

# Quiescent Lockers (Q-Lockers): The First Commercial Active Vibration Isolation for Microgravity Manufacturing

## Condensed Q-Locker Technical Proposal

Russell Boyer | Purdue Aero/Astro '29

[boyer103@purdue.edu](mailto:boyer103@purdue.edu)   [russellhaien@gmail.com](mailto:russellhaien@gmail.com)   +1-714-722-5168

---

### 1. Executive Summary

With upcoming commercial space stations (Vast Haven-1, Axiom, Starlab) increasing access to microgravity environments, microgravity manufacturing and research has been of growing importance. While raw microgravity conditions are suitable for many payloads, high-sensitivity projects like semiconductor crystal growth, protein crystallization, exotic alloy solidification, quantum sensor calibration, etc., all require sustained acceleration environments below 1  $\mu\text{g}$  RMS. These upcoming commercial space stations provide the orbital platform but not the vibrational stillness. Even transient events like crew motion, thruster firings, and structural flexure produce 100–1,000  $\mu\text{g}$  of disturbance at the payload level. Legacy isolation systems (ARIS, MVIS, g-LIMIT, MAIS/MAVIS) were designed for full-rack payloads on government-funded platforms with big budgets, consume excessive volume, and are not commercially available. No locker-level active vibration isolation product exists on the market at any price.

The MLE-AVIS (Middeck Locker Equivalent - Active Vibration Isolation System, a.k.a Q-Locker) described below is a retrofittable, contactless, 6-DOF active vibration isolation platform designed for single and double MLE payloads on commercial space stations. It uses a planar sleeve-in-sleeve architecture with pure Lorentz-force levitation: two concentric thin-walled aluminum frames separated by a 4 mm air gap, with 24 flat electromagnetic actuators across six faces providing fully decoupled force and torque control. No physical contact occurs between the payload-carrying inner frame (flotor) and the station-mounted outer frame (stator) during on-orbit operation. A single hybrid flexible printed circuit (FPC) ribbon carries all power and data across the gap.

Magnetic negative stiffness (with  $K_{neg} \approx -38,200$  N/m being the baseline when permeable stator materials are present) is reduced by 99.7% through non-permeable stator construction to a residual  $< 100$  N/m per axis. Earnshaw's theorem does not allow passive magnetic stability in all axes simultaneously, so rather than partially mitigating with passive magnets (or any other forms of passive systems) at the cost of axis asymmetry and system complexity, the design relies entirely on the active controller, which stabilizes the residual 0.31 Hz unstable pole with  $< 0.2$  A DC bias per coil ( $< 1$  W total, a bandwidth ratio of 64,000 $\times$  relative to the instability). This two-layer architecture (material selection combined with active control) produces symmetric, identical plant dynamics in all translational axes.

A dual-loop skyhook +  $H_\infty$  feed-forward control law achieves active vibration isolation across 0.05–100 Hz ( $-20$  to  $-85$  dB), with passive mechanical decoupling providing partial isolation extending to 0.01 Hz. The sensor noise floor (not the actuator, structure, or controller) is the sole performance-limiting subsystem at every tier, enabling a sensor-swap upgrade path from  $< 3$   $\mu\text{g}$  to sub-100 nano-g without structural modification.

Parameter	Value
Target platform	Vast Haven-1 MLE bays (verified against HUG v0.7)
AVIS system mass	~4.7 kg (with 15% margin)
Customer payload mass	~25.3 kg (of 30 kg slot limit)
AVIS power consumption	~2 W nominal / <7 W peak
Baseline isolation	<3 $\mu\text{g}$ RMS (0.05–100 Hz) design target
Upgrade path (sensor swap)	<1 $\mu\text{g}$ RMS, architecture supports <100 ng
Sway authority	$\pm 3.5$ mm all axes
Actuator force constant	0.461 N/A per coil (48 turns, 26 AWG)
Halbach gap field	$B_{eff} \approx 0.20 T$ (N48SH, range 0.15–0.27, pending COMSOL) $B_{eff} \approx 0.20 T$ (N48SH, range 0.15–0.27, pending COMSOL)
Control bandwidth	20 kHz digital
BOM cost	\$2,800–3,800 (COTS components)

## 2. System Architecture

### 2.1 Two-Frame Design

The outer frame (stator) is a skeletal rectangular frame of Al 6061-T6 extruded L-section rails ( $10 \times 10 \times 1.5$  mm) with 0.3 mm aluminum face panels, conforming exactly to the Haven-1 MLE envelope ( $440 \text{ W} \times 253 \text{ H} \times 513 \text{ D}$  mm). Twelve edge rails carry all launch loads via eight corner-mounted titanium launch lock pins. Six face panels serve as mounting surfaces for coils, gap sensors, and PCB-integrated electronics, and as the thermal path from coils to station air cooling (Al 6061:  $167 \text{ W/m}\cdot\text{K}$ ).

The inner frame (flotor) is a lightweight rectangular shell of Al 6061-T6 corner rails with G10 fiberglass face panels (0.5 mm, stiffened by a  $4 \times 4$  aluminum rib grid on large panels,  $2 \times 2$  on smaller). The flotor exterior carries 24 Halbach magnet units arranged as a  $2 \times 2$  grid per face and the flotor-side accelerometers. The flotor interior is the customer payload volume, accessed via a hinged front face (door).

The nominal air gap is 4 mm on all six faces, providing  $\pm 3.5$  mm sway authority with 0.5 mm clearance to the silicone bumpers at maximum displacement.

## 2.2 Component Inventory Per Face

Each of the six faces carries four stator-side coils (26 AWG, 48-turn rectangular spirals,  $20 \times 30 \times 2$  mm, potted in thermally conductive epoxy), four flotor-side Halbach magnet units (N48SH,  $20 \times 30 \times 5$  mm, 4-segment with 0.5 mm flotor-side back-iron), and gap sensors (2 per lateral face, 1 per front/back face to summate to 10 total). Coil offsets from face center are optimized per face to balance the 6-DOF Jacobian condition number for both force and torque generation.

## 2.3 Launch Locks and Fail-Safe

Eight titanium (Ti-6Al-4V) shear pins at the flotor corners engage conical seats in the stator corner gussets, carrying all launch loads (11 g random vibration at 7 *GRMS* per HUG qualification). Pins are retracted by small solenoids and re-engaged by spring-return mechanisms (fail-closed). A 2.5 V, 10 F ultracapacitor provides ~30 seconds of controller and sensor power after main power loss, enabling orderly recentering and lock engagement. Silicone bumpers (1 mm, Shore A 30) at all eight corners prevent metal-to-metal contact during any transient excursion. *GRMS* per HUG qualification). Pins are retracted by small solenoids and re-engaged by spring-return mechanisms (fail-closed).

## 2.4 Operational Modes

The system operates in three modes: launch-locked (pins engaged, controller off), 6-face isolation (all 24 actuators active, full 6-DOF control), and 5-face crew access (front face de-energized, door open, controller switches to 5-face Jacobian with ~25% reduced Z-authority and ~50% reduced Z-torque authority). A software-only recentering loop (0.01 Hz proportional controller on gap sensors) maintains geometric center without injecting noise into the 0.05–100 Hz isolation band.

