I.C.2E Propulsion Unit White Paper

Reverse-Flow Dual-Spool Ejector-Fan Outrunner Architecture

Design Analysis, Manufacturability, and Validation Roadmap

Prepared for Aerial MECHANICA Inc.

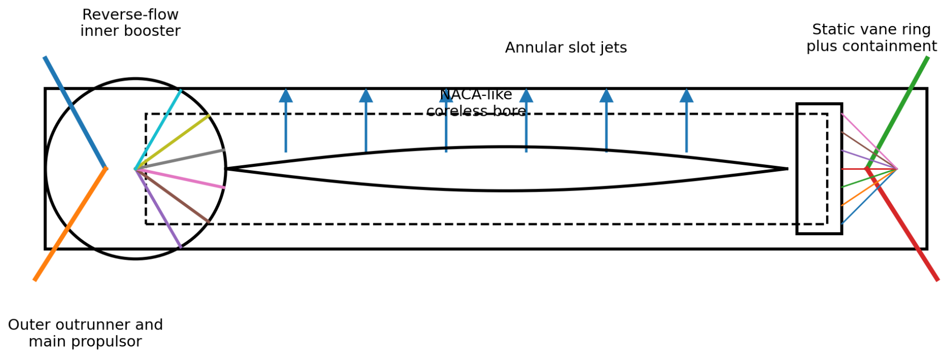

Document type:	Technical white paper
Program:	I.C.2E integrated propulsion unit
Prepared by:	**Arinze N. N. Eze** (technical drafting support)
Status:	Conceptual engineering package for internal design, business, and IP work
Date:	March 30, 2026

Engineering note. This document advances the I.C.2E as a high-efficiency integrated propulsion architecture intended to deliver substantially improved installed thrust density, thermal endurance, and operating robustness relative to conventional direct-drive units in the same integration envelope. Gains, if achieved, must come from effective-area recovery, controlled entrainment, swirl management, electro-thermal efficiency, and airframe integration rather than from a violation of conservation laws.

Contents

1 Executive Summary **1**

2 Context and Design State **3**

3 Requirements and Use Cases **5**

4 System Architecture **7**

5 Baseline Geometry and Internal Data **10**

6 Aerodynamic Foundations **12**

7 Reverse-Flow Inner Booster **14**

8 Coanda Ejector and Coreless Bore **16**

9 Outer Outrunner and Main Propulsor **18**

10 Electromagnetic Architecture **20**

11 High-Voltage Powertrain and Controls **22**

12 Thermal Management **24**

13 Materials and Process Selection **26**

14 Structural, Containment, and Rotordynamics **29**

15 Preliminary Performance Model **31**

16 Noise, Safety, and Failure Management **33**

17 Manufacturing Plan **35**

18 Test and Validation Roadmap **37**

19 Aircraft Integration for eVTOL and Drone Use **39**

20 IP and Program Roadmap **41**

21 Conclusion **43**

A Nomenclature **45**

B Key Equations **47**

B.1 Momentum-Theory Relations . . . 47
B.2 Tip Speed Estimate . . . 47
B.3 Electrical Loss Framing . . . 48
B.4 Hoop Stress Screening . . . 48

C Extracted Datasheet Tables **49**

C.1 Motor Baseline Summary . . . 49
C.2 Propulsor 27.8 x 8.5 LV . . . 49
C.3 Propulsor 27.8 x 8.5 HV . . . 50
C.4 Propulsor 27.8 x 10 LV . . . 50
C.5 Propulsor 27.8 x 10 HV . . . 51
C.6 Propulsor 27.8 x 12 LV . . . 51
C.7 Propulsor 27.8 x 12 HV . . . 52

D Figure and Modeling Notes **53**

D.1 Cross-Section Figure . . . 53
D.2 Airflow Figure . . . 53
D.3 Electromagnetic and Thermal Figures . . . 53
D.4 Containment, Manufacturing, and Validation Figures . . . 54

E Open Questions **55**

E.1 Aerodynamic Questions . . . 55
E.2 Electro-Thermal Questions . . . 55

E.3 Structural and Program Questions 55

F LBNL Sensing and Control Relevance 57

F.1 PRISM as an Advanced Development Diagnostic 57

F.2 DDRM and CALIPR as Methods for Unstable Test Systems 58

F.3 Recommended Positioning Inside the I.C.2E Program 58

G Design Decisions and Baseline Departures 59

G.1 Documented Baseline Versus Recommended Product Direction 59

G.2 Reason for Recommending an 18 Percent Opening 60

G.3 Preferred External Performance Clause 60

List of Figures

4.1 Concept cross-section of the reverse-flow dual-spool I.C.2E architecture. 7

4.2 Representative electrical and control architecture for a dual-spool I.C.2E unit. . . 9

6.1 Representative airflow logic for the bore slots, entrainment zone, and conditioned wake. 13

8.1 Representative radial-flux concept for a coreless winding annulus integrated around the bore. 17

12.1 Representative thermal architecture using liquid cooling, passive spreading, and prop-wash rejection. 25

14.1 Notional containment and energy-absorption stack for the ring and inner-stage boundary. 29

17.1 Representative manufacturable process chain for an instrumented I.C.2E prototype. 35

18.1 Validation ladder from thrust-stand work to flow and endurance testing. 37

List of Tables

3.1 Representative top-level requirements for the next I.C.2E iteration. 6

5.1 Key baseline values extracted from current internal documents. 10

13.1 Representative material allocation for a manufacturable I.C.2E build. 27

15.1 Conceptual gain and loss channels in the I.C.2E architecture. 32

16.1 Representative safety and failure-management logic. 34

20.1 Representative phased roadmap for the I.C.2E program. 42

C.2 Extracted datasheet values for 27.8 x 8.5 LV. 49

C.3 Extracted datasheet values for 27.8 x 8.5 HV. 50

C.4 Extracted datasheet values for 27.8 x 10 LV. 51

C.5 Extracted datasheet values for 27.8 x 10 HV. 51

C.6 Extracted datasheet values for 27.8 x 12 LV. 52

C.7 Extracted datasheet values for 27.8 x 12 HV. 52

Chapter 1

Executive Summary

The I.C.2E propulsion unit is best understood as a next-generation integrated propulsor architecture rather than as a conventional motor with an unusual hub. The current internal documents already show a serious baseline: a 707.16 mm integrated five-blade propulsor built around a 264.16 mm outer hub and a 199.16 mm inner bore, with nominal operation near 3000 rpm, peak short-duration electrical input above 7 kW, and measured static thrust in the 19.6 kgf to 28.6 kgf range depending on speed, pitch, and winding option [1]. That baseline alone is enough to justify a formal engineering package. The purpose of this white paper is to translate that baseline into a coherent next-step architecture for heavy-lift drones, ring-wing aircraft, and distributed eVTOL propulsion arrays.

The core thesis is that the present geometry contains a strategically valuable void. In a conventional propeller the inner region is mostly structural and electromagnetically occupied, with little aerodynamic value beyond hub fairing. In the I.C.2E, the 199.16 mm bore can instead become an active flowpath. A reverse-flow inner booster can ingest air through side and rear passages, impart modest pressure rise, and feed a toroidal plenum that discharges through carefully angled annular slots into a shaped bore wall. Those jets can energize the local boundary layer, entrain ambient air, and reduce internal separation. Downstream, a static vane and containment ring can recover swirl while carrying the mechanical burden of a dual-flow rotor system. The practical result is not a mystical air amplifier. It is a physically grounded attempt to reclaim effective area, reduce losses, and improve the net thrust-per-kilowatt-per-kilogram package.

The white paper takes a disciplined stance on claims. The preferred external formulation is not a headline multiplier. The preferred formulation is that the I.C.2E is being developed to deliver substantially improved installed thrust density, lower effective disc loading, and stronger thermal-operational margin than conventional direct-drive propulsion units occupying the same aircraft integration envelope. That statement remains ambitious, but it is physically grounded and suitable for technical, investment, and partnership audiences.

Several technical conclusions follow from that framing. First, the inner reverse-flow stage should be used as a pressure-assist and flow-conditioning device, not as a tiny primary fan that chases excessive pressure ratio or near-sonic compressor Mach numbers. Second, the outer spool should remain the dominant thrust producer by moving the majority of air at low tip Mach and moderate loading. Third, the concept only becomes commercially credible when the aerodynamic architecture, electromagnetic packaging, containment system, and cooling system are designed as one machine rather than as separate disciplines. Fourth, the migration to an aircraft-level 800 V to 1000 V DC bus is not ornamental. It is central to lowering current, reducing copper and switching losses, and freeing thermal margin for sustained operation.

This paper is written as a stakeholder-ready technical position document. It provides an executive-level narrative, a technical concept definition, a manufacturability and materials strategy, a preliminary performance model, a validation roadmap, and appendices with extracted operating data from the current datasheet. It also includes illustrative figures that communicate the design logic without pretending to be final CAD. The goal is not to freeze the design. The goal is to provide Aerial MECHANICA with a structured foundation for engineering reviews, partner discussions, investor communication, and future patent or grant support.

The most important practical recommendation is straightforward. The next build should preserve the documented outer propulsor scale and performance envelope, but add an instrumented inner cartridge containing the reverse-flow booster, annular plenum, slot geometry, pressure taps, and containment structure. That cartridge should be bench-validated in modes that separate outer-spool-only thrust from combined-mode thrust. If the measured gains appear in pressure recovery, wake regularity, and net thrust at fixed electrical input, the architecture deserves full-scale refinement. If the gains do not materialize, the platform still retains value as a high-performance integrated propulsor. The design therefore has a rational downside and an unusually strong upside if executed with discipline.

Chapter 2

Context and Design State

The current design state is anchored by two useful source documents and one supporting summary document. The Fusion-derived drawing package captures the geometric envelope and naming conventions for the motor-prop unit, while the final datasheet captures the baseline electrical, thermal, and propulsor performance values [1, 2]. The separate summary PDF contributes high-level product framing but comparatively little quantitative design detail. For that reason this white paper relies primarily on the drawing and datasheet for quantitative claims and uses the summary document only as contextual presentation support.

The internal datasheet is already strong enough to define the current product class. It specifies a brushless three-phase unit with 18 pole pairs, a weight target of 2.1 kg for the motor plus integrated propulsor, and two documented winding options at 52 V and 74 V nominal input [1]. The no-load speed of both windings is listed as 3550 rpm, with recommended maximum speed 3600 rpm. Across the three documented pitch variants, the system produces useful static thrust from low speed all the way through short-duration high-power operation above 7 kW. This is not a speculative desktop rotor. It is already a meaningful heavy-lift propulsion module.

The current physical geometry matters even more than the catalog numbers. The drawing identifies a 264.16 mm outer hub and a 199.16 mm inner hub or bore, nested inside a 707.16 mm overall propulsor diameter [2]. That leaves a large central opening relative to the total disk. In a typical open propulsor this would be viewed as a pure penalty because the inner region contributes to disk area but not to thrust. In the I.C.2E concept it becomes the signature opportunity. The white paper therefore treats the central bore not as dead volume but as a candidate secondary-flow device.

The present performance tables are equally valuable because they expose both the capability and the limitations of the baseline architecture. The 27.8 in by 8.5 in configuration reaches approximately 19.6 kgf at 2900 rpm and 28.6 kgf at 3500 rpm on the HV winding, with efficiency falling gradually as speed and loading rise [1]. The 10 in and 12 in pitch versions trade hover efficiency for different forward-flight bias. This suggests that the current motor and blade family are already operating in a balanced, design-worthy region. The task is not to invent thrust from scratch. The task is to see whether a more sophisticated central flow architecture can increase net system value without breaking noise, heat, or manufacturability limits.

That framing also highlights the current evidence gap. The documents do not yet include measured plenum pressure, slot exit velocity, bore-wall static pressure distributions, unsteady acoustic spectra, or endurance results for a secondary high-speed spool. They do not show whether a Coanda-style annular slot system remains attached across the hover-to-transition envelope. They do not yet establish bearing life or containment strategy for a reverse-flow inner stage. These are exactly the missing pieces that a next-phase white paper should expose rather than obscure. The proper design state, then, is “promising and well-bounded,” not “already proven.”

The correct program interpretation is that Aerial MECHANICA already has a credible integrated propulsor baseline and a distinct architectural opening. The logical next step is not to enlarge the propeller indiscriminately or to overmarket the current design. It is to build a disciplined secondary-flow demonstrator inside the existing envelope, generate the missing pressure, temperature, vibration, and acoustic data, and then decide whether the inner-stage concept deserves productization. The rest of this white paper is structured around that decision path.

