestream turbulence level for the second-mode correlation and the classical Mack correlation.

Table 9: Critical N-factor from the two correlations at the turbulence levels of table 8.

Tu_∞ (%)	$N_{\text{crit,SM}}$	$N_{\text{crit,Mack}}$	Difference
0.10	7.903	8.149	-0.245
0.20	5.453	6.485	-1.032
0.40	3.548	4.822	-1.273
0.80	2.416	3.158	-0.742
1.60	1.870	1.494	+0.376

Below $Tu_\infty \approx 1.30\%$ the Mack value is the larger, by up to about 1.3 in the range tabulated; the two curves cross at $Tu_\infty \approx 1.30\%$. They also differ in shape: the Mack expression falls without bound and reaches zero at $Tu_\infty \approx 3.0\%$, while the second-mode expression approaches a lower bound of 1.5. A difference of one N-factor is a factor $e \approx 2.7$ in disturbance amplitude, so the choice of threshold is a material uncertainty for any transition estimate. The Mack expression is the same as Eq. (20) of Barraza et al. [2], who cite Mack’s 1977 JPL report for it; the original report was not consulted here. Both thresholds are compared with computed N-factors in sections 3.7.3 and 3.7.4.

Dependence on the local state. Both thresholds must be evaluated with local conditions, because the boundary layer on each ramp sits behind a different shock. The second-mode frequency scales as $f \sim u_e/\delta$. Using the laminar Blasius form $\delta \sim \sqrt{\nu_e x/u_e}$ with the edge states of table 5, at the same streamwise distance the boundary layer on ramp 2 is thinner than on ramp 1 by a factor of 0.84 (ramp-2 over ramp-1 thickness), and the u_e/δ frequency scale is higher by a factor of 1.14. This is a scaling argument only: it ignores compressibility and wall temperature, gives no absolute frequency, and is not a stability calculation.

3.7.2 Sensitivity of a calibrated N-factor model

To show how much a transition estimate depends on the threshold, a one-parameter amplification model is used,

$$N(x) = c \left(\sqrt{Re_x} - \sqrt{Re_{x0}} \right), \quad Re_x = \frac{\rho_e u_e}{\mu_e} x, \quad (9)$$

with x_{tr} given by $N(x_{tr}) = N_{crit}$. Here $Re_{x0} = 10^5$ is an assumed neutral-point Reynolds number and c is an assumed, dimensionless amplification coefficient. *Neither is derived in this paper:* c is treated as unknown and varied over 0.003–0.008 about a central value of 0.005, and the model contains no stability calculation. Variants with a hand-set neutral point, or with a numerical integration of the same growth law, reduce to this form (the integration converges to it as the grid is refined), so they are not independent methods and are not reported separately.

The case is a stand-alone sharp wedge, not an engine station: $M_\infty = 5$ at 27 km with a 4° weak oblique shock, which gives an edge Mach number of 4.59, edge velocity from conservation of total enthalpy, and $Re/m = 3.73 \times 10^6$ per metre. At $Tu_\infty = 0.4\%$ the second-mode threshold $N_{crit,SM} = 3.55$ gives $x_{tr} = 0.28$ m at $c = 0.005$, and 0.15–0.60 m across the coefficient range. Figure 10 and table 10 show that the coefficient and the threshold each move x_{tr} by a factor of several, far more than the wedge angle does.




Figure 10: Illustrative wedge ($M_\infty = 5$, 27 km). Left: x_{tr} against Tu_∞ for three amplification coefficients with the second-mode threshold, and for the Mack threshold at $c = 0.005$. Right: x_{tr} against wedge angle at $Tu_\infty = 0.4\%$.

Table 10: Transition location x_{tr} (m) of the calibrated model on the illustrative wedge, by turbulence level and amplification coefficient c .

Tu_∞ (%)	$N_{crit,SM}$	x_{tr} , second-mode threshold			$N_{crit,Mack}$	x_{tr} , Mack
		$c = 0.003$	$c = 0.005$	$c = 0.008$		$c = 0.005$
0.10	7.90	2.334	0.964	0.456	8.15	1.015
0.20	5.45	1.220	0.530	0.267	6.49	0.698
0.40	3.55	0.602	0.282	0.155	4.82	0.440
0.80	2.42	0.337	0.171	0.102	3.16	0.241
1.60	1.87	0.237	0.128	0.081	1.49	0.101

These numbers are a sensitivity study. Because c is not calibrated against stability calculations or experiment, the absolute values of x_{tr} carry no predictive weight, and agreement between differently-tuned versions of the same law would not be a validation. What the study does show is that a credible transition estimate needs a computed $N(x)$ and a verified threshold.

