


Cornell Hyperloop

Cornell University

May 31, 2026



Engineering Design

0. Project Overview & Design Objectives

The overarching objective for the 2025–2026 design cycle is to transition the hyperloop pod architecture away from heavy, rigid, off-the-shelf industrial hardware toward a highly optimized, custom-engineered, and modular ecosystem. Subsystem development prioritizes maximizing performance efficiency and operational reliability while aggressively driving down total mass, footprint, and manufacturing complexity.

Key design initiatives across the core subsystems include:

- **Structural Mass Optimization:** Minimizing framework deadweight to enhance acceleration dynamics, achieved through a transition from a 65 lb welded steel frame to a modular aluminum geometry, yielding a 62% mass reduction and 32% length reduction.
- **Custom Electromagnetic Integration:** Engineering tailored propulsion and levitation solutions to eliminate off-the-shelf constraints, utilizing a parametrically optimized Linear Induction Motor (LIM) alongside custom cylindrical electromagnets to establish precise field uniformity.
- **Fail-Safe Deceleration Redundancy:** Implementing a multi-layered braking strategy consisting of a switchable magnetic eddy-current braking subsystem paired with a lightweight, dual-sided pneumatic friction clamping assembly.
- **EMI Management & Signal Integrity:** Isolating sensitive control hardware from high-power emitting electronics via a strict vertical layer separation strategy within an expanded chassis footprint.
- **Modular Command Infrastructure:** Deploying a distributed architecture of edge hubs and standalone firmware structures to support real-time telemetry streaming, low-latency control loops, and secure remote operations.

1. Magnetic Braking Subsystem

1.1 Principles of Function

The magnetic braking system operates based on Lenz's Law; a changing magnetic flux through a conducting surface induces an eddy current that opposes the change. The resulting braking force is expressed mathematically as: $F = \frac{vB^2At}{\rho}$. The system is "switchable," meaning the magnetic field can be toggled between 'on' and 'off' states. While the magnetic field magnitude remains constant, the motion of the magnets creates a spatially changing field that induces the opposing eddy currents in the track.

1.2 System Specifications

The switching mechanism relies on the alignment of magnetic polarities to either contain the

field within the brake or propagate it toward the track. **Mechanical Configuration:** The assembly consists of a base and two rotating disks containing four permanent magnets arranged in a circle. When "off", poles are alternating and the field propagates internally. When "on", the disk is rotated 90 degrees, polarities are antisymmetric, and the field propagates outward toward the track.

Power and Weight Logistics: Disks are actuated by two two-phase stepper motors. Due to current limits, each motor utilizes its own driver, powered by an e-bike battery. This equipment combines to weigh ~2.7 kg.

2. Frictional Braking System

2.1 Overview

The objective of this project was to redesign the Hyperloop pod's pneumatic friction braking system to reduce complexity, weight, and manufacturing cost while maintaining reliable braking performance. The completed design successfully meets these objectives by providing a lightweight, manufacturable, and cost-effective system capable of dual-sided friction braking while maintaining a simple pneumatic architecture.

2.2 Principles of Function & Operation

The final design consists of a C-bracket assembly, linear guide rails, a pneumatic actuator, external return springs, a brake pad backing plate, and brake pads. The linear guide rails and C-bracket serve as the key components that enable dual-pad braking.

Clamping Mechanism: The actuator is mounted rigidly to the upper plate of the C-bracket, with the actuation shaft extending through the top plate. As compressed air enters the pneumatic cylinder, the actuator rod extends downward, pressing the upper brake pad against the top surface of the track's I-beam flange. Once sufficient force is generated, the resulting reaction force lifts the entire C-bracket assembly upward along the linear guide rails. This movement causes the lower brake pad to engage the underside of the flange, creating a dual-sided clamping action that increases braking force and effectiveness.

2.3 Mechanical Design & Component Selection

- **Actuator Configuration:** To resolve a sourcing issue, a double-acting pneumatic cylinder with integrated mounting features was selected and paired with external extension springs.
- **Spring Mechanics:** These external extension springs ($k = 17 \frac{lb}{in}$) generate enough force to overcome the weight of the actuator and C-bracket assembly. This returns the system to its disengaged position when air pressure is removed. This approach allows the brake to function similarly to a single-acting system while avoiding additional pneumatic complexity.
- **Manufacturability:** To minimize manufacturing costs and simplify fabrication, the C-bracket was designed as an assembly of three distinct plates connected using gusset brackets, completely avoiding the need for a single, complex bent or machined component.

3. Aeroshell Subsystem

3.1 Introduction

A scaled mini-pod (23.6% length, 31.2% width) evaluated fabrication failure modes in a controlled, low-cost setting. Adopting a single-piece convex mold reduced assembly error, improved vacuum integrity, and smoothed the outer surface. Fabrication complexity remained concentrated at the nose. To mitigate this, a simplified geometry with a sheared trailing edge was selected; later analysis shows this significantly influences flow separation and wake formation, contrary to initial assumptions.