Chapter 3

Requirements and Use Cases

The I.C.2E architecture should be developed against mission classes, not against abstract laboratory hero numbers. Three mission classes matter most. The first is heavy-lift multirotor and coaxial drone systems, where static thrust, packaging density, and thermal reliability dominate. The second is ring-wing or annular-frame aircraft, where the central bore and side inlets can be integrated into the surrounding structure and the propulsion unit can do double duty as both lift source and flow-control element. The third is distributed eVTOL lift arrays, where the combined value of high specific thrust, graceful degradation, and better electro-thermal performance could materially reduce installed propulsion mass.

These missions imply different design priorities. Heavy-lift drones value thrust density, dirt tolerance, and ease of field replacement. Ring-wing aircraft value inlet management, structural integration, and the ability to route secondary flow through the surrounding frame. Distributed eVTOL platforms impose the harshest standards because every unit must be quiet, fault-tolerant, thermally stable, and suitable for fleet maintenance. The architecture proposed here is deliberately biased toward the second and third categories, because that is where the central-bore ejector concept has the most system-level leverage.

Table 3.1: Representative top-level requirements for the next I.C.2E iteration.

Requirement	Target or design intent
Net static thrust	Preserve current 20 kgf to 30 kgf class in stand-alone mode and pursue higher net thrust only with measured validation
Outer spool role	Primary thrust producer in all safe operating modes
Inner spool role	Pressure assist, bore-flow control, and off-design wake conditioning
Electrical architecture	Path from current low-voltage prototypes to aircraft-level 800 V to 1000 V DC distribution
Thermal durability	Sustained high-power operation with explicit winding, magnet, and inverter thermal margin
Containment	Fragment management for inner stage and retained-failure behavior for outer stage
Serviceability	Replaceable electronics cartridge and inspectable inner flow hardware
Acoustics	Avoid tip-Mach creep and control narrowband tones introduced by slots or stators

A central requirement for credibility is graceful degradation. The aircraft must never depend entirely on the inner booster for basic lift. The outer propulsor has to remain able to carry the mission safely, or at least to preserve controlled descent and hover exit logic, if the booster or its inverter is disabled. This requirement changes the entire control philosophy. It means the architecture should support an outer-only operating mode, an assisted mode in which the inner stage improves efficiency or wake quality, and a fully coupled mode reserved for the flight conditions where it provides real benefit.

Another requirement is discipline in how performance is presented. The comparison baseline must be explicit every time the concept is discussed. Same motor mass is not the same as same propulsor diameter. Same electrical input is not the same as same thermal margin. Same package volume is not the same as same aircraft integration value. The recommended external clause is therefore "substantially improved installed thrust density, operating endurance, and flow-control authority relative to conventional direct-drive units in the same airframe integration envelope." That language is still ambitious, but it is physically interpretable, commercially useful, and testable.

The resulting use-case hierarchy suggests a phased product strategy. Phase one should be a bench-validated stand-alone propulsion unit with optional inner-stage assist. Phase two should be a ring-wing module in which the inlet, bore, and external structure are co-designed. Phase three should be a full aircraft powertrain configuration with high-voltage bus integration, redundancy logic, and vehicle-level control allocation. That staged approach is slower than jumping directly to a flagship eVTOL demonstrator, but it is vastly more likely to produce defensible engineering data and a fundable program narrative.

Chapter 4

System Architecture

The proposed architecture combines three machines inside one assembly: an outer low-tip-Mach main propulsor, an inner reverse-flow booster, and a non-rotating structural ring that acts simultaneously as stator, containment shell, and flow straightener. Figure 4.1 captures the concept in its simplest form. Air enters through side and aft passages, is turned by the inner booster, collected in an annular plenum, and discharged through distributed slots into a shaped bore wall. Around that bore sits the outer outrunner and its open blades, which remain responsible for most of the net thrust.

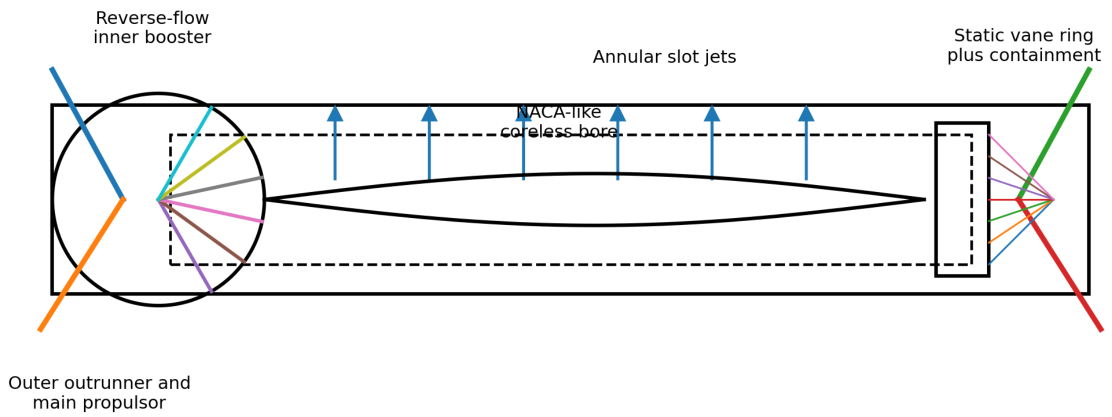

Figure 4.1: Concept cross-section of the reverse-flow dual-spool I.C.2E architecture.

The outer spool should remain conceptually simple. Its job is to process the majority of mass flow with as little induced and profile loss as practical. That means low-to-moderate blade section lift coefficient in hover, carefully managed twist and chord distribution, and a tip speed that stays comfortably subsonic even when the aircraft demands transient over-thrust. The present 27.8 in geometry is already favorable in this respect. At 3500 rpm, the blade tip speed is roughly 129.6 m/s, which is below Mach 0.4 at standard sea-level conditions. For future variants, the more general architectural target should still be to keep hover tip Mach at or below about 0.65 to 0.75. The performance ceiling of the current hardware is therefore set by thermal, acoustic, and structural concerns before it is set by compressibility.

The inner spool is the opposite kind of machine. It should not chase large mass flow or large net thrust directly. It should chase controlled pressure rise and well-directed discharge momentum. In practical terms this means the booster should act like a compact centrifugal or mixed-flow compressor feeding a plenum, not like a miniature turbojet compressor. The reverse-flow packaging logic is still useful, because it lets the system pull air from protected frame inlets and route it toward the bore. But the stage loading should remain conservative enough that surge margin, mechanical life, and acoustic behavior remain manageable. A practical first-pass target is a pressure ratio on the order of 1.10 to 1.30 feeding a thin annular slot with an equivalent gap in roughly the 0.5 to 2.0 mm range, to be tuned by test rather than fixed by styling.

Between the two sits the static ring. This is where the architecture either becomes elegant or becomes heavy. If the ring is treated as dead structure, the concept will struggle. If it is treated as a multifunction component, it can carry a large share of the design burden. The ring can support vanes that recover swirl, house pressure and temperature sensors, spread heat from the inner stage into the surrounding shell, and provide a robust burst barrier around the highest-speed rotating hardware. The same ring can also be a mounting interface into a ring-wing or nacelle structure.

Function stacking is therefore essential. The bore wall is not only an aerodynamic surface. It is also the electromagnetic boundary of the coreless machine, a cooling surface, and a plenum discharge surface. The surrounding frame is not only a structural bracket. It is an inlet diffuser. The containment ring is not only a safety shell. It is part of the aerodynamic nozzle system. This is the type of architecture that can outperform a conventional propulsor in aircraft integration terms, but only if every component is asked to do more than one job.

Figure 4.2 shows the same philosophy from the electrical side. The aircraft high-voltage bus supplies a protected DC link. Two field-oriented inverter channels then feed the outer and inner spools independently under supervisory control. The supervisory layer schedules torque split, booster pressure target, and thermal derating using rotor speed, temperature, plenum pressure, and vibration inputs. That architecture is more complex than a single ESC, but it is still tractable with modern power electronics and is aligned with the kind of control-rich distributed propulsion systems already used in advanced electric aircraft programs.

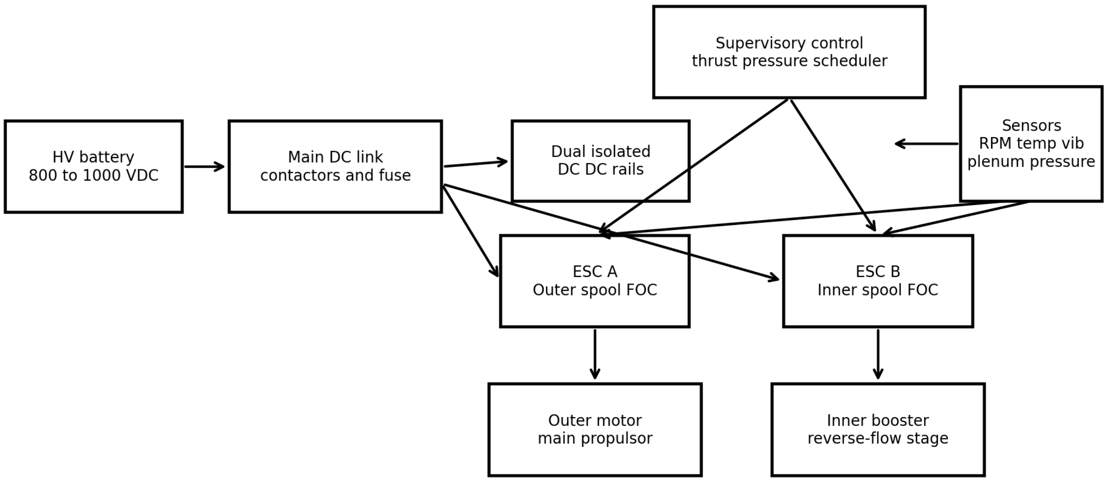

Figure 4.2: Representative electrical and control architecture for a dual-spool I.C.2E unit.

The system recommendation is therefore precise. Preserve the present outer propulsor scale, create a removable inner booster cartridge with its own instrumentation and containment, and design the non-rotating ring as a multifunction structural, thermal, and aerodynamic element. That gives the concept a clean experimental path and avoids the common error of embedding too many unverified assumptions into one inseparable assembly.

Chapter 5

Baseline Geometry and Internal Data

Any serious next-iteration concept must remain tied to the geometry and operating points already documented internally. The current datasheet gives the most reliable baseline numbers for size, mass, electrical constants, and integrated propulsor performance [1]. The drawing contributes the mechanical layout and confirms that the central bore is not a trivial cosmetic feature but a major geometric characteristic [2]. Those baseline numbers define both the opportunity and the engineering boundaries.

Table 5.1: Key baseline values extracted from current internal documents.

Parameter	Internal value
Outer hub diameter	264.16 mm
Inner hub or bore diameter	199.16 mm
Overall integrated propulsor diameter	707.16 mm or 27.84 in
Blade count	5 in the documented family
Nominal operating speed	3000 rpm
Maximum recommended speed	3600 rpm
Nominal voltages	52 V LV and 74 V HV
Max efficiency	92 percent
Weight target	2.1 kg motor plus integrated propulsor
Pitch options	8.5 in, 10 in, and 12 in

The bore-to-diameter ratio is one of the clearest places where the next design iteration should depart from the present baseline. An inner diameter of 199.16 mm inside a 707.16 mm propulsor means the central opening is roughly 28 percent of the total diameter. That is large enough to support a secondary flowpath, but it also removes a substantial amount of effective actuator area near the hub. For the next product-oriented iteration, this paper recommends reducing the active central opening target to approximately 18 percent of propulsor diameter in the current size class, corresponding to roughly 127 mm diameter at a 707 mm outer propulsor. That reduction preserves room for a controlled annular slot system while materially improving disk loading, containment packaging, and structural efficiency. The current 28 percent geometry remains useful as an experimental reference, but it is not treated here as the preferred end-state geometry.

The performance tables also provide an important caution against overstatement. The current hardware is already efficient in the only sense that matters to a vehicle designer: it produces useful thrust across a wide operating range without absurd rotational speed. In the 27.8 in by 8.5 in configuration, the HV winding produces roughly 22.4 kgf at 3100 rpm and 28.6 kgf at 3500 rpm for about 5.0 kW and 7.1 kW input respectively [1]. That means any future inner-stage architecture is competing not with a poor rotor, but with a baseline that is already good.

That baseline leads to two useful engineering interpretations. First, even small percentage gains in overall efficiency are valuable. A 10 percent increase in thrust at fixed electrical input or a 10 percent reduction in power at fixed thrust is strategically important at the aircraft level. Second, the outer stage is already operating in a speed regime that should remain acoustically and structurally manageable. There is no reason to chase a radically faster outer prop. Most of the design ambition should migrate inward, into better pressure management and better electro-thermal efficiency.

