

**MIT Hyperloop
Final Report**

MIT HYPERLOOP FINAL REPORT

*An overview of the design, build,
and testing process for MIT's
entry in the SpaceX Hyperloop
Competition 2015-2017*

August 2017

CLEARED FOR PUBLIC RELEASE

©MIT Hyperloop Team 2017. All rights reserved.

ABSTRACT

The MIT Hyperloop Team competed in the SpaceX Hyperloop competition from June 2015 to January 2017. The goal for this competition is to design and build a scaled Hyperloop pod to test in a 1-mile long test track in Hawthorne, California. In doing so, the teams develop technology that could one day be used in a full-scale Hyperloop system.

The MIT Hyperloop Team won the design component to this competition in January 2016, and went on to unveil the first scaled Hyperloop prototype in May 2016. At the final competition in January 2017, the team's pod demonstrated its stable magnetic levitation in the first-ever Hyperloop run in a vacuum environment, won the Safety & Reliability Award, and got third place in Design & Construction. This report details the road to these accomplishments by providing a detailed explanation of the design of the pod, the construction of the pod, and the tests that were performed – including an analysis of the competition run.

The pod uses an electrodynamic suspension maglev system for levitation, which operates at relatively large gap heights and which becomes more efficient at high speeds. The braking system uses the same physical principle to slow the pod down, only now optimized for maximum drag. The braking system is able to slow the pod down at upwards of $2.4 G$ and is designed such that it can be used as an emergency braking system on a full-scale Hyperloop pod, and is completely fail-safe. The lateral stability is safeguarded by two Lateral Control Modules at the front and rear of the pod, which again use permanent magnets to provide a stabilizing force. Most of the systems on the pod are passive, and only the brakes require active control to keep them open. The majority of the electronics on the pod are therefore used for monitoring, such that the performance of the system can be quickly improved after each run. Finally, all these systems are covered by a lightweight aerodynamic shell.

The levitation system is crucially important for this pod, and therefore a testing campaign was launched to validate the simulations used for the levitation design. A rotary and a linear test rig were constructed, and the results from both test rigs agreed with the simulation data. During the competition run, the levitation design was further validated, as the pod fully lifted off and was stable throughout the flight.

PREFACE

The MIT Hyperloop Team was founded more than two years ago to aid in developing the Hyperloop concept by means of competing in the inaugural SpaceX Hyperloop Competition. We have succeeded in that mission by winning the first design competition for a Hyperloop concept, unveiling the first Hyperloop prototype, and demonstrating the stable magnetic levitation of our pod in the first half-scale Hyperloop run at speed in a low-pressure environment. This report provides a high-level overview of our journey towards these accomplishments. Note that this report focuses on the design of a Hyperloop pod specifically for the SpaceX Hyperloop competition, not necessarily on a design for a full-scale Hyperloop system.


We would not have been able to complete this project without the tremendous support from a lot of people. First, we would like to thank our sponsors for supporting us to build the Hyperloop pod. All team sponsors are listed on page 8. The team was fortunate to have some incredible team advisors: Professor Douglas Hart, Professor John Hansman, Professor Maria Yang, Professor (astronaut candidate) Warren Hoburg, Dr. Bob Shin, and Noel Zamot. The team is indebted to you for sharing your time and advice so generously. The same goes for our industry advisors Tracy Clark from Magnemotion, Dr. Bruce Montgomery from Magplane, and Ben Parker from Whitecap Composites, who helped the team out tremendously with their expertise.

Within MIT, we would like to thank the students' advisors for allowing them to work on this project. In general, the support the team has gotten from students, faculty, the administration, and the alumni has been amazing. In particular, the MIT Edgerton Center – where the team found a home during these past two years – cannot be thanked enough for their support. Specifically, Sandra Lipnoski from the Edgerton Center was a lifesaver by handling all the paperwork and supporting the team, and Patrick McAtamney from the Edgerton Center made sure the team had access to a machine shop and garage space, and supported the team whenever he could.

THE MIT HYPERLOOP TEAM IS GRATEFUL FOR THE GENEROUS SUPPORT FROM ITS SPONSORS

hyperloop | one

DRAPER



LIST OF TEAM MEMBERS

TEAM LEADERSHIP

Max Opgenoord

Team Lead, *Year 2*
Aero/Structures Lead, *Year 1 & 2*

Philippe Kirschen

Team Captain, *Year 1*
Levitation, *Year 1*

Greg Monahan

Levitation Lead, *Year 1*

Shawn Zhang

Vehicle Dynamics Lead, *Year 2*
Vehicle Dynamics, *Year 1*

Chris Merian

Chief Engineer, *Year 1*

Colm O'Rourke

Electronics Lead, *Year 1 & 2*

Derek Paxson

Levitation, *Year 1*

Georgiana Vancea

Business Co-lead, *Year 2*
Business, *Year 1*

Chuan Zhang

Business Lead, *Year 1*

John Mayo

Project Manager, *Year 1*

Greg Izatt

Software Lead, *Year 1 & 2*

Charlie Wheeler

Vehicle Dynamics Lead, *Year 1*

Nargis Sakhibova

Business Co-lead, *Year 2*
Business, *Year 1*

TEAM MEMBERS

Raghav Aggarwal

Vehicle Dynamics, *Year 1*

Philip Caplan

Aero/Structures, *Year 1*

Stephanie Chen

Electronics, *Year 1*

Cory Frontin

Aero/Structures, *Year 2*

Lakshya Jain

Electronics, *Year 1 & 2*

Josh Nation

Aero/Structures, *Year 1*

Sarthak Vaish

Aero/Structures, *Year 1 & 2*

Nick Baladis

Business, *Year 1*

Peter Chamberlain

Vehicle Dynamics, *Year 1*

Rachel Dias Carlson

Aero/Structures, *Year 1*

Abe Gertler

Vehicle Dynamics, *Year 2*

Richard Li

Aero/Structures, *Year 1 & 2*

Aleksandr Rakitin

Business, *Year 1*

Scott Viteri

Software, *Year 1*

Lucy Zhang

Business, *Year 2*

Sabrina Ball

Vehicle Dynamics, *Year 1*

Josh Chen

Business, *Year 1 & 2*

Dan Dorsch

Aero/Structures, *Year 1 & 2*

Yiou He

Electronics, *Year 1 & 2*

Geronimo Mirano

Software, *Year 2*

Nick Schwartz

Levitation, *Year 1 & 2*

Evan Wilson

Software, *Year 1*

CONTENTS

Abstract 5

Preface 7

1 Introduction 15

1.1 Background 16

1.2 SpaceX Hyperloop Competition 16

1.3 About the Team 16

Design & Construction

2 Top-Level Overview 19

2.1 Conceptual Design 19

2.2 Final Design 21

3 Levitation 23

3.1 Feasibility of Air Bearings and Compressor 23

3.2 Lift Ski Magnetic Design and Simulation 24

3.3 Lift Ski Mechanical Design and Manufacturing 29

4 Vehicle Dynamics 33

4.1 Suspension 34

4.2 Lateral Control Modules 40

4.3 Brakes 44

4.4 Low Speed System 47

5 Aerodynamics & Structures 51

5.1 Frame Design 51

5.2 Aerodynamic Design 54

5.3 Shell Structural Design 66

5.4 Shell Manufacturing 67

6 Electronics 70

6.1 Sensing & Communication 70

6.2 Power Systems 78

6.3 Wiring Harness 83

7 Software 87

7.1 Overview of System Modules 87

- 7.2 Communications 90
- 7.3 Guidance, Navigation, and Control 93
- 7.4 UI 98
- 7.5 Data Logging 98
- 7.6 Simulation 99

Testing

- 8 Levitation Testing Campaign 101
 - 8.1 Rotary Test Rig 101
 - 8.2 Linear Test Rig 105
- 9 Vibration Testing Campaign 110
- 10 Competition Log 112
- 11 Competition Data Analysis 115
- Bibliography 118
- List of Abbreviations 121

Appendices

- A Project Schedule 123
- B PCB Schematics 125
- C Final Prototype 129

INTRODUCTION

The Hyperloop is a concept for high speed ground transportation, consisting of passenger pods traveling at transonic speeds in a partially evacuated tube. The concept was originally proposed in a white paper published by SpaceX in 2013¹ as an alternative to the high-speed rail system currently being developed between Los Angeles and San Francisco, which was deemed too expensive and slow.

The Hyperloop concept could fill a growing need for an alternative transportation mode for short-haul travel. For short routes, such as Los Angeles – San Francisco, or Boston – New York, the time spent traveling at the cruise speed is quite low compared to overall end-to-end travel time due to inescapable inefficiencies in air travel (runway taxiing, climb, descent, holding patterns, etc.). The high-frequency throughput of Hyperloop pods could alleviate some of these inefficiencies. Recently, KPMG published a preliminary study commissioned by Hyperloop One – one of the companies commercializing the Hyperloop concept – on the Helsinki–Stockholm corridor where they found that the Hyperloop could cut down end-to-end travel time by 75% to 28 minutes.² Furthermore, the market share for high-speed transport is projected to grow rapidly over the next few decades,³ and the Hyperloop concept could take some pressure off increasingly congested airports and flight routes.

Momentum is growing in the Hyperloop movement, with a number of newly founded companies attempting to commercialize it. In addition, SpaceX is sponsoring a student competition to encourage innovation and to help accelerate the development of a working prototype, starting June 2015^a. Over 1,000 teams submitted their intent to compete, and over 100 teams made it to Design Weekend in January 2016. The student team from the Massachusetts Institute of Technology – the MIT Hyperloop Team – won first place overall in that design weekend.⁴

Academic research into the Hyperloop concept has focused mostly on system integration. A conceptual sizing tool using the OpenMDAO framework⁵ focuses primarily on the aerodynamic and thermodynamic interactions between the pod and tube, with recent work focusing on the energy consumption of the system.⁶ The pods for the SpaceX Hyperloop Competition were the first physical prototypes of the Hyperloop concept. The goal for this report is therefore to shed some light on the design of the MIT Hyperloop pod for this competition.

Section 1.1 discusses some historic background of the Hyperloop concept, Section 1.2 provides an overview of the SpaceX Hyperloop competition, and Section 1.3 gives a brief overview of the team itself. A high-level overview of the pod is provided in Chapter 2. Part I focuses on the design and construction of our pod, while Part II focuses on the testing phase of the project.

^awww.spacex.com/hyperloop



1.1 Background

Concepts for high-speed trains in vacuum or evacuated tubes can be traced back as far as 1909, when rocket pioneer Robert H. Goddard proposed high-speed passenger-carrying pods traveling through evacuated tubes.⁷ Bachelet introduced the core idea behind magnetically levitating trains as early as 1910.⁸ Over the years these ideas have been further refined, for instance by the Rand Corporation in 1972 with their “Very High Speed Transport System”.⁹

The Hyperloop Alpha white paper combined several of these historic concepts¹ and spurred a great deal of public interest in the concept, something the earlier ideas were somewhat lacking. This white paper discusses a Hyperloop pod that travels at 1220 *km/h* in a partially evacuated tube (1/1000th of atmospheric pressure) levitating using air bearings. The use of wheels at these high speeds would be quite problematic because of the massive centripetal forces on them. Air bearings are proposed as a more efficient mechanism, where the pod floats on a thin film of compressed air. In the Hyperloop Alpha concept, this compressed air is supplied by an onboard compressor. Propulsion is provided by a linear induction motor. The benefit of this is that the heavy components are built track-side and the pod only has to carry a rotor which makes the propulsion quite efficient. Furthermore, that same linear induction motor can also be used for braking at the other end of the tube to recover a substantial amount of energy.