3.7.3 Computed second-mode N-factor from linear stability

The assumed coefficient c of the previous subsection can be removed by computing the amplification directly. A locally parallel linear-stability solver for two-dimensional waves on a compressible flat-plate boundary layer was written and verified for this purpose. The mean flow is the compressible similarity solution (perfect gas, Sutherland viscosity, $Pr = 0.72$, adiabatic or isothermal wall). The linearised compressible Navier–Stokes equations in primitive variables, with the Stokes hypothesis and heat conduction proportional to viscosity, are discretised by Chebyshev collocation (100 points, algebraically stretched, domain height $3\delta_{99}$) and solved as a generalised eigenvalue problem. The spatial growth rate $-\alpha_i$ at real frequency is found from the temporal second-mode branch by Newton iteration. The amplification is $N = -\int \alpha_i dx = 2 \int (-\alpha_i L) dR$ with $L = \sqrt{\nu_e x / U_e}$ and $R = \sqrt{Re_x}$, integrated at fixed physical frequency from the lower neutral point; the envelope is the maximum over frequency.

Verification. Table 11 lists the checks, all run in the accompanying notebook. The solver reproduces the Blasius mean flow and the incompressible critical Reynolds number (519.4 against the classical 519.4), which an independent Orr–Sommerfeld solver on the same profile also confirms. It reproduces a published subsonic eigenvalue of Malik [12] (as reproduced by Sharma and Girimaji [15]) and a published Mach 4.5 second-mode eigenvalue of Erlebacher and Hussaini [6]: with $Pr = 0.70$ the real part of the latter agrees to 8×10^{-6} (relative) and the growth rate to 0.06 %. The Prandtl number is not stated in the source values that were available, and the match identifies $Pr = 0.70$; with $Pr = 0.72$ the growth rate differs by 7 %, so growth rates at this level are sensitive to the Prandtl number. The calculations below use $Pr = 0.72$. The baseline second-mode growth rate changes by 0.02 % between 90 and 180 collocation points and with a taller domain.

Table 11: Verification of the linear-stability solver against exact limits and published values.

Check	Reference	Computed	Difference
Blasius $f''(0)$	0.33206	0.33206	2.7×10^{-6}
Blasius δ^*/L	1.7208	1.7208	1.2×10^{-5}
Critical Re_{δ^*} , $M \rightarrow 0$	519.4	519.4	0.00 %
Neutral Re_{δ^*} at $\alpha\delta^* = 0.3011$, independent Orr–Sommerfeld solver	519.4	519.2	0.03 %
Malik (1990): $M = 0.5$, $R = 2000$, $\alpha = 0.1$, $Pr = 0.72$	$0.0290817 + 0.0022441 i$	$0.0290871 + 0.0022479 i$	6.6×10^{-6}
Malik (1990): same, $Pr = 0.70$	$0.0290817 + 0.0022441 i$	$0.0290719 + 0.0022446 i$	9.8×10^{-6}
Erlebacher–Hussaini: $M = 4.5$ second mode, $Pr = 0.70$	$2.046738 + 0.021490 i$	$2.046754 + 0.021477 i$	8.0×10^{-6} (Re), 6.4×10^{-4} (Im)
Erlebacher–Hussaini: same, $Pr = 0.72$	$2.046738 + 0.021490 i$	$2.047640 + 0.023080 i$	4.4×10^{-4} (Re), 7.4×10^{-2} (Im)
Convergence of the baseline second-mode ω_i , $N = 90$ to 180	0.001630	0.001630	1.8×10^{-4} rel.

Result for the illustrative flat plate. The edge state is that of the wedge of section 3.7.2 ($M_e = 4.59$, $T_e = 249$ K, $Re/m = 3.73 \times 10^6$ per metre). Two walls are used: adiabatic, and

isothermal at $T_w = 2T_e$ (≈ 500 K). The wall temperature is an assumption of this study and is not produced by the framework. Figure 11 shows the computed second-mode growth rate and the N-factor envelope; table 12 gives the transition location for the thresholds of tables 8 and 9.




Figure 11: Illustrative flat plate ($M_e = 4.59$, $T_e = 249$ K). Left: second-mode growth rate for the adiabatic wall; the white line is the neutral curve. Right: computed N-factor envelope for the adiabatic wall and for $T_w = 2T_e$, with the calibrated $\sqrt{Re}$ model of section 3.7.2 (dotted) and two thresholds at $Tu_\infty = 0.4\%$ (dashed lines).

Table 12: Transition location x_{tr} (m) at which the computed N-factor envelope reaches the critical value, for the adiabatic wall and for $T_w = 2T_e$. “>” means the threshold is not reached within the computed range.