3.2 Factor of Safety & Structural Analysis

The aeroshell was modeled as a thin-walled pressure vessel under internal pressurization ($P_{int} = 101,325 \text{ Pa}$, $t = 0.01 \text{ m}$).

Geometric and structural properties:

- $r_o = 0.528$, $r_i = 0.508$, $L = 1.867$
- *Hoop Stress*: $\sigma_h = \frac{Pr}{t} = 2.6 \text{ MPa}$; *Longitudinal Stress*: $\sigma_l = \frac{Pr}{2t} = 1.3 \text{ MPa}$
- $W = \frac{\rho g \pi (r_o^2 - r_i^2)}{4} = 279 \frac{\text{N}}{\text{m}}$; $M = WL_m = 121 \text{ Nm}$
- $Z = \frac{\pi (r_o^4 - r_i^4)}{32 r_o} = 0.0021$; *Bending Stress*: $\sigma_b = \frac{M}{Z} = 0.06 \text{ MPa}$

Hoop stress dominates the factor of safety. As $\Delta P \rightarrow 0$, structural loading transitions to bending-dominated behavior. Internal pressurization provides radial stiffening but increases susceptibility to sagging under self-weight



Figure

1: ANSYS deformation (pinned-pinned), $P = 101325 \text{ Pa}$, $E \approx 230 \text{ GPa}$.

3.3 Flow Physics & Aerodynamic Simulation

Upstream of the nose, flow is well approximated by inviscid potential flow over a cylindrical body. Viscous effects become important along the surface, leading to boundary layer growth and separation. Wall shear stress approaches zero at the separation point, where boundary layer approximations break down under an adverse pressure gradient. Beyond separation, the flow enters a recirculation wake region. No closed-form solution exists due to transverse velocity components.

Simulation Parameters: Air approximated residual aerodynamic effects in a low-pressure environment. Wrought aluminum approximated surface roughness $O(10^{-6} m)$ to capture viscous losses. Inlet velocity was set to $480 \frac{m}{s}$. Results show clear flow separation along the trailing edge due to the abrupt geometric transition, highlighting the aerodynamic penalty of simplified geometry.

Mitigation strategies: Include gradual aft curvature or boundary layer control, which increase manufacturing complexity.

Trade-off: Efficiency vs. manufacturability. Manufacturing was prioritized, accepting modest aerodynamic drag penalties for reliability and lower production costs.




Figure 2: ANSYS streamlines and pressure contour

4. Chassis Subsystem

4.1 Structural Geometry & Mass Architecture

The SQ5 chassis is a two-layer rectangular prism measuring 6.1 ft long, 16 in wide, and 10 in high, weighing ~25 lbs. This is a 62% mass reduction and 32% length reduction compared to the previous 65 lb, 9 ft welded steel frame.

Chassis Metric	Previous Iteration (Steel Frame)	SQ5 Iteration (Modular Aluminum)	Performance Delta
Material	Welded Low-Carbon Steel	6063 Aluminum / PA6 + 30% GF	Significant material pivot
Mass / Length	65 lbs / 9.0 feet	25 lbs / 6.1 feet	62% Mass / 32% Length Reduction
Width	12 inches	16 inches	Expanded component packaging capacity
Assembly Tooling	External Welding Shop	Rubber Mallet (2 people, <30 mins)	Eliminated third-party dependencies

4.2 Material Selection & Fastening Mechanics

The structural framework utilizes 1"x1" square aluminum tube extrusions (6063 alloy) joined by press-fit connectors molded from polyamide 6 (PA6) reinforced with 30% glass fiber. These multi-way configurations provide high stiffness, creep resistance, and dimensional stability under static and dynamic loads. Assembly requires only a rubber mallet and can be completed by two people in under 30 minutes, bypassing the external welding shop schedule dependencies of the

previous steel design.

4.3 Electromagnetic Interference (EMI) Management Strategy

The two-layer geometry enables a vertical separation strategy for EMI management, entirely absent in the previous single-plane chassis. Emitting subsystems are physically separated from sensitive subsystems to eliminate risk of electromagnetic coupling:

- **Lower Layer:** Houses EMI-producing subsystems (propulsion electronics, power distribution, motor controllers).
- **Upper Layer:** Houses EMI-sensitive components (sensors, control electronics, communication modules).

Widening the frame from 12 to 16 inches accommodated this vertical architecture while optimizing subsystem packaging efficiency and shrinking the overall pod footprint.