The onboard compressor is also used to improve the efficiency of the pod at higher speeds. Once the pod reaches transonic speeds, the flow around the pod will start to choke, i.e. the flow around the pod will become sonic. At this sonic condition – the so-called *Kantrowitz limit*¹⁰ – the mass flow around the pod is at its maximum. Therefore, when the speed is further increased, not all flow can travel around the pod and is therefore collected in front of the pod. The result is a column of air being pushed by the pod throughout its run. That pressure build-up results in significant additional drag. The Hyperloop Alpha concept therefore introduces the on-board compressor to compress the additional flow and suck it through the pod, while at the same time supplying compressed air to the air bearings.

1.2 SpaceX Hyperloop Competition

Two years after publishing the Hyperloop Alpha white paper, SpaceX announced the Hyperloop student competition to advance the concepts proposed in that white paper to the next level, while at the same time garnering more interest from students, universities and the general public for the concept. In this competition, students are to design and build Hyperloop pods which are then tested on a 1 mile long, 6 *ft* diameter tube built by SpaceX. The tube has a flat concrete base on which aluminium (Al 6101) track plates and an aluminium I-beam are installed.

The competition is intentionally kept very open, allowing for a variety of different designs. Taking the levitation concepts as an example, wheels, air bearings, and magnetic levitation are all allowed. The Hyperloop Alpha concept used a linear induction motor to get up to speed. Because such a system requires a great deal of integration between the pod and track, and it is quite hard to handle this integration with a large number of teams, SpaceX provides a pusher vehicle to bring pods up to speed.

The most important scoring criteria for the competition are the overall run time (therefore favoring high cruise speeds and fast braking), and the potential to scale up the technologies used in the competition vehicle to a full-scale Hyperloop vehicle.

1.3 About the Team

The MIT Hyperloop Team was founded in June 2015 by students from the Aeronautics & Astronautics, and Mechanical Engineering departments at MIT to participate in the SpaceX Hyperloop Competition.

Most students had experience in other design and build competitions, such as the FSAE engineering design competition. Regardless of whether students believed Hyperloop would indeed be the next successful mode of transportation, everyone was excited to design and build a vehicle that no one had designed or built before. Most students on the team were graduate students, and all students worked on the team part-time, while taking classes, performing research, and writing theses. The whole project took well over a year to complete, an overview of the schedule is provided in Appendix A.

At its peak, the team had 35 team members from a variety of departments. Most students were Mechanical Engineering students (15), followed by Electrical Engineering & Computer Science (7), Sloan School of Management (7), and Aeronautical & Astronautical Engineering (6).

PART I

DESIGN & CONSTRUCTION

TOP-LEVEL OVERVIEW

The competition goal is to advance the Hyperloop concept to the next level by exploring technologies that one day could be used in full-scale Hyperloop and by using these technologies in a working Hyperloop prototype. The top-level design of an interpretation of that competition goal is discussed here.

The overall goal of the MIT Hyperloop Team within this competition was to design and build a pod that scores well in both aspects of the competition, i.e speed and scalability of systems used. The speed goal favors a lightweight design, while the scalability goal comes back in every aspect of the pod. There are several approaches that could be taken to design for scalability, for instance to design a full Hyperloop pod with a large passenger compartment and scale it down for the competition. The MIT Hyperloop team took the approach of focusing on the most important technologies for the Hyperloop concept, and developing scalable technologies for those. The pod therefore does not have a dedicated pressurized passenger compartment, as this technology is already available from fuselage design in aircraft, and thus not unique to a Hyperloop. Instead, a major focus throughout the design and build process was the scalability of the levitation design, the braking design, and the aerodynamics.

Other major driving requirements for the design were imposed by the team itself. Firstly, the pod had to be built in four months (February 2016 start of manufacturing – May 2016 fully assembled), which meant that simpler/easier-to-manufacture designs were often favored. Secondly, the pod needed to accelerate at $2.4 G$ – the maximum proposed acceleration offered by the SpaceX pusher. Finally, the pod had to be robust to changes in performance specifications and track tolerances.

2.1 Conceptual Design

A conceptual design tool was developed to aid in the initial sizing of the pod. Due to the Hyperloop Alpha design,¹ this tool was mostly focused on using air bearings and mitigating the Kantrowitz limit, and thereby was instrumental in discounting air bearings once the track specifications were released. Our developed tool is similar in focus to a conceptual sizing tool developed Chin et al. using the OpenMDAO framework,⁵ which focuses primarily on the aerodynamic and thermodynamic interactions between the pod and tube. Our tool, however, extends beyond this work by considering more factors, including, but not limited to, alternative pod configurations and levitation systems, and system weight estimation.

Considering that no Hyperloop system has been designed before, we cannot rely on correlations in the same way as for instance Torenbeek¹¹ or Raymer¹² do. We need to consider physics-based models for the high-level sizing, similar to TASOPT¹³ or SUAVE.¹⁴ The only exception is for secondary components, such as batteries, that can be sized based on historical correlations.

The goal for this conceptual design tool is to size a pod that avoids the Kantrowitz limit and provides enough lift from the air bearings. For that purpose, several configurations were investigated, such as air bearings with compressor, air bearings with pressure vessel with or without compressor for mitigating exceeding the Kantrowitz limit, et cetera.

As a low-fidelity model for the air bearings, we consider an aerostatic bearing by Khatait et al.¹⁵ This allows us to use analytical expressions for the relation between load capacity of the bearings and supply pressure and mass flow rates. The large assumption here is that the air bearing is static, which is an optimistic estimate, because the performance of the air bearings will most likely deteriorate for higher speeds.

This conceptual design tool was used to design a pod with air bearings for this competition. However, it was found that the required volume for the pod to keep sufficient track clearance ($> 2\text{ mm}$) with the track's tolerances^a was simply too large. More information about this tradeoff is provided in Section 3.1. Therefore, the pod instead uses a passive magnetic levitation (maglev) is used to levitate above the track, specifically the pod uses an *Electrodynamic Suspension (EDS)* system. EDS works with larger gap heights (on the order of a few millimeters), and additionally becomes more efficient at higher speeds.

The overall size of the pod was decided on by looking at performance sensitivities. SpaceX provides a heavy pusher that – according to original specifications – outputs a fixed force independent (to first order) of pod mass. Because the pod weighs relatively little compared to the SpaceX pusher, the influence of weight on the overall run time is actually fairly small, as seen in Fig. 2.1a, which is very different for e.g. race cars and aircraft. In the design of for instance race cars and aircraft, the weight optimization in the design is one of the most time-consuming parts of the design process, because it is an iterative process. Considering the tight timeline for this project, we decided to fix the pod target weight at 250 kg , because it seemed like this was an attainable weight goal while the performance gain for going for a lower weight was fairly small. Fixing the weight to 250 kg allows for skipping the whole weight iteration process, and if the eventual pod would be lighter, ballast could be added to avoid upsetting the system dynamics. Note that this design strategy is in line with our design philosophy, to focus on systems critical the Hyperloop concept. Weight optimization is then more of a detail in a final design for a Hyperloop system, once its feasibility is proven.

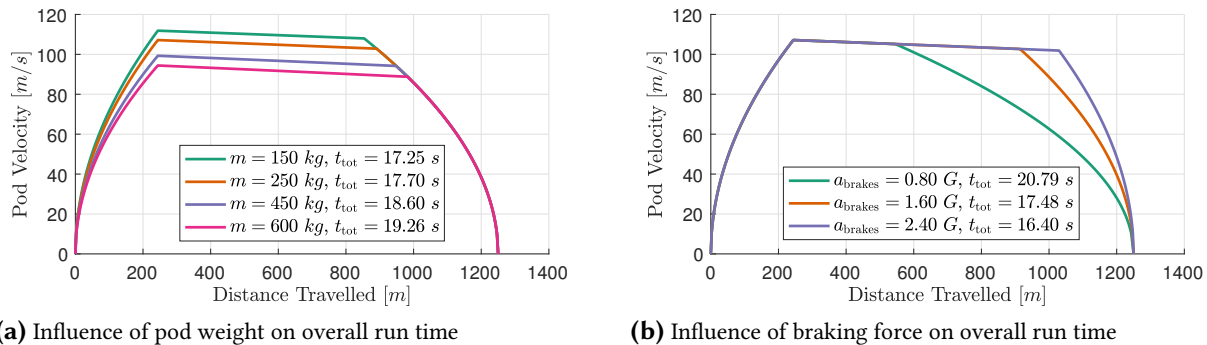


Figure 2.1: Influence of pod weight and braking performance on overall run time.

Whereas the pod weight's influence on the performance is rather slim, the pod's braking force on the other hand has a large influence of the performance, as shown in Fig. 2.1b. Besides the levitation system, the braking system is therefore another major focus in our design.

^aSteps in height between the aluminum plates of the track were expected to be as large as 1 mm , but were found to be higher for the actual track in some places.

2.2 Final Design

An overview of the final MIT Hyperloop design is shown in Fig. 2.3. Photographs of the final prototype are included in Appendix C. The prototype is 2.4 m long, weighs 258 kg, and is designed for a top speed of 400 km/h. The pod has no propulsion on board, as we rely on the SpaceX pusher to get up to cruising speed. An example of a velocity trajectory for a flight of this pod is shown in Fig. 2.2.

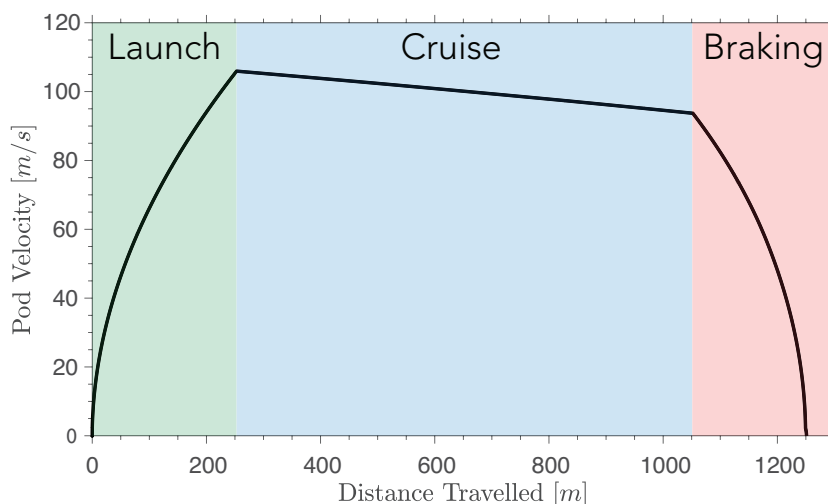


Figure 2.2: Velocity profile for MIT Hyperloop pod run.

As mentioned before, the pod uses an *Electrodynamic Suspension (EDS)* system to be robust to track tolerances and to have a scalable levitation system. The levitation system does not require power, therefore it scores well for scalability because of its safety and efficiency at high speeds. Magnetic levitation is notorious for being underdamped, therefore a suspension is added between the magnet arrays and the pod.

The eddy-current braking system uses the same physical principle as the levitation design, only now the design is optimized for drag rather than lift. The design is similar to an eddy-current brake on a rollercoaster. As discussed earlier, linear induction motors will likely be the best technology for propulsion and deceleration. However, track-side infrastructure is required for that. In the case of an anomaly mid-run (e.g. a problem with a pod further down the track, tube de-pressurization/breach, etc.) the pod still has to brake. Therefore, the braking system on the MIT Hyperloop pod is designed as an emergency braking system for a full-scale Hyperloop, decelerating the pod upwards of 2.4 *G*.

