Multifunctional Leaf-Vein Micro-Channel Spar Architecture for Cryogenic Hydrogen Thermal Management and Structural Load-Bearing in Long-Endurance UAVs

Majd Chalak*

Independent Researcher

San Mateo, CA, USA

ABSTRACT

Long-endurance hydrogen PEM fuel-cell UAVs must pre-heat cryogenic hydrogen from $-150\,^{\circ}$C to 70–85 °C at the fuel cell inlet while simultaneously dissipating kilowatts of motor and stack waste heat. Conventional approaches require separate dedicated subsystems, imposing mass penalties that reduce endurance.

This paper presents a *conceptual design* for a *leaf-vein micro-channel spar architecture* in which the primary wing spars of the MARID UAV simultaneously carry structural bending and torsion loads and serve as cryogenic hydrogen thermal conduits. The architecture uses a *paired-bridging serpentine* channel network: adjacent secondary channels are connected in series by mechanisms of 30 micro-channels each, traversed single-file (no parallel splitting) rather than branching independently to a return line. Two such chains (front and rear of the main spar) run per wing, each a 27.23 m single-file path from root to tip and back, substantially reducing dedicated thermal-management hardware while routing cold hydrogen through thermally critical zones to suppress the vehicle's infrared (IR) signature—a key requirement for covert ISR missions. (This pass covers the main wings only; V-tail chains are deferred to future work.)

A one-dimensional finite-segment thermal model, with an ambient-conductance path (internal convection, Ti wall, CFRP skin, an assumed low-observable coating layer, and external convection in series) and a Darcy-Weisbach pressure-drop calculation, shows that the extended path length yields strong passive pre-heating: with internal waste heat sources (motor, avionics, PEM cooling jacket) included, sea-level ambient heat exchange along the wing raises the outlet from $-85.1\,^{\circ}$C (internal sources alone) to 51.3 °C, and an electric trim heater of only 86 W at sea level (3.3% of total electrical demand) closes the balance to 77.6 °C within the 70–85 °C PEM window. The flow regime is *turbulent* ($Re \approx 42{,}900$), since only 4 parallel chains share the design flow across a comparatively small number of micro-channels; the resulting pressure drop, including minor losses at the 87 hairpin turns within the zigzag routing ($\approx$145 kPa per chain, with turn losses adding $\approx$39% on top of the friction-only baseline), is non-negligible but still yields a small total pump power ($\approx$2.1 W, $\approx$0.1% of electrical demand) given the low volumetric flow. A wing-exit temperature check—relevant to IR signature, since wall temperature there governs external skin emission—shows the hydrogen (and wing wall) closing most, but not all, of the gap to ambient by the time flow exits the wing structure, at every altitude tested, before any downstream trim heating occurs: a residual of a few degrees remains, attributable to the added thermal resistance of the assumed low-observable coating, itself included specifically because it also suppresses IR

emission from the hotter internal components upstream.

I. INTRODUCTION

A. *Motivation*

Hydrogen PEM fuel cells offer energy densities an order of magnitude higher than lithium-ion batteries at the system level, making them the enabling technology for UAV endurance missions exceeding 24 hours [1]. However, PEM stacks operate optimally between 60 and 85 °C; membrane degradation accelerates outside this window [2]. Cryogenic hydrogen storage—whether compressed gaseous (GH_2, $-150\,°C$ at 350 bar) or liquid (LH_2, $-253\,°C$)— therefore requires a pre-heating system before the fuel cell inlet.

Simultaneously, electric motors, the PEM stack, and avionics collectively dissipate several kilowatts of waste heat. On a long-endurance ISR platform, this heat must be managed to prevent structural overtemperature and, critically, to suppress the vehicle's infrared plume. Conventional designs route separate cooling loops through heat exchangers and radiators, adding significant dry mass [3].

B. *Proposed Approach*

This paper proposes eliminating the dedicated H_2 pre-heater and substantially reducing PEM cooling hardware by routing cryogenic hydrogen directly through the wing spars. The spars simultaneously carry aerodynamic bending and torsion loads and act as distributed heat exchangers. As cold H_2 flows from the cryogenic storage vessel outboard through the wing, it absorbs waste heat deposited at motor mounts and the PEM bay, arriving at the fuel cell inlet pre-heated to the required temperature. No dedicated H_2 pre-heater is required, and a portion of the PEM stack's cooling duty is offloaded onto the same hydrogen flow (Section 7.1).

Because the wing chain network permanently holds hydrogen in transit rather than passing it through momentarily, the spar also functions as a distributed buffer volume: at the design operating point, the four wing chains together hold approximately 93 g of H_2 (2.6 L of internal volume) at any instant — a modest $\approx 1.2\%$ of the 48-hour fuel budget, but a real fourth function (structural, thermal, fuel delivery, and buffer) layered onto the same physical structure.

The architecture is inspired by biological leaf-vein networks, in which a hierarchical branching structure simultaneously provides mechanical support and fluid transport—a well-established example of structural multifunctionality in nature [4].

C. *Related Work*

Multifunctional structures that combine load-bearing and fluid transport functions have been explored for satellite thermal control [5] and battery thermal management in electric vehicles [6]. Micro-channel heat exchangers embedded in composite panels have been demonstrated for airborne electronics cooling [7]. However, to the author's knowledge, the application of this concept to cryogenic hydrogen pre-heating within a primary structural spar—where the working fluid is also the fuel—has not been previously reported.

D. Paper Contributions

This paper makes the following contributions:

1. A dual-function wing spar architecture that eliminates the dedicated H_2 pre-heater and substantially reduces waste heat rejection hardware.
2. A one-dimensional finite-segment thermal model validated against the 70–85 °C PEM inlet requirement for the MARID 113.6 kg UAV.
3. Quantification of the recirculation pump requirement (22.6:1 thermal-to-fuel flow ratio) and its small ($< 0.1\%$ of electrical demand) but non-negligible pressure-drop budget under the paired-bridging architecture.
4. A working fluid trade study comparing GH_2 against LH_2 as the two physically viable candidates.
5. Analysis of the IR signature suppression mechanism enabled by selective cold-hydrogen routing.

II. MARID PLATFORM

A. Vehicle Overview

The MARID (Multi-Axis Rotary-wing Integrated Drone) is a fixed-wing UAV with independently rotating main wings and V-tail surfaces that provide full pitch, roll, and yaw authority without conventional control surfaces [12]. Key geometric and mass properties, extracted from the ROS2/Gazebo simulation model, are summarized in Table 1.

B. Propulsion and Energy System

The primary energy source is a PEM hydrogen fuel cell stack supplemented by a lithium battery buffer that serves two roles: (i) transient peak power during takeoff and manoeuvres, and (ii) sole electrical supply during cold-start, before the fuel cell reaches operating temperature. Hydrogen is stored cryogenically and fed to the fuel cell after passing through the micro-channel spar system described in Section 3.

Battery Sizing

At cold-start the trim heater (Section 3.5) must pre-heat the H_2 supply before the PEM stack can sustain output. The battery must therefore cover two sequential energy demands:

1. **Trim heater phase** (PEM offline, ≈10 min): worst-case heater power 2,103 W yields $2{,}103 \times \frac{10}{60} \approx 350\,\mathrm{Wh}$.

Table 1: MARID vehicle properties (from simulation URDF); mass is a provisional structural baseline only—see note below

Parameter	Value	Source
Structural mass (URDF)†	113.6 kg	Link mass summation
Fuselage length	3.0 m	Tail joint at $x = -2.56$ m
Wing semispan (joint to tip)	1.5 m	Inertial CoM at $y = 0.61$ m
Wing root offset from CL	0.50 m	Wing joint y-position
Wing mass (each)	3.662 kg	URDF inertial
Thruster configuration	1 (single)	Top-mounted aft-fuselage pusher
Wing rotation limit	$\pm 8.6°$	Joint limits ± 0.15 rad
Design endurance	>48 hr	Design objective
Design mission	Covert ISR	—

† The URDF mass of 113.6 kg represents the structural airframe as modeled in simulation only. It excludes: the single high-power thruster system (motor + propeller + ESC, estimated 7–16 kg); cryogenic hydrogen storage and tank (fuel + LH_2 tank: estimated 43–53 kg); startup and contingency battery pack (960 Wh at 1.30× safety, estimated 3.8–4.8 kg); PEM fuel cell stack; and sensor payload. Accounting for these, real MTOW is estimated at **175–210 kg**. All thermal analyses in this paper scale proportionally with gross mass and remain architecturally valid at the corrected MTOW.

2. **Takeoff transient** (≈3 min, PEM at 50% capacity): net battery demand $\approx 7{,}800\,\mathrm{W} \times \frac{3}{60} \approx 390\,\mathrm{Wh}$.

The combined baseline energy requirement is therefore ≈740 Wh. Applying a 1.30 safety and redundancy factor gives a design capacity of

$$E_{\text{bat}} = 740 \times 1.30 \approx 960\,\mathrm{Wh}. \tag{1}$$

At a cell-level specific energy of 200–250 Wh kg^{-1} (LiPo to NMC Li-ion) the pack mass is estimated at **3.8–4.8 kg**. Once the PEM is at full output the battery is trickle-charged from excess fuel-cell capacity and held at ≈80% state of charge as a contingency reserve for the remainder of the mission.

C. Heat Sources

Table 2 lists the primary waste heat sources and their locations along the hydrogen flow path, as derived from the power budget in Section 4.