4.4 Reconfigurability and Shipping Logistics

Press-fit construction allows individual tube segments to be removed or replaced without affecting the surrounding structure. The frame can be fully disassembled for flat-pack transport and reassembled on-site, optimizing competition logistics. New frame configurations can be tested with zero material waste, adapting dynamically as internal spatial requirements mature throughout the development timeline.




Figure 3: Modular Chassis Design

5. Propulsion Subsystem

5.1 Custom LIM Design & Computational Framework

To overcome the limitations of off-the-shelf linear induction motors, a custom scaled mini-LIM was engineered, utilizing ANSYS Maxwell simulations for parameter optimization. The stator creates a traveling magnetic field that induces currents in the track, creating a Lorentz force that drives the pod. Stator dimensions were constrained to a length of ~50 cm and a width of 20 cm. Initial modeling explicitly represented each coil winding, resulting in simulation times exceeding five hours per run. To optimize iteration speed, coils were simplified as solid stranded conductor blocks, preserving electromagnetic fidelity while reducing runtime to 15 minutes to facilitate a fully parametric design framework.

5.2 Optimization Sweeps & Results

Parametric sweeps evaluated variables over a five-second transient period, tracking

position-time curves as the primary metric for overall thrust generation. Current amplitude was swept from 5 A to 40 A in 5 A increments, revealing maximum thrust at 40 A.

Optimized Mini-LIM Specifications: Stator core measures 47 cm long, 18 cm wide, and 2 cm high, containing 24 coils arranged in a three-phase configuration. Core teeth have heights of 3 cm and widths of 1 cm. Pole pitch is set to 6 cm, with each coil containing 200 turns of 25-gauge wire operating at 10 Hz. This setup generates ~120 N of thrust in transient simulations, providing a verified baseline for physical fabrication.

6. Levitation Subsystem

6.1 Core Mechanics and Winding Optimization

To eliminate contact resistance at high operating speeds, magnetic levitation is utilized. Previous small-scale mini-pods relied on four modified transformers for static levitation, which limited control over field uniformity and introduced high variability. To establish precise field uniformity, custom electromagnets with cylindrical cross-sections wound on high-purity iron cores were designed. Force generation scales quadratically with wire turns; however, excessive turns increase coil length and resistance, driving up thermal dissipation and power consumption. An optimized target of ~1,500 turns per electromagnet was established to balance force output against efficiency and controllability.

6.2 Coil Winding Automation Rig

Because manual winding is impractical, an automated coil winding machine was engineered. It utilizes a synchronized dual stepper motor system: one motor guides a wire feeder traversing the linear length of the core, while the second rotates the core itself to achieve precise, repeatable layer packing. Current upgrades will allow multiple winding pattern configurations to further optimize field intensity while maintaining safe, low coil resistance and stable thermal limits.

7. Power Systems

The Power Systems team developed an integrated architecture consisting of a custom battery pack, a battery management system (BMS), an emergency power-off (EPO) circuit, and a gate-driver-based DC-to-3-phase AC inverter. The goal for this cycle was to transition away from bulky off-the-shelf hardware toward compact, modular, and controllable custom electronics. KiCad was used for schematic capture and PCB layout, while LTspice and Tinkercad simulations validated circuit behavior prior to fabrication.

7.1 Battery Pack and Battery Management System (BMS)

The previous pod used a standard automotive battery to power the linear induction motor (LIM). To improve control over voltage, current capability, and pack geometry, the team began development of a custom battery pack composed of individual 3.7 V lithium-ion cells in series. A 12-cell test pack (3.3 V nominal per cell) was used to validate the Orion 2 BMS. The BMS provides real-time monitoring of cell voltage, temperature, and state of charge (SOC), and

enforces strict protection limits to prevent overcharging, over-discharging, and cell imbalance critical for lithium-ion safety.

The battery pack is designed to be modular. By adjusting the number of series cells, the output voltage can be tuned to match LIM and inverter requirements. To avoid the risks of spot-welding, the team plans to use Vruzend CelCapz battery kits, which improve safety, reconfigurability, and repairability.




Figure 4: Battery state of charge monitoring through the BMS software.

7.2 Emergency Power-Off (EPO) System

The previous EPO relied heavily on external wiring and multiple relays, making it difficult to integrate and prone to wiring faults. The redesigned EPO PCB consolidates functionality into a compact board featuring: flyback protection, a single latching relay, MOSFET-driven relay control, remote and physical shutdown inputs, and LIM thermistor routing through the motherboard. Remote shutdown is now handled by an N-channel MOSFET instead of a relay. A signal from the motherboard pulls the MOSFET gate low, opening the path between the +12 V rail, the LIM contactor, and ground. A momentary input triggers the latching relay, which then maintains power delivery until a shutdown signal is received. Thermal monitoring is integrated through an additional MOSFET that interrupts the remote-stop signal when an over-temperature condition is detected. Fuses and pull-down resistors provide additional protection and ensure predictable behavior under fault conditions.