The current low-voltage and high-voltage windings also reinforce the case for future bus-voltage escalation. The HV option produces essentially the same speed and useful thrust with noticeably lower current than the LV option, which is exactly what should happen when copper loss is a material fraction of the thermal budget. This trend is important because the eventual move to 800 V to 1000 V distribution is not about raw performance theater. It is about making a high-power distributed propulsion system thermally and electrically civilized.

Finally, the current pitch variants already encode a product strategy. The 8.5 in pitch version is hover-biased, the 10 in pitch version is balanced, and the 12 in pitch version is cruise-biased. Any next-generation central-flow feature has to preserve that flexibility. The right design is therefore schedulable and modular, not a one-condition fixed gimmick. The extracted datasheet tables in Appendix C are included precisely so the white paper can keep referring back to the baseline rather than drifting into speculation.

Chapter 6

Aerodynamic Foundations

The right way to analyze the I.C.2E is to begin with momentum theory and then layer secondary-flow phenomena on top. For a hovering actuator disk the ideal induced velocity is

$$v_i = \sqrt{\frac{T}{2\rho A}}, \tag{6.1}$$

where T is thrust, ρ is air density, and A is the effective disk area. The ideal induced power is then

$$P_i = T v_i = \frac{T^{3/2}}{\sqrt{2\rho A}}. \tag{6.2}$$

These relationships matter because they immediately explain what the concept must do to win. It must either increase effective area, reduce losses, or both. It cannot get around the energy balance.

The central bore therefore has to be reframed as a controllable effective-area recovery mechanism. If the bore were merely open, it would reduce disk area and hurt hover efficiency. If, however, the bore is turned into an entraining and pressure-conditioned flowpath, the aircraft may recover part of that lost area and also reduce the recirculatory losses that normally accompany a large hub. This is why the architecture is better described as a controlled entrainment system than as a raw thrust multiplier.

Figure 6.1 shows the qualitative flow logic. Secondary air enters through side and aft inlets, the inner booster turns and energizes it, and the resulting plenum discharges through annular slots on a shaped bore wall. Those sheet jets are intended to remain attached long enough to entrain surrounding fluid and pull more mass flow through the bore. Downstream, static vanes remove swirl and regularize the merged wake before it leaves the center of the propulsor.

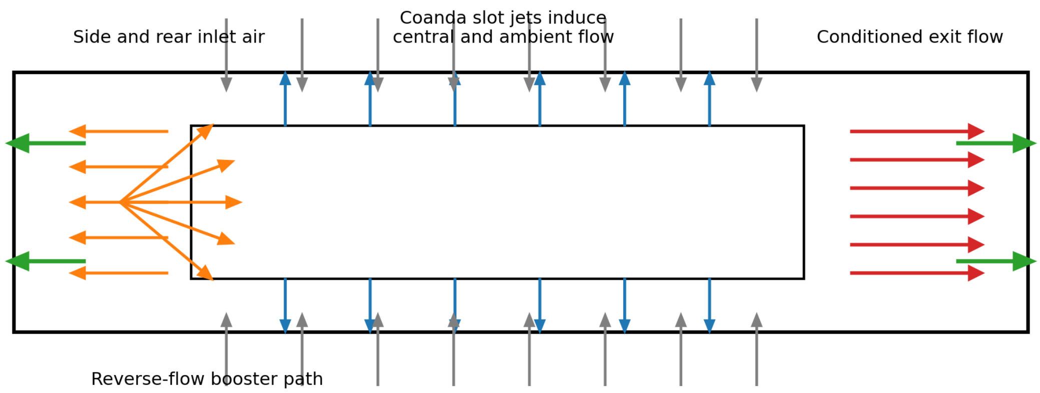

Figure 6.1: Representative airflow logic for the bore slots, entrainment zone, and conditioned wake.

Several aerodynamic mechanisms may contribute to net benefit. Boundary-layer energization can delay separation on the bore wall. Jet-pump entrainment can pull additional ambient flow into the center region. Swirl recovery can convert circumferential velocity into more useful axial momentum. Pressure recovery in the bore and stator region can reduce wake disorder. The challenge is that every one of these mechanisms has an associated failure mode. If the slot jet is underpowered, it does little. If it is overpowered, it adds mixing loss and noise. If the slot angle is wrong, attachment is lost. If the vane ring is too aggressive, blockage and tonal noise increase.

That is why the term "Coanda" must be used carefully. The Coanda effect is not a magic adhesion law. It is the result of a surface jet interacting with curvature and pressure gradients in a specific Reynolds-number regime. The shaped bore therefore needs to be designed for an internal flow problem, not for the aesthetics of a classical external airfoil. A NACA 0012-like section can be a useful starting point because it provides smooth curvature and acceptable thickness, but the final internal contour may need a blunter lip, a flatter recovery region, and a slot shoulder tuned specifically to annular jet attachment.

The biggest aerodynamic mistake would be to make the inner stage chase dramatic pressure ratio or near-sonic local conditions. That path usually ends in poor efficiency and ugly acoustics. The correct target is an optimal slot momentum ratio: enough jet energy to improve the net mass-flow field, but not so much that the secondary system spends more power than it recovers. In benchmark terms, the concept should therefore be compared against a same-diameter open propulsor or another high-disc-loading baseline rather than against an ideal actuator disk. The rest of the white paper treats that optimum as a measurable design target and structures the hardware, controls, and validation plan around finding it.

Chapter 7

Reverse-Flow Inner Booster

The inner booster is the most novel part of the I.C.2E concept and also the part most likely to fail if it is overdesigned. A reverse-flow impeller around 16 000 rpm is entirely compatible with a practical development program, but 16 000 rpm alone does not imply near-sonic operation. Tip speed depends on radius. A booster with a 75 mm to 90 mm radius at 16 000 rpm has tip speed on the order of 125 m/s to 150 m/s, which is well below sonic conditions. That is good news. The design does not need near-sonic absolute speed to be useful. It needs stable pressure rise and well-directed discharge momentum. The correct near-term framing is a compact transonic booster with modest pressure ratio, not a shock-heavy miniature "supersonic compressor."

The recommended booster concept for the next prototype is therefore a contained centrifugal or mixed-flow stage driving a modest pressure ratio into a toroidal plenum. A centrifugal architecture is attractive because it is easier to route into an annular plenum, more tolerant of manufacturing variation than a tiny axial compressor, and better aligned with the reverse-flow packaging logic. It can ingest air from peripheral frame inlets, throw it outward into the plenum, and let the plenum feed the circumferential slot system with relatively uniform static pressure. A realistic design point is a booster that comfortably supports a 1.10 to 1.30 pressure ratio while preserving efficiency and surge margin.

Several design rules follow from that choice. The impeller blades should be short, thick enough for manufacturability, and designed for controlled diffusion. The plenum should have enough volume to smooth unsteady inflow and avoid large circumferential pressure gradients at the slot exits. The exit slot area should be chosen from the desired slot momentum and expected pressure rise, not from arbitrary aesthetic spacing. Because the booster is a secondary machine, total pressure stability and flow uniformity matter more than peak compressor map bragging rights.

The reverse-flow packaging logic still matters even with a centrifugal stage. Protected side and rear inlets can be integrated into a ring-wing or nacelle frame, reducing ingestion of the hottest recirculated air and allowing the inner system to operate in a cleaner pressure environment than a front-facing exposed inlet would provide. Reverse-flow routing also shortens the path between the booster, the plenum, and the bore discharge slots. That is valuable because every bend, contraction, and sudden expansion in the secondary path costs pressure and therefore costs the very benefit the inner stage is supposed to create.

The booster should be treated as its own rotordynamic program. A 16 000 rpm secondary spool is not extreme by turbomachinery standards, but it is demanding in an electric aircraft environment where cost, maintainability, and safety matter. Bearing DN limits, shaft or rim stiffness, and synchronous and asynchronous vibration modes have to be analyzed early. If a rim-driven implementation is used, the rotor shell and magnet retention become part of the burst-risk picture. If a shafted cartridge is used, the support bearings and seals become the maintenance driver. Hybrid ceramic bearings are the minimum credible baseline for an aggressive prototype, while air or magnetic bearings belong in the future-options category rather than in the first build.

From a controls standpoint the booster should be commanded by pressure or slot-momentum demand rather than by open-loop speed ratio alone. That means the next prototype should include plenum pressure sensing and ideally several circumferential taps to confirm slot uniformity. The outer rotor demand can then schedule a target inner-stage pressure rise, while thermal and vibration logic set the hard limits. That control approach keeps the secondary stage tied to a measurable aerodynamic objective and reduces the temptation to run it harder than the flow actually needs.

The conclusion of this chapter is intentionally conservative. The reverse-flow booster is worth building, but only as a modest-pressure, high-control, heavily instrumented device. If the engineering team tries to make it a heroic miniature compressor, the concept will likely lose more power to heat, noise, and complexity than it gains in useful thrust. If it is built as a disciplined pressure-assist stage, it has a credible chance of improving the whole system.

Chapter 8

Coanda Ejector and Coreless Bore

The bore is where the I.C.2E moves from "interesting propeller" to "architecturally distinct propulsion system." It is also where the mechanical, electromagnetic, and aerodynamic design teams must stop working in isolation. The bore wall carries the slot geometry, defines the entrainment surface, separates the inner and outer flow systems, and bounds the coreless electromagnetic region. The design will only work if those functions are deliberately co-designed.

The most useful starting point is a symmetric airfoil-derived contour, such as a NACA 0012 or nearby thickness distribution, adapted for internal flow. The value of such a profile is not that it is a perfect answer. The value is that it provides smooth curvature, manufacturable wall thickness, and a defensible place to locate the slot lips. In practice the final section may depart from the textbook profile to support local thickening for magnets, cooling passages, or plenum walls. That is acceptable. What matters is that the surface seen by the slot jet remain smooth enough to sustain attachment over the required arc length.

The slot geometry deserves more attention than almost any other feature. A thin annular slit will usually provide a higher discharge velocity for a given plenum pressure than a broad slot, but it is also more sensitive to blockage, manufacturing tolerance, and contamination. A broader slot is easier to manufacture but may diffuse too quickly and lose attachment. The design target should therefore come from desired momentum coefficient and acceptable pressure drop, not from stylistic imitation of a Dyson fan. The correct physical analogy is a controlled annular wall jet inside a propulsion duct.

The discharge angle is equally important. A slot aligned too close to the surface may remain attached but entrain insufficient core flow. A slot angled too far inward can detach or create a shear layer that wastes energy in mixing. Because the flow is annular and the outer propulsor is simultaneously pulling on the same region, the optimal angle will depend on the main rotor loading. This is one of the strongest arguments for building the slot system into a removable experimental cartridge. Slot angle, width, and lip radius should be treated as rapid-iteration parameters in early testing.

The coreless aspect of the bore architecture is also central. Removing an iron core reduces eddy current and hysteresis losses, eliminates one major source of thermal mass and saturation behavior, and opens more freedom in shaping the bore region. The penalty is that the magnetic circuit must be built more carefully around air gaps, magnet retention, and support structure. The bore wall therefore becomes a shared aerodynamic and electromagnetic boundary. Figure 8.1 illustrates the notional radial-flux arrangement in which embedded magnets on the rotor shell interact with a coreless winding annulus supported by a lightweight structural carrier.

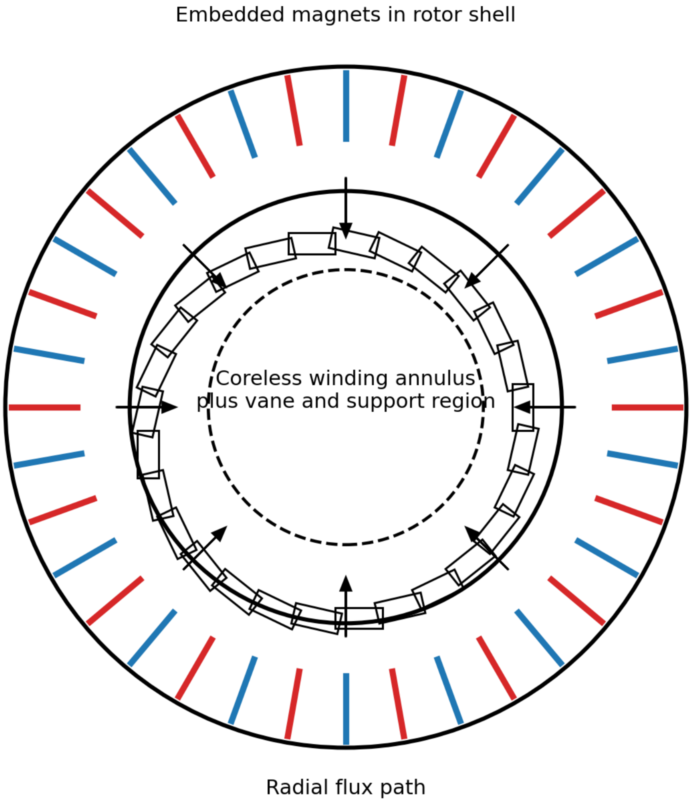

Figure 8.1: Representative radial-flux concept for a coreless winding annulus integrated around the bore.