III. LEAF-VEIN MICRO-CHANNEL SPAR ARCHITECTURE

A. Hierarchical Channel Network

Figure 1 illustrates the general engineering concept underlying the leaf-vein thermal architecture: a closed hydrogen circuit integrated into the wing spar structure, in which cold hydrogen delivered from the fuselage circulates through the wing, absorbs waste heat, and returns pre-heated to the fuel

Table 2: Waste heat sources and spatial distribution

Zone	URDF location	Cruise (W)	Takeoff (W)
Aft motor mount (single)	top-mounted, $x \approx -2.0\,\text{m}$	221	1,000
Fuselage avionics	$x \approx -1.5\,\text{m}$	80	160
PEM fuel cell stack	CoG at $x = -1.876\,\text{m}$	2611	7833
Total		**2,912**	**8,993**

cell inlet, while the same spar structure simultaneously carries aerodynamic bending and torsion loads. The specific channel routing shown is illustrative of this underlying principle rather than an as-built rendering of the paired-bridging architecture detailed in this section; a CAD rendering specific to the paired-bridging architecture is intentionally withheld from this version of the paper pending prosecution of a corresponding non-provisional patent application, with the architecture instead fully specified in text and in Table 3 below.

The main spar runs root to tip and divides the wing chordwise into a front side and a rear side; each side independently carries its own closed hydrogen circuit, so a wing has two parallel, symmetric *chains*. On each side, cold hydrogen travels a short 150 mm stub (structural bore, $D_i = 27\,\text{mm}$) from the rotary union to the first of four *secondary channels*, connected in series. Between each adjacent pair of secondaries, a *paired-bridging mechanism* of 30 micro-channels ($D_i = 1.5\,\text{mm}$, 276 mm each) carries the flow single-file—no parallel splitting—in a staggered, back-and-forth pattern that maximises heat-exchange surface area within a compact chordwise span. With 4 secondaries in series, a chain crosses 3 such mechanisms (90 micro-channels total), connected by 2 short secondary-channel legs ($D_i = 3.7\,\text{mm}$, 220 mm each). At the last secondary, flow turns and returns along a 1,600 mm perimeter pipe ($D_i = 19\,\text{mm}$; see Section 3.2 for why Murray's Law does not govern this bore) back to the wing root, into an L-connector feeding the outer annulus of the rotary union.

Each secondary channel functions as a pressure-relief checkpoint: with no parallel redundancy inside a chain, a blockage anywhere along the 27 m single-file run has nowhere else to go, so under overpressure conditions hydrogen in the micro-channels is redirected, within the sealed circuit, back into the adjoining secondary and onward to either the main delivery stub or the return pipe, without dedicated relief-valve hardware. This is an internal flow-redistribution path within the closed hydrogen circuit, not an external leak or loss of containment. This is also why pressure relief appears at the main spar itself (the sole parallel-split point) as well as at every secondary.

Cryogenic hydrogen enters the wing circuit through a *dual-channel rotary union* mounted at the wing-root pivot, using a concentric-tube arrangement (inner bore: cold delivery; outer annulus: warm return), with the bearing housing simultaneously transferring wing bending and torsion loads into the fuselage. Per wing, this gives 2 chains $\times$ 3 mechanisms $\times$ 30 micro-channels = 180 micro-channels; across both wings, 360 micro-channels total. V-tail chains are not yet re-derived under this architecture and are treated as future work (Section 7.4).

B. Channel Geometry and Path Length

The dimensions presented below reflect the current preliminary design iteration and will be refined as structural optimisation and detailed thermal analysis mature toward a flight-ready configuration.

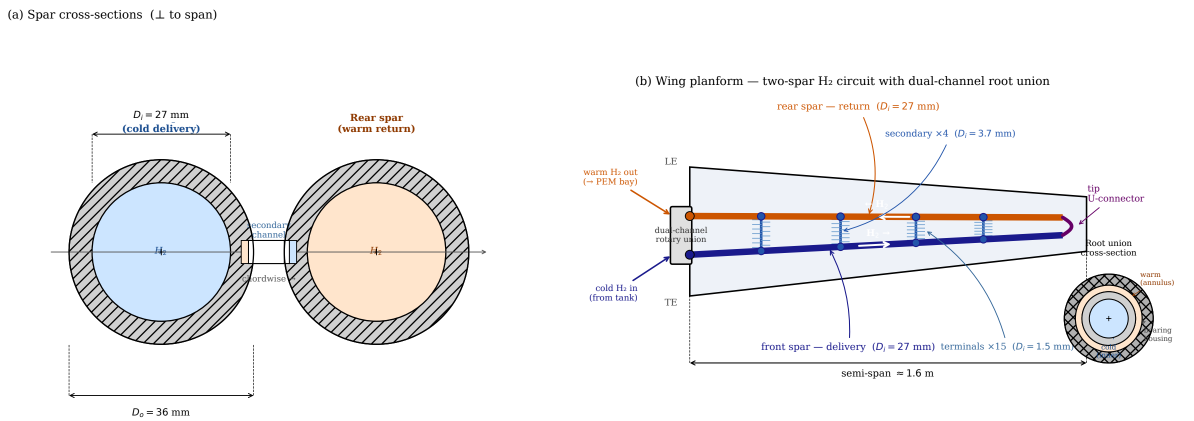

Figure 1: **Conceptual illustration of the leaf-vein thermal circuit.** A representative dual-channel closed circuit: front spar (cold delivery, blue) and rear spar (warm return, orange), with parallel chordwise secondaries branching to terminal micro-channels, illustrating the general principle of routing cryogenic hydrogen through the wing structure for combined structural and thermal duty. The paired-bridging architecture described in the text refines this general concept into a single-file, series-connected chain per side of the main spar, detailed in Table 3; a CAD rendering specific to that refined architecture is withheld pending patent prosecution.

Murray's Law ($D^3_{\text{parent}} = \sum_i D^3_{\text{child},i}$) minimises pumping power by maintaining constant wall shear stress at each parent/child split in a branching network. The paired-bridging architecture eliminates branching almost entirely: within a chain, flow travels single-file from the root stub, through all 90 series micro-channels, to the return pipe, with no point at which one conduit splits into multiple daughters or multiple conduits merge into one. Murray's Law is defined at branch points; a chain simply has none downstream of the root stub, so the law does not apply to any of the micro-channel, connector, or return-pipe bores—all three are set at reasonable starting values pending a topology-specific derivation (e.g., a frictional-loss-minimisation trade against the pump power budget of Section 6). The only genuine branch point remaining in the architecture is the root-level split into the 4 parallel wing chains (Section 3), which is not itself Murray's-Law-derived in this pass—the 4 chains are simply assumed to share the design flow equally by symmetry. Table 3 summarises the per-chain path segmentation.

Common parameters: wall thickness = 0.75 mm; material = Ti-6Al-4V ($k = 7\,\mathrm{W\,m^{-1}K^{-1}}$); operating pressure = 350 bar; channel wall temperature range -150 to $+100$ °C.

C. Structural Spar Sizing

The main spar is a hollow Ti-6Al-4V shaft that rotates relative to the fuselage through cross-roller bearings and is driven by a high-torque brushless motor (harmonic drive) to provide full pitch/roll/yaw authority. It must simultaneously carry:

1. Wing root bending moment from aerodynamic lift
2. Motor torsion transmitted to the wing
3. Provide adequate shaft diameter for the motor clamp interface

Table 3: Paired-bridging chain geometry (one representative chain; by symmetry all 4 wing chains are identical under this model)

Segment	Length	Count	D_i	Sizing driver
Root stub (root $\rightarrow$ 1st secondary)	150 mm	1	27.0 mm	Structural (bending + torsion)
Zigzag (3 mech. $\times$ 30 micro-ch, series)	24,840 mm	90	1.5 mm	Heat transfer / manufacturability
Connecting legs (between mechanisms)	440 mm	2	3.7 mm	Preliminary sizing
Return pipe (perimeter, tip $\rightarrow$ root)	1,600 mm	1	19.0 mm	Preliminary sizing
Fuselage routing	200 mm	—	—	Downstream of wing
Total path length, one chain	**27,230 mm (27.23 m)**			

Parallel chains: 2 per wing (front/rear of main spar) $\times$ 2 wings = 4 total, sharing the wing design flow rate; each chain is single-file with no further internal parallel splitting. Aircraft-wide (wings only, this pass): 4 chains $\times$ 90 micro-channels = 360 total.

At the 2.5g limit load factor, each front wing generates a root bending moment of approximately 1,575 N·m (ultimate). Combined with the 700 N·m torsion limit (URDF joint specification), the equivalent moment is:

$$M_{\text{equiv}} = \sqrt{M_b^2 + T^2} = \sqrt{1575^2 + 700^2} \approx 1{,}723\,\text{N·m} \tag{2}$$

Using Ti-6Al-4V allowable stress $\sigma_{\text{allow}} = 587\,\text{MPa}$ (yield $\div$ 1.5) and a tube wall ratio $k = D_i/D_o = 0.75$, the required outer diameter is $D_o \approx 36\,\text{mm}$, yielding $D_i \approx 27\,\text{mm}$, since the root stub carries the full bending and torsion load regardless of the downstream channel topology. For context, a thermal-only bore estimated using Murray's Law for a generic, purely-branching (non-series) channel topology—the kind of sizing rule that would apply to a conventional branching design, not this paper's paired-bridging architecture—works out to only ≈5.9 mm. This figure is offered only as a rough sense of scale, not a bore requirement derived for the paired-bridging architecture presented here and not a valid apples-to-apples comparison: it has not been derived for the paired-bridging topology (Section 3.2), which does not use Murray's Law for its downstream bores at all, and is retained here only as an indication that the structural requirement dominates rather than as a precise thermal-minimum bore for this architecture. A rigorous thermal-minimum bore would need to be re-derived from the per-chain flow rate and turbulent heat-transfer requirement, but given the wide margin already evident at this looser estimate, the qualitative conclusion—the structural bore comfortably accommodates the H_2 flow path—is not expected to change. The V-tail spar (500 N·m limit, lower aerodynamic loads) sizes to $D_i \approx 18\,\text{mm}$ by the same procedure.