7.3 Gate-Driver DC to 3-Phase AC Inverter

The original propulsion system used an AIMS Power DC-to-AC inverter paired with a 7.5 kW A2-series variable-frequency drive (VFD). Although functional, this configuration was heavy, bulky, and difficult to integrate into the pod. The new inverter design replaces these components with a custom PCB built around the Texas Instruments DRV8302 gate-driver IC. The system uses three pairs of N-channel MOSFETs to generate a 3-phase AC output from a DC supply. Each MOSFET pair is switched with a 120° phase offset, producing the rotating magnetic field required to drive the LIM.

A microcontroller sets the switching frequency by configuring the INH and INL pins on the gate driver, which in turn control the GH and GL outputs connected to the MOSFET gates. The prototype board currently implements a single phase for validation before scaling to the full

3-phase system.

Thermal management was a major design consideration. The PCB incorporates thermal vias around the gate driver, large ground planes on both layers, and high-temperature-rated materials to ensure reliable operation under high switching loads.




Figure 5: Phase 1 design of the gate driver inverter.

8. Electronic, Command, and Control (ECC) Subsystem

The system is built around a distributed architecture of Raspberry Pi hubs, an AWS-hosted API, and a browser-based Command & Control GUI. This platform enables real-time monitoring, high-frequency control loops, and robust remote operations for an Electromagnetic Levitation (EML) pod. A core design principle is modularity: firmware, cloud services, and the GUI evolve independently through standardized I/O contracts and opaque data payloads.

8.1 Telemetry and Command Infrastructure (GUI)

The Command & Control GUI serves as the operational interface for the pod, supporting low-latency data streaming, secure authentication, and remote hardware management. Telemetry originates on Arduino microcontrollers and is forwarded over USB serial to Raspberry Pi hubs. These hubs communicate with a Pi-based edge gateway, which provides a cellular uplink to an AWS API. The API fans out data to the GUI via WebSockets.

- **Opaque Payloads:** Serial data is preserved end-to-end, allowing firmware teams to modify sensor formats or control logic without requiring GUI or cloud schema updates.
- **Data Framing:** Newline-delimited framing with time-sliced fallback ensures data integrity over unstable cellular connections.

8.2 Secure Operations and OTA Management

Authentication: A Caddy reverse proxy integrates Cornell Single Sign-On (Shibboleth) and Duo MFA for secure access. **Over-the-Air Flashing:** The GUI interfaces with a build service that allows operators to flash firmware to any microcontroller remotely, using either raw code or validated scripts.

Session Health Monitoring: A state machine tracks up to 12 concurrent hardware sessions, displaying connectivity and fault states through diagnostic badges.

8.3 Motherboard and Levitation Control Hardware

The motherboard acts as the pod's central nervous system, coordinating high-I/O subsystems, power electronics, and levitation control. The Arduino Nano R4 provides a 32-bit control platform, but its native 22 digital pins are insufficient for the pod's requirements, which include stepper drivers, PWM fans, and magnetic brake actuation. To address this:

- **I/O Expansion:** MCP23017 I²C expanders and cascaded 74HC595 shift registers provide scalable digital I/O.
- **Communication Bridging:** A MAX3232 RS232-to-TTL converter supports industrial stepper drivers, while a TCA9548A multiplexer resolves PC address conflicts.

8.4 Motion Control and EML Implementation

- **PWM Control:** Electromagnets are driven using IRF520/IRLZ44N MOSFETs with flyback protection to handle inductive loads.
- **PID Stability:** A high-frequency PID loop maintains consistent preload and reduces oscillatory "wavy motion" during levitation.
- **Fault Detection:** Temperature and current thresholds are continuously monitored, triggering automatic shutdown under unsafe conditions.

8.5 Interfacing and Robust Hardware Integration

Molex Micro-Fit 3.0: Provides secure, vibration-resistant connections for all high-current and mission-critical interfaces. I²C Harnessing: An 8-pin harness includes dedicated interrupt lines, enabling sensors such as the MCP9808 to signal threshold violations immediately.

8.6 Power Management and Signal Integrity

The pod's 24 V supply is stepped down through a series of LM2596 buck converters to 12 V, 5 V, and 3.3 V rails.

- **Thermal Efficiency:** Cascaded conversion reduces heat generation across regulators.
- **Noise Mitigation:** Decoupling capacitors at each device node stabilize power delivery and reduce EMI from high-current magnetic brake lines.




Figure 6: The diagrams collectively map the system's software architecture, detailing the Pi Hub data ingestion pipeline, the cloud service orchestration layer, and the GUI's real-time command and telemetry framework.