## 2.5 Mass Budget

The system mass totals approximately 4.7 kg with 15% margin, leaving 25.3 kg for customer payload within the 30 kg MLE slot limit. Stator: ~2.1 kg (rails 362 g, panels 240 g, coils 288 g, electronics 235 g, locks 160 g, misc 315 g). Flotor: ~1.0 kg (rails 480 g, panels/ribs 420 g, Halbach segments 540 g, back-iron 60 g, sensors 50 g, FPC 5 g). Margin: ~0.6 kg.

Component	Qty	Unit Mass	Total Mass
<b>Stator structure</b>			
Corner rails (Al 6061, 10×10×1.5mm L)	12	varies	362 g
Corner gussets (lock pin mounts)	8	15 g	120 g
Face panels (0.3 mm Al 6061)	6	varies	757 g
<b>Flotor structure</b>			
Corner rails (Al 6061, 8×8×1.2mm L)	12	varies	227 g

Corner brackets (lock pin receivers)	8	8 g	64 g
Face panels (0.3 mm G10 fiberglass)	6	varies	484 g
Panel stiffening ribs (Al, 3×3 mm)	lot	—	318 g
<b>Electromagnetic</b>			
Halbach segments (20×30×5 mm N48SH)	24	22.5 g	540 g
Stator coils (26 AWG, potted)	24	12 g	288 g
Back-iron shims (flotor side, 0.5 mm)	24	4.5 g	108 g
<b>Electronics</b>			
FPGA (SmartFusion2)	1	15 g	15 g
H-bridge drivers (DRV8870)	24	1.5 g	36 g
Accelerometers (SF3600A flotor, SF1500S stator)	4	10 g	40 g
Gap sensor ASICs (AD7746)	10	2 g	20 g
Temperature sensors (NTC)	4	1 g	4 g
PCBs (stator main + flotor)	2	60 g	120 g
Voltage regulators, passives	lot	—	30 g
<b>Mechanical</b>			
Launch lock solenoids + pins	8	20 g	160 g
Silicone bumpers	8	3 g	24 g
Ultra-capacitor backup	1	5 g	5 g
Door limit switch	1	2 g	2 g
<b>Interconnects</b>			

FPC umbilical	1	5 g	5 g
Door flex cable	1	3 g	3 g
Internal wiring harness	lot	—	40 g
<b>Thermal</b>			
Potting compound (coils)	24	3 g	72 g
Thermal interface material	lot	—	15 g
Fasteners & adhesive	lot	—	60 g
<b>AVIS TOTAL</b>			<b>3,919 g</b>
Margin (~15-17%)			700 g
<b>AVIS TOTAL WITH MARGIN</b>			<b>~4.7 kg</b>
<b>Available for payload</b>			<b>~25.3 kg</b>

## 2.6 Panel Structural Design

Flotor panels use G10 fiberglass (0.5 mm) rather than aluminum to maintain the non-permeable stator principle and minimize eddy-current interactions. Stiffening ribs (Al 6061,  $3 \times 5$  mm cross-section) raise the panel first bending mode above 50 Hz (target  $>5 \times$  the  $H\infty$  crossover frequency). Simply-supported first-mode estimates for bare 0.5 mm G10 panels range from 4–10 Hz depending on span; the 4x4 rib grid reduces effective span to ~50 mm, raising modes above 100 Hz with a mass penalty of ~25 g per large panel. Rib-panel bonding uses structural epoxy with peel stops at rib ends to arrest disbond propagation under thermal cycling. Note, there are Plan Bs and fallbacks.

## 3. Electromagnetic Design

### 3.1 Halbach Field Derivation

Each flotor face carries four Halbach-array magnet units (N48SH,  $B_r = 1.38$  T,  $H_{cj} = 20$  kOe,  $T_{max} = 150^\circ$  C, selected over N52 ( $B_r = 1.45$  T,  $H_{cj} = 11$  kOe,  $T_{max} = 60^\circ$  C) for superior coercivity margin under thermal cycling.). Each unit is a 4-sub-segment linear Halbach array with  $90^\circ$  progressive magnetization rotation (spatial period  $\lambda = 4 \times 5 = 20$  mm). The gap-face field for an ideal infinite Halbach array of thickness  $h$  and spatial period  $\lambda$  at distance  $g$  is:

$$B(g) = B_r (1 - e^{-2\pi h/\lambda}) e^{-2\pi g/\lambda}$$

Substituting array parameters ( $h = 5$  mm,  $\lambda = 20$  mm,  $g = 4$  mm,  $B_r = 1.38$  T for N48SH):  
Thickness factor:

$$1 - e^{-2\pi(5)/20} = 1 - e^{-\pi/2} = 1 - 0.2079 = 0.7921$$

Gap decay factor:

$$e^{-2\pi(4)/20} = e^{-2\pi/5} = e^{-1.2566} = 0.2840$$

Raw field (infinite array, no back-iron):

$$B_{raw} = 1.38 \times 0.7921 \times 0.2840 = 0.310 \text{ T}$$

It is important to note that the spatial period  $\lambda$  is determined by the full magnetization rotation cycle of the Halbach array. With four sub-segments of 5 mm width each, one complete rotation spans  $4 \times 5 = 20$  mm. A common error is to use the array footprint width (20 mm) or half the footprint as the spatial period. The correct value is the magnetization cycle length, which for a 4-segment array equals the total array width along the magnetization direction.

The formula above for  $B_{raw}$  assumes an infinite periodic array. A single 4-sub-segment unit is the minimum repeating unit. The field at the center of the unit is close to the infinite-array prediction, but falls off near the edges due to fringe effects. For a single-period array (4 sub-segments), the center-field reduction relative to infinite is typically 70–85%. I apply a conservative finite-array correction factor (working estimate) of 0.75:

$$B_{finite} = 0.310 \times 0.75 = 0.233 \text{ T}$$

A soft-iron shim of 0.5mm on the flotor-side behind the Halbach array provides a low-reluctance return path for flux leaking from the imperfectly cancelled back face, redirecting a portion of it to reinforce the front-face field. For a finite 4-segment array with ~10–20% back-face leakage, the enhancement is estimated at 5–15%. A factor of 1.15 $\times$  is applied below as a working assumption. This factor is not derived from the specific array geometry and is not independent of the finite-array correction (both depend on fringe-field behavior at the array boundaries). The combined product of 0.75  $\times$  1.15 may overestimate the net effect. COMSOL simulation will resolve the actual  $B_{eff}$  including back-iron in a single integrated calculation, replacing both correction factors with a validated result. The 1.15 $\times$  factor is retained in the analytical model to establish a working design point; if COMSOL yields a lower  $B_{eff}$ , the architecture accommodates values as low as 0.15 T without redesign.

$$B_{backiron} = 0.233 \times 1.15 = 0.268 \text{ T}$$

Gap-averaged field: The coil occupies a 2 mm thickness starting at 4 mm from the magnet face. The field decays across this thickness. The field at the far edge of the coil (6 mm from magnet face):

$$B(6 \text{ mm}) = B_r \times 0.7921 \times e^{-2\pi(6)/20} = 1.38 \times 0.7921 \times e^{-1.885} = 1.38 \times 0.7921 \times 0.1520 = 0.166 \text{ T (raw)}$$

Applying the same finite-array ( $\times 0.75$ ) and back-iron ( $\times 1.15$ ) corrections:  $0.166 \times 0.75 \times 1.15 = 0.143 \text{ T}$

Gap-averaged over the 2 mm coil thickness:

$$B_{eff} = \frac{B(4 \text{ mm}) + B(6 \text{ mm})}{2} = \frac{0.268 + 0.143}{2} = 0.206 \text{ T}$$

Summary of  $B_{eff}$  estimates:

Condition	$B_{eff}$ (T)	Notes
Raw infinite array, no back-iron, at gap face	0.310	Upper theoretical bound
Finite-array corrected ( $\times 0.75$ ), no back-iron	0.233	Realistic without back-iron
With flotor-side back-iron ( $\times 1.15$ )	0.268	At near edge of coil
Gap-averaged across coil thickness	0.206	Working design value
Pessimistic (finite correction $\times 0.65$ , back-iron $\times 1.10$ )	0.155	Lower bound
Optimistic (finite correction $\times 0.85$ , back-iron $\times 1.25$ )	0.272	Upper bound

The adopted working value is  $B_{eff} = 0.20$  T for all subsequent force calculations, with a sensitivity range of 0.15 - 0.27 T. COMSOL simulation is required to resolve the finite-array correction factor and back-iron boost, which dominate the uncertainty.