The pod does not have an onboard compressor, both because of the lower speeds in the competition and because the pod does not have air bearings that need to be supplied with air. The Kantrowitz limit will be further discussed in Section 5.2.3.

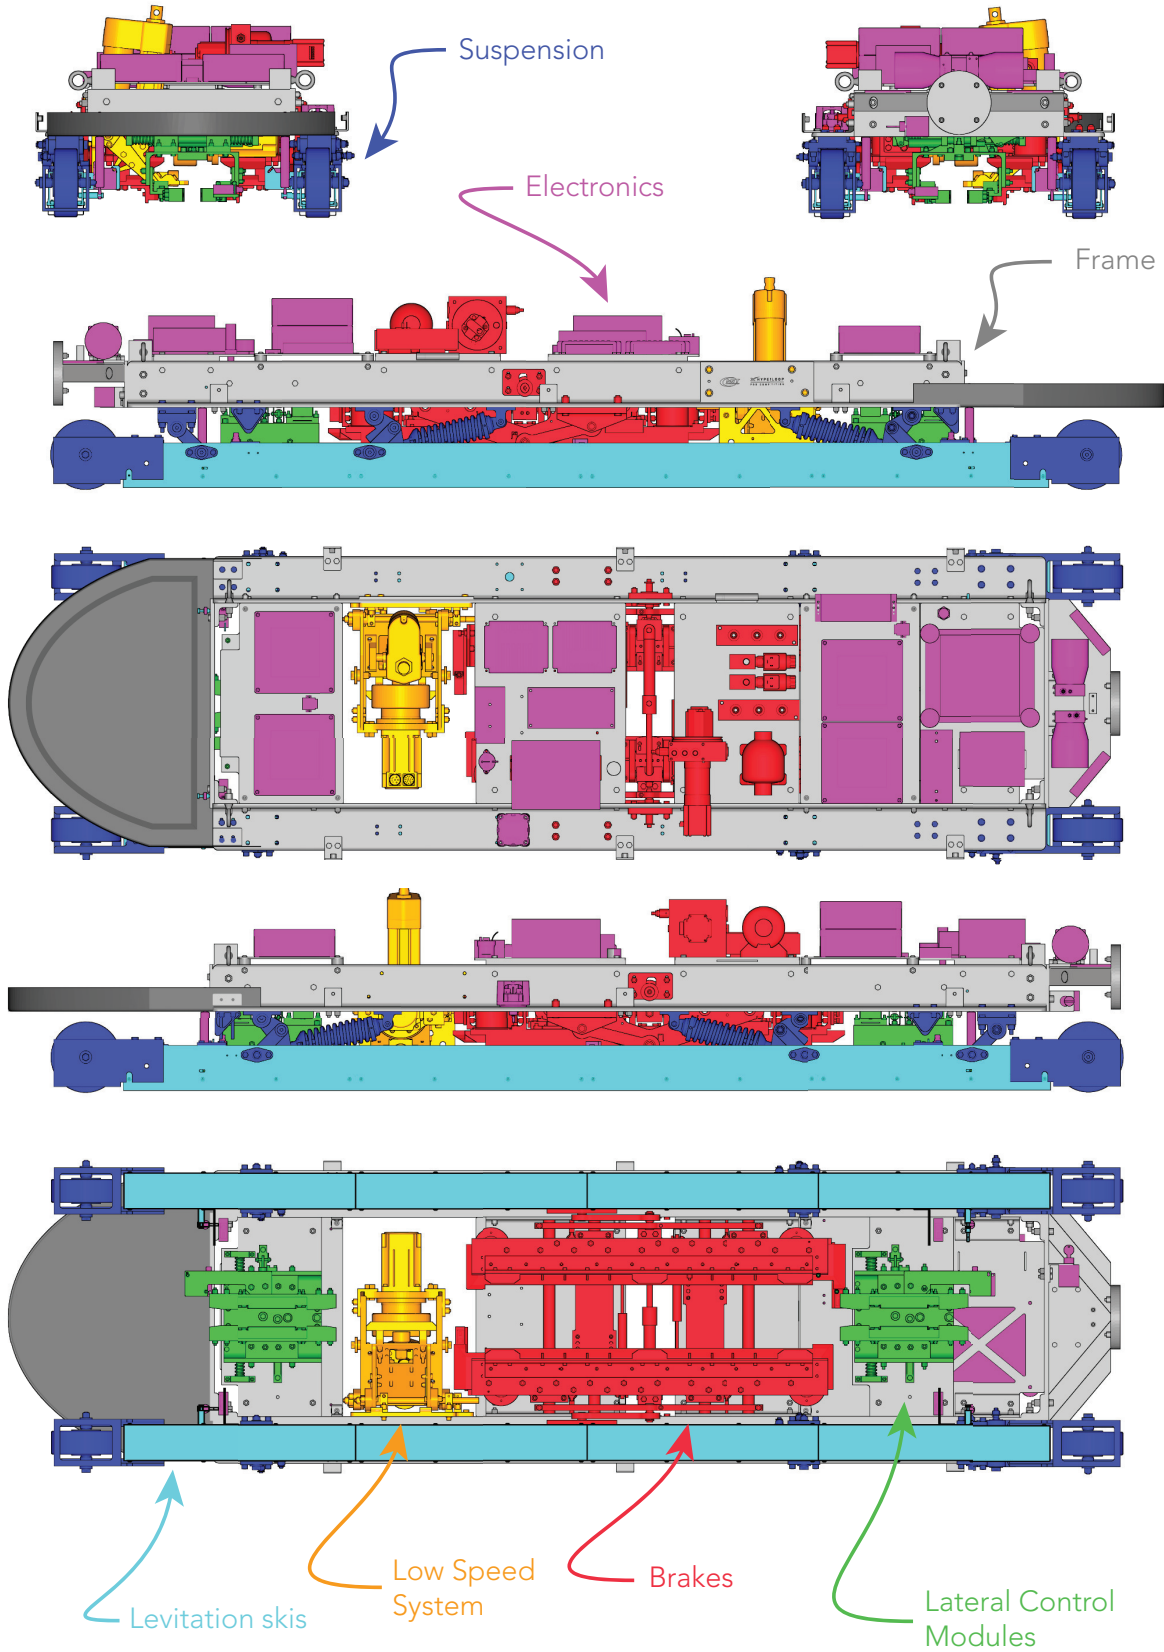


Figure 2.3: Overview of the MIT Hyperloop pod. For clarity the aerodynamic shell is not shown here.

LEVITATION

A number of potential options for a levitation system were investigated in an iterative process as the rules of the competition were slowly developed. This is not necessarily an optimized method for levitation but fits within the context of the competition (prototype version), the tight timeline, and the simplicity of applying this technology.

Air bearings, while investigated early on in the process, were determined to be impractical due to the large gap heights (1 *mm*) between plates of the track. Wheels were found not to be scalable and a selection was made within the team to compete in the non-wheeled vehicle class for more relevance to an actual Hyperloop system. An Electromagnetic Suspension (EMS) maglev system was not possible without the necessary infrastructure within the tube. While the necessary control system is complicated, the system has many positives that make it a serious option in the event of future tube changes. Ultimately, an Electrodynamic Suspension (EDS) maglev system was chosen due to it being passive, scalable with increasing L/D with speed, large gap heights, compatible with the track, and relatively straightforward to implement.

The vast majority of the design effort was spent in characterizing the physics behind EDS maglev, determining the relevant magnetic array parameters and their effects, and simulating various array geometries in an attempt to optimize L/D while minimizing mass, cost, and maintaining acceptable gap height. The main parameter affecting L/D was determined to be array wavelength which drove the majority of design decisions.

3.1 Feasibility of Air Bearings and Compressor

During the early phases of conceptual design, air bearings were considered due to the initial suggestions of the Hyperloop Alpha paper¹ as well as the lack of details of the tube and track specifications. In a situation where the track is simply the bottom surface of the tube, air bearings or wheels become some of the only feasible options. In this early phase, a general pod sizing code – which was already discussed in Section 2.1 – was developed that looked at air bearing levitation requirements in terms of size, and air flow to lift a pod given the pod dimensions and mass. For a relatively small test pod with a small area ratio and at speeds well below $M_\infty = 1$, no compressor is needed to divert air flow, as also discussed in Section 5.2.3.

In that case, a compressed air tank is the simplest method to supply air to the air bearings under a short run time. The feasibility study covers what are basic sizing requirements for air bearing area, pressure vessel size and pressure, and gap heights.

The air bearing gap height is found to be the driving parameter for the entire system since the required air flow rate scales with the cube of the gap height for a static air bearing. Flow rates then increase due to the additional losses of a dynamic air bearing. As more competition rules and track specifications were developed and shared, namely the 1 *mm* plate steps between track sections, air bearings quickly became

infeasible without a massive risk of making contact with the track and putting enormous shear loads on the air bearings and suspension.




Figure 3.1: Required tank volume as a function of air bearings pad area and gap height.

A sample of the required tank volume as a function of gap height and pad area with a tank pressure at 5000 *psi* in Fig. 3.1 shows that the tank quickly becomes unfeasibly large at gap heights approaching 1 *mm* or higher. However, the addition of the aluminum track in the SpaceX track allows for the possibility of maglev systems to be implemented.

3.2 Lift Ski Magnetic Design and Simulation

Once air bearings were determined to be infeasible due to the large step heights between plates, other options had to be considered. From the beginning, the team made the philosophical decision to construct a pod in the non-wheeled vehicle class due to the poor scalability of wheels in a theoretical “full scale” Hyperloop. This left two types of maglev systems as potential options. The first is an Electromagnetic Suspension (EMS) maglev system which is a maglev system that works by attraction between magnets on a pod and ferritic material. Powerful permanent or electro magnets are attracted to the bottom of a ferritic rail or track, lifting up the vehicle. Since this is an inherently unstable system that wants to contact, electromagnetic control systems are required to modulate the magnetic field and hold it in at a fixed gap height. This system has extremely high lift to drag ratios which are a strong benefit of this type of system. For the competition, however, the lack of a dedicated ferritic track or rail prevented this from being a viable option.

Therefore, the Electrodynamic Suspension (EDS) system is chosen as the mode of levitation for this pod. The trade-off of these systems is summarized in Table 3.1.

2D Simulations

2D simulations were used heavily in the design process for the levitation magnet arrays due to their relatively low computational cost compared to 3D simulations while still capturing the majority of the

Table 3.1: Three levitation systems were considered.

| | Air bearings | Electromagnetic suspension | Electrodynamic suspension |
|--|-----------------------------|--|--|
| Sources of lift | Thin film of compressed air | Magnetic attraction to a ferromagnetic surface | Repulsive magnetic force produced by relative motion to a conducting surface |
| Track suitability | No | No | Yes |
| Lift to drag ratio | High | High | Moderate |
| Gap height | Very small | Large | Large |
| Stiffness | High | High when controlled | Low |
| Stability | Moderate | Inherently unstable | Low |
| Proven technology at required conditions | No | Yes | Yes |
| Power requirements | Sensitive to gap height | Moderate | Low |
| Static levitation | Yes | Yes | No |
| Complexity | Moderate | High | Low |

physics involved. After some basic sizing and design intuition was gained from our analytical models, the 2D simulations were used for sizing optimization and variable sweeps. Here, we use the FEMM¹⁶ package for the 2D simulations.