This shared boundary creates a manufacturing implication. The bore cannot be designed as a thin decorative shell. It needs enough section thickness and stiffness to hold aerodynamic tolerances, protect the magnetic gap, and carry the slot lips without deformation. That likely means a hybrid wall with a metallic liner or insert, a composite or polymer structural overbuild, and local thermal paths to prevent the slot region from becoming a heat trap.

The bore chapter therefore ends with a simple rule: the bore is a system component, not an airfoil graphic. Its geometry must be selected by the coupled requirements of slot attachment, structural stability, thermal path continuity, and electromagnetic clearance. If those four demands are aligned, the bore can become the signature advantage of the I.C.2E rather than its most delicate weakness.

Chapter 9

Outer Outrunner and Main Propulsor

The outer spool remains the primary lift generator in every credible version of the I.C.2E. This point is strategically important because it preserves a rational fallback mode and keeps the whole concept attached to efficient hover physics. Large thrust in hover comes from moving a large mass of air with a modest velocity increment. The existing 27.8 in integrated propulsor already does that well enough to serve as the stable backbone of the architecture [1].

The present blade family uses five integrated blades across three pitch options. That is a sensible baseline for heavy-lift and distributed-lift operation because higher blade count raises solidity and allows the rotor to absorb useful power at relatively low rotational speed. A broader future family spanning roughly five to ten blades is also plausible if quieter ring-duct variants or more compact high-solidity lift modules are pursued. What should not change is the low-tip-Mach philosophy: even as solidity evolves, hover operation should remain well below severe compressibility and ideally within a practical target band no higher than roughly Mach 0.65 to 0.75.

The next iteration should preserve several geometric virtues of the current outer rotor. The diameter is large enough to avoid extreme disk loading for the documented thrust levels. The tip speed at normal operating range stays comfortably below compressibility trouble. The integrated hub provides mechanical continuity and a clean platform for balancing. Most important, the current propulsor data already show differentiated mission tuning through pitch variation. The dual-spool architecture should not erase that flexibility. Instead, the inner-stage schedule should complement the outer-blade family.

Blade section design deserves renewed scrutiny in light of the bore-flow concept. The inner section of each blade sees a more complex inflow than the outer section because it interacts with the bore discharge and the static vane ring. It may therefore benefit from a different local camber, twist law, or chord taper than a blade designed for a purely passive hub. Likewise, the root structural design must account for any unsteady pressure fluctuations created by the slot system or the secondary flow in transition flight. None of this requires a radical reinvention of the outer rotor, but it does require acknowledging that the inflow field is no longer canonical.

The outer rotor also becomes part of the acoustic story. One reason to keep the inner-stage pressure ratio modest is that the outer rotor already has a strong low-frequency and blade-passing signature. If the inner stage adds its own tonal system on top of that, the unit risks becoming acoustically objectionable even if the net thrust improves. The outer rotor should therefore be optimized not only for thrust and efficiency, but also for a clean interaction with the stator ring and any central-slot discharge.

There is also a control-system implication. Because the outer rotor is the primary thrust machine, vehicle demand should be mapped primarily to outer-spool torque or speed, with the inner stage scheduled around that demand. This seems obvious, but it prevents the supervisory controller from using the booster as a crutch to compensate for poor outer-rotor matching. In a healthy architecture, the outer spool alone still defines the basic thrust curve and the inner system sharpens, stabilizes, or improves it.

The right design stance is therefore to leave the outer rotor recognizably continuous with the current datasheet hardware. That protects the baseline product, simplifies A-B testing, and ensures that any measured performance gain can be attributed to the new central-flow features rather than to a wholesale change in the main propeller. From both an engineering and a business perspective, that continuity is a major advantage.

Chapter 10

Electromagnetic Architecture

The I.C.2E concept depends on a low-loss electromagnetic package as much as it depends on clever aerodynamics. A coreless stator approach remains attractive because it offers lower eddy-current and hysteresis losses, reduced mass, no ferromagnetic saturation in the classical tooth-core sense, and greater freedom in shaping the central structure. These are all real advantages if the magnetic circuit and thermal path are managed properly.

Coreless machines trade one set of problems for another. Iron removal reduces classic core losses, but it also weakens flux concentration and can reduce torque density if the winding geometry, magnet arrangement, and air-gap control are not carefully designed. This means the rotor shell, magnet segmentation, and support carrier become more important than in a standard slotted machine. The overall architecture must therefore be treated as a carefully engineered air-gap machine rather than a conventional BLDC with its teeth removed.

For the outer machine, the current integrated outrunner logic remains attractive. Embedded magnets in the inner face of the rotating shell can interact with a lightweight coreless winding ring fixed to the static structure. A Halbach-like magnet segmentation strategy is worth studying in this context because it can strengthen useful air-gap flux while reducing dependence on heavy back iron. The result is a large mean radius for torque production and a clean mechanical path into the blades. Figure 8.1 illustrates that radial-flux framing. In this arrangement, torque density comes from mean radius and active circumference rather than from ferromagnetic tooth concentration.

The inner machine is a separate electromagnetic problem. If the booster is rim-driven, it will likely require its own compact radial-flux arrangement with a smaller diameter and much higher electrical frequency. If it is shaft-driven with an external motor cartridge, the electromagnetic and fluid machines can be partially separated. The white paper favors partial separation in early prototypes because it reduces coupling risk. The concept can still preserve the reverse-flow aerodynamic logic without forcing both rotating machines into one unserviceable electromagnetic package.

High electrical frequency introduces nontrivial inverter and winding design consequences. Skin effect, proximity effect, switching ripple, and common-mode EMI become increasingly relevant as pole count and speed rise. Coreless windings help with some of the classic loss channels, but they do not remove the need for careful conductor choice, litz or multi-strand strategy where needed, and disciplined inverter layout. The electromagnetic architecture therefore has to be developed alongside the powertrain, not upstream of it.

Magnet retention is another central issue. In a high-radius outrunner, the magnets see large centrifugal loads. If the machine also lives near warm internal airflow and transient overtemperature, adhesive-only retention becomes a poor long-term bet. Mechanical capture features, nonmagnetic sleeves, or hybrid bonded-and-captured construction should be assumed from the start. The same is true of the inner booster if a rim-driven rotor is used. Burst energy scales quickly, and the electromagnetic design must assume the structural design rules rather than treating them as secondary.

The practical conclusion is that a coreless architecture is still the right strategic direction, but it must be justified as a whole-system choice. It is not enough to say "coreless equals efficient." The complete machine must show stable air gaps, magnet security, manageable electrical frequency, controlled thermal gradients, and manufacturable winding support. If those are delivered, the aerodynamic concept gains the electrical headroom it needs. If they are not, the secondary-flow benefits will be consumed by electrical loss and maintenance burden.

Chapter 11

High-Voltage Powertrain and Controls

The current datasheet records two low-voltage configurations at 52 V and 74 V. Those values make sense for present bench hardware and early propulsion modules, but they are not the end state for a serious distributed electric aircraft system. The design brief correctly points toward an 800 V to 1000 V DC architecture. That direction is not cosmetic. It is one of the few ways to reduce current, cable mass, copper loss, and inverter stress at the vehicle level while preserving the ability to push multiple propulsors simultaneously.

The transition to high voltage does not require the motors themselves to see 1000 V directly at the phase terminals. The cleanest architecture is aircraft-level high-voltage distribution with local protected conversion and inversion inside or near each propulsion unit. This allows the wiring backbone to remain light and efficient while the local ESC modules are tailored to the voltage, current, and switching-frequency needs of the outer and inner spools. It also improves fault isolation, which matters in any distributed-lift platform.

The dual-spool arrangement naturally maps to dual inverter channels under one supervisory controller rather than two independent control brains. One channel feeds the outer spool, which is the primary thrust machine. The second feeds the inner booster, which is a pressure-assist machine. The supervisory controller sits above both and maps commanded vehicle thrust into outer-spool demand plus a scheduled secondary-pressure or jet-momentum target. That target can depend on rotor speed, current flight mode, plenum pressure, inlet distortion, acoustic constraints, and measured thermal state. In other words, the booster should be asked to do only the amount of work that is useful in the current aerodynamic condition.

Sensoring is therefore a first-order design requirement. Rotor speed and phase current are necessary but not sufficient. The next prototype should include winding and magnet temperature estimation, plenum pressure sensing, at least one differential pressure or slot-pressure indicator, and vibration sensing on the static ring. These are the signals that allow the controller to protect the hardware and also to learn whether the booster is genuinely improving system performance. A dual-spool propulsion unit without such measurements would be harder to tune and easier to damage.

There is also a reasonable case for an advanced development-only sensing layer built around

LBNL-origin methods such as PRISM and DDRM-style differential learning. PRISM is presented as a magnetometry scheme with strong common-mode rejection against vibration, temperature and strain variation, pulse imperfections, and background fields, while retaining broadband magnetic sensitivity [3]. That does not make it a replacement for standard propulsion sensors, but it does make it interesting as a non-contact magnetic-state observer for development rigs: for example, tracking motor magnetic signatures, current-harmonic disturbances, rotor-stator asymmetry, or EMI in the presence of severe vibration and thermal drift. In the same spirit, DDRM and CALIPR-style control logic are directly relevant as learning frameworks for unstable systems where absolute labels drift too quickly to be useful [4, 5]. That is a closer fit to propulsion development than to certified flight control, and it should be framed that way.

Field-oriented control is the obvious starting point for both machines, but the inverter design goals are different. The outer spool is relatively slow and torque-dominant, with large inertia and comparatively low electrical frequency. The booster is faster, lighter, and more sensitive to surge-like behavior or sudden load changes. That means the inner inverter may need faster current control, tighter current ripple management, and more aggressive derating rules, while its command loop should be tied to pressure or jet-momentum objectives rather than blunt speed ratio alone. A single generic ESC architecture could work, but two tailored inverter boards or firmware modes are more likely to produce good results.

Electromagnetic compatibility also has to be taken seriously because the architecture places sensors, cooling hardware, and structural electronics in a compact annular package. The temptation to route everything through the central shell must be resisted unless shielding, grounding, and isolation are engineered deliberately. High-voltage distributed propulsion units live or die on wiring discipline. The right powertrain architecture therefore looks less like hobby ESC packaging and more like aerospace power electronics in miniature.

The design recommendation is to use the current datasheet hardware as the aerodynamic and mechanical baseline, but build the next electronics around the future aircraft architecture rather than around a temporary low-voltage bench stack. That ensures the work done on controls, sensing, thermal design, and safety cases is directly reusable when the concept scales into a true aircraft module.

Chapter 12

Thermal Management

Thermal management is one of the few subsystems that can make or break every other promise in the concept simultaneously. If the windings overheat, copper loss rises and magnet temperature margin disappears. If the inner-stage inverter overheats, the control authority that makes the secondary-flow concept valuable disappears. If the bore wall overheats locally, slot geometry can distort, adhesives can weaken, and the entire aerodynamic story degrades. The thermal architecture therefore has to be deliberate from the first serious prototype.

The internal datasheet already provides useful thermal context. It lists winding-to-housing thermal resistance of about 0.5 to 0.55 K/W and housing-to-ambient thermal resistance of about 0.11 to 0.12 K/W, with winding thermal time constants in the 25 to 27 second range and overall motor thermal time constants around 60 to 65 seconds [1]. Those are not bad values for a compact propulsion unit, but they are also not enough to support a more complex dual-spool architecture without improvement. The inner booster and its electronics will raise local heat flux, not lower it.

The preferred response is an integrated thermal stack built from three layers. The first layer is a conductive path from windings, inverter switches, and bearing seats into the surrounding shell. The second is an active liquid-cooling loop using microchannels in the stationary structure. The third is passive spreading and rejection via heat pipes or vapor-chamber-like paths into the ring shell, which is then washed by ambient airflow. Figure 12.1 captures that layered strategy.