### 3.2 Lorentz Force Model

On any face, the Halbach field  $B_{eff}$  is directed normal to the face (across the gap). Current-carrying coil segments parallel to the face produce force  $dF = I \cdot dl \times B$ . In a rectangular spiral winding, each turn  $k$  (counting from  $k = 0$  at the outermost turn to  $k = (N_{per\,layer} - 1)$  at the innermost) has progressively smaller dimensions:

$$W_z(k) = W_{z,outer} - 2k \cdot p$$

Where  $p$  (0.5 mm) is the winding pitch (wire diameter + clearance). The first turn is offset from the coil edge by approximately  $p/2$ , so:

$$W_{z,outer} = W_{z,coil} - p \approx 30 - 0.50 = 29.5 \text{ mm}$$

The primary-axis (X) force from the entire coil (all layers, all turns) is:

$$F_x = N_{layers} \cdot 2 \cdot B_{eff} \cdot I \cdot \sum_{k=0}^{N_{tp}-1} W_z(k)$$

where  $N_{tpl}$  is the number of turns per layer.

Evaluating the sum for  $N_{tpl} = 12$  turns per layer,  $p = 0.50$  mm:

$$\begin{aligned} \sum_{k=0}^{11} W_z(k) &= \sum_{k=0}^{11} (29.5 - 2 \times 0.50 \times k) \times 10^{-3} = \sum_{k=0}^{11} (29.5 - k) \times 10^{-3} \\ &= 10^{-3} \times (29.5 + 28.5 + 27.5 + \dots + 18.5) = 10^{-3} \times 12 \times \frac{29.5 + 18.5}{2} = 10^{-3} \times 12 \times 24.0 = 0.288 \text{ m} \end{aligned}$$

With  $N_{layers} = 4$ :

$$F_x = 4 \times 2 \times B_{eff} \times I \times 0.288 = 2.304 \times B_{eff} \times I$$

Force constant (primary axis):

$$k_{f,x} = \frac{F_x}{I} = 2.304 \times B_{eff}$$

At  $B_{eff} = 0.20$  T, the force constant is  $k_{f,x} = 0.461$  N/A. Even at the pessimistic  $B_{eff} = 0.15$  T, the primary-axis force constant of 0.346 N/A per coil provides adequate authority.

Similarly, the cross-axis (Z) force constant:

$$\sum_{k=0}^{11} W_x(k) = 10^{-3} \times 12 \times \frac{19.5 + 8.5}{2} = 10^{-3} \times 12 \times 14.0 = 0.168 \text{ m}$$

$$k_{f,z} = 4 \times 2 \times B_{eff} \times 0.168 = 1.344 \times B_{eff}$$

At  $B_{eff} = 0.20$  T, the cross-axis force constant is  $k_{f,z} = 0.269$  N/A.

### 3.3 Coil Electrical Properties

Mean turn perimeter: 76.0 mm. Total wire length: 3.648 m. DC resistance: 0.489  $\Omega$  (20°C), 0.56  $\Omega$  (50°C).

Inductance (modified Wheeler formula): ~45  $\mu$ H. At 20 kHz PWM, the inductive impedance ( $2\pi \times 20000 \times 45 \times 10^{-6} = 5.65 \Omega$ ) dominates the resistive impedance, ensuring rapid current response.

The time constant of ~80  $\mu$ s is well below the 50  $\mu$ s control period.

### 3.4 Thermal Limits

Maximum continuous current per coil:  $I_{max} = \sqrt{P_{max}/R} = \sqrt{2.9/0.56} = 2.28$  A, where  $P_{max} = 2.9$  W is derived from the 95°C temperature rise limit (coil at 120°C, station air at 25°C) through the thermal path (epoxy potting  $\rightarrow$  Al panel  $\rightarrow$  station air convection at  $h \approx 5$  W/m<sup>2</sup>·K). The operational current limit is set conservatively at 1.5 A per coil.

---

## 4. Magnetic Negative Stiffness and Two-Layer Mitigation

### 4.1 The Problem

Baseline  $K_{neg}$  represents the worst case scenario where all stator-face materials are permeable (e.g., steel coil formers, ferrous fasteners, back-iron behind stator coils). The Q-Locker design explicitly excludes all permeable stator materials, but this baseline establishes the reference point from which mitigation is measured.

To start, we calculate the attractive magnetic pressure between a Halbach face and a soft-iron surface at gap  $g$ :

$$P_{mag} = \frac{B_{eff}^2}{2\mu_0}$$

At  $B_{eff} = 0.20$  T:

$$P_{mag} = \frac{(0.20)^2}{2 \times 4\pi \times 10^{-7}} = \frac{0.040}{2.513 \times 10^{-6}} = 15,915 \text{ Pa}$$

Attractive force per Halbach segment (face area  $A = 20 \times 30 \text{ mm} = 6.0 \times 10^{-4} \text{ m}^2$ ):

$$F_{attr} = P_{mag} \times A = 15,915 \times 6.0 \times 10^{-4} = 9.55 \text{ N per segment}$$

Negative spring constant per segment (for attraction scaling as  $F \propto 1/g^2$ , the stiffness is  $dF/dg = -2F/g$ ):

$$K_{neg,segment} = -\frac{2F_{attr}}{g} = -\frac{2 \times 9.55}{0.004} = -4,775 \text{ N/m per segment}$$

Net translational  $K_{neg}$  per axis (baseline, with back-iron on stator): For translation along an axis (e.g., Y), the two opposing faces ( $\pm Y$ ) each have 4 segments. The near face's attraction increases while the far face's attraction decreases, but nonlinearly ( $1/g^2$  on the near side vs  $1/g^2$  on the far side). For small displacements  $\delta$  from the centered position:

$$K_{neg,axis} = -2 \times 4 \times K_{neg,segment} = -8 \times 4,775 = -38,200 \text{ N/m per axis (baseline)}$$

The factor of 2 comes from the two opposing faces' contributions adding (near face becomes more attractive, far face becomes less attractive, thus both changes are destabilizing).

Baseline unstable pole frequency ( $m = 26 \text{ kg}$ , fully loaded flotor):

$$\omega_u = \sqrt{\frac{|K_{neg}|}{m_{flotor}}} = \sqrt{\frac{38,200}{26}} = \sqrt{1,469} = 38.3 \text{ rad/s} = 6.1 \text{ Hz}$$

### 4.2 Layer 1: Non-Permeable Stator Materials

All stator-face components are constructed from non-magnetic, non-ferrous materials: G10 coil formers, Al 6061 face panels, FR4 PCB substrates (without copper ground planes within 3 mm of the gap face), and titanium

fasteners. No ferromagnetic back-iron exists on the stator side (back-iron is only on the flotor side, behind the Halbach arrays).