The basic model consists of an array of alternating north and south permanent magnets and a permeable back iron over an aluminum plate. An example of such a configuration is shown in Fig. 3.2.




Figure 3.2: Example configuration of an EDS system.

The model is parameterized to allow for sweeps of magnet thickness, back iron thickness, array wavelength, number of periods, gap height, magnet grade, and speed. Transient analysis is chosen in order to capture the generation of eddy currents due to motion which creates the levitation forces. Table 3.2 is a summary of the results of the variable sweeps and the contribution of each parameter on lift force and L/D . These results generally follow the expected trends based on the analytical models.

These 2D simulations were used for fast parametric sweeps. The influence of the wavelength of the array is shown in Fig. 3.3, where it is seen that a large wavelength is the most efficient, although it does reduce the L/W of the magnet array. However, for overall run time weight is less important than efficiency during cruise (as already explained in Section 2.1), and therefore at the design point we have a large L/D but a smaller L/W .

2D simulations were also used to characterize end effects - the decreased lift-per-length of a finite length compared with an infinite array. In general, a single period generates approximately 30% less lift per unit length than an infinite array. End effects also result in lower lift at the ends of the array, which makes for lower pitch stiffness. This means that having fewer periods means having lower pitch stiffness.

Table 3.2: Effect of design parameters on magnet array performance.

| Design parameter | Effect on lift | Effect of L/D | Effect on weight |
|---------------------|-------------------------|----------------------|------------------|
| Array wavelength | Slight decrease | Square root increase | Increase |
| Number of periods | Linear increase | No effect | Increase |
| Array width | Linear increase | No effect | Increase |
| Array thickness | Squared increase | Negligible | Increase |
| Magnet grade | Increase | No effect | No effect |
| Back iron thickness | Slight increase | Negligible | Increase |
| Nominal gap height | Inverse square decrease | No effect | No effect |





Figure 3.3: Effect of magnet array wavelength. L/D ratio increases with wavelength, but longer wavelengths also mean lower specific lift.

3D Simulations


3D simulations were used to assess the effect of varying the width of the array (“edge effects”) and to corroborate the 2D simulations. It is expected that an infinitely wide array performs better than a finite width array, therefore these simulations are used to quantify how much lift is lost by using a finite width array.

Interestingly, the 3D simulations give higher lift to drag ratios than the 2D simulations. This is due to the fact that the 2D simulations cannot capture the eddy currents generated along the edge of the magnets (shown in Fig. 3.4). These eddy currents along the edge of the magnet in axial direction generate an additional lift force.

Some parameter sweeps for the magnet array width for the 3D simulations are shown in Fig. 3.5. The number of 3D simulations that were run during the design phase were unfortunately limited due to their high computational cost.



(a) Eddy currents in aluminum track.



(b) Eddy currents exert forces along the edge of the magnet arrays.

Figure 3.4: Eddy currents in the track and forces along the magnet arrays.




Figure 3.5: Results of parameter sweeps for 3D simulations. The magnet width for the final design is 50 mm.

3.2.1 Magnet Bridge (Back Iron) Design

Because the magnetic array is an up-down configuration (as opposed to a halbach array) a back iron is required to complete the so-called “magnetic circuit”, channeling the flux from one pole to the next. This has the possibility of adding significant weight to the array since the back iron height can be more than the thickness of the magnets while having a similar density. Because of the extremely high aspect ratio of our magnet arrays (20” pole length vs. 3/4” magnet thickness), the flux is concentrated mostly at the junctions between poles. This can be seen in the baseline 2D simulation of the magnets in Fig. 3.6.

To minimize weight, the back iron shape was changed from a constant thickness to a sinusoidal shape. The back iron height $h(x)$ was of the form,

$$h(x) = h_{\max} \left(\frac{\cos^n \left(\frac{2\pi x}{\lambda} \right) + 1}{2} \right) \quad (3.1)$$

where λ and h_{\max} are the wavelength of the array and maximum back iron height, respectively and n is a parameter which can be varied to control the width of the sinusoidal “hump” at the magnet junctions. A parametric 2D simulation in FEMM¹⁶ was used to observe the effect of the sinusoidal back iron on the magnetic field at the intended track height below the magnets. It was found that the back iron weight

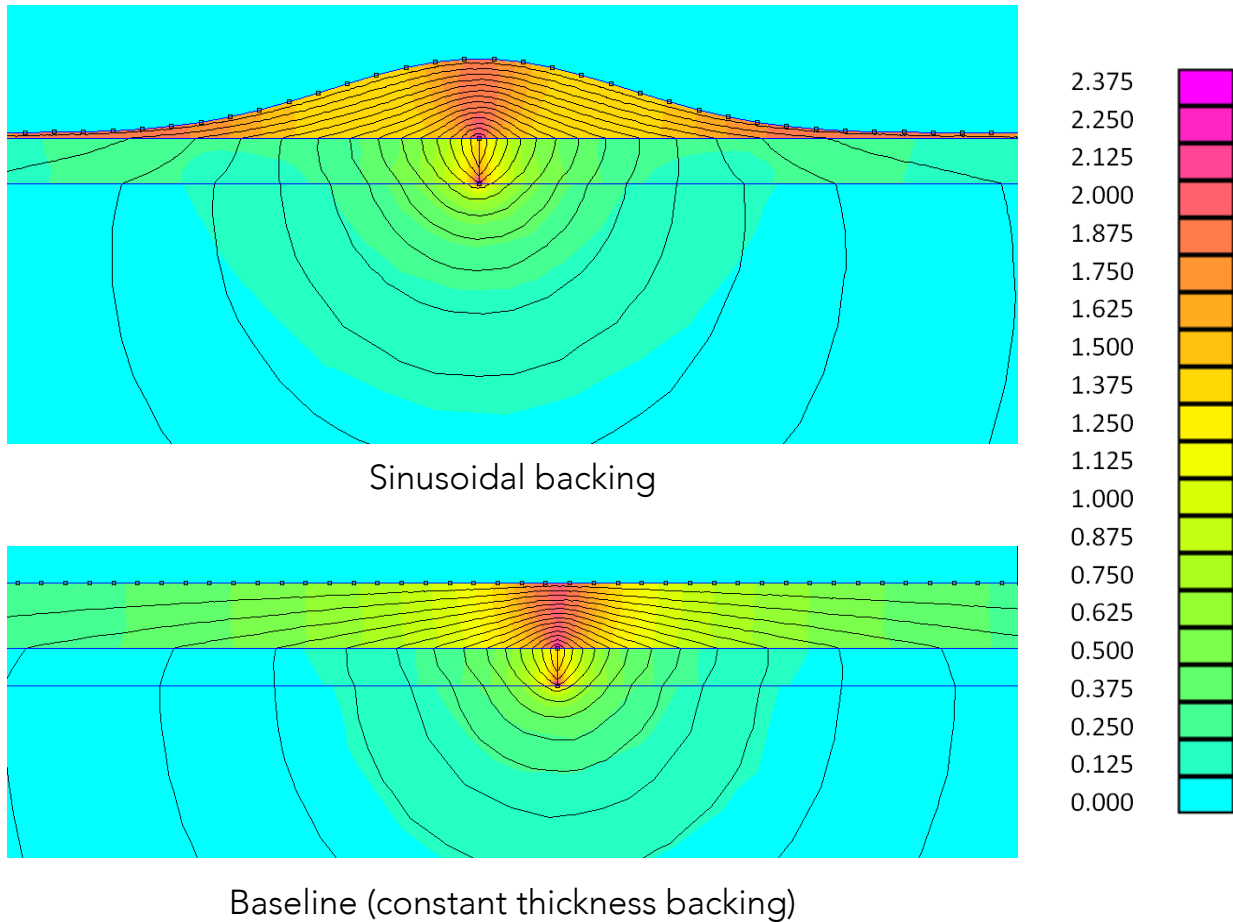


Figure 3.6: Comparison between magnetic field strength for constant iron backing or sinusoidal shaped iron backing.

could be reduced by more than 70% with only a 1% reduction in peak magnetic field at the inspection plane, which is shown in Fig. 3.7.

3.2.2 Final Array Design

The final parameters of the magnet array are shown in Table 3.3.

Table 3.3: Final design parameters for levitation arrays.

| Design Parameter | Value |
|---------------------|----------------------|
| Array wavelength | 1 m |
| Number of periods | 2 |
| Array width | 50 mm |
| Array thickness | 19.05 mm |
| Magnet material | NdFeB, Grade N42 |
| Back iron thickness | 1 – 20 mm (stepwise) |
| Nominal gap height | 10 mm |




Figure 3.7: Comparison between field magnitude at plate for constant thickness backing iron and sinusoidally shaped backing iron.

Note that demagnetization can occur with closely interacting powerful magnets. N52 magnets were originally chosen but showed 30% field reduction over time due to demagnetization. We switched to higher coercivity and weaker N42 magnets to avoid static loss of field strength. This meant a 10% weaker field than the nominal N52 array. FEMM¹⁶ was used to investigate the demagnetization. The comparison in BH-curves for these two grades of magnets is shown in Fig. 3.8.

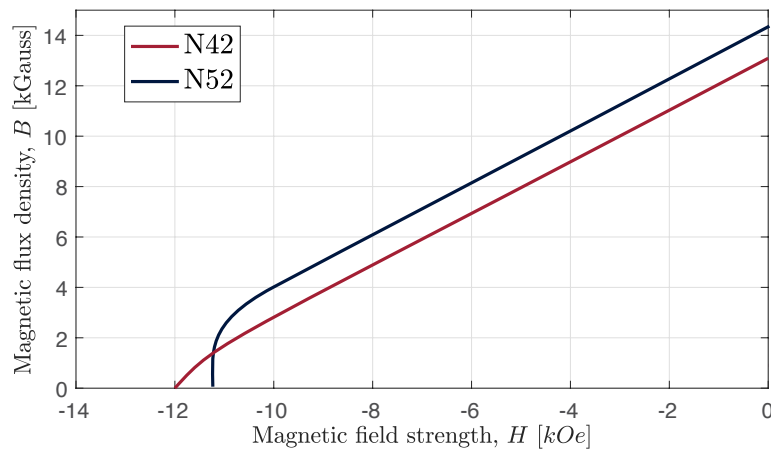


Figure 3.8: BH curves for Neodymium PH grade N42 and grade N52.

3.3 Lift Ski Mechanical Design and Manufacturing

The lift skis are the mechanical connection between the magnets and the pod suspension. Apart from the functional requirements laid out below, the main design drivers were manufacturability and ease of assembly.



The main functional requirements for the skis are,

1. Provide mounting locations for magnets, back iron, wheels and suspension
2. Minimize deflection under load (less than 1 *mm* at the center under maximum load).

The interfaces for the skis are,

1. Magnets: two threaded holes per magnet
2. Suspension: bolted on bearing plates (3 pairs of 2)
3. Low Speed Wheels: two per ski.

Basic Design

The ski consists of four main sets of components: the ski box section, the continuous back iron, the magnet bridges, and the magnets themselves. The box section is a $3 \times 3 \times 1/8$ " wall square 6061 aluminum tube into which various cutouts were made to accommodate the dampers, magnet bridges, and various other mounting points for the interfaces listed above. A continuous piece of 19 gauge 1018 steel is used to establish a constant magnetic path for flux between poles while stacked $1/4$ " steel pieces were used to make an approximation of the sinusoidal magnet bridges at the poles outlined in the magnet design section in Section 3.2.