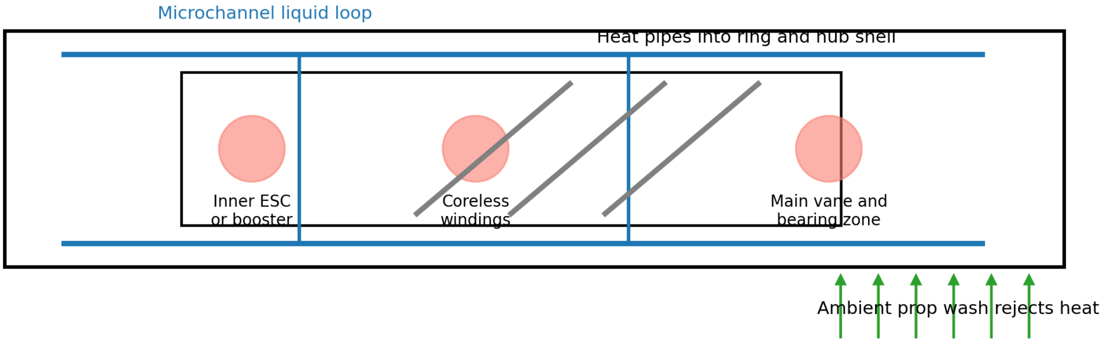

Figure 12.1: Representative thermal architecture using liquid cooling, passive spreading, and prop-wash rejection.

Microchannel liquid cooling is justified here not because it is fashionable but because the design has concentrated thermal zones in a narrow annulus. A modest coolant loop can hold the stationary shell and the inverter cold plate near a much tighter temperature band than air alone can manage. That, in turn, stabilizes magnet temperature, air-gap growth, and adhesive life. The coolant path does not need to be elaborate. It needs to be inspectable, serviceable, and tolerant of vibration.

Heat pipes or equivalent passive spreaders are especially attractive in the ring and bore region because they can move heat away from the slot discharge area without requiring liquid channels in the most delicate aerodynamic surfaces. This matters because the slot lips and bore wall need tight geometric control. Embedding every cooling function directly in those surfaces would complicate manufacturing and raise leak risk. A hybrid approach with conductive inserts, passive spreaders, and remote liquid channels is therefore more credible than an all-liquid concept.

The thermal-control logic should be as explicit as the hardware. Outer-spool torque demand can be limited by winding temperature estimator and coolant state. Booster demand can be limited by plenum benefit relative to electrical and thermal cost. If the secondary stage becomes thermally expensive at a given flight point, the controller should simply back it off. This is another reason the booster must remain an enhancement device rather than a single point of failure. A thermally intelligent schedule is part of the architecture, not a software afterthought.

The main conclusion is that thermal design should be treated as a source of performance, not just protection. Better cooling directly enables lower resistance, tighter clearances, more stable magnets, and longer usable high-power duration. In a propulsion unit whose differentiation depends on subtle improvements in net system efficiency, that thermal headroom is one of the most valuable forms of real engineering margin.

Chapter 13

Materials and Process Selection

The I.C.2E concept is only interesting if it can be built repeatably. That immediately turns materials selection into a system problem. The outer blades want high specific stiffness and fatigue resistance. The bore and slot region want dimensional stability, local heat tolerance, and fine-feature manufacturability. The ring wants burst containment and thermal spreading. The electromagnetic carriers want electrical insulation, structural accuracy, and predictable thermal expansion. No single material family does all of that well.

The most credible answer is a hybrid stack. The outer blades should remain carbon-composite dominant because mass at radius is expensive and the current product already uses carbon-composite propulsors. For the housing and duct shell, carbon-fiber-reinforced PEEK or PAEK deserves specific attention because it offers heat tolerance, stiffness, and electrical friendliness in a compact propulsion package. The bore insert and slot lips should likely include a metallic substructure or liner, particularly if they also carry magnetic or cooling features. Titanium and aluminum both remain credible depending on local temperature, corrosion environment, and cost target, while Ti-6Al-4V is an especially strong candidate for the inner impeller. If the booster is ultimately run hotter or harder than planned, Inconel 718 becomes a defensible fallback at the expense of mass. Aramid or glass reinforcement becomes attractive where impact toughness and electrical isolation matter more than peak stiffness.

Table 13.1: Representative material allocation for a manufacturable I.C.2E build.

Subsystem	Preferred material family	Rationale
Outer blades	Carbon composite with toughened epoxy	High specific stiffness, low rotating mass, established process path
Containment shell	Carbon plus aramid hybrid laminate over metallic liner	Combines stiffness, fragment capture, and thermal stability
Bore insert and slot lips	Machined or printed aluminum or titanium insert	Fine geometry, cooling compatibility, dimensional control
Housing and duct shell	CF-PEEK or CF-PAEK	Heat-tolerant structural shell with strong electrical and thermal compatibility
Coreless winding support	Glass-filled or high-temperature polymer composite carrier	Electrical isolation with moderate structural duty
Stator ring and vanes	Metallic or hybrid metallic-composite assembly	Precision, heat spreading, and vane-root strength

Additive manufacturing has a particular role in the next iteration, especially in the bore, plenum, and vane support region. It enables the team to explore integrated ducts, cooling passages, and mounting features that would be awkward to machine conventionally. However, the final product should not assume that every component remains additively manufactured. The right program approach is to use additive processes for early design convergence and then migrate stable parts toward more scalable machining, molding, or laminate tooling as geometry matures.

Bonding strategy deserves its own discipline. The concept contains rotating composites, magnets, inserts, and likely multiple dissimilar materials in thermally active regions. Adhesives must therefore be chosen not only for strength but for temperature capability, peel resistance, and long-term environmental durability. Where practical, adhesives should be backed up by geometric retention features so the design does not rely on pristine bond-line behavior after repeated thermal cycles.

Surface finish also matters more than usual because several key surfaces are fluid-dynamic. The bore wall, slot shoulders, and plenum flowpath directly affect attachment and mixing losses. A rough additive surface may be acceptable in early flow-visualization work, but a serious performance demonstrator should include post-processing on all aerodynamic passages. The same is true of the outer blade surfaces if low-noise claims are important.

From a manufacturing-readiness perspective, the materials plan should deliberately separate “proof-of-physics” builds from “pilot-manufacturable” builds. A proof-of-physics build can accept slightly heavier inserts, modular cartridges, and visible fasteners if that accelerates learning. A pilot-manufacturable build then uses the data from those prototypes to collapse part count, shift to molded or co-cured assemblies where justified, and lock in containment and balance features. This staged material strategy makes the architecture more believable to both engineers and investors because it acknowledges the difference between prototype convenience and product discipline.

Chapter 14

Structural, Containment, and Rotordynamics

The structural problem in the I.C.2E is more demanding than the current open-prop baseline because the concept adds a second rotating system and a nontrivial containment requirement. The outer rotor already stores meaningful rotational energy, particularly because it carries multiple blades at large radius. The inner booster may store less total energy because of its smaller radius, but it rotates faster and can fail more abruptly. The architecture is only credible if both rotating systems have explicit retained-failure logic.

The first rule is to treat the static ring as a true containment component. It should not merely hold vanes. It should be sized for hoop load, fragment strike, and thermal distortion. A hybrid laminate with an inner metallic liner and an outer carbon-plus-aramid shell is attractive because it combines a hard inner impact surface with a tougher energy-absorbing wrap. Figure 14.1 illustrates the concept schematically. Final sizing would still require fragment-energy analysis and spin testing, but the design intent should be obvious from the first CAD revision. The same ring should also carry sensors and cooling passages wherever possible so the containment mass contributes directly to system value.

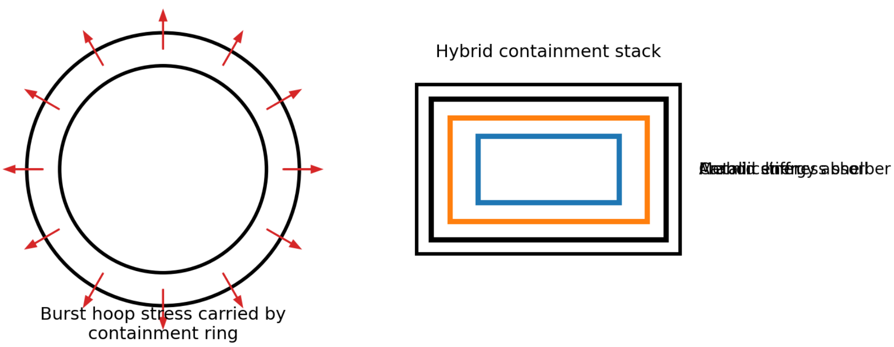

Figure 14.1: Notional containment and energy-absorption stack for the ring and inner-stage boundary.

Rotordynamics should be treated separately for the outer and inner spools before any coupled model is trusted. The outer rotor is dominated by blade and shell stiffness, balance quality, and support compliance through the main bearings and frame. The inner booster adds higher-frequency modes and potentially more sensitivity to plenum or seal asymmetry. Campbell diagrams, critical speed estimates, and unbalance response studies should therefore be developed early, even at reduced fidelity. The goal is not to perfect the model on day one. The goal is to avoid obvious mode-placement mistakes before the first serious spin test.

Bearing selection is likely to be one of the hardest tradeoffs in the concept. Large-diameter slow outrunners like robust bearings that tolerate contamination and static loading. High-speed boosters like smaller precision bearings with lower DN burden and tighter internal geometry. A shared bearing architecture across both stages may look elegant, but it is rarely the best engineering answer. A cartridge-based inner stage with its own bearing strategy is therefore more attractive for early development than a fully shared support system. Hybrid ceramic rolling bearings are the sensible baseline, while air or magnetic bearings should be treated as later-stage options only if the rest of the system proves worthy of that complexity.

Structural stiffness also feeds directly into aerodynamic performance. The bore wall has to hold slot geometry under pressure and temperature. The vane ring has to preserve throat area and turning angle under centrifugal reaction loads transmitted through the frame. The outer blade roots have to tolerate any unsteady pressure field created by the booster. These are not secondary refinements. If the structure moves too much, the aerodynamic model becomes meaningless.

The recommended verification sequence is therefore structural before aerodynamic in the highest-risk areas. First prove overspeed capability, balance behavior, and retained-failure logic for the inner cartridge. Then prove outer-spool and ring stiffness under representative pressure and centrifugal loads. Only then should the team attempt aggressive coupled tests where both spools operate together near the intended combined envelope. This order may feel conservative, but it is the fastest route to useful data because it avoids destroying expensive hardware while chasing flow effects.

In short, the structural and rotordynamic program should be framed not as a compliance burden but as a design-enabling discipline. The moment the ring becomes credible as a containment and stiffness element, the rest of the architecture becomes much easier to defend.

Chapter 15

Preliminary Performance Model

The performance model in this white paper is intentionally preliminary and explicitly bounded. It does not claim to predict final thrust exactly. It exists to show where real gains could come from and where inflated expectations are likely to fail. The starting point is the measured baseline from the current datasheet. The HV 27.8 in by 8.5 in configuration produces about 19.6 kgf at 2900 rpm, 22.4 kgf at 3100 rpm, and 28.6 kgf in short-duration operation at 3500 rpm [1]. Those are the numbers any secondary-flow concept has to beat on a net basis, not on a local-jet basis.

The most credible gain channels are fivefold. First, the bore-flow system may recover effective area by pulling useful throughflow where a large hub would otherwise be dead. Second, the slot system may reduce center-hub recirculation and improve the wake seen by the inner portions of the outer blades. Third, the vane ring may recover swirl and reduce kinetic energy that would otherwise leave the system in unhelpful circumferential motion. Fourth, high-voltage architecture and better thermal control may reduce electrical losses, letting more of the same stored battery energy reach the airflow. Fifth, the integrated ring-wing installation may reduce external interference losses compared with a more conventional packaged propulsor. The fairest comparison case for all of this is therefore a same-diameter open propulsor or another high-disc-loading baseline, not an ideal disk that already has the effective area the concept is trying to recover.

The main loss channels are equally important. The inner booster consumes power and adds its own inefficiency. The slot system adds mixing loss. The vane ring adds blockage and potential tonal noise. The larger structural shell adds mass. The electronics become more complex. Any claim of net benefit is therefore really a claim that the gain channels will outweigh the loss channels at the operating points that matter most. That is why the booster must be scheduled, not simply turned on at all times.

Table 15.1: Conceptual gain and loss channels in the I.C.2E architecture.

Channel	Expected effect on net performance
Effective-area recovery through bore entrainment	Lowers equivalent induced velocity for the same aircraft-level thrust target
Swirl recovery in vane ring	Converts rotational kinetic energy into more useful axial momentum
Improved cooling and higher bus voltage	Reduces copper loss and extends sustained high-power operation
Secondary-flow mixing	Penalizes efficiency if slot momentum is excessive or poorly directed
Added booster and containment mass	Penalizes thrust-to-weight ratio unless multifunction integration offsets it
Acoustic and structural margins	May force the system to operate below theoretical peak points

This framing also clarifies why the program should move away from multiplier language. A factor near three may be conceivable for a narrow system-level comparison if the baseline is a smaller, same-mass, thermally constrained conventional unit installed in the same aircraft bay. It is not credible as a pure aerodynamic statement against an ideal rotor disk of equal effective area and equal shaft power. The recommended positioning is therefore comparative and architectural rather than multiplicative.