Tu_∞ (%)	Second-mode threshold			Mack threshold		
	N_{crit}	adiabatic	$T_w = 2T_e$	N_{crit}	adiabatic	$T_w = 2T_e$
0.10	7.90	> 5.4	2.80	8.15	> 5.4	2.94
0.20	5.45	> 5.4	1.56	6.49	> 5.4	2.04
0.40	3.55	3.85	0.82	4.82	> 5.4	1.29
0.80	2.42	2.23	0.48	3.16	3.25	0.69
1.60	1.87	1.57	0.34	1.49	1.18	0.26

At the station $x = 0.28$ m where the calibrated model of section 3.7.2 gave $N = 3.5$, the computed envelope is $N = 0.39$ for the adiabatic wall and $N = 1.61$ for the cold wall, so the assumed coefficient overestimates the amplification by a factor of about 9 and 2.2 respectively. The computed N-factor grows roughly linearly in R after the neutral point, not as $\sqrt{Re_x}$. The dominant frequency at that station is about 213 kHz (adiabatic) and 248 kHz (cold wall). With the second-mode threshold at $Tu_\infty = 0.4\%$ the computed transition location is 3.8 m for the adiabatic wall and 0.82 m for the cold wall, against 0.28 m from the calibrated model; the Mack threshold is not reached on the adiabatic wall within 5.4 m (maximum $N = 4.5$) and is reached at 1.29 m on the cold wall. Cooling the wall raises the amplification several-fold, consistent with the known destabilisation of the second mode by wall cooling.

Limits of this result. Only two-dimensional waves are computed; oblique first-mode waves can be more amplified at these Mach numbers and would move the transition location upstream, so the values above apply to the two-dimensional second-mode family only. The flow is treated as locally parallel, on a flat plate at constant edge conditions, without pressure gradient, nose bluntness, entropy layer or real-gas effects, and the wall temperatures are assumed. The critical N-factors are the correlations of sections 3.7 and 3.7.1, which the authors fitted to their own cases and which are used here outside that range. The engine ramps are treated in section 3.7.4. The numerical uncertainty from the integration is about 1–2% in N . The values are an illustration of what a computed amplification gives, not a prediction for the engine.

3.7.4 Second-mode N-factor at the engine ramp states

The solver of section 3.7.3 is now applied to the edge states of the engine’s own two ramps (baseline mission). Ramp 1 has $M_e = 4.30$, $T_e = 343$ K and $Re/m = 124 \times 10^6$ per metre; ramp 2 has $M_e = 3.66$, $T_e = 438$ K and $Re/m = 175 \times 10^6$ per metre. At the corner ($x = 0.299$ m) the Reynolds number is $\sqrt{Re_x} = 6097$, far beyond the range of the illustrative case above, so new stability tables were computed (100 collocation points, $Pr = 0.72$, about 13 minutes per table on two cores).

Method. Each ramp is a locally parallel flat plate at its edge state. Ramp 1 starts at the nose. Ramp 2 starts from a virtual origin placed so that its similarity momentum thickness equals that of ramp 1 at the corner (the momentum thickness is carried through the corner; $x_0 = 0.40$ m on ramp 2’s own origin for the adiabatic wall). Waves keep their physical frequency across the corner, so the non-dimensional frequency changes by $F_2/F_1 = 0.737$, and N is carried over and not allowed below zero. A second bound is a boundary layer restarted at the corner (“fresh”). The wall temperature is not an output of the framework and is assumed: adiabatic, 700 K, and 450 K, the same on both ramps. The turbulence level Tu_∞ is the freestream value used in section 3.7.




Figure 12: Computed second-mode N-factor envelope along the engine’s two ramps for three assumed wall temperatures (left; solid: ramp 1, dashed: ramp 2 with momentum-thickness matching), and the resulting transition location against turbulence level for the two thresholds (right).

Table 13: Computed second-mode N -factor on the engine ramps. N at three stations of ramp 1 and at the corner, at the end of ramp 2 (carried through the corner, and for a boundary layer restarted at the corner), and the dominant frequency at the corner.

Wall	T_w/T_e (ramp 1)	$N(0.05 \text{ m})$	$N(0.10 \text{ m})$	$N(0.20 \text{ m})$	$N(\text{corner})$	$N(\text{end r2}), \text{carried}$	$N(\text{end r2}), \text{fresh}$	$f \text{ (kHz)}$
adiabatic	4.10	1.3	2.3	3.9	5.2	6.8	2.0	1572
700 K	2.04	4.8	7.7	11.9	15.2	41.2	14.7	1788
450 K	1.31	6.7	10.5	16.1	20.5	57.2	22.3	1991

Table 14: Transition location x (m from the nose) at which the computed envelope reaches the critical N -factor of each threshold. “r2”: reached on ramp 2 (carried-through bound); “> range”: not reached within the two ramps.