Magnets

Several options were considered for mounting the magnets to the skis. One such option is to clamp the magnets from the sides and retain them in vertical direction by a small step on the clamp plates. The advantage of this method is that off-the-shelf magnets could be used to minimize the lead time. However, this option is abandoned due to risks during both magnet installation and operation. Another option is to bond the magnets together, but this requires significant analysis and prevents any repair or maintenance of the arrays. Furthermore, with this option there is no adjustability in the array rearranging the magnets later, and was therefore not seriously considered.

Mechanical retention with threaded fasteners was decided as the attachment method of choice when it was found that custom magnets could be procured fast enough to fit within our production timeline. The final attachment strategy was to divide each 20" pole of the array into 5 separate $2 \times 4 \times 3/4$ " pieces with two counterbored holes for $1/4$ -20 button head cap screws. Counterbores were put in both sides of the magnets to allow for magnets to be mounted with either the north or south pole facing up, requiring only one type of lift magnet on the pod. To minimize interference between adjacent magnets, the final magnet length was undersized from the nominal mounting spacing by 0.01 " – 3.99 " as opposed to 4.00 ".

The magnets are installed using a custom jig. A single magnet is incredibly strong and should therefore be handled with care. To make matters worse, during the installation the magnets had to be placed against each other with NP–NP and SP–SP. The jig allows for placing the magnets using a linear actuator, by clamping the ski down and then pressing the magnet in.

Ski Box Section

For simplicity, the box sections are restricted to have holes and profiles that could be achieved on a CNC tube laser cutter without the need for precision machined features. Due to the thin wall thickness of the beam, threaded riv-nuts are installed to provide threaded hard points for the $1/4$ -20 fasteners used to secure the magnets. A large number of 10-24 holes tapped directly into the side of the beam are used to secure the






Figure 3.9: Tool to install magnets on magnet skis.

bent aluminum magnet covers. The magnet bridges, low speed wheels, suspension mounts, and suspension bump stops are all installed using clearance holes to limit the required manufacturing precision.



(a) The backing iron doubles as nut-plate and is therefore longer than required, therefore the backing features rectangular cut-outs.



(b) Section view of magnet bridges.

Figure 3.10: Magnet bridges used to channel flux from one pole to the next.

Magnet Bridges

The magnet bridges are constructed from stacked pieces of $\frac{1}{4}$ " 1018 steel. This method was chosen to greatly simplify machining. The pieces are all 2D extrusions allowing for them to be cut with a waterjet



with minimal post-machining (threaded and countersunk holes). The bridge pieces are held together with recessed flat head screws through the bottom plate and threaded into the top plate and held to the box section by a series of 10-24 fasteners installed from the bottom. Because the magnet bridges double as a nut-plate, the second piece in the stack was much longer than that which would be required to establish the necessary sinusoidal shape determined before. This led to square cut outs on the ends where the thick steel was unnecessary, as shown in Fig. 3.10a.

Continuous Back Iron

While the majority of the flux is channeled through the magnet bridges directly over the poles, a small thickness of steel is required to channel all the flux. Because of this, a thin piece of 19 gauge 1018 steel spans the entire length of the ski – shown in Fig. 3.10b – and is fastened both by the magnet clamping forces and by the fasteners used to attach the magnet bridges.

Suspension Mounts

The suspension components are attached to the lift skis using two bolt flanges with precision bores, as shown in Fig. 4.9. This satisfied the requirement of no precision holes directly machined in the ski and allowed for the shoulder bolts on which the suspension members pivot to have greater contact area. These pieces also relieve the potential over constraint that could be caused by a lack of concentricity between the shoulder bolt bores on either side of the skis. The clearance holes through which the bolts are installed allow the flanges to float and self align before tightening.

VEHICLE DYNAMICS

Vehicle Dynamics comprises all components that are (not necessarily continuously) in contact with the track. These components ensure pod is stable during flight and can brake at the end of the run. An overview of the vehicle dynamics components is shown in Fig. 4.1.




Figure 4.1: Components of the Vehicle Dynamics system. The rest of the pod is hidden here.

A key point for stability was to separate the degrees of freedom between different contact groups, such that each degree of freedom is controlled only by *one* contact group. Table 4.1 shows which degree of freedom is controlled by which contact group.

Table 4.1: Each degree of freedom on the pod is controlled by only *one* contact group.

| | Translation | Rotation |
|-----|-------------------------|-------------------------|
| x | Brakes | Lift skis |
| y | Lateral Control Modules | Lift skis |
| z | Lift skis | Lateral Control Modules |



4.1 Suspension

The pod has a suspension system for two primary reasons: to improve the dynamics of the pod during cruise, and to allow for ride height adjustability. SpaceX intended to judge teams based on their pod's vibration profile throughout the run, and therefore improving the dynamics of the pod was quite important. The levitation magnets provide high stiffness and low damping during cruise, which would result in large, under-damped oscillations in the pod and skis if no damping is introduced into the system. By isolating the mass of the skis and the pod with a suspension, we can provide damping to the system and limit the effects of track disturbances. This way we can respect the clearances around the rail of the lateral control modules and brakes, and at the same time limit the vibrations of the pod during cruise. Many designs were evaluated to accomplish these goals, and target values for geometry, stiffness, and damping were established with simulation and calculations.

4.1.1 Functional Requirements

The suspension is responsible for controlling motion of the pod in heave, pitch, and roll. The skis will contact the track surface with wheels at rest and low speed and will levitate off the track surface at speeds above roughly 10 m/s . With the suspension responsible for these degrees of freedom, the pod would be capable of compensating for uneven track surfaces and track disturbances. For example, if the track was sloped downward on one side and upward on the other, the suspension should allow the levitation magnets to maintain their ideal gap height.

As per the SpaceX specifications for the track, a "step" of up to 1 mm might be present due to placement tolerance of the aluminum plate sections. Any step encountered while levitating would introduce a force disturbance to the skis, which would cause ski oscillations in pitch and heave. This poses an issue for the maintaining the tolerance of the LCMs and brakes. It would also introduce vibrations into the system, which could have adverse effects on electronics, bolts, and any imagined passenger. In keeping with transportation industry standards for passenger comfort, the suspension should keep the root mean squared vertical acceleration below 0.1 Gs for the range of frequencies we expect to encounter. In addition, the oscillations at the front and rear end of the skis should be kept well below 5 mm , as this is the expected distance of the wheels to the track during cruise. Ideally, the system would be critically damped. In order to do so, the heave and pitch degrees of freedom should be largely de-coupled.

During the launch phase of travel, the pod will pitch forward due to the pusher interface being located above the center of mass. The pitch suspension will need to be stiff enough to ensure that excessive pitching does not occur, which would crash the front end of the pod into the rail or cause other interference issues with the brakes or LCMs.

During braking, the inertial force of the skis will be in direct competition with the braking acceleration of the pod. It is important to design for this phase from a robustness standpoint, as the dynamics of braking at up to 4 Gs will likely already be undesirable. The design of a hard-stop for unpredictable loading cases could be necessary.

Due to the unknown effects of the magnets when the pod is oscillating laterally around the rail, the suspension must be capable of handling side loading. This includes eliminating slop in any linkages or connections between the pod and the skis.

The system must be designed to have repeatable installation and performance. Unwanted friction should be mitigated and sufficient, but not excessive, points of adjustability should be designed. As with almost every component on the pod, the system should be designed to be of minimal weight while maintaining manufacturability and robustness.

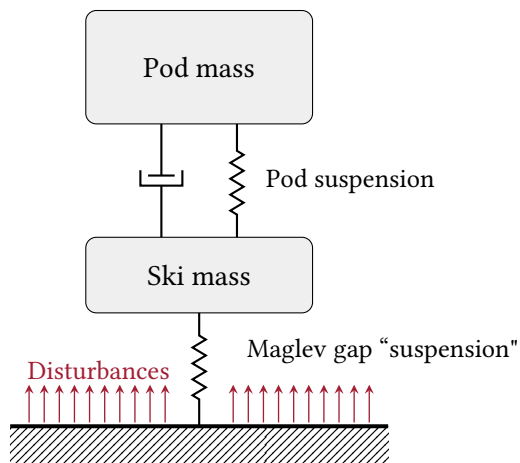


Figure 4.2: Simple two-mass system to describe the pod dynamics in pure-heave setting.

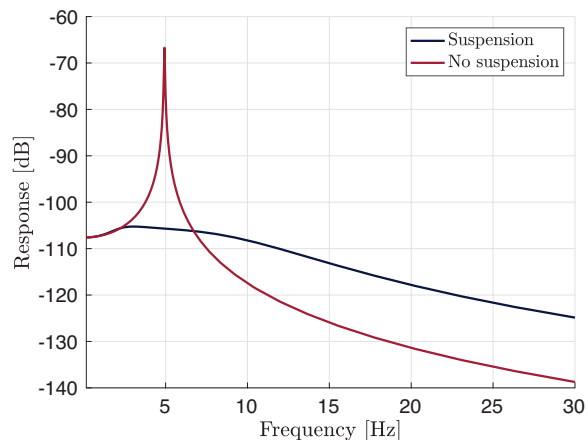


Figure 4.3: Comparison of system dynamics with and without pod suspension.

4.1.2 Preliminary Modeling

To determine the relative values of stiffness and damping for the suspension, simple two-mass systems (Fig. 4.2) were proposed as below to describe the behavior of the system in a pure-heave setting, with and without a suspension.

The frequency peak in the “no suspension” model in Fig. 4.3 results in oscillations in excess of the maximum 5 mm when a 1 mm track disturbance is introduced. Alternatively, with the suspension, the system has a low gain. With this model, various spring stiffness and damping values can be explored to find the ideal dynamic behavior.

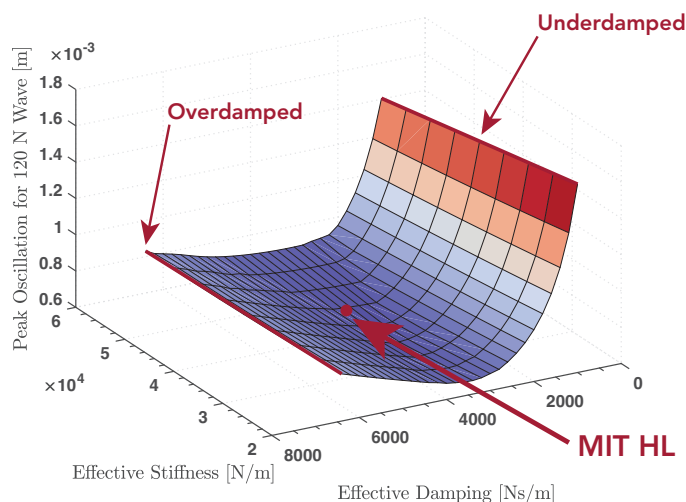


Figure 4.4: Heave-mode suspension optimization.

If a 120 N wave is introduced into the system (a simulation of the 1 mm track disturbances given the magnet gap stiffness), the peak oscillation can be minimized by changing the damper value, see Fig. 4.4. For this system, it was found that a damping value of around 3000 Ns/m would result in critical damping. These preliminary values informed future design decisions and provided a baseline for advanced modeling.