A rotary union at the wing root allows pressurised hydrogen to pass from the fixed fuselage plumbing into the rotating spar; such unions are commercially available for high-pressure applications and represent a known engineering solution for this interface.

D. Material Selection

Ti-6Al-4V is selected for the spar and channel walls on the basis of its high specific strength, good fatigue resistance under thermal cycling, and good resistance to hydrogen embrittlement under anticipated operating conditions (expected hydrogen partial pressures) [8]. 316L stainless steel is an alternative for the channel inserts, offering higher thermal conductivity ($k = 16\,\mathrm{W\,m^{-1}K^{-1}}$) at the cost of higher density. CFRP outer skins are bonded over the metallic spar core, providing aerodynamic surface and additional bending stiffness while maintaining radar-cross-section discipline through controlled edge alignment. The specific CFRP layup, resin system, and through-thickness conductivity used in the ambient heat-exchange model (Section 4.2) remain a placeholder pending dedicated material selection research and are flagged in Section 7.4.

E. Electric Trim Heater

The extended paired-bridging path substantially reduces, but does not eliminate, the need for supplemental heating: passive sources alone (internal waste heat + ambient) are insufficient to guarantee 70–85 °C at the PEM inlet across all altitudes. The architecture therefore retains a compact inline *electric trim heater*, now positioned within the final 200 mm fuselage-routing segment (downstream of the wing structure entirely), between the avionics zone and the PEM cooling jacket.

Operating Modes

The trim heater operates in three modes (Table 4):

1. **Startup** (PEM cold, no waste heat available): sized separately from the cruise analysis below; the startup transient has not yet been re-derived for this architecture, so a conservative worst-case figure (2,103 W) is retained pending re-analysis (Section 7.4).
2. **Altitude / reduced load:** the trim heater compensates for reduced ambient pre-heating at altitude. Required power by ISA altitude is given in Table 4.
3. **Nominal steady-state cruise (primary mode):** at sea level, with internal waste heat sources (motor, avionics, PEM cooling jacket) included, ambient heat exchange along the wing raises the outlet to 51.3 °C, so the trim heater's cruise duty is modest.

Table 4: Trim heater power by ISA altitude (wing chains only, this pass)

Altitude	Ambient	T_{out} (no trim)	P_{trim} (% of elec. demand)
Sea level	+15.0 °C	51.3 °C	86 W (3.3%)
2,000 m	+2.0 °C	39.3 °C	126 W (4.8%)
5,000 m	−17.5 °C	21.2 °C	185 W (7.1%)
Tropopause	−56.5 °C	−15.0 °C	303 W (11.6%)
No ambient (worst case)	—	−85.1 °C	527 W (20.2%)

Sizing and Control

The cruise-mode trim heater is sized for the worst credible case (no ambient contribution at all, e.g., cold-soaked ground operation): 527 W. Passive internal sources alone (motor, avionics, PEM cooling jacket, without any ambient contribution) bring the outlet to only $-85.1\,^{\circ}\mathrm{C}$, so in this worst case the trim heater must close most of the remaining gap to the 70–85 °C target itself; it is ambient heat exchange along the extended path, not internal waste heat alone, that does most of the passive pre-heating work under normal (ambient-present) operation. At sea-level cruise the nominal trim heater power is only 86 W (3.3% of total electrical demand). Implementation (resistance element, closed-loop PID on PEM-inlet thermocouple) follows standard practice for inline trim heaters.

An *electric* trim heater, rather than a gated PEM-to-hydrogen heat path, is retained specifically because of the startup case: PEM waste heat does not exist until the stack is producing power, so it cannot substitute for the trim heater during the cold-start transient (Table 4 Startup mode). At steady-state cruise, by contrast, PEM's own uncaptured waste heat ($\approx$2,111 W, Section 7.1) is far larger than the trim heater's entire duty range (86–527 W across all conditions modelled), so a suitably gated PEM-to-hydrogen interface could in principle substitute for the electric trim heater once the stack is running; retaining an independent electric element keeps the architecture functional across the full startup-to-cruise envelope without relying on a heat source unavailable at the moment it is first needed.

F. IR Signature Suppression

The single aft-fuselage motor mount and the PEM bay exhaust ports are the primary IR emitters on the vehicle. The motor is deliberately positioned adjacent to the LH_2 storage tank at the aft fuselage, where the cryogenic tank surface (20–120 K) and cold H_2 vapour boil-off provide a passive cold shroud around the dominant heat source. This co-location suppresses the motor mount IR plume without any additional hardware.

The extended paired-bridging path introduces a second, quantifiable IR-relevant effect at the *wing* structure itself. The simulation tracks fluid and wall temperature at the wing-exit checkpoint—the transition from the return pipe into the fuselage-routing segment, i.e. after the hydrogen has traversed the full wing network but before reaching any heat-generating equipment. With the low-observable coating's added thermal resistance included in the ambient-conductance network (Section 4.2), the fluid temperature at this checkpoint no longer converges to within a fraction of a degree of local ambient; instead it closes between 94.2% and 94.4% of the total temperature gap between the flow immediately downstream of the root stub and local ambient, at every altitude tested, leaving a modest residual (Table 5) rather than a near-exact match.

Table 5: Wing-exit fluid and wall temperature residual relative to local ambient

Altitude	Ambient	Residual (fluid)	Residual (wall)
Sea level	$+15.0\,^{\circ}\mathrm{C}$	$-8.1\,^{\circ}\mathrm{C}$	$-5.9\,^{\circ}\mathrm{C}$
2,000 m	$+2.0\,^{\circ}\mathrm{C}$	$-7.3\,^{\circ}\mathrm{C}$	$-5.3\,^{\circ}\mathrm{C}$
5,000 m	$-17.5\,^{\circ}\mathrm{C}$	$-6.2\,^{\circ}\mathrm{C}$	$-4.5\,^{\circ}\mathrm{C}$
Tropopause	$-56.5\,^{\circ}\mathrm{C}$	$-4.0\,^{\circ}\mathrm{C}$	$-2.9\,^{\circ}\mathrm{C}$

Practically, this means the wing skin runs a few degrees cooler than the surrounding air rather than presenting no thermal contrast at all: a substantially smaller signature than the tens-to-hundreds of degrees implied by the internal-only case ($-85.1\,^{\circ}$C at the outlet), but a measurable one rather than an exact null. The residual IR signature is therefore reduced at the wing but not eliminated there, with the dominant remaining contribution still concentrated at the aft fuselage (motor mount, PEM bay, trim heater segment), where the passive cryogenic shrouding described above applies.

This incomplete convergence is a direct, and in this case deliberate, consequence of the low-observable coating included in the ambient-conductance network. Fluid temperature approaches ambient with a characteristic length set by $\dot{m}c_p/UA$; the coating's own conduction resistance dominates the combined external-side resistance, reducing the effective external-side conductance from ≈80 to $\approx 11.4\,\mathrm{W\,m^{-2}K^{-1}}$ at the terminal micro-channel bore (Section 4.2) and correspondingly lengthening the characteristic relaxation length relative to an uncoated skin. Where an uncoated wing chain would close 90% of its approach to ambient by $x \approx 8$ m and 99% by $x \approx 16$ m—well inside the 27.03 m wing-internal path—the coated chain modelled here reaches 90% only by $x \approx 25$ m, and does not reach 99% at all within the wing-internal path, capping out at the 94.2–94.4% figures above by the time flow exits the wing. In this sense the same coating included to suppress radar and infrared observability of the airframe's hotter internal components (motor, PEM bay) necessarily also insulates the wing skin somewhat from the ambient airstream, trading a small, deliberate residual thermal contrast at the wing for a lower-conductivity, presumably lower-observability skin overall—a physically coherent trade-off rather than a modelling shortfall. Qualitatively, the radiative power from a surface scales as T^4; reducing a motor mount from 80 °C (353 K) to 40 °C (313 K) reduces emitted power by approximately 38%. The present model predicts surface temperature, not detector response; translating these surface temperatures into a predicted signal at a given sensor band and range requires a radiometric detector model, which is outside the scope of this paper. Quantitative IR signature modelling, including of the coating's combined radar- and IR-observability trade-off at the wing, is deferred to future work integrating a Gazebo thermal plugin.

IV. THERMAL MODEL

For simplicity, the present 1-D model deposits motor waste heat at the wing root stub rather than at its physical aft-fuselage location (Section 7.4); this affects only where along the 27.23 m path the motor's contribution is added, not the total amount of waste heat modelled.

A. Power Budget

Table 6 summarises the power budget for the MARID at the cruise design point. The aerodynamic drag is estimated assuming a lift-to-drag ratio of 14, representative of a clean medium-aspect-ratio fixed-wing platform at cruise.