A realistic near-term target would be much more useful: demonstrate that with the outer spool held near a known thrust level, a modestly powered booster and bore-slot system can produce a repeatable increase in net thrust or a repeatable reduction in power for the same thrust. If that improvement is then paired with higher-voltage and better cooling at the unit level, the aircraft-level gain may become large even if the isolated aerodynamic gain is moderate. That is precisely the sort of result that can justify the architecture to both engineers and investors.

The correct performance question is therefore not “can the slot jets multiply thrust by a headline factor?” The correct question is “can the integrated architecture produce enough net benefit, at acceptable mass, noise, and risk, to outperform conventional units in the aircraft contexts that matter?” That is a much harder question, but it is also the right one, and it is the one this concept can realistically answer in its favor.

Chapter 16

Noise, Safety, and Failure Management

Any propulsion concept aimed at eVTOL or urban-adjacent drone use has to confront noise and failure management early. The I.C.2E introduces additional tonal sources relative to a conventional open propulsor: the inner booster, the slot jets, and the static vane ring. Even if each source is individually modest, their interaction can produce acoustic signatures that are disproportionately annoying. That means acoustic design is not a final tuning problem. It is embedded in spool speed, slot geometry, vane count, and control strategy from the beginning.

The strongest noise-control lever is restraint. The outer propulsor should remain comfortably subsonic, the booster should avoid excessive discharge velocity, and the vane count should be selected with both aerodynamic turning and blade-passing interaction in mind. Narrowband tones can be especially problematic in distributed arrays because many units can phase-lock perceptually even if they do not phase-lock dynamically. A small acoustic error in one propulsor becomes a fleet-level problem when multiplied across dozens of units on an aircraft.

Failure management follows the same logic. The architecture has more parts than a conventional direct-drive unit, so it must be better organized, not merely more complicated. The outer spool must remain a viable degraded mode. The booster should fail safe by descheduling or isolating rather than by becoming a hard obstruction. The ring and vane assembly must contain or at least redirect credible fragments from the inner stage. The electronics must detect overtemperature, overcurrent, over-vibration, and implausible pressure behavior quickly enough to protect the hardware without driving nuisance shutdowns.

Table 16.1: Representative safety and failure-management logic.

Failure mode	Required design response
Inner-stage overspeed or high vibration	Immediate deschedule of booster torque and controlled shutdown of secondary path
Plenum blockage or slot fouling	Detect through pressure deviation and revert to outer-spool-only mode
Outer winding overtemperature	Torque derating with aircraft-level thrust redistribution
Inverter fault in one channel	Preserve independent operation of the remaining spool where safe
Bearing degradation indication	Maintenance flag and restricted operating envelope before hard failure

Safety case logic also benefits from the architecture's natural modularity. A removable inner cartridge can be inspected, balanced, and replaced without disturbing the main outer rotor. That reduces maintenance burden and makes the concept easier to certify internally even before any formal certification path is defined. If the team tries to fuse everything into one inaccessible assembly too early, both reliability and serviceability will suffer.

Acoustic and safety validation should be coupled during testing. Overspeed tests, burst tests, and endurance runs should be instrumented acoustically because tonal emergence often tracks structural or flow problems before any catastrophic failure occurs. Likewise, flow-visualization tests should include a basic microphone array or at least spectral logging. It is much cheaper to understand where the tones come from during bench work than after the architecture is deeply embedded in an aircraft.

The architectural message is simple. The I.C.2E will never win by ignoring noise and safety until late in the program. It can win only if the same design discipline that enables thrust and efficiency also produces predictable failure behavior and an acoustically tolerable operating envelope.

Chapter 17

Manufacturing Plan

Manufacturability is a competitive advantage for this concept only if it is planned explicitly. The architecture contains more subsystems than a conventional propulsor, but many of those subsystems can be organized into repeatable modules rather than bespoke one-off parts. The recommended manufacturing plan is built around a staged convergence process: proof-of-physics builds first, pilot-manufacturable modules second, and product-oriented part consolidation only after the data justify it.

Figure 17.1 shows a representative process chain. The bore insert and static ring are strong candidates for early additive manufacturing because they contain ducts, slot geometry, and mounting features that benefit from geometric freedom. The blades and outer shell are better aligned with established composite tooling and balancing workflows. Electronics and sensors should be packaged as swappable modules from the outset so testing and rework do not require destructive disassembly.

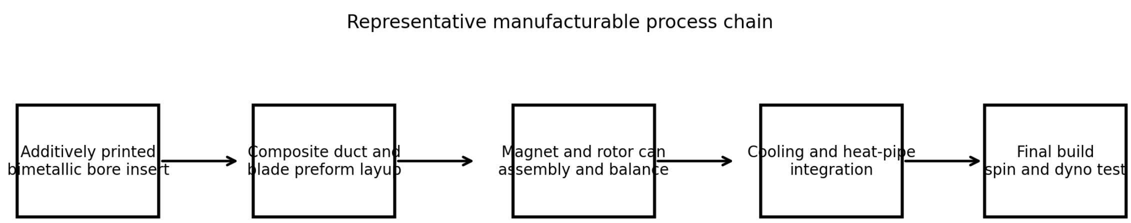

Figure 17.1: Representative manufacturable process chain for an instrumented I.C.2E prototype.

The first production-like build should probably be split into five modules: outer rotor with blades, static ring with vane carrier, inner booster cartridge, power-electronics cartridge, and cooling manifold. That modular split allows the team to iterate the high-risk parts while keeping the low-risk parts stable. It also aligns with serviceability. A failed inverter should not require destruction of the aerodynamic core. A slot-geometry experiment should not require a new full composite rotor.

Balancing and traceability deserve special handling. Because the concept combines a large outer rotor and a smaller high-speed inner rotor, the manufacturing process should include separate balance operations and separate serial traceability for each module. The temptation to do a single final balance and call the problem solved should be resisted. Inner-stage balance drift, adhesive cure variability, or minor plenum eccentricity can all matter at high speed even when the outer rotor still appears healthy.

The cooling system also needs product-minded manufacturing choices. Fittings, seals, and service ports should be standardized early. The worst outcome would be a beautifully integrated cooling scheme that is impossible to leak-test or repair. Likewise, electrical connectors should be chosen for vibration endurance and inspectability rather than for compactness alone. Compactness is useful, but a dead propulsion unit in service is not.

The manufacturing plan should also include a deliberate transition path from additive and machined parts into scalable processes. That means designing critical geometries with eventual casting, molding, or laminate layup in mind even if the first prototypes are printed or milled. It is easier to preserve scale-up freedom early than to retrofit it after the geometry has become dependent on prototype-only features.

Overall, the manufacturing message is that the I.C.2E does not need to be simple to be manufacturable. It needs to be modular, inspectable, and disciplined in how prototype freedom is turned into product discipline. If that happens, the architecture can move quickly without becoming a maintenance or quality-control trap.

Chapter 18

Test and Validation Roadmap

The concept should be validated in a ladder, not in one dramatic demonstration. That ladder starts with structural and electrical sanity checks, moves into isolated booster characterization, then into integrated static-thrust and flow-visualization work, and only later into endurance and aircraft integration. The purpose of the ladder is to preserve causality. If the team tries to validate everything at once, it will be impossible to know whether a gain or failure came from the booster, the bore, the vanes, the controls, or the outer rotor.

Figure 18.1 summarizes the recommended validation sequence. The first station is a thrust stand instrumented with load cell, voltage, current, RPM, and thermal sensing. The second is a flow-visualization rig with pressure taps and either smoke or particle-imaging capability around the bore. The third is a durability and safety station focused on overspeed, containment, bearing life, and thermal cycling.

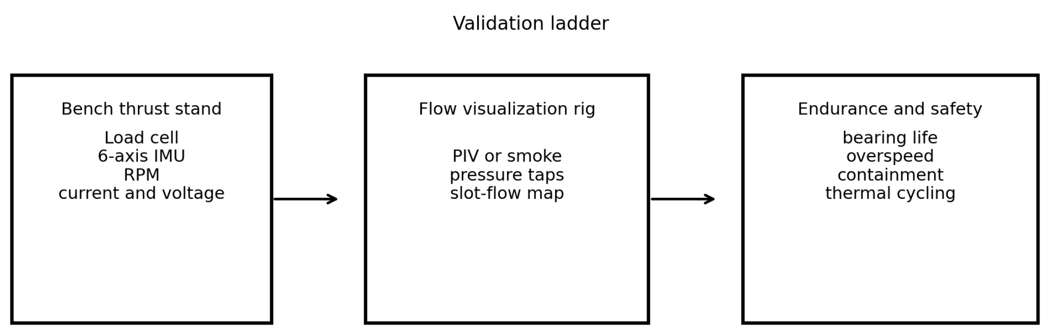

Figure 18.1: Validation ladder from thrust-stand work to flow and endurance testing.

Phase zero should be documentation and instrumentation readiness. Before spinning any hardware, the team should finalize data acquisition channels, thermocouple and pressure-tap placement, safe overspeed procedures, and inspection criteria for each module. At minimum that instrumentation set should include rotor speed for both spools, plenum pressure, duct static taps, inferred or calibrated slot mass flow, inverter and winding temperatures, and an IMU or dedicated vibration package on the containment ring. This is boring work, but it is what separates interpretable test data from anecdote.

Phase one should isolate the booster. Characterize its pressure rise, current draw, speed response, and vibration behavior in a safe rig. The objective is to understand whether the secondary machine itself is healthy before it is coupled to the full propulsor. Phase two should couple the booster to the plenum and slot system with the outer rotor inactive or removed where practical, simply to visualize whether the annular jet remains attached and whether the bore sees meaningful entrainment.

Phase three is the first integrated propulsion phase. Measure outer-only thrust, power, and temperature over the known datasheet operating range. Then add scheduled booster assist and repeat the same points. The critical outputs are not just thrust; they are change in net thrust at fixed electrical input, change in power at fixed thrust, plenum pressure behavior, and acoustic consequences. The recommended A-B-C comparison set is explicit: the same outer fan with no duct, the same outer fan inside the duct but with slot jets inactive, and the full duct plus slot-jet plus vane-ring configuration. If those do not move in the right direction, the geometry should be iterated before any endurance program begins.

An optional parallel workstream can materially strengthen stakeholder interest if framed correctly. PRISM-derived quantum magnetometry can be trialed as a development diagnostic for non-contact magnetic-state monitoring around the inner and outer machines, especially where standard probes are polluted by vibration, thermal drift, and broadband electromagnetic background [3]. Separately, DDRM and CALIPR-style differential learning can be used on the test stand to learn mappings from sensor-pattern stacks such as pressure taps, microphone arrays, accelerometers, current harmonics, and camera data to corrective settings even while the rig drifts thermally and mechanically [4, 5]. These technologies should be positioned as accelerants for development, calibration, and fault diagnosis, not as prerequisites for the basic propulsion architecture.

Phase four is endurance and abuse testing. This includes thermal soak, repeated transient cycling, dirty-inlet tolerance, overspeed margin, and retained-failure demonstrations. Only once those phases are complete does it make sense to integrate the unit into a ring-wing frame or aircraft subsystem. This order is not conservative for its own sake. It is the fastest route to trustworthy evidence because it lets the team prune weak design branches early.

Success criteria should also be explicit. The booster should demonstrate repeatable pressure behavior without unstable surge-like events. The combined system should show a measurable net benefit relative to outer-only operation at the points that matter to the target aircraft. Thermal and vibration margins should remain acceptable. And the safety case should improve with each design turn, not deteriorate as performance is pushed upward. The cleanest first comparison set is also explicit: same outer fan without duct, ducted outer fan without slot jets, and full duct plus slot jets plus vane set. If those criteria are met, the concept will have crossed from interesting idea into serious propulsion program.

Chapter 19

Aircraft Integration for eVTOL and Drone Use

The I.C.2E should ultimately be judged not as a standalone rotor but as an aircraft subsystem. That is where several of its unusual features become valuable. A ring-wing or annular frame can provide protected side and aft inlets for the reverse-flow booster. It can also provide structure that turns secondary air smoothly into the plenum while shielding the central system from direct debris ingestion. In that context, the propulsion unit is not fighting its installation. It is using it.

For heavy-lift drones the integration story is slightly different. The unit can still add value through high thrust density and improved thermal behavior, but the side-inlet and frame-coupling advantages may be weaker if the propulsor is mounted on a simple arm or pylon. That suggests the earliest commercial expression of the architecture may be a stand-alone module whose inner-stage assist is optional rather than essential. The same hardware can then evolve into a more integrated ring-wing version where the secondary-flow advantages grow.