Wall	Tu_∞ (%)	N_{crit} (second-mode)	x_{tr}	N_{crit} (Mack)	x_{tr}
adiabatic	0.1	7.90	> range	8.15	> range
adiabatic	0.2	5.45	0.454r ²	6.49	0.505r ²
adiabatic	0.4	3.55	0.178	4.82	0.271
adiabatic	0.8	2.42	0.107	3.16	0.152
adiabatic	1.6	1.87	0.078	1.49	0.060
700 K	0.1	7.90	0.104	8.15	0.109
700 K	0.2	5.45	0.059	6.49	0.077
700 K	0.4	3.55	0.032	4.82	0.050
700 K	0.8	2.42	0.019	3.16	0.027
700 K	1.6	1.87	0.014	1.49	0.011
450 K	0.1	7.90	0.064	8.15	0.067
450 K	0.2	5.45	0.037	6.49	0.048
450 K	0.4	3.55	0.020	4.82	0.031
450 K	0.8	2.42	0.012	3.16	0.017
450 K	1.6	1.87	0.009	1.49	0.007

Result. For the adiabatic wall the computed N reaches 5.2 at the corner. At $Tu_\infty = 0.4\%$ the thresholds ($N_{\text{crit}} = 3.55$ and 4.82) are reached at $x = 0.18 \text{ m}$ and 0.27 m , that is 18 and 27 cm along ramp 1 (length 0.299 m). For the cooled walls the amplification is several times larger: $N = 15.2$ at the corner for 700 K and 20.5 for 450 K, and the same thresholds are reached at 3–5 cm (700 K) and 2–3 cm (450 K), a few centimetres from the nose. Wall temperature is therefore the dominant uncertainty, and it is an assumption of this study. The dominant frequencies are about 1.57, 1.79 and 1.99 MHz, because the boundary layers are thin at these high unit Reynolds numbers. Through the corner the adiabatic-wall N rises from 5.2 to 6.8 at the end of ramp 2 (carried through) and would reach only 2.0 for a restarted boundary layer. The adiabatic-wall envelope first falls after the corner (fig. 12), because the wave that dominates at the end of ramp 1 has, at the new Reynolds number and frequency scaling, a frequency below the ramp-2 unstable band; it decays until it re-enters the band and other frequencies take over the envelope. For the cooled walls the values at the end of ramp 2 are far above any threshold (41 and 57 carried through, 15 and 22 restarted) and are meaningful only as evidence that the thresholds are exceeded.

Consequence for the reduced-order model. Unless the walls are close to adiabatic and the freestream turbulence is low, these results place transition on ramp 1, well upstream of the isolator. The reduced-order model uses laminar (Blasius) friction on the ramps and isolator. For the cooled walls that treatment is inconsistent with the computed state, and the laminar friction should be read as a lower bound; this was not quantified.

Limits. The assumptions listed at the end of section 3.7.3 apply, and in addition: the wall temperature is assumed; the momentum-thickness matching ignores the shock–boundary-layer interaction at the corner; the edge state is the inviscid state behind each shock, with no entropy-layer or bluntness effect; only two-dimensional waves are computed, so oblique first-mode waves, which can be more amplified and would move transition upstream, are not included; and the thresholds were fitted by their authors to Mach-6 cases and are applied here at Mach 4.30 and 3.66. At N above roughly 10 the linear theory is far outside the range in which it relates to transition, so the large values are not predictions of transition location. The numerical uncertainty in N from the integration is 1–2% (collocation points, domain height, and Reynolds-number grid were varied).

Apart from the calibrated model and the linear-stability computations above, the framework does not compute the streamwise amplification $N_{\text{SM}}(x)$ for other configurations. The transport-equation route of Barraza et al. would require the trained neural-network weights, which are not bound to the present framework. **The transition locations of section 3.7.4 are model results under the stated assumptions; they are not predictions for the vehicle.**

4 Ignition Feasibility of the Combustor Entry State

The calibrated combustor of section 2.5 releases heat by an energy balance; it contains no ignition model. Whether the entry state can ignite hydrogen at all was therefore checked separately. For each mission, the isolator-exit state (T_3, p_3) was used as the initial condition of a homogeneous, constant-pressure reactor in Cantera 3.1 [8] with the mechanism of Li et al. [10] and a hydrogen–air equivalence ratio of 0.85. The ignition delay is the time to a 400 K temperature rise; “none” means no such rise within 0.5 s. The combustor residence time is $2H_c/U_3$, the combustor length of the model divided by the entry velocity.

Table 15: Homogeneous ignition delay at the isolator-exit state (times in ms). Delays at fixed 1000 K and 1200 K use the same pressure p_3 .

Mission	T_3 [K]	p_3 [kPa]	t_{res}	$\tau_{\text{ign}}(T_3)$	$\tau_{\text{ign}}(1000 \text{ K})$	$\tau_{\text{ign}}(1200 \text{ K})$
$M_\infty = 4$, 2 km, 15 kN	449	397.5	0.073	none	11.6	0.0164
$M_\infty = 5$, 3 km, 30 kN (baseline)	438	350.5	0.077	none	11.2	0.0173
$M_\infty = 6.5$, 25 km, 8 kN	362	12.6	0.145	none	1.29	0.542

At none of the three entry states does hydrogen autoignite. The entry temperatures of 360 K to 450 K are far below the 1000 K to 1200 K range in which the delay of the two lower-altitude cases falls from about 11 ms to below the residence time, and the 25 km case is still too slow at 1200 K (0.542 versus 0.145 ms). The heat release that the thrust closure assumes is therefore conditional on an ignition source (pilot flame, flameholder recirculation or a heated surface) that the model does not represent, and every performance number in this paper should be read as “if ignited”. The homogeneous delay is also a lower bound on the time needed in a real combustor, because mixing is not instantaneous.