B. One-Dimensional Finite-Segment Formulation

The 27.23 m per-chain path (root stub + 90 series micro-channels across three paired-bridging mechanisms + connecting legs + return pipe + fuselage routing, Section 3.2) is discretised into

Table 6: MARID power budget at cruise (25 m/s, $L/D = 14$)

Quantity	Value	Notes
Gross mass	113.6 kg	URDF
Aerodynamic drag	79.6 N	$W/(L/D)$
Required shaft power	1,990 W	$F_D \times V$
Motor input power (single thruster)	2,211 W	$\eta_{\text{motor}} = 0.90$; cruise sizing
Avionics + actuators	400 W	Estimate
Total electrical demand	2,611 W	
PEM chemical power input	5,222 W	$\eta_{\text{PEM}} = 0.50$
H_2 consumption (cruise)	0.044 g/s	7.52 kg per 48 hr
Motor waste heat	221 W	
PEM stack waste heat	2,611 W	
Miscellaneous waste heat	80 W	
Total cruise waste heat	**2,912 W**	
Total takeoff waste heat	**8,993 W**	single high-power motor at takeoff thrust

$N = 2{,}000$ equal segments of length $\Delta x = L/N \approx 13.6$ mm. Segment diameter and exposure (wing-internal vs. fuselage-internal) are looked up per segment from the geometry table. At each segment i the fluid temperature T_i is updated by an energy balance:

$$T_{i+1} = T_i + \frac{\dot{q}_i \, \Delta x}{\dot{m}_{\text{chain}} \, c_p(T_i)} \tag{3}$$

where $\dot{q}_i$ (W m^{-1}) is the distributed heat flux at position x_i and $c_p(T_i)$ is the temperature-dependent specific heat of gaseous hydrogen. Because the four wing chains (two sides $\times$ two wings) are geometrically and thermally identical, the per-segment mass flow uses the per-chain rate $\dot{m}_{\text{chain}}$ rather than the aircraft-total rate, and the results below apply uniformly to all four chains. The channel wall temperature is:

$$T_{\text{wall},i} = T_i + \frac{\dot{q}_i}{h_i \, \pi D_i} \tag{4}$$

where h_i is the convective heat transfer coefficient derived from the Nusselt number correlation appropriate to the local flow regime and segment diameter (Section 4.4). This locally-computed h_i is also used directly in the ambient heat-exchange term (rather than a separate fixed coefficient): the ambient conductance per unit length is built from a five-resistance series network—internal convection (h_i), titanium channel-wall conduction, CFRP outer-skin conduction, a low-observable (radar/IR) coating layer, and external convection—so that the wall-temperature and ambient-exchange calculations remain mutually consistent at every segment.

The external convection coefficient is a flat-plate-turbulent order-of-magnitude estimate, $H_{\text{ext}} \approx 80$ W m^{-2}K^{-1}, consistent with the 25 m/s, $L/D = 14$ cruise condition used throughout the Power Budget (Table 6); this is deliberately kept consistent with that reference speed rather than independently corrected to the higher $\approx$60 m/s speed used in separate CFD runs elsewhere in this work, since reconciling the whole power/drag/fuel budget to that higher speed is a larger undertaking left to future work (Section 7.4). The coating layer represents an assumed low-observable (radar/IR signature reduction) treatment on the outer skin: because the specific coating design is a classified topic outside the scope of this paper, its thickness ($t_{\text{coat}} = 2.0\,\text{mm}$) and through-thickness conductivity ($k_{\text{coat}} = 0.035\,\text{W}\,\text{m}^{-1}\text{K}^{-1}$) are an engineering guesstimate rather than a specification, deliberately biased toward a thick, low-conductivity, foam-like broadband layer as a conservative placeholder. Combined in series with the external convection term, this coating resistance reduces the effective external-side conductance from ≈80 to $\approx 11.4\,\text{W}\,\text{m}^{-2}\text{K}^{-1}$ at the terminal micro-channel bore—a net ≈7× reduction. The CFRP conduction, coating, and external convection terms all use placeholder property values pending final material selection and any disclosable coating data (Section 7.4).

C. Hydrogen Thermophysical Properties

Thermophysical properties of gaseous hydrogen at 350 bar are approximated by polynomial fits to NIST WebBook data over the range 100–420 K:

$$c_p(T) = 12{,}200 + 8.0(T - 100) - 0.012(T - 100)^2 \quad [\text{J}\,\text{kg}^{-1}\text{K}^{-1}] \tag{5}$$

$$k(T) = 0.065 + 5.83 \times 10^{-4}(T - 100) \quad [\text{W}\,\text{m}^{-1}\text{K}^{-1}] \tag{6}$$

$$\mu(T) = [4.2 + 0.0163(T - 100)] \times 10^{-6} \quad [\text{Pa}\,\text{s}] \tag{7}$$

The mean specific heat over the operating range (123–353 K) is approximately 13,500 $\text{J}\,\text{kg}^{-1}\text{K}^{-1}$—nearly 3.2× that of water and 6.8× that of thermal oil.

Density, used throughout (Darcy-Weisbach pressure drop and pump volumetric flow, Section 6; buffer-volume mass, Section 1.2), is computed from the ideal-gas relation

$$\rho(T, P) = \frac{P}{R_{\text{H}_2} T}, \qquad R_{\text{H}_2} = 4{,}124\ \text{J}\,\text{kg}^{-1}\text{K}^{-1} \tag{8}$$

evaluated at the local fluid temperature and the 350 bar storage pressure. Real hydrogen at 350 bar deviates from ideal-gas behaviour (compressibility factor $Z \gtrsim 1$ in this pressure/temperature range), which would push actual density, and therefore the buffer-volume mass estimate and the exact pressure-drop and pump-power figures, measurably below this ideal-gas estimate. This is flagged as a simplification pending a real-gas (compressibility-factor) correction, consistent with the other engineering-estimate placeholders in this paper.

D. Nusselt Number Correlation

The Reynolds number per channel is:

$$Re = \frac{4\,\dot{m}_{\text{ch}}}{\pi D_i \mu} \tag{9}$$

For $Re < 2{,}300$ (laminar), uniform-heat-flux internal pipe flow: $Nu = 4.36$. In the transition region ($2{,}300 < Re < 10{,}000$), the Gnielinski correlation is applied [9]. For fully turbulent flow ($Re > 10{,}000$), Dittus-Boelter is used: $Nu = 0.023\,Re^{0.8}\,Pr^{0.4}$. All three regimes are evaluated per-segment using the local diameter, so the correlation switches automatically as the flow transitions between the narrow (1.5 mm) micro-channel bore and the wider connector/return/fuselage bores along the path.

Because the extended path length increases per-channel Reynolds number well past 10,000 for the design flow rate (Section 4.5.1), pressure drop is non-negligible and must be carried alongside the thermal solution. The Darcy friction factor is evaluated per segment as $f = 64/Re$ (laminar) or the Blasius correlation $f = 0.316\,Re^{-0.25}$ (turbulent, smooth-wall, valid to $Re \lesssim 10^5$), and the segment pressure drop follows the Darcy-Weisbach relation:

$$\Delta p_i = f_i\,\frac{\Delta x}{D_i}\,\frac{\rho v_i^2}{2} \tag{10}$$

summed over all N segments to give the friction (major loss) component of the total per-chain pressure drop. This alone does not capture the full picture: the staggered, back-and-forth routing of the zigzag section (Section 3) requires the flow to execute a series of tight hairpin turns, each of which contributes an additional *minor loss* not represented by straight-pipe friction. Section 6 quantifies this contribution and combines it with the friction total above.

E. Results

4.5.1 Flow Regime

At the design cruise flow rate of 0.984 g/s distributed across the four wing chains (two sides × two wings), the per-chain mass flow is 245.9 mg/s, because each chain carries the full delivery flow through a single-file series path rather than splitting it across many parallel terminal channels. Reynolds number therefore varies substantially with local bore diameter along the path: in the 1.5 mm-bore micro-channels (root of the zigzag section)

$$Re_{\text{zigzag}} = \frac{4\,\dot{m}_{\text{chain}}}{\pi D_{\text{micro}}\,\mu} \approx 42{,}900 \tag{11}$$

placing the micro-channels and the 3.7 mm connectors ($Re \approx 17{,}400$) firmly in the *turbulent* regime (Dittus-Boelter, Section 4.4). The wider root stub (27 mm) sits close to the laminar/transitional boundary ($Re \approx 2{,}460$), while the 19 mm return pipe and fuselage segment operate modestly into the transitional regime ($Re \approx 3{,}380$); both are evaluated with the Gnielinski correlation. The paired-bridging architecture achieves substantially higher convective heat transfer coefficients in the micro-channel section as a result, at the cost of a non-trivial recirculation pump power budget (Section 6).

4.5.2 Temperature Profiles and Ambient Pre-Heating

Figure 2(a) shows the H_2 temperature profile along the 27.23 m per-chain path (0.15 m root stub + 24.84 m zigzag micro-channel section + 0.44 m connecting legs + 1.60 m return pipe + 0.20 m fuselage routing) for three cases at the design per-chain flow rate of 245.9 mg/s (22.6:1 recirculation). Heat zones are anchored to absolute position rather than path fraction. There is exactly one motor, one avionics suite, and one PEM stack aircraft-wide; since each of the 4 parallel, symmetric chains carries its proportional 1/4 share of these single aircraft-wide sources, the motor mounts (55.3 W cruise, i.e. 221.1 W aircraft-wide ÷ 4) sit at the root stub (0–0.15 m), while the avionics zone (20.0 W, i.e. 80 W÷4), electric trim heater, and PEM cooling jacket (125.0 W, i.e. 500 W÷4) are all located within the final fuselage-routing segment, downstream of the entire wing structure. The 500 W aircraft-wide PEM cooling jacket figure is only a fraction of the PEM stack's full 2,611 W waste-heat duty (Table 6); it is the portion of PEM waste heat assumed to be routed through the hydrogen loop, not the stack's total heat rejection, and the implications of that choice for claimed mass savings are discussed in Section 7.1.