This eliminates the dominant magnet-to-back-iron attraction. The residual  $K_{neg}$  sources are: (a) Halbach fringe fields interacting with non-magnetic conductors, which produces only velocity-dependent eddy-current damping, not position-dependent stiffness ( $\approx 0$  N/m static); (b) image forces from Halbach arrays reflecting in opposing conductive panels, which is actually a weak *repulsive* (positive) force of  $\sim 1$ – $10$  N/m; (c) Halbach-to-Halbach interactions between opposing flotor faces, which produces  $0$  N/m to gap-crossing stiffness because they are internal forces within a single rigid body.

Residual  $K_{neg}$  after Layer 1:  $< -100$  N/m per axis (down from  $-38,200$ ). A 99.7% reduction.

Caveat: eddy currents induced in the aluminum stator panels during flotor motion at the control loop frequency may (heavy emphasis) produce an effective position-dependent stiffness that the static analysis does not capture. COMSOL time-harmonic simulation (Phase 1) is required to quantify this. The two-layer architecture accommodates  $K_{neg}$  values up to  $-790$  N/m without exceeding actuator current limits.

#### 4.4 Layer 2: Active Controller Stabilization

With  $K_{neg} \approx -100$  N/m per axis and no passive cancellation, all three translational axes have identical unstable poles at:

$$f_u = \frac{1}{2\pi} \sqrt{\frac{100}{26}} = 0.312 \text{ Hz}$$

The 20 kHz controller has a bandwidth ratio of  $64,100\times$  relative to this instability. The standard magnetic bearing systems routinely stabilize poles at  $50$ – $200$  Hz with  $10$ – $20$  kHz loops ( $50$ – $200\times$  ratio). The DC stabilization current at maximum displacement ( $\delta = 3.5$  mm) per dominant-face coil:

$$I_{stab} = \frac{|K_{neg}| \times \delta}{4 \times k_f} = \frac{100 \times 0.0035}{4 \times 0.461} = 0.190 \text{ A}$$

Total stabilization power (worst case, all 3 axes at max displacement):  $0.95$  W. Typical single-axis:  $0.32$  W.

Composite stiffness with controller ( $K_p = 500$  N/m):

$$K_{total} = K_p + K_{neg} = 500 + (-100) = 400 \text{ N/m (all axes, symmetric)}$$

$K_p > |K_{neg}|$  is required for stability. At  $K_{neg} = -100$  N/m, the minimum usable  $K_p \approx 150$  N/m (with margin).

---

## 5. Control Architecture and Transfer Function

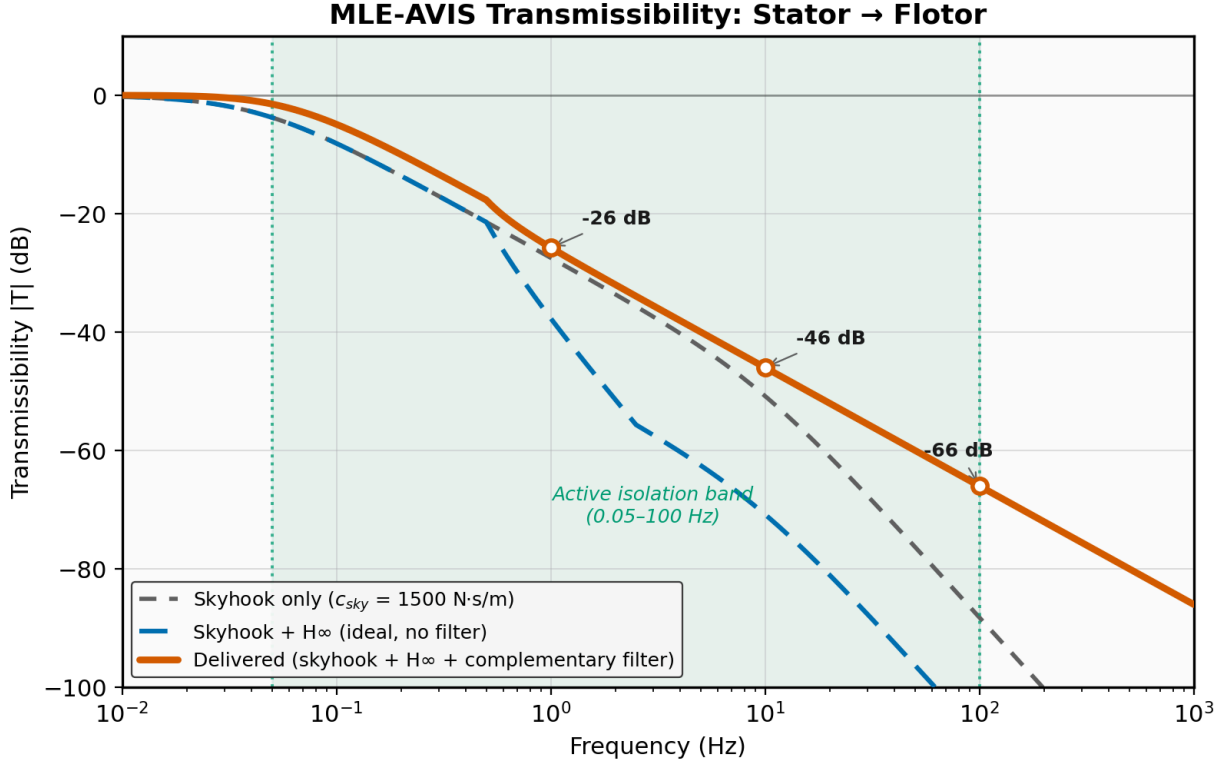


Figure 1: analytically estimated transmissibility graph from DC to 1000 Hz

### 5.1 Single-Axis Plant Model

The flotor (mass  $m$ ) is connected to the stator by the net passive stiffness  $K_{neg}$  (destabilizing) and the active control force  $F_c$ . With relative displacement  $z = x_f - x_s$ :

$$m\ddot{z} + c_{sky}\dot{z} + K_{total} \cdot z = -m\ddot{x}_s$$

where  $c_{sky}$  is the skyhook damping coefficient and  $K_{total} = K_p + K_{neg}$ . The closed-loop transmissibility (ratio of flotor to stator acceleration) is:

$$T(s) = \frac{K_{total} + c_{sky}s}{m s^2 + c_{sky}s + K_{total}}$$

with natural frequency  $\omega_n = \sqrt{K_{total}/m}$  and damping ratio  $\zeta = c_{sky}/(2\sqrt{m \cdot K_{total}})$ .

### 5.2 Parameter Selection

For  $m = 26$  kg (fully loaded),  $K_{total} = 400$  N/m:  $\omega_n = 3.92$  rad/s = 0.624 Hz. Heavy overdamping ( $\zeta \gg 1$ ) eliminates resonance peaks. At  $c_{sky} = 1500$  N·s/m:  $\zeta = 7.36$ . Closed-loop poles:  $s_1 = -57.4$

rad/s,  $s_2 = -0.274$  rad/s. Both real and negative, monotonic roll-off with no resonance. Gain margin  $> 12$  dB (limited by anti-alias filter, not plant). Phase margin  $> 90^\circ$  (heavily overdamped).