5 CFD Cross-Check with Eilmer 4

5.1 Set-up

The notebook writes planar Eilmer 4 [7] cases of the solved flowpath (Cell 16e, non-reacting; Cell 16g, reacting). All runs are inviscid (Euler), run for eight vehicle-length flow-throughs, and

use 8990 cells unless stated. Section 5.5 reports a four-grid study of the diverging-combustor jobs; the constant-area jobs were not repeated on other grids. Station values are plain averages over the duct height, compared with the reduced-order model at the same stations.

5.2 Non-reacting flow

The non-reacting case (ideal air, $\gamma = 1.4$, 8990 cells, heat release off) tests the inlet and isolator models.

Table 16: Two-dimensional Eilmer 4 against the reduced-order model for the baseline mission, non-reacting. Entries are CFD / model.

Station	M	p [kPa]	T [K]	$\dot{m}'$ [kg/s/m]
Capture plane (inlet exit)	3.63 / 3.66	345.1 / 350.5	444 / 438	244.9 / 251.8
Isolator exit	3.33 / 3.66	431.2 / 350.5	514 / 438	245.5 / 251.8

At the capture plane the CFD reproduces the model’s Mach number (-0.8%), pressure (-1.5%) and temperature (1.4%); the mass flow is 2.7% lower. At the isolator exit the CFD Mach number is 9% lower and the pressure 23% higher than the model. This is a model-form difference, not a numerical one: the model’s isolator is a straight friction duct whose inflow is the ramp-2 exit state, whereas the CFD retains the cowl-lip shock and the reflected waves in the isolator (section 7). The model therefore overstates the isolator-exit Mach number that the combustor would see and understates its pressure.

5.3 Reacting flow

Cell 16g converts the job to nine-species hydrogen–air chemistry with the 18-reaction Bittker–Scullin scheme [4], built with Eilmer’s `prep-gas` and `prep-chem`. The mechanism as transcribed was checked on Eilmer’s own Bittker–Scullin test case within that case’s tolerances. Because the solver is inviscid there is no molecular mixing, so hydrogen enters as a perfectly mixed volumetric source in a strip at the combustor entry. Three settings limit what the result means:

- In that constant-area job the fuel flow was capped so that complete combustion released at most 85% of the heat that would thermally choke a constant-area duct (section 2.5); this gives $\phi = 0.26$, well below the model’s $\phi = 0.85$.
- Since hydrogen does not autoignite at this state (section 4), a numerical ignition source relaxes the gas in the forward half of the lower duct to 1700 K . It adds energy: at most 114 MW per metre of depth, which is 51% of the fuel energy. It is not a flameholder model.
- The run uses 20 cells across the capture height (2737 cells) and is not grid-converged.

Table 17: Coarse reacting Eilmer 4 run, baseline mission, with the numerical ignition source. These are feasibility results, not predictions for an unaided engine.

Station	M	p [kPa]	T [K]	η_{H_2}	$\eta_{\text{H}_2\text{O}}$
Isolator exit	3.35	411.4	511	–	–
Combustor exit	1.33	1979.9	1762	0.80	0.74
Nozzle-lip plane	1.78	989.9	1522	0.80	0.74

Here η_{H_2} is the fraction of the injected hydrogen consumed and $\eta_{\text{H}_2\text{O}}$ the fraction converted to water; the difference is hydrogen still held as OH, H, O and HO_2 at these temperatures. Neither is a heat-release efficiency. The flow stays supersonic through the combustor (exit Mach number 1.33) and the isolator-exit Mach number (3.35) matches the non-reacting run (table 16). This run does *not* test the model’s combustor: its fuel flow is 30 % of the model’s, and up to half of the fuel energy is supplied by the numerical ignition source. What it shows is that, once ignited, hydrogen at this equivalence ratio burns to a large extent within the duct and the flow does not unstart. The nine-species chemistry, the ignition source and the mixing assumption each need separate study before the result is used quantitatively.

5.4 Diverging-combustor jobs (notebook version 16)

The constant-area jobs above choke thermally: a first non-reacting run at the default resolution (8990 cells), heated to $T_t = 2501$ K, reached $M = 0.96$ at the combustor exit (2.66 MPa), and a first reacting run at that resolution stopped in the chemistry integrator during the start-up transient. The current export (Cells 16e and 16g) therefore lets the body-side wall of the combustor drop linearly to an area ratio of 1.29, the largest the nozzle geometry allows, and limits the heat, or the hydrogen, to what a one-dimensional integration of flow with area change and heat addition keeps at or above $M = 1.8$ at the combustor exit. The margin is large because the earlier constant-area run reached an exit Mach number about 0.4 below its one-dimensional estimate. The fuel and ignition source are ramped up over six flow-throughs and the reacting run uses a CFL number of 0.3.