The dash-dot grey curve (internal heat sources only, no ambient, no trim heater) reveals the thermal path structure: a small temperature rise at 0–0.15 m (motor mounts), a long flat plateau across the 24.84 m zigzag section and 0.44 m connecting legs and 1.60 m return pipe (no internal sources along the wing itself), and a step at the very end of the path where the avionics zone and PEM cooling jacket add heat. Without ambient or trim heating the outlet reaches $-85.1\,°C$, well below the 70–85 °C target: the modest per-chain share of internal waste heat (200.3 W total) is not, on its own, enough to substantially pre-heat the flow. This internal-only case isolates the contribution of motor, avionics, and PEM waste heat alone, providing the baseline against which the ambient and trim heater contributions below are added.

The solid blue curve adds ambient pre-heating through the CFRP-skinned, coated titanium channel walls along the entire wing-internal path (root stub through return pipe). Cryogenic hydrogen at $-150\,°C$ is separated from ambient air at $+15\,°C$ (sea level ISA) by the five-resistance conductance network described in Section 4.2 (internal convection, titanium wall, CFRP skin, low-observable coating, external convection), evaluated locally at each segment using the actual local convective coefficient. Because the path is long and largely turbulent (Section 4.5.1), ambient heat exchange—combined with the same internal waste heat sources above—is still effective despite the added coating resistance: at the design flow rate the outlet reaches $51.3\,°C$—short of, but approaching, the target window without any trim heater contribution. It is this combination of ambient exchange along the extended path plus internal waste heat, not ambient alone, that closes most of the energy balance.

The heavy red curve adds the electric trim heater (Section 3.5) at its sea-level design point of 86 W. H_2 exits at $77.6\,°C$—within the 70–85 °C PEM operating window. The grey shaded band in Figure 2(a) marks the final 200 mm fuselage segment, where the trim heater is now located; because this segment is only ≈0.7% of the 27.23 m path, panel (b) zooms into it directly, showing the avionics, trim heater, and PEM cooling jacket zones individually.

4.5.3 Outlet Temperature vs. Flow Rate — Altitude Parametric

Figure 2(c) shows outlet temperature as a function of per-chain flow rate for four ISA conditions: sea level (+15 °C), 2,000 m (+2 °C), 5,000 m ($-17.5\,°C$), and tropopause ($-56.5\,°C$). No trim heater is applied in these sweeps. The green band marks the 70–85 °C target. Every altitude curve crosses into and through the target band as flow rate is reduced below the 0.984 g/s (22.6:1 recirculation) design point: at ∼75% of design flow, sea level already exceeds 70 °C, and the remaining altitudes cross the band at successively lower flow fractions as ambient assistance falls off. This is the direct consequence of the long, largely turbulent ambient-exposed path: because outlet temperature is highly sensitive to flow rate, *flow rate modulation* becomes a viable complementary (or, in principle, alternative) means of PEM-inlet temperature control alongside the electric trim heater. The present design retains the fixed 0.984 g/s / 22.6:1 recirculation operating point with the trim heater as primary control, since flow-rate-based control introduces coupled transient dynamics (residence time changes) that are left to future work (Section 7.4).

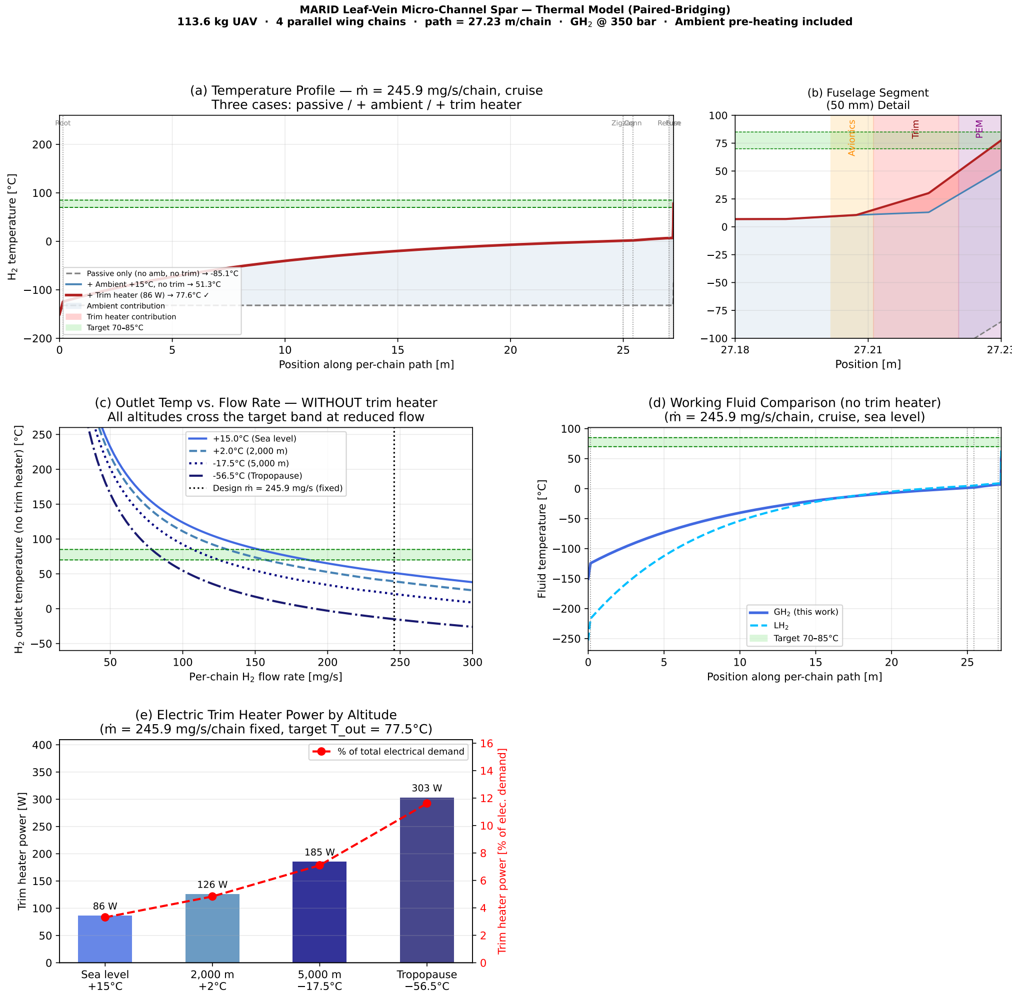

Figure 2: Thermal model results (paired-bridging architecture, 27.23 m per-chain path, design flow 245.9 mg/s/chain, 22.6:1 recirculation). (a) H_2 temperature profile along the per-chain path for three cases: dash-dot grey = internal sources only (exits $-85.1\,°C$); solid blue = +sea-level ambient pre-heating (exits $51.3\,°C$); heavy red = +electric trim heater 86 W (exits $77.6\,°C$ ✓); red fill = trim heater contribution; blue fill = ambient contribution; grey shaded band = final 200 mm fuselage segment, the final 50 mm of which is detailed in (b); green band = 70–85 °C PEM operating window. (b) Zoomed detail of the final 50 mm of the fuselage segment from (a), showing the avionics zone, trim heater zone, and PEM cooling jacket (orange/red/purple shading respectively), invisible at the full 27.23 m scale of (a). (c) Outlet temperature vs. per-chain flow rate for four altitude conditions (no trim heater); all curves cross the target band at reduced flow. (d) Working fluid comparison at design flow, sea level. (e) Required trim heater power by ISA altitude (bars) and fraction of total electrical demand (red line); sea level 86 W (3.3%), tropopause 303 W (11.6%).

4.5.4 Maximum Wall Temperature

The peak channel wall temperature occurs in the zigzag micro-channel section near the PEM-side end of the path, where turbulent internal convection ($Re \approx 42{,}900$) drives the wall temperature close to the local fluid temperature. Because the internal-only baseline outlet is $-85.1\,°\mathrm{C}$ (well below the trim-heater-assisted $77.6\,°\mathrm{C}$ design case), the peak wall temperature across the full thermal duty cycle remains well below both the yield temperature of the Ti-6Al-4V channel liner (>900 °C) and the CFRP skin's approximate epoxy service ceiling (≈150 °C, pending final material selection, Section 7.4). In the wider root stub and return-pipe segments ($D_i = 27\,\mathrm{mm}$ and 19 mm respectively), the larger bore and correspondingly lower local convective coefficient keep the wall temperature closer to the ambient/CFRP-skin side of the five-resistance network rather than the fluid side. The channel network is structurally safe across the full thermal duty cycle.

V. WORKING FLUID TRADE STUDY

Figure 2(d) compares the two physically viable candidate working fluids, each at its own correct storage inlet temperature, at the design per-chain flow rate (245.9 mg/s) under cruise heat loading, sea level, no trim heater. Table 7 summarises the results.

Table 7: Working fluid comparison at design per-chain flow (245.9 mg/s, sea level)

Fluid	c_p (J kg^{-1}K^{-1})	ΔT (K)	Notes
GH_2 (this work)	13,500	201 K	Fuel = coolant; no added mass; exits 51.3 °C
LH_2	9,700	323 K	Own $-253\,°\mathrm{C}$ storage inlet; exits 69.5 °C; phase-change complexity

Gaseous hydrogen is the fuel itself: no additional working fluid, pump, or reservoir is required, since the recirculation pump circulates fuel that is being consumed anyway. Non-cryogenic surrogates such as thermal oil ($c_p \approx 2{,}000\,\mathrm{J\,kg^{-1}K^{-1}}$) and water ($c_p \approx 4{,}186\,\mathrm{J\,kg^{-1}K^{-1}}$) are not evaluated as candidates here: both would be frozen solid at cryogenic hydrogen-line temperatures (typical thermal-oil pour points are no lower than about $-90\,°\mathrm{C}$, and water freezes at $0\,°\mathrm{C}$) and so could not actually be pumped through the micro-channels, making a simulated temperature-profile curve for either fluid physically meaningless at this inlet condition. Even setting that aside, introducing a genuine second fluid loop would add mass, complexity, and a heat-exchanger interface to transfer heat back into the actual H_2 fuel stream before the PEM inlet, defeating the single-fluid simplicity that motivates the architecture.