4.1.3 Mechanical Design

A modified 4-bar linkage system (see Fig. 4.5) was designed to allow for control of the heave and pitch degrees of freedom on both sides of the pod. When the pod is purely moving in heave, the suspension will move up and forward, compressing only the heave spring in the middle. When moving in pitch, the pivot on the right will change position by compressing the pitch spring, with minimal displacement of the heave spring. In a Solidworks motion simulation, a 0.1° rotation of the ski caused a 103 N reaction from the pitch spring and a 2 N reaction from the heave spring.




Figure 4.5: Modified 4-bar linkage system for suspension.




Figure 4.6: Effect of geometry on effective stiffness.

The effect of the geometry can be condensed down to a single equation that relates the stiffness of the suspension element to the effective stiffness of the system in heave,

$$\frac{K_e}{K_s} = \frac{l_2^2}{l_1^2} \frac{2 \sin \theta (\sin (\varphi + \epsilon) \cos \delta + \cos (\varphi + \epsilon) \sin \delta)^2}{\left(\frac{\sin \varphi}{\tan \theta} + \cos \varphi \right) (\sin \varphi \cos \theta + \sin \theta \cos \varphi)} \quad (4.1)$$

where the parameters are defined in Fig. 4.6.

As the angles are changed, the stiffness ratio changes as well. A design point was chosen so that the stiffness ratio remains constant with respect to changes in geometry. At this point, where $\theta = 45^\circ$ and $\varphi = 30^\circ$, the stiffness of the system should remain constant at about $3 \times$ the stiffness of the suspension

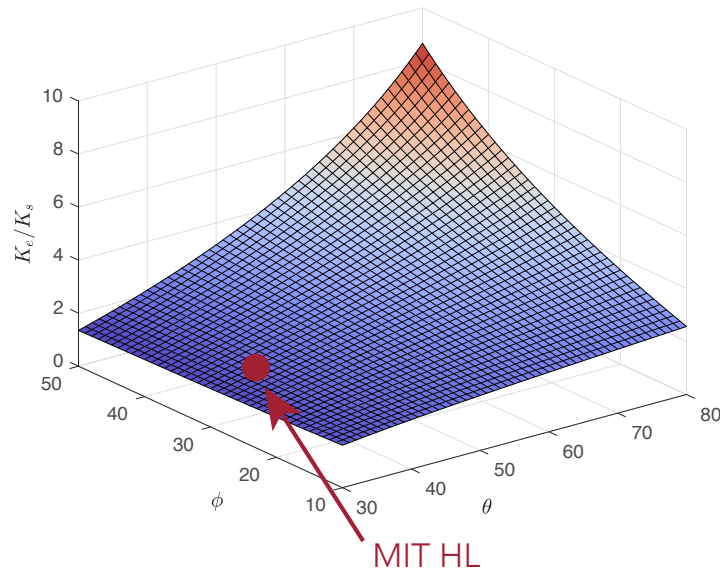


Figure 4.7: Stiffness ratio. The design point is indicated in red, it is chosen such that the stiffness ratio remains approximately constant with respect to changes in geometry.


element. All other dimensions and angles were parameterized as well to obtain a linear system within the bounds of the design.

These values were further validated with full 2D modeling in Solidworks. Using Eq. (4.1) at $\theta = 60^\circ$ and $\varphi = 30^\circ$, the resulting stiffness ratio should be 3.3. After this was established, the Solidworks model of a ski was configured to the “frame”, or mount points of the suspension, allowing the ski to dangle. A purely vertical upward load of 1250 N was applied to the ski to represent the weight of the pod pre-loading the suspension. Then a 465 N force was applied in the same fashion, as to match our worst-case loading scenario. The resulting displacement and reaction force yielded an effective stiffness of 3.2, almost perfectly matching the hand calculations.


The mechanical design of the linkages is shown in Fig. 4.8. To reduce friction, Oilite brass bushings and thrust bearings are used on all contact points. The preload in each bolt stack-up is limited by the spring, which ensures that the nut can be tightened onto the shoulder of the bolt/shaft while keeping friction low in the stack-up. The flanged bushings are press-fit into the machines holes in the linkages.

The bottom shaft (connecting the suspension to the skis) goes through two mounting blocks that provide a thick, precision surface. A stack-up is inserted in between the two ski walls to reduce slop. This bolt is then preloaded to a low torque to ensure that there is no lateral slop in the connection between the skis and the suspension. This is only necessary on this joint because it is not constrained by a precision external surface (such as is the case with the middle and upper linkage).

The bottom U-link was subjected to an FEA to ensure that it could handle a torsional load of 500 N , which is an upper estimate of the magnetic force from a small lateral movement. The factor of safety for this part was above 50. Hard stops were designed for each ski on both the front and back to the limit the relative motion between the pod and the skis. The stops are shimmed with rubber sheet to change the allowed motion. Due to the pitching analysis seen in the next section, the gap should be between 5 and 10 mm to the hard stops, such that the hard stops do not engage during launch, but do engage during braking.



(a) Linkage



(b) Exploded view of linkage

Figure 4.8: Mechanical design of linkage.




Figure 4.9: Close-up of connection from suspension linkage to magnet ski.

4.1.4 3D Dynamics Analysis

In order to choose the proper spring and damper values for the suspension elements, a 3D model of the pod was constructed in SimMechanic (Simulink) to determine which stiffness and damping values best mitigated the known track disturbances.

The model assumed that the magnet gap could be represented as an array of linear springs with a collective stiffness equal to the linearization of the force versus displacement results from magnetic levitation simulations in Maxwell. Therefore, each ski has eight springs attached to the bottom to represent

the magnets during cruise. Springs and dampers are placed in their exact location as well as the suspension linkages with the proper constraints. Thus, this model was capable of moving in roll, pitch, and heave.

To simulate the track disturbances, force inputs as a pulse wave were applied to each spring sequentially with a time delay based on the simulation pod velocity. The force pulses were then repeated every 12.5' (the length of one track plate). A square pulse wave represents the worst-case scenario, since the force is applied and removed instantaneously resulting in large upward accelerations. However, magnetic levitation data shows that this is a decent approximation of how the track disturbances would be seen by the magnets.

The displacement oscillations of the skis at the front and back end were measured as well as the vertical acceleration and displacement oscillations of the pod. For many different stiffness and damping values, an experimental sweep was done with simulations at many different pod velocities. The result was a sort of frequency response plot that was dependent on the suspension values. It also allowed for validation of the choice to include a pitching degree of freedom, as oscillations and accelerations were greatly reduced with this degree of freedom included. These results are shown in Fig. 4.10.

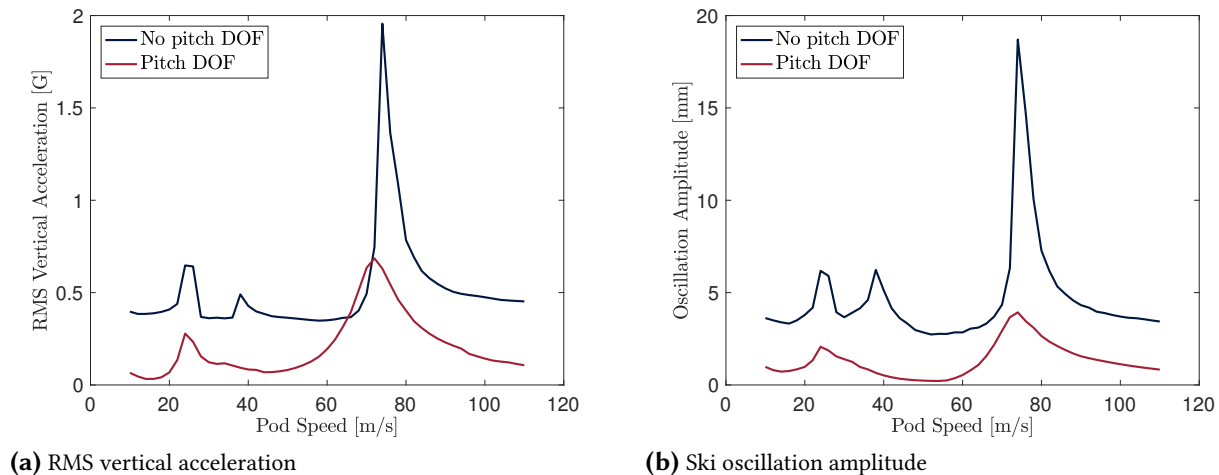


Figure 4.10: Effect of 1 mm track steps on ski dynamics.

While the RMS vertical accelerations were higher than the “comfortable” level for many velocities, with the pitch DOF they stay below the “tolerable” level for 1 G for all velocities. Similar trends exist for the ski and pod oscillations. With a suspension that included pitch, the extreme oscillation values are kept to about 2 mm, fulfilling our functional requirement.

This model also allowed us to study the effect of pitching during the Launch phase. Fig. 4.11 shows the vertical (z) displacement of the front of the pod when the 2.4 G pushing force is applied at 4 in above the center of mass of the pod. The displacement is about 10 mm.

Using this model, the optimal spring stiffnesses and damping values were chosen. A heave spring stiffness of 15 kN/m and a damping of 1 kNs/m were chosen, while a pitch spring stiffness of 30 kN/m and a damping of 2 kNs/m were chosen. Given the previous analysis, the effective stiffness in heave was calculated. These stiffness and damping values are listed in Table 4.2.

Comparison of these values to those calculated in the transfer function for heave motion suggests these values will produce a critically damped system at high speeds. It also gives confirmation that the full 3D model agrees with the hand calculations of the two-mass system. The results from both these analysis methods show that the oscillations of the front/rear of the skis will remain below 3 mm in amplitude for all speeds. The RMS vertical acceleration of the pod will remain below 1 G in worst-case scenario, which is the tolerable level for humans.




Figure 4.11: Vertical displacement of the front of the pod when a 2.4 G pushing force is applied at 4 in above the center of mass of the pod.

Table 4.2: Stiffness and damping for suspension element, ski, and pod.

| | Heave | Pitch |
|-------------------------------|------------|-----------|
| Suspension element stiffness | 15 kN/m | 30 kN/m |
| Suspension element damping | 1 kNs/m | 2 kNs/m |
| Effective stiffness – one ski | 44 kN/m | n/a |
| Effective damping – one ski | 3 kNs/m | n/a |
| Effective stiffness – pod | 88 kNs/m | n/a |
| Effective damping – pod | 6 kNs/m | n/a |

4.2 Lateral Control Modules

The purpose of the Lateral Control Modules (LCMs) is to keep the pod in the center of the track during the competition. This is specifically designed within the constraints of the competition, and for the SpaceX track. The track has an aluminum I-beam down the center of the tube resting on a flat concrete bed with two flat aluminum plates running on either side of the beam. Lateral control modules are especially important for magnetic levitation vehicles, because of their low resistance to side or yaw motion (unlike wheeled vehicles). These lateral control modules are specifically designed for the SpaceX track (i.e. no corners), and would require a redesign for sharper corners, for instance by allowing the LCMs to pivot around the z -axis to allow for yaw in turns. The resulting design, however, is scalable and could be used for a full-scale Hyperloop if there is an I-beam located in the tube. The lateral control modules placed as far in the front and rear of the pod as possible to ensure they have a large moment arm to the center of gravity of the pod such that they can effectively control the yaw mode, as already seen in Fig. 4.1.

4.2.1 Design Process

The lateral control modules were designed with manufacturability in mind. All custom parts are constructed from easily obtainable stock materials and with simple geometric features.