Table 18: Diverging-combustor Eilmer 4 runs, baseline mission, 8990 cells, inviscid. Entries are CFD, with the model value for the non-reacting stations.

Case	Station	M	p [kPa]	T [K]	$\dot{m}'$ [kg/s/m]
Non-reacting, T_t rise 75 % of model	Isolator exit	3.34 (3.66)	429.2 (350.5)	512 (438)	245.3 (251.8)
	Combustor exit	1.49	1225.8	1678	246.4
Reacting, hydrogen 10 % of the model’s, ignition source on	Isolator exit	3.38	420.9	505	242.9
	Combustor exit	1.94	912.4	1389	245.9
	Nozzle-lip plane	1.86	918.8	1420	246.8

Both jobs run to the end and the combustor-exit flow stays supersonic. In the reacting run the fuel and the ignition source together are limited to the heat of the non-reacting case; the hydrogen is 10 % of the model’s flow ($\phi = 0.08$), and the ignition source is allowed up to 155 % of the fuel energy, so the burning here is dominated by the source and is a feasibility result, not a prediction. The fraction of hydrogen consumed is 0.62 from the species at the combustor exit and 0.72 from the water formed (section 5.5 gives the values on four grids); the two estimates differ because they are plain averages along one line. A first version of the reacting job that counted only the fuel against the heat limit (25 % of the model’s hydrogen) failed in the chemistry integrator at both resolutions, once the fuel reached full flow, and the earlier job without the wider combustor failed during start-up; none of these failed runs is used here. The heat release assumed by the reduced-order model, 2.97 MJ per kilogram of air, is about four times the heat these duct shapes take while the flow stays supersonic at flight Mach 5. With the model’s own heat release these inviscid duct shapes would choke thermally; closing that gap needs a larger area ratio (a different nozzle layout), a momentum-consistent reduced-order combustor, or a higher flight Mach number.

5.5 Grid convergence

The diverging-combustor jobs were repeated at four resolutions, 2737, 5592, 8990 and 12 168 cells, by changing only the number of cells across the capture height (20, 29, 36 and 42); the geometry, fuel, ignition source and run time are identical. Station values are the averages over the duct height used above, and the impulse flux is $\int (p + \rho u^2) dy$ per metre of span. Tables 19 and 20 give the values, the change from the third to the fourth grid (“last step”), the spread over all four grids relative to the finest, and, where the four grids converge monotonically, a grid-convergence index [14] on the finest grid (safety factor 1.25, observed order capped at 2, cell size taken as $N^{-1/2}$). Because the flow contains shocks, the formal order is only an estimate.

Table 19: Non-reacting grid study (ideal air, inviscid). A dash means the four grids do not converge monotonically, so no index is given.

Quantity	2737	5592	8990	12 168	Last step [%]	Spread [%]	GCI [%]
Isolator exit M	3.304	3.327	3.339	3.342	+0.1	1.1	0.3
Isolator exit p [kPa]	421.8	427.3	429.2	430.4	+0.3	2.0	1.2
Isolator exit impulse [kN/m]	376.3	383.1	384.9	385.9	+0.3	2.5	0.9
Combustor exit M	1.493	1.523	1.489	1.481	-0.5	2.8	–
Combustor exit p [kPa]	1193	1181	1226	1243	+1.4	5.0	–
Combustor exit T [K]	1676	1646	1678	1684	+0.3	2.2	–
Combustor exit impulse [kN/m]	388.4	394.8	395.2	396.2	+0.3	2.0	–
Nozzle-lip plane p [kPa]	1180	1173	1172	1172	+0.0	0.7	–
Nozzle-lip plane impulse [kN/m]	386.8	393.7	395.5	396.3	+0.2	2.4	0.7

Table 20: Reacting grid study (fuel at 10 % of the model flow, numerical ignition source). A dash means the four grids do not converge monotonically, so no index is given.

Quantity	2737	5592	8990	12 168	Last step [%]	Spread [%]	GCI [%]
Combustor exit M	1.89	2.02	1.94	1.95	+0.6	6.5	–
Combustor exit p [kPa]	923	853	912	905	-0.8	7.8	–
Combustor exit T [K]	1402	1331	1389	1390	+0.1	5.1	–
Combustor exit impulse [kN/m]	391.8	397.7	400.4	400.9	+0.1	2.3	0.4
Nozzle-lip plane impulse [kN/m]	392.5	398.8	401.6	402.2	+0.2	2.4	0.5
Combustor exit $\eta_{\text{H}_2\text{O}}$	0.72	0.62	0.71	0.60	-19.3	20.5	–