### 5.3 $H_\infty$ Feed-Forward Layer

The stator accelerometers measure the incoming disturbance  $\ddot{x}_s$  directly. An  $H_\infty$ -synthesized feed-forward controller generates a cancellation force on the flotor proportional to the estimated disturbance, providing 20 dB of additional rejection above 1 Hz without amplifying flotor sensor noise (because it operates on stator sensors in an open-loop feed-forward path). The combined skyhook + H transmissibility is:

$$|T_{combined}| \approx |T_{skyhook}|/G_{H_\infty}$$

where  $G_{H_\infty} \approx 10$  (20 dB) above the  $H_\infty$  crossover frequency ( $\sim 1$  Hz). Below 1 Hz, the  $H_\infty$  contribution rolls off and the skyhook operates alone.

The 20 dB  $H_\infty$  contribution is a design target contingent on controller synthesis; the exact achievable rejection depends on the stator-to-flotor disturbance path coherence and the sensor noise floor. This is validated in Phase 3 (Simulink) and identified as Risk R8.

### 5.4 Complementary Filter and Isolation Bandwidth

The skyhook controller requires flotor velocity, obtained by integrating the flotor accelerometer signal. This integration amplifies sensor noise as  $1/f$ , diverging at DC. A first-order complementary filter blends the accelerometer (high-pass, above  $f_c = 0.05$  Hz) with the gap sensors (low-pass, below  $f_c$ ):

Below  $f_c$ , the gap sensors provide position control (flotor tracks stator, no isolation). Above  $f_c$ , the accelerometers provide inertial velocity feedback (full isolation). The system provides active isolation from 0.05–100 Hz, which is better than ARIS (0.1 Hz lower bound) and far superior to any passive system. Below 0.05 Hz, the passive mechanical decoupling of the air gap provides partial isolation without active control.

### 5.5 6-DOF Jacobian

The 24 coils (4 per face  $\times$  6 faces) produce a wrench vector  $\mathbf{w} = [F_x, F_y, F_z, \tau_x, \tau_y, \tau_z]^T$  related to the current vector  $\mathbf{i}$  by the Jacobian  $\mathbf{J}$ :

$$\mathbf{w} = \mathbf{J} \cdot \mathbf{i}, \quad \mathbf{J} \in \mathbb{R}^{6 \times 24}$$

Each column of  $\mathbf{J}$  encodes the force and torque produced by one coil at unit current, computed from the Lorentz force model (Section 3.2) and the coil's position and orientation. The controller computes the minimum-norm current allocation via the weighted pseudoinverse:

$$\mathbf{i} = \mathbf{J}^T(\mathbf{J}\mathbf{J}^T)^{-1}\mathbf{w}$$

Gain scheduling corrects for displacement-dependent coil-magnet overlap variation (17.5% at  $\pm 3.5$  mm, a single multiply per axis per control cycle). In 5-face mode (door open), the Jacobian reduces to  $6 \times 20$ ; full 6-DOF control is maintained with reduced Z-axis and Z-torque authority.

---

## 6. Sensor Noise Floor and Performance Tiers

### 6.1 Why Sensor Noise Dominates: The Integration Problem

The skyhook control law commands a force proportional to the flotor's absolute velocity:  $F = -c_{sky} \cdot \dot{x}_f$ . But accelerometers measure acceleration, not velocity. This means that the controller must integrate the accelerometer output to estimate velocity:

$$\hat{x}_f(t) = \int \ddot{x}_f(t), dt$$

This integration is the origin of the fundamental noise problem. If the accelerometer reports a noise signal  $n(t)$  superimposed on the true acceleration, the controller cannot distinguish noise from signal. It integrates the noise alongside the real acceleration and commands a spurious force:

$$F_{noise}(t) = c_{sky} \int n(t), dt$$

In the frequency domain, integration multiplies by  $1/(2\pi f)$ , effectively amplifying low-frequency noise and suppressing high-frequency noise. The resulting force noise spectral density on the flotor is:

$$S_F^{noise}(f) = \frac{c_{sky}^2 \cdot S_n(f)}{(2\pi f)^2}$$

where  $S_n(f)$  is the accelerometer's noise power spectral density (in  $(\text{m/s}^2)^2/\text{Hz}$ ). Converting to acceleration noise on the flotor ( $a_{noise} = F_{noise}/m$ ):

$$S_a^{noise}(f) = \frac{c_{sky}^2 \cdot S_n(f)}{m^2 \cdot (2\pi f)^2}$$

If the sensor noise were white instead ( $S_n(f) = n_0^2 = \text{constant}$ ), the injected noise PSD would scale as  $n_0^2/f^2$ , which already diverges at low frequencies. This is a fundamental property of *any* accelerometer-based velocity-feedback controller, not specific to this design.

### 6.2 The Complementary Filter: Bounding the Divergence

Without mitigation, integrating  $S_a^{noise} \propto 1/f^2$  from DC to 100 Hz produces infinite RMS noise, rendering the controller unusable. The complementary filter (Section 5.4) solves this by removing the accelerometer from the loop below the crossover frequency  $f_c = 0.05$  Hz. Below  $f_c$ , the gap sensors provide position information directly (no integration needed, no  $1/f$  amplification). Above  $f_c$ , the accelerometer drives the skyhook.

The high-pass filter contribution to the injected noise PSD is:

$$|H_{HP}(f)|^2 = \frac{f^2}{f^2 + f_c^2}$$

Multiplying into the noise PSD:

$$S_a^{inj}(f) = \frac{c_{sky}^2 \cdot S_n(f)}{m^2 \cdot (2\pi f)^2} \cdot \frac{f^2}{f^2 + f_c^2} = \frac{c_{sky}^2 \cdot S_n(f)}{m^2 \cdot (2\pi)^2 \cdot (f^2 + f_c^2)}$$

The  $f^2$  from the high-pass filter *cancels* the  $1/f^2$  from the integration, producing a PSD that is bounded at DC:  $S_a^{inj}(0) = c_{sky}^2 \cdot S_n(0)/(m^2 \cdot (2\pi)^2 \cdot f_c^2)$ . The integral converges, and the RMS residual is finite.

For white sensor noise ( $S_n(f) = n_0^2$ ), the RMS integral evaluates analytically:

$$a_{white,RMS} = \frac{c_{sky} \cdot n_0}{m} \cdot \frac{1}{2\sqrt{2\pi}f_c}$$

This is only true *if* the noise is truly white down to DC.

### 6.3 The 1/f Correction: When Sensor Noise Isn't White

All MEMS and force-balance accelerometers exhibit flicker (1/f) noise at low frequencies. The true noise PSD is:

$$S_n(f) = n_0^2 \left( 1 + \frac{f_{1/f}}{f} \right)$$

where  $f_{1/f}$  is the sensor's 1/f corner frequency a.k.a the frequency below which flicker noise exceeds the white noise floor. This is important here for a specific reason: the 1/f sensor noise rises in exactly the frequency range where the controller is most sensitive. Substituting into the injected noise PSD:

$$S_a^{inj}(f) = \frac{c_{sky}^2 \cdot n_0^2}{m^2 \cdot (2\pi)^2} \cdot \frac{1}{f^2 + f_c^2} \cdot \left( 1 + \frac{f_{1/f}}{f} \right)$$

The first factor  $1/(f^2 + f_c^2)$  is bounded, so the complementary filter does its job. But the second factor  $(1 + f_{1/f}/f)$  adds a  $1/f$  divergence at low frequencies that the complementary filter does not fully suppress.