One of the most important questions for the lateral control modules was whether or not a contact or non-contact system should be used to center the pod within the tube. The two main considerations for this decision are the drag produced and the design complexity. Due to the overall goal of the team to focus on minimizing the travel time in the tube, we needed to reduce the amount of drag on the pod during the cruise phase of flight, including the drag from the LCMs. The four main concepts we looked at were

using skids, wheels, permanent magnets, or electromagnets. Skids are obviously easy to design and build. However, it is not a sustainable method, requiring a large number of skid swaps. Furthermore, the rail is aluminium, which is a soft material and therefore the skids had to be made from an even softer material to not damage the rail. This option would result in a lot of friction during the run.

Wheels are also easy to implement, and do not have to be replaced after each run. However, the high velocity would require either a large wheel, which is poor choice from a packaging standpoint, or at least custom wheels or bearings, which can be quite expensive. Additionally, these wheels are running against the rail the whole run which adds drag. Moreover, using wheels was against the team's design philosophy.

Permanent magnets then are complete passive, and do not have to be replaced after each run. Furthermore, they are relatively easy to understand. However, their effectiveness is low at low speeds, thereby requiring a backup method. Furthermore, maglev is notoriously underdamped which is not ideal for a stability providing mechanism.

Finally, electromagnets can control the damping and position of the system accurately. However, this option is difficult to implement, requiring substantial testing. Additionally, they have a considerable power consumption, and the system is more error-prone.

In the end, we chose to use pairs of permanent magnets at the front and back of the pod with backup Jabroc skids for low speeds. We also toyed with the idea of using permanent magnets with supplemental electromagnets to provide damping, but due to limited resources in the electronics and software team, it was decided to add a secondary horizontal suspension to the LCMs instead, shown in Fig. 4.12.




Figure 4.12: Different components of the Lateral Control Module (LCM).

4.2.2 Final Design

The lateral control module features a completely aluminum frame constructed out of readily available sized stock. Most of the holes (with the exception of the 1D accelerometer holes) are clearance sized for AN-3 and AN-4 nuts and bolts. Both vertical and horizontal slots exist between the upper C-channel frame and

the lower 90° angle bracket which holds the magnet stack and skids. This allows for adjustability in the vertical and horizontal directions with respect to the I-beam. The different components from the LCM system are shown in Fig. 4.12, and an isometric view with labeled components is shown in Fig. 4.14. The LCMs in the final prototype are shown in Fig. C.4.

Magnet Stack

The magnet “stack” refers to the 1” × 2” × 1/4” magnets, a threaded steel back-iron piece, an aluminum spacer used to get the correct magnet separation, and the magnet hat covers which protect the magnets from debris. On each LCM there are four magnets, on each side there is a north and a south facing magnet, which mirror the same pole on the other side as shown in Fig. 4.13.




Figure 4.13: Magnet configuration for Lateral Control Modules.

Skids

The skids are a mechanical failsafe system on the lateral control modules. At low speeds, the magnet centering force is essentially zero, therefore, we anticipate riding on the skids at the beginning of the run and during all operation of the low speed system. There is also a large skid that covers the bottom of the C-channel frame which prevents the pod from going into the bottom “no-go” zone on the SpaceX track. The skids are made of the jabroc wood composite used in F1 cars as a skid plate.

Shafts, supports, and bearings

There are two 3/4” diameter precision ground steel rods per lateral control module. The forward shaft has two bearings evenly spaced apart, while the aft shaft has one centered bearing. This design ensures that the bearings do not bind given the geometry of the part regardless of the loading condition. The bearings on the double bearing shaft are 2 3/4” long with 0.0015” clearance and self-alignment capabilities. The single shaft bearing is an oil-embedded, 7/8” wide sleeve bearing with 5° misorientation capability. These were chosen for their dimensions, availability, and misorientation tolerance. They are mounted on aluminum spacers to ensure that the rods are at the same height and to provide enough clearance for the rod flange mounts. The rod is attached to the ladder frame of the pod through off-the-shelf flanged mounts that clamp tightly to the ends of the rod using socket head cap screws. The shaft mounts on the shorter rod serve an additional role as hard stops to limit the displacement between the lateral control modules and the pod.

Springs and dampers

The ideal spring and damper values were determined by using two separate Simulink models: one for strictly lateral motion and one for rotational motion using a mass moment of inertia calculated by estimating each subsystem as a separate simple geometry.




Figure 4.14: Isometric view of the Lateral Control Module.

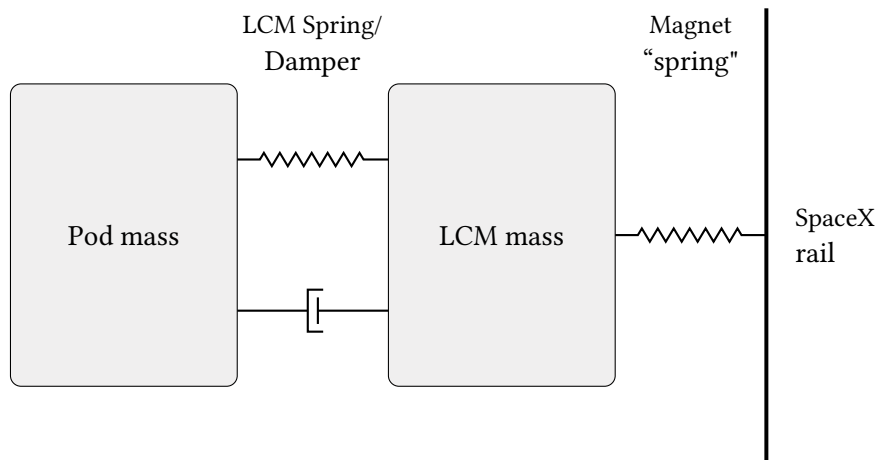


Figure 4.15: Model for LCM dynamics.

Fig. 4.15 shows a diagram for the model that was used in Simulink. The spring and damper were sized such that they damp out the lateral motion of the pod within a reasonable amount of time. Both the spring and the damper are off-the-shelf products, with the spring requiring small modifications.

Electronics

The electronics were packaged to the left side of the pod for ease of wiring. The electronics package on the LCMs consists of a gap height sensor (measures LCM to rail), a 1D accelerometer (measures impact and loads), and a potentiometer (measures LCM position relative to pod). The gap height sensor is mounted on a plate which attaches to the bottom of the LCM because this sensor has a limited operating range. The 1D



accelerometer is mounted directly to the side of the LCM to get accurate data without geometry induced measurement errors. Finally, the potentiometer is mounted with an aluminum spacer between the top of the LCM and the pod's ladder frame. All of the electronics are mounted to the same side of the LCM for ease of installation and wiring.

4.3 Brakes

The pod uses eddy-current brakes to come to a safe stop. Our brakes stop the pod by moving a large array of permanent magnets into close proximity with the central I-beam in the track. The relative motion between the magnets and the I-beam induces eddy-currents in the I-beam, which produce a magnetic field which retards the motion of the pod. Eddy current brakes allow us to bring the pod to a stop without requiring contact with the I-beam, thereby allaying concerns of damage to the beam or our pod due to high contact forces. In this section, we will discuss the functional requirements and design of the brakes.

4.3.1 Functional Requirements

The goal here is to design a pod that can run down the track in the smallest amount of time possible. To this end, we need a lightweight pod which can be accelerated to a high speed, and powerful brakes, so that we can coast for as long as possible before stopping. This design objective has dictated the functional requirements of the braking system.

The brakes are required to hit the following performance targets:

1. **Stopping rate:** 2.4 G of braking force, the maximum amount allowed
2. **Engagement time:** 0.5 second engagement time
3. **Modulation time:** We need to be able to modulate the braking force so that we can stop within a 50 ft section at the end of the track.
4. **Mass:** $\sim 20\%$ of the pod weight
5. **Rail clearance during cruise:** The brakes should not influence the dynamics of the pod when they are not engaged. They should not make contact with the rail even under maximum lateral motion during cruise.
6. **DOF constrained:** The brakes should have minimal effect on the lateral and vertical degrees of the freedom of the pod.

4.3.2 Eddy Current Brakes vs. Friction Brakes

An earlier version of the pod design had used contact based friction brakes. These brakes used a hydraulic power cylinder to clamp brass pads onto the central I-beam on the hyperloop track. This system was rejected in favor of eddy-current brakes for the following reasons:

- **Potential for damage to track:** The eddy current brakes have an inherently lower risk of damaging the I-beam since they do not contact the I-beam.
- **Difficulty of modeling:** We attempted to model the interaction between the brass brake pads and the aluminum I-beam. This included modeling of high speed impacts between the brass pads and the 1 mm steps in the aluminum I-beam.
- **Difficulty of experimental validation:** Experimentally simulating the extreme conditions present at the interface of the brass friction pads and the I-beam would have required significant design work.

4.3.3 Design Overview

The brakes can be decomposed into three subsystems: The magnet arrays, the motion control system, and the suspension. The magnet arrays allow us to generate the eddy-currents in the I-beam. The motion control system moves the magnet array relative to the I-beam, allowing us to modulate the braking force. The suspension isolates the brakes from the rest of the pod in all degrees of freedom except for the direction along the track. This is necessary to prevent contact between the brake magnets and the I-beam when the pod pitches during braking.

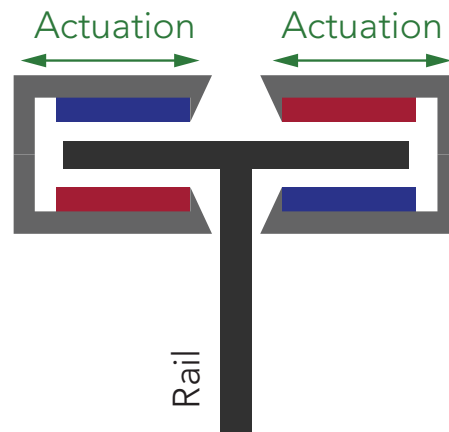


Figure 4.16: Brakes orientation with respect to rail while brakes are engaged.

To actuate the brakes, two C-channels containing magnet arrays slide laterally across the rail, as shown in Fig. 4.16. The reason for actuating the brakes laterally is that the magnets generate massive lift and drag forces, and sliding the channels orthogonal to both lift and drag requires the least force. Sliding the brake channels over the rail angled would require the actuator to move against that massive drag force, requiring a bulky actuator. The magnets are encapsulated in a sturdy C-channel, because the magnets also create a large lift force. These repellant forces from the brakes are now contained in the C-channel.

One of the two C-channels with magnets is shown in Fig. 4.17. Whereas for maximum L/D of a magnet array the wavelength is to be maximized, for maximum drag the wavelength needs to be minimized. These magnets are therefore only $1.5'' \times 1/2'' \times 1/4''$ resulting in a very short wavelength in this design, and over 200 magnets in the entire brakes system.

The brakes are suspended in their own frame – as shown in Fig. 4.18 – which allows them to freely slide laterally. The brakes therefore do not control the lateral or yaw motion of the pod, even when they are engaged. The brakes are designed to be fail-safe, such that if any system on board fails (including power), the brakes engage and the pod comes to a full stop. For this reason, the brake halves are spring loaded and a hydraulic actuator is used to open the brakes. The hydraulic system is designed to fail depressurized such that the springs close the brakes, as will be discussed in the next section. The skirts on the brakes prevent any metal or magnets from the brakes hitting the rail while they are engaged. During worst case pod motions, they also aid in sliding the channels onto the rail.