In the non-reacting case the isolator-exit state, the impulse flux and the nozzle-lip pressure change by 0.1 to 0.4 % between the two finest grids, with indices of about 1 % or less. The combustor-exit Mach number and pressure are less settled: they do not converge monotonically, the pressure changes by 1.4 % in the last step and by 5 % over the four grids, because the position of the waves on the diverging wall moves from grid to grid. In the reacting case the impulse flux is robust (a spread of about 2 %), but the combustor-exit Mach number, pressure and temperature scatter by 5 to 8 % without a trend, and the fraction of hydrogen converted, taken from the water formed, is 0.72, 0.62, 0.71 and 0.60 on the four grids. The 5592-cell reacting case was run twice, once from the start and once restarted from a saved state, and the two agree exactly, so the scatter is a property of the grids and not of the restart. With a numerical ignition source and no molecular mixing, the burning is set by where the source and the strip of fuel meet the waves in the duct, and that changes with the mesh. The two-dimensional result therefore supports the impulse flux to a few percent and the combustor-exit state to about 10 %, but it does not fix the combustion efficiency better than about ± 0.1 .

5.6 Three-dimensional case

A three-dimensional half-vehicle job is provided with the notebook (Cell 16f). Only a coarse smoke test of that job has been run, and no vehicle three-dimensional result is reported in this paper.

A second, smaller job (Cell 16h) extrudes the reacting two-dimensional job across the engine width: the same inlet, isolator, combustor and nozzle between side walls, with no vehicle body, nine-species chemistry, the same fuel cap and the same numerical ignition source. It is a half-span structured mesh (116 870 cells at the default resolution, in 38 blocks), estimated at about 26 core-hours, a figure scaled from the measured run below. Only a 19 159-cell version (20 cells across the capture height, seven across the half span, uniform fuel) was run to the end. It reproduces the two-dimensional job on the same cell size to the digits of table 20: combustor-exit Mach number 1.89, pressure 923 kPa, temperature 1402 K, impulse flux 391.8 kN m^{-1} and $\eta_{\text{H}_2\text{O}} = 0.72$, with no variation across the span. That is the expected result with inviscid slip side walls and uniform fuel, so the run tests the three-dimensional set-up and nothing more. It does not show side-wall boundary layers or corner flows, and so says nothing yet about their effect on starting or burning; that needs viscous flow with wall cells of about $1 \mu\text{m}$ at this Reynolds number, which is beyond a single-machine run. An option to inject the same hydrogen through spanwise columns, which does give spanwise variation, starts and sources correctly in a short test but has not been run to a result.

6 Validation Summary

Table 21: Sub-model validation summary.

Sub-model	Test	Result
θ - β - M	NACA 1135, $M_\infty = 2$, 4 cases	PASS (max $\Delta\beta$ 0.112° , max rel. Δp $1.6\text{e-}03$)
θ - β - M	Root-find residual, 5 cases	PASS (max $1.1\text{e-}12^\circ$)
Isolator	Closed form vs quadrature and vs $\bar{c}_f$ form (solved design)	PASS (rel 10^{-15})
Combustor	T_4 monotonic in ϕ	PASS (5 values)
Nozzle	$F_{\text{gross}} = \dot{m}_{\text{gas}}U_e - \dot{m}_{\text{air}}V_\infty$ ($p_e = p_\infty$)	PASS (rel 10^{-15})
Inlet-isolator	Isolator inflow = ramp-2 exit state	PASS (rel 10^{-15})
Geometry	Shock 2 lands on lip	PASS (resid $2.8\text{e-}17$ m)
Geometry	Lip gap $\times \cos \theta_2 = A_{\text{eff}}/W$ (continuity)	PASS (rel 10^{-15})
Ramp search	Optimum $\geq$ brute-force grid (0.5° steps)	PASS (π 0.900606 vs grid 0.897059)
Ramp search	Equals Oswatitsch equal-strength optimum	PASS ($ \Delta\pi $ $2.2\text{e-}16$)
Inverse solve	Thrust closure	PASS (rel $9.3\text{e-}13$)
Constraints	π_{total} , p_2/p_∞ , H_c within limits	PASS (3 of 3)

These sub-model checks test internal consistency and the tabulated shock relations. The only comparison against an independent flow solver is section 5; it is inviscid and two-dimensional, and its grid dependence is quantified in section 5.5.