The integral still converges (because the  $1/f$  from the sensor times the bounded  $1/(f^2 + f_c^2)$  from the filter is integrable), but it is larger than the white-noise estimate.

Separating the integral into white and flicker terms and evaluating (with a practical low-frequency cutoff at  $f_{min} = 1/T_{obs}$  for observation time  $T_{obs}$ ):

$$a_{sensor,RMS} = a_{white,RMS} \times \sqrt{\Gamma}$$

where the correction factor is:

$$\Gamma = 1 + \frac{2f_{1/f}}{\pi f_c} \ln(f_c \cdot T_{obs})$$

The logarithmic dependence on  $T_{obs}$  is weak. Changing  $T_{obs}$  by 10× changes  $\Gamma$  by only ~30%. At  $f_c = 0.05$  Hz and  $T_{obs} = 100$ s:

Sensor	$f_{1/f}$ (Hz)	$\Gamma$	$\sqrt{\Gamma}$	RMS Penalty
Sercel DSU1-508	0.005	1.10	1.05	+5%
SF3600A (optimistic)	0.03	1.61	1.27	+27%
SF3600A (pessimistic)	0.10	3.05	1.75	+75%
Open-loop MEMS	0.20	5.10	2.26	+126%

The 1/f correction matters only for the Standard tier (MEMS sensors with  $f_{1/f}$  near  $f_c$ ). For the Advanced tier and above,  $f_{1/f} \ll f_c$  and  $\Gamma \approx 1$ . The sensor upgrade path inherently solves the 1/f problem. Higher-grade sensors have lower 1/f corners, which is precisely why they are higher-grade.

The SF3600A achieves the 3  $\mu$ g target with  $H_\infty$  feed-forward at the optimistic 1/f corner, but degrades to ~4  $\mu$ g at the pessimistic corner. This does not invalidate the sensor choice, but it means the 3  $\mu$ g headline number carries  $\pm 30\%$  uncertainty that is resolved by measuring  $f_{1/f}$  in Phase 2 bench testing via Allan deviation over  $\geq 1000$  s integration time.

## 6.4 Optimal Skyhook Gain: The Fundamental Tradeoff

The skyhook gain  $c_{sky}$  controls a direct tradeoff between two noise sources, transmitted vibration and sensor noise injection. Transmitted vibration decreases with  $c_{sky}$ . The skyhook force  $F = -c_{sky}\dot{x}$  acts as an ideal viscous damper anchored to inertial space. A stronger damper (higher  $c_{sky}$ ) resists flotor motion more aggressively, reducing the transmitted vibration. From Section 5.1, the transmissibility for frequencies well above the isolation corner ( $f \gg \omega_n/2\pi$ ) is:

$$|T(f)| \approx \frac{K_{total}}{c_{sky} \cdot 2\pi f}$$

The transmitted acceleration scales proportional to  $1/c_{sky}$ , so double the gain, halve the transmitted vibration. For a representative station input of  $a_{station} \approx 300$   $\mu$ g RMS concentrated near 1–10 Hz:

$$a_{transmitted} \approx \frac{a_{station} \cdot K_{total}}{c_{sky} \cdot 2\pi f_{dominant}}$$

But sensor noise injection increases with  $c_{sky}$ . From Section 6.1, the spurious force is  $F_{noise} = c_{sky} \int n, dt$ . Higher  $c_{sky}$  amplifies the noise proportionally. The injected RMS (from Section 6.2) is:

$$a_{sensor} = \frac{c_{sky} \cdot n_0}{m \cdot 2\sqrt{2\pi} f_c}$$

This scales linearly with  $c_{sky}$ , so double the gain, double the noise injection.

The total residual is the RSS of these two independent contributions:

$$a_{total}^2 = a_{transmitted}^2 + a_{sensor}^2 = \left(\frac{A}{c_{sky}}\right)^2 + (B \cdot c_{sky})^2$$

where  $A$  and  $B$  collect the constants from each term. This is a sum of a decreasing hyperbola and an increasing line, meaning it has a unique minimum. Differentiating and setting to zero:

$$\frac{d}{dc_{sky}} \left( \frac{A^2}{c_{sky}^2} + B^2 c_{sky}^2 \right) = 0 \quad \Rightarrow \quad c_{sky,opt} = \sqrt{A/B}$$

At the optimum, the two contributions are exactly equal:  $a_{transmitted} = a_{sensor}$ . Basically, the best you can do is the point where you are equally limited by the vibration you haven't rejected and the noise you've injected. You cannot improve one without worsening the other.

Substituting the expressions for  $A$  and  $B$  with the system parameters ( $m = 26$  kg,  $K_{total} = 400$  N/m,  $f_c = 0.05$  Hz,  $f_{dominant} \approx 2$  Hz,  $a_{station} = 300$   $\mu$ g):

$$c_{sky,opt} = \frac{m \cdot 2\sqrt{2\pi} f_c \cdot a_{station} \cdot K_{total}}{n_0 \cdot 2\pi f_{dominant}} \cdot \frac{1}{\sqrt{\text{bandwidth}}}$$

This evaluates to  $c_{sky,opt} \approx 332/\sqrt{n_0}$  (with  $n_0$  in  $\mu$ g/ $\sqrt{\text{Hz}}$ ,  $c_{sky}$  in N·s/m). The minimum achievable residual at the optimum:

$$a_{min} = 2\sqrt{A \cdot B} \propto \sqrt{n_0}$$

Observe that the minimum residual scales as the square root of the sensor noise density. This is why the sensor is the bottleneck: a 10 $\times$  better sensor (0.03 vs 0.3  $\mu$ g/ $\sqrt{\text{Hz}}$ ) yields only a 3.2 $\times$  better residual. It also means the performance tier roadmap (Section 6.5) is fundamentally set by sensor selection.

With the 1/f correction, the noise term grows by  $\sqrt{\Gamma}$ , shifting the optimum:

$$c_{sky,opt}^{corrected} = c_{sky,opt}^{white} / \Gamma^{1/4}, \quad a_{min}^{corrected} = a_{min}^{white} \times \Gamma^{1/4}$$

This is accurate to within 15% when  $f_{1/f} < 0.3f_c$ . The SF3600A is near this boundary; the Sercel DSU1-508 satisfies it comfortably.

## 6.5 Performance Tier Roadmap

Tier	Target	Sensor	$n_0$ ( $\mu\text{g}/\sqrt{\text{Hz}}$ )	$f_{1/f}$ (Hz)	With $H_\infty$ ( $\mu\text{g}$ )	Notes
Standard	<3 $\mu\text{g}$	Colibrys SF3600A	0.3	0.03 est.	~3.0	Baseline product
Standard (pessimistic)	<12 $\mu\text{g}$	SF3600A	0.3	0.10	~4.0	Still meets commercial spec
Advanced	<1 $\mu\text{g}$	Sercel DSU1-508	0.05	0.005	~0.9	Sensor swap only
Ultra	<0.4 $\mu\text{g}$	Research MET	0.005	<0.001	~0.3	Sensor + electronics swap
Nested	<0.05 $\mu\text{g}$	MET + 2-stage	0.001	<0.001	~0.05	Second-stage platform

## 6.6 Actuator Force Noise

Lorentz actuator force noise from H-bridge driver current noise ( $S_i \approx 1 \mu\text{A}/\sqrt{\text{Hz}}$ ) produces ~0.09  $\mu\text{g}$  RMS integrated over 0.01–100 Hz across all 24 coils, which is negligible at the Standard and Advanced tiers, comparable to the sensor floor only at the Ultra tier. Precision current drivers ( $< 0.1 \mu\text{A}/\sqrt{\text{Hz}}$ ) reduce this to ~0.009  $\mu\text{g}$  for the upper tiers.