Liquid hydrogen is evaluated at its own, physically correct storage inlet of $-253\,°\mathrm{C}$ (saturated liquid at 1 bar), not the GH_2 dense-gas inlet of $-150\,°\mathrm{C}$ used for the other rows—the two are different storage states and must not share an inlet temperature. Despite starting $103\,°\mathrm{C}$ colder than GH_2, LH_2 exits *warmer*, and by more than its lower specific heat alone would suggest—two compounding mechanisms are responsible. First, its assumed constant $c_p = 9{,}700\,\mathrm{J\,kg^{-1}K^{-1}}$ is about 25% lower than GH_2's actual temperature-dependent c_p averaged over its own path (≈12,960 J kg^{-1}K^{-1}), so for a given amount of heat picked up along the path, LH_2's lower specific heat converts that heat into a larger temperature rise ($\Delta T = Q/\dot{m}c_p$)—a factor of ≈1.34× on its own. Second, because

LH_2 starts and remains colder than GH_2 over most of the path, the ambient driving temperature difference ($T_{\text{amb}} - T_{\text{fluid}}$) is larger at every point, so LH_2 genuinely absorbs ≈20% *more* total heat along the same path than GH_2 does, not merely converting the same heat more efficiently. Combined, these two effects (≈1.34× from c_p, ≈1.20× from the larger driving ΔT) give ≈1.60× the temperature rise of GH_2, matching the simulated ΔT ratio closely and, together, more than offsetting LH_2's colder start. Over the additional $\sim 100\,°\text{C}$ of warming this implies, LH_2's outlet temperature (69.5 °C) falls just short of the 70–85 °C target window at this flow rate, by less than a degree—a consequence of the long ambient-exposed path being just barely insufficient once the added coating resistance in the ambient path (Section 4.2) is included. Unlike GH_2, whose shortfall is closed by the electric trim heater in the architecture as modelled, an LH_2-fuelled variant would need either a small trim-heater contribution of its own or a modest flow-rate reduction to close this last degree; this is a follow-on design detail rather than an architecture-level obstacle, given how close the constant-c_p estimate already comes. This result should be read with the same caveat noted throughout: the constant-c_p model used here does not capture LH_2's phase-change physics (boiling onset at $-253\,°\text{C}$ at 1 bar), the associated latent heat, or potential two-phase flow instability in micro-channels, all of which GH_2 avoids entirely by remaining single-phase across the full path. The present analysis recommends GH_2 for initial prototype development, with LH_2 as a follow-on option once single-phase GH_2 behaviour is validated and a two-phase model is developed for the LH_2 case.

VI. RECIRCULATION PUMP ANALYSIS

A. Required Flow Rate vs. Fuel Consumption

At cruise, the fuel cell consumes hydrogen at $\dot{m}_{\text{fuel}} = 0.044$ g/s. The thermal management system operates at a fixed total wing-chain flow of 0.984 g/s (245.9 mg/s across each of the four parallel chains), giving an aircraft-wide *recirculation ratio* of 22.6:1, set by the cruise waste-heat energy balance (Table 6). The flow rate is held constant across altitudes; the electric trim heater (Section 3.5) adjusts its power output to maintain the 70–85 °C PEM inlet temperature rather than varying the flow rate, although Section 4.5 notes that flow-rate modulation is now also a physically viable control channel. V-tail chains are not yet included in this design flow (deferred, Section 7.4); the 0.984 g/s figure covers the wing chains only.

The fuel supply system must incorporate a recirculation pump that circulates 22.6 volumes of hydrogen for every volume consumed by the fuel cell, split four ways across the parallel wing chains. The surplus hydrogen is recirculated to the cryo storage vessel or through a bypass loop.

B. Pump Power

The paired-bridging chain is a single-file series path (Section 3.2) with no parallel splitting within a chain, so pressure drop is computed as a straightforward sum of Darcy-Weisbach losses along the path (Eq. 10). Because the micro-channel and connector segments operate in the turbulent regime (Section 4.5.1), the Blasius friction factor dominates over most of the path length; only the wider root-stub, return-pipe, and fuselage segments contribute laminar/transitional losses.

Integrating Eq. 10 over all $N = 2{,}000$ segments at the design per-chain flow rate (245.9 mg/s) gives a friction (major loss) pressure drop of

$$\Delta P_{\text{friction}} \approx 104{,}232 \text{ Pa}, \tag{12}$$

the direct consequence of routing the full chain flow through a 1.5 mm bore in series over 24.84 m, rather than splitting it across many parallel terminal channels. This is lower than an otherwise-identical uncoated case would give, because the added coating resistance (Section 4.2) leaves the flow colder, and therefore denser and slower at fixed mass flow, over more of the path.

This friction term alone understates the true pressure drop, because it treats the zigzag section as a straight pipe. In reality, the staggered, back-and-forth routing that connects the 30 series micro-channels of each paired-bridging mechanism requires the flow to execute $N_{\text{turns/mech}} = 29$ internal 180° hairpin turns per mechanism; across three mechanisms this gives $N_{\text{turns}} = 3 \times 29 = 87$ turns per chain (the two mechanism-to-mechanism transitions route through the wider connector legs rather than a sharp hairpin, and are excluded from this count). Each turn contributes a minor loss

$$\Delta p_{\text{turn}} = K_{\text{turn}} \frac{\rho v^2}{2}, \tag{13}$$

evaluated at the local fluid density and micro-channel velocity at that turn's position along the thermal profile, using a representative hairpin coefficient $K_{\text{turn}} \approx 1.75$ (literature range ≈1.5–2.2 for a tight-radius 180° return bend in small-bore tubing; not independently verified for this specific geometry). Summing Eq. 13 over all 87 turns at the design flow rate gives a total minor-loss contribution of

$$\Delta P_{\text{minor}} = \sum_{j=1}^{87} K_{\text{turn}} \frac{\rho_j v_j^2}{2} \approx 40{,}680 \text{ Pa}, \tag{14}$$

an addition of ≈39% on top of the friction-only total. The combined per-chain pressure drop is therefore

$$\Delta P_{\text{total}} = \Delta P_{\text{friction}} + \Delta P_{\text{minor}} \approx 144{,}912 \text{ Pa} \approx 21.0 \text{ psi per chain}. \tag{15}$$

The pump power per chain, at 60% pump efficiency and volumetric flow $\dot{Q} = \dot{m}_{\text{chain}}/\rho(T_{\text{in}})$:

$$P_{\text{pump,chain}} = \frac{\Delta P_{\text{total}} \times \dot{Q}_{\text{chain}}}{\eta_{\text{pump}}} \approx 516.5 \text{ mW}, \tag{16}$$

and for all four wing chains combined:

$$P_{\text{pump,total}} \approx 2{,}065.8 \text{ mW} \approx 2.1 \text{ W}. \tag{17}$$

This remains a small fraction (≈0.1%) of the 2,611 W total electrical demand, so the pump power budget is not a design driver. The pressure drop itself, at ∼21 psi per chain (more than a quarter of it from hairpin-turn minor losses rather than straight-pipe friction), is a non-negligible design consideration for recirculation pump sizing (a miniature cryogenic centrifugal or gear pump capable of that head is required), though it is a minor addition ($< 0.5\%$) on top of the 350 bar storage pressure the micro-channel wall must already be sized to hold, so it does not materially change the wall-thickness design driver. Because K_{turn} is an engineering estimate rather than a measured or CFD-verified value for this exact hairpin geometry, the pump and wall-thickness sizing should carry margin against this uncertainty; a bench measurement or CFD study of the actual turn geometry is flagged as an item for structural follow-up (Section 7.4).

C. Implications for 48-Hour Endurance

The 48-hour design endurance requires 7.52 kg of H_2 at the computed fuel-cell consumption rate (Table 6). The recirculation pump itself adds no hydrogen mass consumption. The electric trim heater draws electrical power from the fuel cell rather than consuming additional H_2 directly; however, increased electrical demand reduces overall fuel cell efficiency.

In the *nominal cruise* operating mode the trim heater operates at a modest design point (86 W at sea level, Section 3.5), because passive ambient pre-heating along the 27.23 m path supplies most of the required energy. The endurance penalty of continuous cruise-mode trim heater operation is correspondingly small.

If the trim heater were operated continuously at the sea-level design power of 86 W, the additional H_2 consumption would be approximately $86/(0.5 \times 120 \times 10^6) \approx 0.00144$ g/s, adding ≈0.25 kg of H_2 over 48 hours—a ≈3.3% increase in the H_2 budget. The trim heater's operating profile:

- **Startup** (0–10 min): sized separately; the startup transient has not yet been re-derived for this architecture (Section 3.5), so a conservative figure ($\lesssim$2,100 W) is retained pending re-analysis.
- **Altitude excursions / reduced PEM load**: variable supplement per Table 4, up to 527 W worst case (no ambient contribution at all).
- **Nominal cruise**: 86 W at sea level, falling further as PEM waste heat and ambient pre-heating combine to close more of the energy balance.

Under this operating profile the 48-hour H_2 budget remains close to 7.52 kg, with the trim-heater contribution small even under continuous worst-case cruise operation.