Distributed eVTOL aircraft are the most demanding but also the most potentially rewarding use case. In such aircraft, propulsion is not only about thrust generation. It is about installation volume, wiring mass, cooling strategy, maintenance interval, acoustic quality, and redundancy behavior. A propulsion unit that can offer better installed thrust density, lower current at the aircraft bus, and cleaner degraded-mode behavior can be worth much more than one that merely posts a higher static-thrust number on a bench.

The architecture also invites interesting vehicle-level control possibilities. Because the booster is a secondary machine, it may be used to shape local wake or inflow in regimes where the main rotor should not be modulated aggressively. This could help with disturbance rejection, transition tuning, or local interference mitigation between neighboring propulsors in a distributed array. Such uses should remain secondary to the primary thrust mission, but they are another reason the concept may be more valuable in an aircraft than on a stand-alone test stand.

Thermal and electrical integration at the vehicle level further support the concept. A high-voltage bus with localized inversion fits naturally into advanced airframes that already require distributed electrical power management. Shared coolant loops or modular cooling manifolds are also more practical at the aircraft level than in an isolated propulsor experiment. In short, some of the architecture's apparent complexity becomes normal once the whole aircraft is considered.

The aircraft-integration chapter therefore supports a phased business narrative. Sell and validate the propulsion unit first where it can stand on its own. Then show how the same unit becomes more valuable when embedded into a ring-wing or distributed-lift system specifically designed to exploit its secondary flowpath. That is a much stronger story than leaning on a single generalized headline number independent of installation.

Chapter 20

IP and Program Roadmap

The IP position should be built around a coordinated claims structure: coreless stator architecture, induction-based airflow mechanisms, Coanda-guided surfaces, entrainment features, high-voltage powertrain logic, and integrated thermal management working together in a manufacturable and controllable system. In patent language, the differentiator is the coordinated system: reverse-flow booster, annular plenum, shaped bore, slot discharge, static vane containment ring, and dual-spool electrical control.

That framing also makes the program roadmap cleaner. The first milestone should not be filing around a loosely stated headline performance claim. It should be documenting the specific geometry, flowpath arrangement, and control logic that turn the central bore into a productive secondary-flow machine. Test evidence that the bore slots alter pressure distribution, maintain attachment, and improve net performance will make every subsequent claim stronger.

From a program-management standpoint, three tracks should run in parallel. The first is hardware and test, focused on the removable inner cartridge and instrumented static ring. The second is modeling, focused on reduced-order performance prediction, CFD of the bore and slot region, and structural/rotordynamic screening. The third is documentation, focused on IP support, investor communication, and partner-ready white-paper outputs like the present package. None of these tracks replaces the others. They reinforce each other.

A fourth, optional advanced-instrumentation track is defensible if Aerial MECHANICA wants to signal technical depth to stakeholders without overcommitting the base product. That track would pair LBNL-origin PRISM sensing as a robust magnetic-state diagnostic under harsh vibration and drift with DDRM/CALIPR-style differential learning for unstable multi-sensor calibration and control-assist on the rig [3–5]. The value proposition is not that these methods are required to make the propulsion unit work. The value proposition is that they can shorten learning cycles, improve observability of failure precursors, and create a stronger partnership narrative around advanced development tooling.

Table 20.1: Representative phased roadmap for the I.C.2E program.

Phase	Primary deliverable	Decision objective
Phase 1	Instrumented inner cartridge and static ring bench hardware	Determine whether the secondary-flow concept produces measurable aerodynamic benefit
Phase 2	Integrated propulsion demonstrator with validated thermal and vibration envelope	Decide whether to lock in geometry for pilot-manufacturable build
Phase 3	Ring-wing or airframe-integrated demonstrator with high-voltage distribution logic	Determine aircraft-level value proposition and partner readiness

The white paper also suggests a specific messaging discipline for external audiences. Use the present datasheet to demonstrate that the baseline propulsion unit already has real thrust and power credibility. Use the reverse-flow concept to explain the next technical step. Use language centered on installed thrust density, controllability, thermal endurance, and integration value, with benchmark conditions made explicit whenever comparative claims are used. That messaging is both more rigorous and more persuasive than a blanket multiplier statement, because it reads like engineering rather than hype.

Finally, the program roadmap should be honest about what is still open. The spec-summary PDF is not strong evidence. The concept still needs slot-pressure measurements, acoustic characterization, and endurance testing. The booster still needs a contained rotordynamic solution. None of those gaps are fatal. They are simply the work that converts an attractive concept into a credible company asset. The advantage of the I.C.2E is that it already has a real geometric and performance baseline from which to do that work.

Chapter 21

Conclusion

The I.C.2E propulsion unit already exists in a form that is technically worth taking seriously. The current internal documents describe a large-diameter integrated propulsor with meaningful thrust, coherent geometry, and a distinct central bore that conventional designs would largely waste. That alone makes the concept worth further engineering. What makes it especially interesting is that the same geometry appears compatible with a disciplined reverse-flow secondary system that could turn the bore into a productive flowpath.

This white paper has taken a deliberately rigorous position. It has not assumed that slot jets or entrainment magically multiply thrust. It has instead identified the actual mechanisms by which the architecture could outperform a conventional unit: effective-area recovery, swirl management, improved pressure distribution, lower electrical loss through higher-voltage operation, and better thermal control. It has also identified the corresponding risks: added mixing loss, tonal noise, structural and containment complexity, and the danger of asking the inner stage to do too much.

The central engineering recommendation is clear. Build the next prototype around a removable, instrumented inner cartridge and a multifunction static ring. Preserve the documented outer propulsor scale and use outer-spool-only operation as the reference condition in every test. Measure plenum pressure, wake quality, net thrust, thermal state, vibration, and acoustic response before making stronger claims. In other words, convert the concept from a persuasive narrative into a measured system.

If the combined tests show that the secondary-flow architecture produces a repeatable net benefit at acceptable mass, noise, and safety cost, then Aerial MECHANICA will have a legitimately differentiated propulsion platform for ring-wing drones and distributed eVTOL aircraft. If the gain is smaller than hoped, the company still retains a high-performance integrated propulsor baseline with clear next-step improvements in voltage architecture and thermal management. That asymmetry makes the program attractive: the downside is manageable, while the upside is strategically meaningful.

The value of the present package is therefore twofold. It provides a full-length technical narrative with figures, references, and appendices suitable for internal, partner, and investor use. And it provides a disciplined engineering frame for the next design cycle. The I.C.2E should now be advanced as a measured, modular propulsion program with a strong baseline and a clear experimental path. That is how it becomes credible enough to build, fund, and eventually fly.

Appendix A

Nomenclature

Symbol or term	Meaning
A	Effective propulsor disk area or local flow area, depending on context
C_μ	Momentum coefficient of a slot or jet-driven surface flow
I^2R **loss**	Copper resistive loss in electrical conductors
BLDC	Brushless direct-current machine, used here in the common industry sense
Bore	Central opening inside the I.C.2E hub, treated as an active flowpath in the proposed concept
Coanda attachment	Tendency of a surface jet to remain attached to a curved wall under suitable pressure-gradient conditions
Coreless stator	Stator winding arrangement without a conventional iron tooth core
DN	Bearing speed factor used to screen bearing operating severity
Distributed eVTOL	Electric vertical takeoff and landing aircraft using multiple propulsion units distributed across the airframe
Effective area recovery	Practical increase in useful induced-flow area achieved by changing the flowfield rather than by enlarging the rotor diameter alone
ESC	Electronic speed controller or inverter stage
FOC	Field-oriented control
HV / LV	High-voltage and low-voltage winding or supply options in current internal documents
I.C.2E	Induction-Coanda-Entrainment architecture as framed in this white paper
Inner booster	Reverse-flow secondary machine that feeds the annular slot system
Plenum	Circumferential pressure volume between booster discharge and slot outlets
Ring-wing	Annular or semi-annular airframe structure that can serve as both vehicle frame and inlet guide architecture
T	Thrust

Symbol or term	Meaning
v_i	Ideal induced velocity in momentum-theory treatment

Appendix B

Key Equations

This appendix records the reduced-order relationships used throughout the white paper. They are not offered as a full design model; they are included to make the conceptual logic explicit.

B.1 Momentum-Theory Relations

For a hovering actuator disk,

$$v_i = \sqrt{\frac{T}{2\rho A}}, \tag{B.1}$$

and the corresponding ideal induced power is

$$P_i = T v_i = \frac{T^{3/2}}{\sqrt{2\rho A}}. \tag{B.2}$$

These expressions show why any strong performance claim must eventually map back to either larger effective area or lower loss.

B.2 Tip Speed Estimate

The blade or impeller tip speed is

$$V_{\text{tip}} = \omega r = 2\pi \frac{N}{60} r, \tag{B.3}$$

where N is rotational speed in rpm and r is the local radius. This is why an inner booster at 16 000 rpm is not automatically near sonic: the tip radius may be small enough that absolute speed remains moderate.

B.3 Electrical Loss Framing

Copper loss scales with current as

$$P_{\mathrm{cu}} = I^2 R. \tag{B.4}$$

This is the primary reason higher aircraft-bus voltage is attractive: for the same power transfer, current decreases and so do conductor and switching burdens.

B.4 Hoop Stress Screening

For a thin rotating ring, a first-pass centrifugal hoop-stress screening relation can be written as

$$\sigma_\theta \approx \rho_m \, \omega^2 r^2, \tag{B.5}$$

where ρ_m is material density. This is not sufficient for final containment sizing, but it is useful for early scaling arguments and for identifying when fragment retention must dominate the structural concept.

Appendix C

Extracted Datasheet Tables

The following tables reproduce the most useful performance data extracted from the current datasheet text. They are included here so the main white paper can stay focused on design interpretation while preserving the numerical baseline in one place.

C.1 Motor Baseline Summary

Parameter	Value
Nominal voltage	52 V / 74 V
No-load speed	3550 rpm
Nominal speed	3000 rpm
Nominal torque	15 N.m continuous
Nominal current	110 A LV / 75 A HV
Stall torque	95 N.m LV / 140 N.m HV
Max efficiency	92 percent
Pole pairs	18
Weight	2.1 kg motor plus integrated propulsor
Outer hub / inner hub	264.16 mm / 199.16 mm
Integrated propulsor diameter	707.16 mm

C.2 Propulsor 27.8 x 8.5 LV

Table C.2: Extracted datasheet values for 27.8 x 8.5 LV.

Speed [rpm]	Current [A]	Torque [Nm]	Thrust [g]	Power [W]	Efficiency [g/W]
900	2.4	1.18	1890	127	14.9
1100	4.5	1.77	2823	232	12.2

Speed [rpm]	Current [A]	Torque [Nm]	Thrust [g]	Power [W]	Efficiency [g/W]
1300	7.4	2.47	3943	382	10.3
1500	11.3	3.29	5250	588	8.9
1700	16.4	4.23	6743	855	7.9
1900	23.0	5.28	8423	1194	7.1
2100	31.0	6.45	10290	1612	6.4
2300	40.7	7.74	12343	2118	5.8
2500	52.3	9.14	14583	2720	5.4
2700	65.9	10.66	17010	3426	5.0
2900	81.6	12.30	19623	4245	4.6
3100	99.7	14.06	22423	5186	4.3
3300	120.3	15.93	25410	6256	4.1
3500	143.5	17.92	28583	7463	3.8

C.3 Propulsor 27.8 x 8.5 HV

Table C.3: Extracted datasheet values for 27.8 x 8.5 HV.

Speed [rpm]	Current [A]	Torque [Nm]	Thrust [g]	Power [W]	Efficiency [g/W]
900	1.6	1.13	1890	121	15.6
1100	3.0	1.69	2823	222	12.7
1300	4.9	2.37	3943	366	10.8
1500	7.6	3.15	5250	562	9.3
1700	11.1	4.05	6743	819	8.2
1900	15.4	5.06	8423	1143	7.4
2100	20.9	6.18	10290	1543	6.7
2300	27.4	7.41	12343	2028	6.1
2500	35.2	8.75	14583	2604	5.6
2700	44.3	10.21	17010	3281	5.2
2900	54.9	11.78	19623	4065	4.8
3100	67.1	13.46	22423	4965	4.5
3300	80.9	15.25	25410	5990	4.2
3500	96.6	17.16	28583	7146	4.0

C.4 Propulsor 27.8 x 10 LV

Table C.4: Extracted datasheet values for 27.8 x 10 LV.