---

## 7. Sensing, Power, and the Umbilical Cord

### 7.1 Sensor Suite

The control law requires three measurement categories: flotor acceleration (skyhook feedback), stator acceleration ( $H\infty$  feed-forward), and relative gap displacement (DC position hold + complementary filter low-frequency component).

Two Colibrays SF3600A triaxial accelerometers ( $n_0 = 0.3 \mu\text{g}/\sqrt{\text{Hz}}$ , ~\$500 each), are mounted at diagonal corners of the flotor, providing 6-DOF translation and angular acceleration via differential measurement across a 678mm baseline. Two Colibrays SF1500S units ( $n_0 = 1.0 \mu\text{g}/\sqrt{\text{Hz}}$ , ~\$300 each) serve the stator feed-forward path at a relaxed noise specification (feed-forward does not amplify noise through integration). Total accelerometer cost: ~\$1,600.

Ten PCB-integrated capacitive gap sensors (10 mm circular electrodes with guard rings, read by AD7746 24-bit CDCs) provide 6-DOF position reconstruction with  $< 0.1 \mu\text{m}$  resolution, which is three orders of magnitude better than the  $1 \mu\text{m}$  requirement. Capacitance at 4 mm nominal gap: 0.174 pF; dynamic range over  $\pm 3.5 \text{ mm}$ : 15:1. Four NTC thermistors on representative Halbach segments provide thermal correction for  $B_r$  variation ( $\alpha_{B_r} = -0.12\%/^{\circ}\text{C}$  for N48SH).

### 7.2 Power Architecture

The system operates from the Haven-1 28.5 VDC supply. Buck converters provide 5 V (FPGA, sensors, logic) and 12 V (H-bridge coil drivers). Total nominal power: ~2 W ( $< 2\%$  of the 100 W MLE allocation). Peak power during full-displacement transients:  $< 7 \text{ W}$ . A SmartFusion2 FPGA runs the 20 kHz control loop, 6-DOF Jacobian computation, complementary filter, and all mode-switching logic.

### 7.3 Hybrid FPC Umbilical

A single 4-layer FPC ribbon (20 mm wide  $\times$  0.22 mm thick  $\times$  350 mm long) carries all power and data across the gap: two 70  $\mu\text{m}$  copper power traces (3 mm wide, 1.8 A each), a 35  $\mu\text{m}$  ground/shield plane, two Ethernet differential pairs plus two RS-422 pairs for sensor data, and two power return traces. Mass: ~5 g.

The FPC routes as an  $\Omega$ -loop (80 mm radius) on the stator back face, providing 91 mm of geometric slack. This slack is 26 $\times$  the maximum flotor displacement. Bending stiffness analysis bounds the parasitic stiffness between the monolithic limit ( $K_{mono} = 3.27 \text{ N/m}$ , layers rigidly bonded) and the free-slip limit ( $K_{slip} = 4.6 \times 10^{-4} \text{ N/m}$ , layers independently bending). The true value lies somewhere in this 7,000 $\times$  range depending on adhesive shear compliance, and as shown below, the exact value does not affect isolation performance. The FPC creates a mechanical path that bypasses the controller. At any frequency, the bypass transmissibility is:

$$T_{bypass}(f) = \frac{K_{umb}}{m(2\pi f)^2}$$

The controller transmissibility at the same frequency (well above the isolation corner) is:

$$T_{controller}(f) \approx \frac{K_{total}}{m(2\pi f)^2}$$

Both roll off as  $\frac{1}{f^2}$  with the same slope. The bypass margin is constant across all frequencies:

$$Margin = 20 \log_{10} \left( \frac{K_{total}}{K_{umb}} \right)$$

Even at the monolithic worst case ( $K_{umb} = 3.27$ ,  $K_{total} = 400$  N/m): margin = 42 dB. The bypass changes total transmissibility by 0.07 dB, which is practically unmeasurable. At 1 Hz, the bypass path is 22 dB below the controller floor; at 10 Hz, 39 dB below; at 100 Hz, 42 dB below. Running the full closed-loop transmissibility with  $K_{umb}$  included in the plant model across the entire range from free-slip to monolithic produces curves that are visually indistinguishable. A 7,000× uncertainty in FPC stiffness produces <0.1 dB variation in delivered isolation.

As for fatigue, the outer-fiber bending strain at maximum displacement is:

$\epsilon = t/(2R) = 0.22/(2 \times 80) = 0.14\%$ . Annealed copper fatigue endurance at this strain exceeds  $10^9$  cycles (~30 years at 1 Hz dominant frequency). This is independent of the stiffness estimate.

An alternative fallback design would be to have a discrete 40 AWG wire bundle (8 wires, service loop at 80 mm radius,  $K_{umb} \approx 0.007$  N/m) be a drop-in replacement if FPC fatigue testing in Phase 2 reveals trace degradation. The wire bundle has 44 dB of bypass margin against  $K_{total}$ , which should be equally negligible.

## 8. Interfaces and Haven-1 Compliance

All compliance is verified against the Haven-1 User's Guide (HUG v0.7, February 2025) and ISS heritage documents (NSTS-21000-IDD-MDK Rev B, SSP-52000-IDD-ERP). I have a separate document for this, available upon request.

## 9. Validation Roadmap and Risk Register

### 9.1 Current Status

Current status is TRL 2, as this is just a concept, although I believe I have established here the analytical foundation of first-principles electromagnetic, control, and structural models. No hardware has been fabricated or tested as of yet. My target is to achieve TRL 5 within 10–14 months at \$105K–322K, ideally in time to achieve TRL 7 onboard the Purdue-1 suborbital launch and the subsequent Haven-1 launch in late 2027.

### 9.2 Validation Phases

Phase 1: COMSOL Magnetostatic Simulation (2–4 weeks, \$0–2K). 3D FEA of a single Halbach unit at 4 mm gap resolves  $B_{eff}$  to  $\pm 5\%$  and quantifies  $K_{neg}$  per axis with non-permeable stator materials, confirming < 100 N/m. Success criteria:  $B_{eff}$  within sensitivity range;  $k_f$  validated within  $\pm 10\%$ .

Phase 2: Single-Axis Bench Demonstrator (6–10 weeks, \$5–15K). One Halbach unit on a linear air bearing with one stator coil, gap sensor, and accelerometer running a 1-DOF skyhook law. Measures force constant, closed-loop transmissibility (0.1–100 Hz), and sensor noise floor. Additional success criterion: SF3600A 1/f corner frequency characterized via Allan deviation over  $\geq 1000$  s integration time (required  $f_{1/f} < 0.1$  Hz; target  $< 0.03$  Hz).

Phase 3: 6-DOF Simulink Model (8–12 weeks, \$0–5K, overlaps Phase 2). Full multi-body simulation with 24-channel Lorentz model, Jacobian,  $K_{neg} = -100$  N/m per axis, sensor noise PSD  $n_0^2(1 + f_{1/f}/f)$  parameterized from 0.01–0.10 Hz, skyhook +  $H\infty$  dual-loop controller, and representative Haven-1 disturbance spectra. Success criteria: transmissibility within  $\pm 3$  dB of analytical model across 0.05–100 Hz; RMS residual  $< 3$   $\mu\text{g}$  at  $f_{1/f} = 0.03$  Hz and  $< 4$   $\mu\text{g}$  at  $f_{1/f} = 0.10$  Hz; stable recovery from 1 N·s crew impulse within 2 seconds; 5-face mode maintains stable 6-DOF control.