7 Limitations and Future Work

Not modelled: transition at the engine stations beyond the flat-plate linear-stability estimate of section 3.7.4 (two-dimensional second-mode waves only, assumed wall temperature, no shock-boundary-layer interaction at the corner, no oblique first-mode waves, thresholds used outside the range of their fit, and the trained neural-network weights of the Barraza transport equation are not bound to this framework); the effect of the predicted early transition on the laminar friction

and heating used by the reduced-order model; turbulent boundary layer; shock-boundary-layer interaction; finite-rate chemistry in the reduced-order model (the Eilmer reacting job of section 5 is a separate check whose combustion efficiency is not grid-converged, section 5.5); ignition and flameholding; fuel–air mixing; three-dimensional effects (no three-dimensional result is reported); vehicle-level thrust–drag balance, lift, mass and range; cooling and film; off-design operation (apart from the Mach march of the accompanying notebook); the SERN and vehicle layout of section 3.6 are planar, inviscid concept-level models with assumed body outlines; spillage and cowl-lip bluntness (the capture streamtube is assumed to enter exactly at the lip); the turning of the ramp-2 flow onto the isolator axis and any isolator shock train (the isolator inflow is set to the ramp-2 exit state); and combustor total-pressure loss other than through the calibrated static-pressure ratio α .

Future work: grid-converged, viscous, mixing-resolved CFD on the emitted geometry, with and without the numerical ignition source; an ignition model in the reduced-order combustor; a shock-train isolator model; applying the linear-stability solver of section 3.7.3 to the engine ramps with their computed wall temperature and adding oblique first-mode waves; binding the Barraza NN weights; replacing the calibrated combustor with finite-rate chemistry; selecting a preferred member of the isolator-length family.

8 Conclusion

RPDS-X is a reduced-order, inverse-first framework for two-ramp external-compression scramjet inlet design. Three contributions: the complete open-source implementation; the inverse-thrust architecture; and the constrained max-recovery two-ramp search, verified against the Oswatitsch equal-strength solution. The baseline mission at $M_\infty = 5.00$, $h = 3000$ m, $F_{\text{req}} = 30$ kN produces $\pi_{\text{total}} = 0.9006$, $H_c = 5.908$ cm, $\dot{m}_{\text{H}_2} = 2653.8787$ kg h^{−1}, and a total axial engine length of 128.8 cm, and a reference design at $L_{\text{iso}}/D_h = 6.0$ with axial engine length 144.6 cm. The framework is offered as a fast, reproducible starting point for high-fidelity CFD studies.

References

- [1] Ames Research Staff. Equations, tables, and charts for compressible flow. Technical Report NACA-TR-1135, National Advisory Committee for Aeronautics, Washington, D.C., 1953.
- [2] B. Barraza, A. Gross, C. Hader, and H. F. Fasel. Hypersonic transition model for second mode and crossflow instabilities. In *AIAA AVIATION 2024 Forum*, 2024. AIAA paper number to be completed; equations 20–22 checked against the author’s copy.
- [3] F. S. Billig. Propulsion systems from takeoff to high-speed flight. In S. N. B. Murthy and E. T. Curran, editors, *High-Speed Flight Propulsion Systems*, volume 137 of *Progress in Astronautics and Aeronautics*, pages 21–102. AIAA, Washington, D.C., 1991.
- [4] D. A. Bittker and V. J. Scullin. General chemical kinetics computer program for static and flow reactions, with application to combustion and shock-tube kinetics. Technical Report TN D-6586, NASA, 1972.
- [5] E. T. Curran and S. N. B. Murthy, editors. *Scramjet Propulsion*, volume 189 of *Progress in Astronautics and Aeronautics*. AIAA, Reston, VA, 2001.
- [6] G. Erlebacher and M. Y. Hussaini. Non-linear evolution of a second mode wave in supersonic boundary layers. Technical Report 89-15, ICASE, 1989.

- [7] N. N. Gibbons, K. A. Damm, P. A. Jacobs, and R. J. Gollan. Eilmer: an open-source multi-physics hypersonic flow solver. *Computer Physics Communications*, 282:108551, 2023.
- [8] D. G. Goodwin, H. K. Moffat, I. Schoegl, R. L. Speth, and B. W. Weber. Cantera: An object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes. <https://www.cantera.org>. Version 3.1.0.
- [9] W. H. Heiser and D. T. Pratt. *Hypersonic Airbreathing Propulsion*. AIAA Education Series, Reston, VA, 1994.
- [10] J. Li, Z. Zhao, A. Kazakov, and F. L. Dryer. An updated comprehensive kinetic model of hydrogen combustion. *International Journal of Chemical Kinetics*, 36(10):566–575, 2004.
- [11] L. M. Mack. Transition prediction and linear stability theory. In *Laminar-Turbulent Transition*, number 224 in AGARD Conference Proceedings, pages 1–1–22, 1977.
- [12] M. R. Malik. Numerical methods for hypersonic boundary layer stability. *Journal of Computational Physics*, 86(2):376–413, 1990.
- [13] K. Oswatitsch. Der druckrückgewinn bei überschallgeschwindigkeit. *ZAMM*, 24:218–222, 1944.
- [14] P. J. Roache. *Verification and Validation in Computational Science and Engineering*. Hermosa Publishers, Albuquerque, 1998.
- [15] B. Sharma and S. S. Girimaji. Prandtl number effects on the hydrodynamic stability of compressible boundary layers: flow-thermodynamic interactions. arXiv:2205.05742 (unpublished manuscript).