Whether the airframe can accommodate 7.52 kg of compressed hydrogen within the structural and volumetric budget is a separate system-level trade (GH_2 at 350 bar requires ≈313 L; LH_2 at 70 kg m^{-3} requires ≈107 L) and is outside the scope of this paper. The thermal architecture presented here is compatible with either storage approach.

The numerical values reported here should be interpreted as first-order engineering predictions rather than validated performance metrics. While absolute outlet temperatures and trim-heater requirements depend on assumed convection coefficients, coating properties, and material parameters, the architectural conclusions are more robust: extending the hydrogen flow path through a multifunctional structural spar substantially increases passive heat recovery while requiring only modest recirculation power. Future CFD and experimental testing will refine the quantitative values reported here.

VII. DISCUSSION

A. Mass Savings

A conventional thermal management system for a UAV of this class would require: (i) a H_2 pre-heater heat exchanger, estimated at 0.5–2 kg, sized to bring the full consumed H_2 flow from cryogenic storage temperature to the PEM inlet target; (ii) a coolant pump and reservoir for PEM stack cooling, estimated at 1–3 kg, sized for the full 2,611 W PEM waste-heat duty (Table 6); and (iii) associated plumbing, fittings, and insulation for both subsystems.

The leaf-vein spar architecture fully eliminates item (i): all consumed H_2 is pre-heated via the combination of ambient exchange, internal waste heat, and the electric trim heater, with no separate heat exchanger required. Item (ii) is only partially addressed. The "PEM cooling jacket" heat zone captures 500 W of the PEM stack's 2,611 W total waste heat (19.2%) — this is a design choice about how much PEM waste heat is routed through the hydrogen loop at the fuselage end of the path (Section 4.5.2), not a finding that the remaining ≈2,111 W requires no rejection. That remaining heat still requires a separate liquid PEM-cooling subsystem external to this architecture; the leaf-vein spar reduces, but does not eliminate, that subsystem's duty, and this paper does not claim its mass is eliminated.

The net mass saving is therefore dominated by item (i)'s full elimination (0.5–2 kg) plus a partial reduction in item (iii)'s plumbing tied to the eliminated pre-heater, less the recirculation pump addition (< 0.1 kg): approximately 0.4–1.9 kg, translating to a modest but real increase in H_2 carrying capacity and endurance. Whether the reduced PEM-cooling duty allows any downsizing of the external coolant pump and reservoir — and therefore any additional mass credit beyond the figure above — is left to a dedicated thermal-system trade study.

B. Manufacturing Considerations

Embedding 360 micro-channel segments of 1.5 mm diameter in titanium spars (90 micro-channels in series per chain, across four parallel wing chains; wings only in this pass, V-tail deferred—see Section 7.4) is achievable by two established processes: (i) *electron beam melting (EBM)* additive manufacturing, which can produce complex internal channel networks in Ti-6Al-4V directly from CAD; and (ii) *diffusion bonding* of photo-chemically etched titanium foils, a mature process used in compact heat exchangers for aerospace applications. Channel tolerances of ± 0.05 mm are achievable with both methods [10, 11]. The paired-bridging zigzag pattern is geometrically more complex than a set of straight parallel terminal channels—each of the 90 series micro-channels per chain requires a U-turn bend connecting it to the next—which favours EBM over diffusion bonding for this architecture, since EBM natively produces arbitrary 3-D internal void geometry whereas diffusion-bonded foil stacks are best suited to channels that remain in a single etched layer.

C. Turbulent Flow Regime

The operating Reynolds number varies substantially with local bore diameter (Section 4.5.1): the 1.5 mm micro-channels and 3.7 mm connectors operate in the turbulent regime ($Re \approx 17{,}400$–$42{,}900$), while the wider root stub, return pipe, and fuselage segments sit near the laminar/transitional boundary ($Re \approx 2{,}460$–$3{,}380$). This has two design consequences. First, the turbulent convective coefficient is substantially higher than the constant $Nu = 4.36$ laminar value, which is the primary reason ambient pre-heating (combined with internal waste heat) can bring the outlet to within a modest trim-heater margin of the 70–85 °C target (Section 4.5.2). Second, the pressure drop is non-negligible: at ≈145 kPa/chain including hairpin-turn minor losses (Section 6), it is a non-negligible design consideration for recirculation pump sizing, though a minor addition against the 350 bar storage pressure the channel wall must already withstand, a trade-off the architecture accepts in exchange for the pre-heating benefit above. Should future iterations find the pressure drop excessive, mitigation options include a larger micro-channel bore (trading some heat transfer coefficient for lower ΔP, since $\Delta P \propto D^{-5}$ roughly in the turbulent regime) or reintroducing limited parallel splitting within each paired-bridging mechanism, at the cost of the pressure-relief-checkpoint benefit that the current series topology provides at each secondary (Section 3).

D. Limitations and Future Work

The present analysis makes several simplifying assumptions that will be addressed in future work:

- **1-D model**: Radial temperature gradients within the spar cross-section and circumferential conduction between channels are neglected. A full 3-D conjugate heat transfer model (e.g., in OpenFOAM or ANSYS Fluent) will be developed.
- **Cruise speed sensitivity**: The present analysis is evaluated at a single 25 m/s, $L/D = 14$ cruise condition. Future work will analyse and simulate the architecture across a range of cruise speeds, including the higher ≈60 m/s condition evaluated separately by CFD elsewhere in this work. Since external convective heat transfer scales with airspeed, faster cruise is expected to improve, not degrade, the thermal architecture's performance: stronger forced convection over the coated skin would further reduce the required trim heater power for a given ambient temperature, so the trim heater figures reported here should be read as conservative relative to higher-speed flight.
- **Constant heat loads**: Heat loads are assumed steady at the cruise and takeoff values. A transient model capturing the thermal time constants of the spar and fuel cell will be needed for control system design.
- **Structural analysis**: A detailed FEA of the spar and wing structure with 360 micro-channel voids under combined bending, torsion, and thermal loading is in progress. This must now also verify the channel wall against the additional ≈145 kPa/chain recirculation pump pressure drop (Section 6) on top of the 350 bar storage pressure.
- **Gazebo thermal simulation**: A lumped-parameter ROS2 node coupled to a Gazebo system plugin will simulate in-flight temperature evolution and IR signature maps, including the wing-exit thermal residual and coating trade-off identified in Section 3.6 (IR Signature Suppression), enabling real-time thermal management in the flight stack.
- **Sensor-driven adaptive heat routing**: The present architecture uses a single fixed heat-exchange topology at a fixed design flow rate. A future extension could use distributed temperature sensor data (wing skin, PEM inlet, ambient) to actively bias the recirculation flow, or a gated PEM-to-hydrogen interface (Section 7.1), toward whichever duty — body/structure

cooling or PEM pre-heating — most needs it at a given moment, rather than the fixed proportional split assumed here. This builds on the flow-rate-modulation control channel already identified in Section 4.5.3 and is left for future control-system work.

- **GH_2 vs LH_2 transition**: Phase-change effects for liquid hydrogen storage will be incorporated in a follow-on study.

- **Incomplete mass budget**: The 113.6 kg gross mass used throughout is the URDF structural baseline and understates the real MTOW. Missing contributions include the single high-power thruster system (motor + propeller + ESC, 7–16 kg), LH_2 fuel and cryogenic tank (43–53 kg), the 960 Wh startup and contingency battery pack (3.8–4.8 kg at 1.30× safety factor), PEM stack, and sensor payload, yielding an estimated real MTOW of **175–210 kg**. Power and drag figures scale proportionally; a full mass-budget closure is left for the system-level design study.

- **Motor heat-source location**: the thermal model deposits motor waste heat at the start of the series path (wing root stub, $x = 0$–0.15 m). With the single aft-fuselage motor, this heat source physically belongs at the fuselage end of the path instead. Updated 1-D and 3-D models reflecting the correct motor location will be presented in future work; this pass focused on re-deriving the wing routing geometry itself and carried this simplification forward.

- **V-tail chains not yet modelled**: This analysis covers only the four wing chains (front/rear side × two wings). The V-tail's two boom chains are not yet re-derived for the paired-bridging architecture and are excluded from the design flow rate, power budget, and pressure-drop figures reported here. Adding them is the most immediate next step for this model.

- **CFRP skin, coating, and external convection properties are placeholders**: The outer-skin conduction term in the ambient heat-exchange model (Section 4.2) uses assumed values for CFRP through-thickness conductivity ($K_{\text{CFRP}} = 0.6$ W/m K) and skin thickness (1.5 mm), pending dedicated material selection research. The external forced- convection coefficient is a flat-plate order-of-magnitude estimate at the 25 m/s cruise condition used throughout this paper ($H_{\text{ext}} \approx 80$ W/m^{2}K before the coating layer), not a boundary-layer analysis. A low-observable coating layer ($t_{\text{coat}} = 2.0$ mm, $k_{\text{coat}} = 0.035$ W/m K) is included as a deliberately conservative engineering guesstimate, since the actual coating design is a classified topic outside the scope of this paper; combined with the external convection term it reduces the effective external-side conductance to ≈11.4 W/m^{2}K, roughly a 7× reduction relative to the original uncoated placeholder. All three terms should be revisited once the CFRP material, coating, and external aerodynamic surface finish are finalised; the qualitative finding that ambient plus internal waste heat together substantially pre-heat the flow is not expected to reverse under reasonable parameter variation, but the exact trim heater power figures in Table 4, and the degree to which the wing-exit temperature approaches ambient (Section 3.6), will shift with these placeholders.