Phase 4: Multi-Axis Prototype (16–24 weeks, \$50–150K). Full 6-DOF Q-Locker at single-MLE scale: stator frame, flotor frame, 24 Halbach units, 24 coils, 8 launch locks, full electronics suite, FPC umbilical, complete firmware. Ground testing includes individual axis force constant measurement, gap sensor calibration, launch vibration qualification (7  $g_{RMS}$ ), and EMC characterization.

Phase 5: Reduced-Gravity Demonstration (4–8 weeks, \$50–150K). Primary opportunity: Purdue 1 Mission (Virgin Galactic, 2027), providing 4 minutes of microgravity in an MLE-compatible payload locker with no adapter hardware required. Test sequence: lock release  $\rightarrow$  levitation acquisition ( $< 2$  s)  $\rightarrow$  steady-state isolation measurement  $\rightarrow$  impulse response  $\rightarrow$  5-face mode switch  $\rightarrow$  safe mode reengagement. Success criteria: measured transmissibility within  $\pm 6$  dB of Simulink predictions; flotor remains within  $\pm 3.0$  mm; no frame contact.

### 9.3 Risk Register

R1:  $B_{eff}$  Lower Than Working Estimate. If COMSOL reveals  $B_{eff} < 0.15 T$ , force constant drops below useful levels. Severity: High. Likelihood: Low. Mitigation: architecture tolerates  $B_{eff}$  as low as 0.15 T without redesign; below that, magnet thickness increase from 5 to 7 mm boosts  $B$  by  $\sim 30\%$ . Residual risk: Low.

R2: Station Deformation Exceeds Gap Budget. Local Haven-1 deformation during thermal cycling exceeds  $\pm 3.5$  mm. Severity: High. Likelihood: Low-Moderate. Mitigation: pursue access to Vast structural FEA; if unavailable, increase gap from 4 to 5 mm or larger if necessary (15%  $B_{eff}$  reduction, within sensitivity range). Residual risk: Low.

R3: Sensor Noise Higher Than Specified. In-situ SF3600A noise floor exceeds  $0.3 \mu\text{g}/\sqrt{\text{Hz}}$  due to EMI or integration effects. Severity: Moderate (performance degrades as  $\sqrt{n_0}$ ; at  $1.0 \mu\text{g}/\sqrt{\text{Hz}}$ , Standard tier achieves  $\sim 5.6 \mu\text{g}$  — degraded but within  $< 12 \mu\text{g}$  spec). Likelihood: Moderate. Mitigation: sensor-swap upgrade path to Sercel DSU1-508; careful PCB layout with isolated analog grounds; Phase 2 bench characterization. Residual risk: Low.

R3A: Sensor 1/f Corner Above 0.1 Hz. SF3600A exhibits  $f_{1/f} > 0.10$  Hz in situ, pushing Standard tier above 4  $\mu\text{g}$ . Severity: Moderate (still within  $< 12 \mu\text{g}$  spec). Likelihood: Low-Moderate. Mitigation: sensor swap to Sercel DSU1-508 ( $f_{1/f} < 0.01$  Hz); increase complementary filter  $f_c$  to 0.1 Hz (firmware change); second-order complementary filter (firmware change). Residual risk: Low.

R4: FPC Fatigue. Copper trace cracking or adhesive delamination under cyclic bending. Severity: Moderate (system inoperable without umbilical). Likelihood: Low (0.14% strain is well below copper endurance limit). Mitigation: discrete 40 AWG wire bundle fallback ( $K_{umb} \approx 0.007$  N/m); Phase 2 bench fatigue testing to  $10^7$  cycles before Phase 4 commit. Residual risk: Low.

R5: Launch Lock Failure to Release. One or more solenoid pins fail to retract on orbit. Severity: High (system cannot operate). Likelihood: Low (spring-return mechanism is fail-safe for engagement, but solenoid energization is required for release). Mitigation: redundant release mechanism (dual solenoid or pyrotechnic backup); ground reliability testing to  $> 100$  cycles; manual crew override via hex bolt access. Residual risk: Low.

R6: Haven-1 Payload Rack Geometry Incompatible with Double MLE. The double-height MLE envelope differs from the HUG v0.7 specification. Severity: Moderate (limits single-MLE only). Likelihood: Low. Mitigation: single-MLE design is fully self-contained and does not depend on double-MLE availability. Residual risk: Low.

R7: Controller Computational Delay. FPGA processing exceeds  $50 \mu\text{s}$  per control cycle. Severity: Moderate (reduces phase margin). Likelihood: Low (SmartFusion2 at 100 MHz handles the 6-DOF Jacobian pseudoinverse and complementary filter in  $< 20 \mu\text{s}$ ). Mitigation: algorithmic simplification; FPGA upgrade. Residual risk: Low.

R8:  $H\infty$  Feed-Forward Below 20 dB. The  $H\infty$  controller, when synthesized and implemented, achieves less than the assumed 20 dB of feed-forward rejection. Severity: High (Standard tier degrades from  $\sim 3 \mu\text{g}$  to  $\sim 30 \mu\text{g}$  without  $H\infty$ ). Likelihood: Moderate. Mitigation: the skyhook-only performance is characterized first (Phase 2); if  $H\infty$  underperforms, the sensor-swap upgrade to Sercel DSU1-508 recovers  $< 3 \mu\text{g}$  without any feed-forward ( $a_{min} = 9.2 \mu\text{g}$  skyhook-only at  $n_0 = 0.05 \mu\text{g}/\sqrt{\text{Hz}}$ ). Residual risk: Moderate.

R9: Residual  $K_{neg}$  Exceeds 200 N/m. COMSOL or bench testing reveals eddy-current-induced stiffness beyond the static estimate. Severity: Low (controller can stabilize up to  $\sim 790$  N/m at maximum displacement within coil current limits). Mitigation: increase  $K_p$  proportionally; no hardware change required. Residual risk: Low.

---

## 10. Conclusion and Next Steps

The Q-Locker attempts to address a gap in the commercial space station ecosystem: no locker-level active vibration isolation product exists. The two-layer architecture (non-permeable stator materials reducing  $K_{neg}$  by 99.7%, plus a 20 kHz Lorentz actuator controller stabilizing the 0.31 Hz residual) is a system that is mechanically simple, commercially manufacturable, and provides isolation performance ( $< 3 \mu\text{g}$  at Standard tier,  $< 100 \text{ ng}$  at upper tiers) competitive with or exceeding prior government-funded systems at 1/100th the development cost.

The most consequential technical risk is the  $H\infty$  feed-forward performance (R8), which determines whether the  $< 3 \mu\text{g}$  target is achievable with the baseline MEMS sensor (SF3600A) or requires the more expensive force-balance accelerometer (Sercel DSU1-508). This risk is resolved in Phase 3 simulation and Phase 2 bench testing, well before the Phase 4 prototype commitment.

---

*This document summarizes the complete Q-Locker technical proposal. Supporting derivations, component specifications, detailed compliance matrices, and extended sensitivity analyses are available upon request.*