Speed [rpm]	Current [A]	Torque [Nm]	Thrust [g]	Power [W]	Efficiency [g/W]
900	2.5	1.24	1800	132	13.6
1100	4.6	1.85	2689	242	11.1
1300	7.7	2.58	3756	399	9.4
1500	11.8	3.43	5000	612	8.2
1700	17.1	4.41	6422	892	7.2
1900	23.9	5.51	8022	1245	6.4
2100	32.3	6.73	9800	1681	5.8
2300	42.5	8.07	11756	2208	5.3
2500	54.5	9.53	13889	2836	4.9
2700	68.7	11.12	16200	3572	4.5
2900	85.1	12.83	18689	4426	4.2
3100	104.0	14.66	21356	5407	3.9
3300	125.4	16.61	24200	6522	3.7
3500	149.6	18.68	27222	7781	3.5

C.5 Propulsor 27.8 x 10 HV

Table C.5: Extracted datasheet values for 27.8 x 10 HV.

Speed [rpm]	Current [A]	Torque [Nm]	Thrust [g]	Power [W]	Efficiency [g/W]
900	1.7	1.17	1800	126	14.3
1100	3.1	1.75	2689	229	11.7
1300	5.1	2.45	3756	378	9.9
1500	7.9	3.26	5000	581	8.6
1700	11.4	4.18	6422	846	7.6
1900	16.0	5.22	8022	1181	6.8
2100	21.6	6.38	9800	1595	6.1
2300	28.3	7.66	11756	2095	5.6
2500	36.4	9.05	13889	2691	5.2
2700	45.8	10.55	16200	3390	4.8
2900	56.8	12.17	18689	4200	4.4
3100	69.3	13.91	21356	5131	4.2
3300	83.6	15.76	24200	6189	3.9
3500	99.8	17.73	27222	7384	3.7

C.6 Propulsor 27.8 x 12 LV

Table C.6: Extracted datasheet values for 27.8 x 12 LV.

Speed [rpm]	Current [A]	Torque [Nm]	Thrust [g]	Power [W]	Efficiency [g/W]
900	2.6	1.27	1665	136	12.2
1100	4.8	1.90	2487	249	10.0
1300	7.9	2.66	3474	411	8.5
1500	12.1	3.54	4625	631	7.3
1700	17.7	4.54	5941	919	6.5
1900	24.7	5.67	7421	1283	5.8
2100	33.3	6.93	9065	1732	5.2
2300	43.8	8.31	10874	2276	4.8
2500	56.2	9.82	12847	2922	4.4
2700	70.8	11.46	14985	3681	4.1
2900	87.7	13.22	17287	4562	3.8
3100	107.2	15.10	19754	5572	3.5
3300	129.3	17.12	22385	6722	3.3
3500	154.2	19.25	25181	8019	3.1

C.7 Propulsor 27.8 x 12 HV

Table C.7: Extracted datasheet values for 27.8 x 12 HV.

Speed [rpm]	Current [A]	Torque [Nm]	Thrust [g]	Power [W]	Efficiency [g/W]
900	1.8	1.21	1665	130	12.8
1100	3.2	1.81	2487	237	10.5
1300	5.3	2.52	3474	391	8.9
1500	8.1	3.36	4625	600	7.7
1700	11.8	4.32	5941	873	6.8
1900	16.5	5.39	7421	1219	6.1
2100	22.2	6.59	9065	1646	5.5
2300	29.2	7.90	10874	2163	5.0
2500	37.5	9.34	12847	2778	4.6
2700	47.3	10.89	14985	3499	4.3
2900	58.6	12.56	17287	4336	4.0
3100	71.6	14.36	19754	5296	3.7
3300	86.3	16.27	22385	6389	3.5
3500	103.0	18.30	25181	7622	3.3

Appendix D

Figure and Modeling Notes

The figures bundled with this project are intentionally illustrative. They are not reverse-engineered from proprietary internal CAD and should not be treated as final production geometry. Their purpose is to communicate system logic clearly in the white paper and in supporting presentations.

D.1 Cross-Section Figure

Figure 1 shows the outer propulsor, the reverse-flow booster, the shaped bore, the annular slot jets, and the static vane and containment ring. It should be interpreted as an architecture sketch rather than a dimensioned design.

D.2 Airflow Figure

Figure 2 illustrates the intended secondary-flow path: side and rear inlet ingestion, reverse-flow compression, slot discharge, entrainment, and conditioned exit flow. It is especially useful in proposal and investor contexts because it explains where the claimed value is supposed to come from.

D.3 Electromagnetic and Thermal Figures

Figures 3 through 5 are conceptual schematics used to show radial-flux packaging, heat rejection paths, and dual-spool electrical architecture. They are not substitutes for electromagnetic finite-element analysis or detailed thermal network models.

D.4 Containment, Manufacturing, and Validation Figures

Figures 6 through 8 communicate program logic. They show how the concept can be made safe, manufacturable, and testable without assuming that the first prototype already solves every aerodynamic or structural problem.

Appendix E

Open Questions

This appendix records the highest-value unanswered questions that remain after the present white paper package.

E.1 Aerodynamic Questions

1. What slot width, lip radius, and discharge angle maximize net thrust benefit rather than local jet momentum?
2. How stable is bore-wall attachment across hover, transition inflow, and crosswind disturbance?
3. How much swirl actually remains downstream of the outer rotor, and what turning should the static vane ring be designed to recover?

E.2 Electro-Thermal Questions

1. What winding and inverter topology best support a high-speed inner stage without excessive switching loss or EMI burden?
2. How much temperature rise occurs in the bore and ring under sustained dual-spool operation?
3. Which parts of the cooling architecture should be liquid-cooled and which should remain passive spreaders?

E.3 Structural and Program Questions

1. What is the most credible retained-failure strategy for the inner stage at target overspeed conditions?

2. Can the concept be prototyped in a removable cartridge form without sacrificing the final product path?
3. Which benchmark definition should Aerial MECHANICA use when speaking externally about “3x” improvement so that the claim remains defensible?

Appendix F

LBNL Sensing and Control Relevance

This appendix records the rationale for referencing Lawrence Berkeley National Laboratory sensing and machine-learning control work in the I.C.2E development program. The purpose is not to claim that propulsion validation suddenly becomes a quantum-sensing problem. The purpose is to identify where those inventions can add credible value in development, diagnostics, and stakeholder positioning.

F.1 PRISM as an Advanced Development Diagnostic

PRISM is presented as a robust magnetometry approach in which differential readout between long-lived prethermal orbit axes suppresses common-mode disturbances including vibration, temperature and strain variation, pulse imperfections, bias drift, and background fields while preserving broadband magnetic sensitivity [3]. That profile is unusually relevant to propulsion development because electric propulsion rigs are full of exactly those disturbance channels: magnetic background from power electronics, temperature drift during thermal soak, structural vibration during spin tests, and control-pulse or current ripple imperfections in the inverter chain.

The strongest propulsion use case is therefore not direct aerodynamic measurement. It is non-contact magnetic-state observation. A PRISM-inspired diagnostic layer could, in principle, be used to monitor stray-field signatures from the inner and outer machines, observe changing harmonic content associated with rotor-stator asymmetry or current imbalance, and distinguish true electromechanical changes from common-mode environmental disturbances on a harsh test stand. If implemented, it should be positioned as an advanced test and health-monitoring instrument rather than as a replacement for standard tachometry, thermocouples, pressure taps, or accelerometers.

F.2 DDRM and CALIPR as Methods for Unstable Test Systems

The DDRM and CALIPR papers are more directly transferable at the algorithmic level. Their central contribution is not laser-specific optics. It is a control-learning method for systems whose absolute state drifts during training, making direct labeled learning unreliable. DDRM addresses this by learning differential action from pairs of measurements separated by known perturbations, while CALIPR-style pattern recognition uses rich sensor patterns rather than a single scalar detector to drive rapid correction [4, 5]. In the laser experiments, these methods reduced noise, tolerated drift during learning, and converged much faster than stochastic parallel gradient descent.

That logic fits propulsion development surprisingly well. Inner-booster and integrated-propulsor rigs are thermally drifting, mechanically noisy, and often under-instrumented during early development. A DDRM-like framework could be used to learn the relationship between sensor-pattern stacks and corrective actions even while the rig drifts. The patterns do not need to be optical interference fringes; they could be assembled from pressure-tap vectors, microphone spectra, accelerometer arrays, current harmonics, magnetic probes, thermal maps, or high-speed camera features. The key transferable idea is differential remapping under drift, not the specific optical hardware.

F.3 Recommended Positioning Inside the I.C.2E Program

1. Use PRISM only where its disturbance rejection is actually valuable: advanced motor-field diagnostics, EMI discrimination, and development-time health monitoring in vibration-heavy or temperature-drifting rigs.
2. Use DDRM and CALIPR only as development and calibration accelerants: multi-sensor rig tuning, unstable-system learning, anomaly classification, and rapid controller prototyping.
3. Do not position either technology as necessary for first-flight propulsion control. The base propulsion architecture should remain credible with conventional sensing and deterministic control.
4. Do position them as high-value differentiators for stakeholder interest, especially where Aerial MECHANICA wants to emphasize rigorous development infrastructure, observability, and collaboration with LBNL-origin inventions.

In summary, there is a valid argument for incorporating both technology lines into the white paper, but only on bounded terms. PRISM is a strong candidate for advanced non-contact diagnostics under harsh noise and drift. DDRM and CALIPR are strong candidates for development-stage learning and multi-sensor control assistance in unstable rigs. Used that way, they strengthen the technical story without weakening the physical credibility of the propulsion concept.

Appendix G

Design Decisions and Baseline Departures

This appendix records the principal places where the white paper deliberately advances the design beyond the currently documented baseline. The aim is to make those decisions explicit for reviewers, partners, and future design teams.

G.1 Documented Baseline Versus Recommended Product Direction

Design element	Documented baseline	Recommended direction in this paper
Central opening ratio	About 28 percent of propulsor diameter from the present 199.16 mm bore inside a 707.16 mm propulsor	Reduce to about 18 percent in the next product-oriented iteration to recover effective disk area while retaining an active central flowpath
Performance language	Historic "thrust multiplier" framing and occasional generalized multiplier language	Replace with language centered on installed thrust density, thermal endurance, controllability, and integration value
Inner-stage intent	Large central opening invites aggressive secondary-flow interpretation	Keep booster role bounded: modest pressure-ratio plenum feed and flow conditioning rather than heroic compressor behavior

Design element	Documented baseline	Recommended direction in this paper
Stakeholder narrative	Geometry-first concept emphasis	Product-development emphasis with explicit validation ladder, design decisions, and measurable comparison cases

G.2 Reason for Recommending an 18 Percent Opening

The current bore is valuable as a development reference because it makes the secondary-flow physics easy to visualize and instrument. It is not, however, the preferred end-state geometry in this paper. A reduction toward roughly 18 percent of propulsor diameter offers three advantages simultaneously: more effective actuator area for the main outer propulsor, a more favorable structural and containment package in the center, and a still-useful annular region for slot-jet driven entrainment and pressure recovery. In other words, the design keeps the aerodynamic function of the central system while reducing the disk-area penalty associated with an oversized opening.

G.3 Preferred External Performance Clause

The recommended clause for partner and investor-facing use is:

> The I.C.2E propulsion architecture is being developed to deliver substantially improved installed thrust density, thermal endurance, and flow-control authority relative to conventional direct-drive propulsion units occupying the same aircraft integration envelope.

This wording is more durable than any simple multiplier statement because it remains physically grounded and naturally accommodates future benchmark data.

Bibliography

[1] Aerial MECHANICA / AVIUS. I-C-2E Thrust Multiplier Datasheet Final, 2026. Internal datasheet, Feb. 2026 edition / subject to change.

[2] Aerial MECHANICA / AVIUS. I.C.2E Thrust Multiplier Motor-Prop Unit Drawing AM-ICE-1000-AXFX, 2026. Internal drawing package, Rev. 01, dated 16 Feb 2026.

[3] Enrico Daniel Richter, Ryan J. Smith, Brayden Glockzin, Emanuel Druga, Thomas Schenkel, and Ashok Ajoy. Robust quantum sensing via prethermal spin orbits. *arXiv preprint arXiv:2603.21057*, 2026. URL `https://arxiv.org/html/2603.21057v1`.

[4] Dan Wang, Qiang Du, Tong Zhou, Derun Li, and Russell Wilcox. Stabilization of the 81-channel coherent beam combination using machine learning. *Optics Express*, 29(4):5694–5709, 2021. doi: 10.1364/OE.414985.

[5] Dan Wang, Qiang Du, Tong Zhou, Antonio Gilardi, Mariam Kiran, Bashir Mohammed, Derun Li, and Russell Wilcox. Machine learning pattern recognition algorithm with applications to coherent laser combination. *IEEE Journal of Quantum Electronics*, 58(6), 2022. doi: 10.1109/JQE.2022.3204437.