- **Hairpin-turn loss coefficient is an engineering estimate**: the minor-loss contribution from the 87 hairpin turns in the zigzag section (Section 6) uses a representative $K_{\text{turn}} \approx 1.75$ drawn from the general literature range for tight-radius 180° return bends in small-bore tubing, not a value measured or CFD-verified for this specific turn geometry (bend radius, surface finish, and any local flow separation are all unresolved). Since this term adds ≈39% on top of the friction-only pressure drop (roughly a quarter of the combined total), a bench test or CFD study of the actual turn geometry is a priority follow-up; the pump and channel-wall sizing should carry margin against this uncertainty in the meantime.

- **Murray's Law no longer governs the wing chain**: a chain has no branch points downstream of the root stub—it is single-file from root stub through 90 series micro-channels to the return pipe—so Murray's Law, which is defined at parent/child splits, does not apply to the micro-channel, connector, *or* return-pipe bore sizing (Section 3.2). All three are set at reasonable starting values, not yet re-derived for this topology, leaving room for a dedicated bore-optimisation pass (e.g., a frictional-loss vs. pump-power trade) in future work. The only remaining branch point in the architecture—the root-level 1-to-4 split into parallel wing chains—is likewise not yet Murray's-Law-derived; the 4 chains are simply assumed symmetric.

VIII. CONCLUSION

A paired-bridging leaf-vein micro-channel spar architecture has been proposed and analysed for the MARID hydrogen fuel-cell UAV (113.6 kg structural baseline; estimated real MTOW 175–210 kg), using a peer-to-peer network of 90 micro-channels in series per chain across four parallel wing chains (360 micro-channels total, wings only in this pass; V-tail deferred). The architecture embeds a 27.23 m per-chain path ($D_i = 1.5$ mm micro-channels) within the primary wing spar structure, allowing cryogenic hydrogen to serve simultaneously as the propellant and the thermal management working fluid, while the network's own internal volume (≈2.6 L, holding ≈93 g of H_2 in transit at any instant) gives the same structure a modest fourth role as a distributed buffer volume.

The one-dimensional thermal model predicts:

1. **Ambient and internal heat sources together do most of the work**: At the fixed design flow of 0.984 g/s (245.9 mg/s/chain, 22.6:1 recirculation), internal waste heat alone yields -85.1 °C at the outlet, since the modest per-chain share of internal waste heat (200.3 W total) is not, on its own, enough to substantially pre-heat the flow; adding sea-level ambient heat exchange along the long, largely turbulent path raises this to 51.3 °C—it is this combination of ambient exchange plus internal waste heat, not ambient alone, that closes most of the energy balance, though still short of the 70–85 °C target without a trim heater contribution.

2. **Trim heater requirement is modest**: A small electric trim heater of only 86 W (sea level) to 527 W (worst case, no ambient at all) closes the remaining gap to 77.6 °C at the design point, a direct consequence of the extended ambient-exposed path.

3. **Turbulent flow, non-negligible pump power**: The micro-channel and connector sections operate turbulently ($Re \approx$ 17,400–42,900), which drives the improved convective heat transfer above but also produces a pressure drop of ≈145 kPa/chain—minor losses at 87 hairpin turns in the zigzag routing add ≈39% on top of straight-pipe friction alone—and ≈2.1 W total pump power across four chains—still a small fraction of the 2,611 W electrical budget,

but a non-negligible design consideration for recirculation pump sizing (a minor addition against the 350 bar channel wall design pressure).

4. **Wing-exit temperature approaches, but does not exactly match, ambient (IR-relevant finding)**: At the checkpoint where flow leaves the wing structure and enters the fuselage, fluid temperature closes 94.2–94.4% of the gap to local ambient at every altitude tested, leaving a residual of 4–8 °C (Table 5)—substantially reduced from the tens-to-hundreds-of-degrees internal-only case, but not an exact null. The residual is attributable to the added thermal resistance of the assumed low-observable coating (Section 4.2), the same coating included specifically to suppress radar/IR observability of the hotter internal components; the dominant remaining IR signature is concentrated at the aft fuselage, where passive cryogenic shrouding already applies (Section 3.6).

5. **Flow-rate modulation is a viable control channel**: Because outlet temperature is highly sensitive to flow rate, reducing flow below the design point allows ambient pre-heating alone to reach the target band at every altitude tested, a control degree of freedom left for future control-system work.

6. **Structural spar**: The root-stub bore ($D_i = 27$ mm) is governed by structural bending and torsion loads. Because the paired-bridging chain has no branch points downstream of the root stub, Murray's Law does not govern *any* of the remaining bores—connector, micro-channel, or return pipe—all of which are set at reasonable starting values rather than derived for this topology, and are flagged for a dedicated bore-optimisation pass.

7. **Endurance impact remains small**: Continuous nominal-cruise trim heater operation adds only ≈0.25 kg (≈3.3%) to the 48-hour, 7.52 kg H_2 budget.

This analysis carries forward several simplifications unchanged from earlier design iterations (motor heat-source location, incomplete mass budget) and introduces open items specific to the paired-bridging architecture (V-tail chains not yet modelled, CFRP skin properties pending material research, channel pressure containment under the new ΔP), enumerated in Section 7.4. Future work will extend the analysis to full 3-D conjugate heat transfer, structural FEA with channel voids under combined thermal and pressure loading, transient thermal modelling, V-tail chain integration, and integration with the Gazebo-based flight simulation environment.

ACKNOWLEDGEMENTS

The author thanks the open-source ROS2, Gazebo, and FAST-LIO communities whose tools underpin the simulation environment used in this work.

NOMENCLATURE

c_p	specific heat at constant pressure, $\mathrm{J\,kg^{-1}K^{-1}}$
D_i	channel inner diameter, m
h	convective heat transfer coefficient, $\mathrm{W\,m^{-2}K^{-1}}$
k	thermal conductivity, $\mathrm{W\,m^{-1}K^{-1}}$
L	channel path length, m
$\dot{m}$	mass flow rate, $\mathrm{kg\,s^{-1}}$
N_{ch}	number of micro-channels
N_{chains}	number of parallel chains (4, wings only, this pass)
Nu	Nusselt number
f	Darcy friction factor
K_{turn}	hairpin-turn minor-loss coefficient
N_{turns}	number of hairpin turns per chain (87)
t_{coat}	low-observable coating thickness, m
k_{coat}	low-observable coating thermal conductivity, $\mathrm{W\,m^{-1}K^{-1}}$
H_{ext}	external forced-convection coefficient, $\mathrm{W\,m^{-2}K^{-1}}$
ΔP	pressure drop, Pa
$\dot{Q}$	volumetric flow rate, $\mathrm{m^3\,s^{-1}}$
$\dot{q}$	heat flux per unit length, $\mathrm{W\,m^{-1}}$
Re	Reynolds number
T	temperature, K or °C (stated in context)
μ	dynamic viscosity, Pa s
η	efficiency
Subscripts	
ch	per channel
chain	per parallel chain
fuel	hydrogen fuel consumption
thermal	thermal management circuit
wall	channel wall surface

REFERENCES

[1] Bradley, T. H., Moffitt, B. A., Parekh, D. E., and Mavris, D., "Development and Experimental Characterization of a Fuel Cell Powered Aircraft," *Journal of Power Sources*, Vol. 171, No. 2, 2007, pp. 793–801.

[2] Larminie, J., and Dicks, A., *Fuel Cell Systems Explained*, 2nd ed., Wiley, Chichester, UK, 2003.

[3] Gong, A., and Verstraete, D., "Fuel Cell Propulsion in Small Fixed-Wing Unmanned Aerial Vehicles: Current Status and Research Needs," *International Journal of Hydrogen Energy*, Vol. 42, No. 33, 2017, pp. 21311–21333.

[4] Katifori, E., Szöllősi, G. J., and Magnasco, M. O., "Damage and Fluctuations Induce Loops in Optimal Transport Networks," *Physical Review Letters*, Vol. 104, 2010, 048704.

[5] Silverman, E. M., Griese, R. W., and Frei, W., "Using Multifunctional Structure to Reduce Spacecraft Cost," *Acta Astronautica*, Vol. 52, Nos. 9–12, 2003, pp. 785–791.

[6] Heenan, T. M. M., et al., "An Advanced Microstructural and Electrochemical Datasheet on 18650 Li-Ion Batteries with Nickel-Rich NMC811 Cathodes and Graphite-Silicon Anodes," *Journal of the Electrochemical Society*, Vol. 167, 2020, 140530.

[7] Thomas, J. P., and Qidwai, M. A., "The Design and Application of Multifunctional Structure-Battery Materials Systems," *JOM*, Vol. 57, No. 3, 2005, pp. 18–24.

[8] San Marchi, C., and Somerday, B. P., *Technical Reference for Hydrogen Compatibility of Materials*, Sandia National Laboratories, SAND2012-7321, 2012.

[9] Gnielinski, V., "New Equations for Heat and Mass Transfer in Turbulent Pipe and Channel Flow," *International Chemical Engineering*, Vol. 16, No. 2, 1976, pp. 359–368.

[10] Murr, L. E., et al., "Metal Fabrication by Additive Manufacturing Using Laser and Electron Beam Melting Technologies," *Journal of Materials Science & Technology*, Vol. 28, No. 1, 2012, pp. 1–14.

[11] Martindale, B., and Bhave, R., "Perforated Plate Support for Micro-Channel Heat Exchangers," in *Compact Heat Exchangers for the Process Industries*, Begell House, 1997.

[12] Chalak, M., "Multi-Axis Rotary-Wing Integrated Drone (MARID)," *MChalak Aerospace Blog*, July 2024, `https://mchalakaerospace.blogspot.com/2024/07/multi-axis-rotary-wing-integrated.html` [retrieved April 2026]; source code and simulation model available at `https://github.com/MChalak1/MARID_UAV` [retrieved April 2026].