The hydraulic system is designed to open the brakes. The schematic for normal operation (fully open) is shown in Fig. 4.19a. During normal operation, the hydraulic pump and proportional valve control the position of the brakes (and opening and closing speed), where the springs are back-driven by the pump. The overflow bladder in the hydraulic accumulator remains empty. There is therefore a direct input-output relationship between the hydraulic pump and the brakes.




Figure 4.17: Magnet orientation for brakes. There are 224 magnets in both brake halves combined.




Figure 4.18: Eddy-current brakes in open and closed position. Note that the springs are not extended in this figure.

During an emergency situation (Fig. 4.19b) the on/off valve opens, the springs drive the brakes shut and thereby back-drive the hydraulic fluid into the overflow bladder.

The brakes' frame is also suspended vertically such that it does not exert a lift force to the pod once it engages. When the brakes are not engaged, the brakes rest on a hard stop on the side of the pod's L-frame, as shown in Fig. 4.20. This hard stop is adjustable such that the vertical position of the brakes inside the pod can be altered. However, when the brakes engage, the brakes move up relative to the pod and that pin inside the hard stop is now free-floating. Instead the brake force is transferred to the pod through a suspension link at the bottom of the brakes into the underside of the pod's L-frame.

The final key parameters for the brakes are listed below. The overall length of the brakes was decided by investigating the sensitivity of the brake length to the stopping distance. As shown in Fig. 4.21, an overall brake length of 24" results in a good trade-off between performance and weight.

1. **Magnets:** Neodymium N42 1.5" × 1/2" × 1/4"
2. **Brake gap:** 2 mm
3. **Overall length:** 24"
4. **Actuation time:** ~ 0.5 s
5. **Brake engagement on flange:** 1.4"
6. **Engagement force:** ~ 1000 N

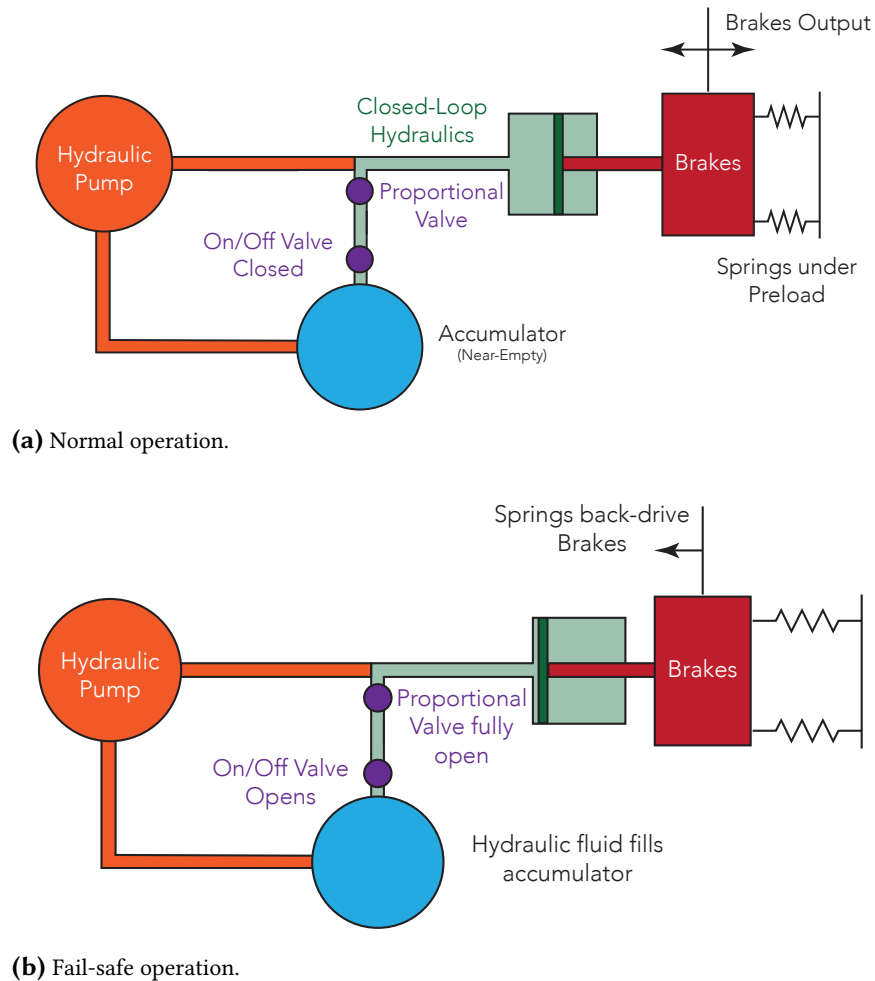


Figure 4.19: Brake hydraulics.

4.4 Low Speed System

The low speed system locomotes the pod in the tube before and after the launch. Pre-launch it may be necessary to position to pod relative to the pusher vehicle. After launch, the pod will coast for a period, then rapidly decelerate in the braking phase. After arriving safely to a stop, the low speed system is tasked with moving the pod from its stationary position to the end of the tube for extraction. The primary design challenge of this low speed system are the high magnetic drag loads experienced at low speed (and low gap




Figure 4.20: Load transfer from brakes into frame.




Figure 4.21: Sensitivity of stopping distance to overall brake length.

heights) resulting from the levitation system. The second design challenge is the rail clearance during high speed travel.

The wheel for the low speed system has to be preloaded, as the terminal velocity due to the large magnetic drag is well below acceptable values if the wheel is not preloaded. As a result, a simply wheeled drive mechanism is inappropriate, because it could not get enough grip. Higher self-propelled velocities could have been attained without clamping by simply retracting the vehicle’s lift magnets away from the track surface. However, a mid-flight magnet retraction mechanism failure would likely result in a catastrophic crash of the pod. Thus, although this approach would elegantly solve the low-speed drag problem, the potential risk posed to high-speed flight caused the team to opt for fixed magnets and a clamping low-speed drive design.

The low speed system as a mechanism is designed to pinch the flange of the rail at the point of contact with the drive wheel, and supports itself with a reaction/dummy wheel from below the flange. The low speed system extends from beside the rail. The two wheels are cantilevered above and below on lever arms, clamped together at a lesser radius by a low backlash acme screw electric linear actuator, providing the


low speed system the ability to produce great normal forces at a single point, without altering the pod suspension. Because of the tremendous torque requirements at low speed, a powerful compact motor is employed with a single stage harmonic gearbox reduction, making the single point of contact of further advantage.

4.4.1 Functional Requirements

The main requirement of the low speed system in our design is to reliably locomote the pod at the beginning and ends of travel. Because the primary objective of our team is high speed travel and rapid, safe deceleration, a light and compact design that would still fulfil the most basic requirement, was a key metric. As such, the primary requirements of the low speed system were as follows:

1. Achieve self-locomotion up to 0.5 m/s
2. Clearance during high speed travel
3. Lightweight
4. Compact

Basic performance specifications apply to self-locomotion. We want to move ourselves out of the tube reliably, and within a reasonable time scale. The low speed system is designed to drive the pod at up to 0.5 m/s for final extraction at the end of travel. The design space including drag faced and requisite power over a range of low speeds predicted from the levitation scheme are shown in Fig. 4.22.



(a) Magnetic drag force experienced vs relative speed at varying magnet gap separations. Magnetic drag falls of with increasing magnet gap and increases with magnet speed relative to the track. This is the drag experienced by the pod as a whole under expected pod weight. We expect a magnet gap of approximately 5 mm at low speed non-levitating travel.

(b) Requisite power versus magnet track relative speed at varying magnet gap separations. Power required increases sharply after 1 m/s for small separations. The system was sized to deliver $\sim 500\text{ W}$ continuous power at a speed 0.5 m/s (at a gap height of 5 mm); it is slightly over-powered to account for any potential changes to the gap/magnet configuration.

Figure 4.22: Design space for low speed system.

The low speed system must fully engage the rail in order to drive in the low speed regime, but must release contact from the rail and position itself in a safe-from-contact area dictated by suspension hard stops and the lateral control modules during high speed travel. This positional safety window for the reaction wheel below the flange is narrow because of the small space between the two flanges of the rail including the safe zones prescribed by SpaceX, and the expected dynamics of the pod during high speed

travel, including heave, pitch and sway displacement from nominal position. The low speed system must furthermore sufficiently resist self displacement from inertial loading at high speed travel.

4.4.2 Design Overview

The low speed system is most fundamentally a drive motor cantilevered on a lever arm pinching mechanism as shown in Fig. 4.23. An isometric view of the Low Speed System is shown in Fig. 4.24. The electric linear actuator produces a clamping force between the drive and reaction wheel, allowing the drive motor enough normal force to drive the pod. The Thomson DA24-05B65M was used as linear actuator for this system. The Torque Systems T0851-X0 motor is used in the low speed system. This is compact motor is mated to a Harmonic Drive gearbox (CSF-25-50-2A-GR), leading to a massive single-stage reduction.




Figure 4.23: Different components of the Low Speed System.




Figure 4.24: Isometric view of the Low Speed System.

The low speed system is mounted to the pod from the side of the frame and projects onto the rail, as shown in Figures 2.3 and 4.1. It is positioned in the front of the brakes, mostly for packaging reasons with the electronics and also to keep the center of gravity in the center of the skis.

AERODYNAMICS & STRUCTURES

One of the earliest decisions made for the Aero/Structures subsystem was to split the functions of the structural frame and the aerodynamic shell into two separate components. This allowed us to reduce the design and manufacturing risk substantially.

The structural frame consists of an aluminum ladder frame. This is an appropriate structural concept because all the major loads are in one direction – the pusher loads and the braking loads. Furthermore, the aluminum ladder frame allows for a great deal of flexibility in the design as well as during the manufacturing and testing phase, if components need to change. The aerodynamic shell on the other hand is made from carbon fiber and covers all components of the pod to reduce the aerodynamic drag as much as possible.

This chapter discusses the frame design in Section 5.1, as well as the aerodynamic design methodology and final aerodynamic design in Section 5.2, and the structural design for the aerodynamic shell in Section 5.3, and its manufacturing process in Section 5.4.

5.1 Frame Design

The main challenge in the frame design is to fit all the components in the frame while staying within weight and center of gravity requirements. Furthermore, the redesign work from Vehicle Dynamics due to rules changes also affected the design substantially.

The entire frame weighs in around 37 *kg* and the packaging is such that the center of gravity is at the center of the skis.

5.1.1 Conceptual Design

Several frame concepts were compared against the functional requirements for the system to determine which would provide the best stiffness and strength for the lowest weight, cost, and manufacturing complexity.

Several concepts were considered for the frame. First, a carbon fiber monocoque (integrated shell and frame) was considered, which would be quite a lightweight design. However, such a design increases the manufacturing time considerably, and at the same time is quite expensive. Furthermore, the manufacturing risk would be increased substantially as well. The project was already on a tight design deadline, and if a carbon fiber monocoque had been chosen, the monocoque interfaces had to be finalized quite early to allow for mold manufacturing time.

Second, a steel or composite space frame was considered. This frame can be built to place support only in areas where loading occurs, making it lighter than other designs. A big manufacturing challenge is the careful alignment of each of the tubes - many jigs would need to be made in order to do this successfully. Connecting composite tubes requires either custom tube joiners or wrapping the joints in more carbon fiber (with this option, the alignment of fibers in the joint is critical for stiffness, and not easily tested).

A box frame was also discussed as the main frame structure. In this concept, three or four large beams run along the length of the pod and provide the attachment points for the outer shell. Supporting structural beams would connect these and provide support in areas of loading. This design is heavier than the others, although it would provide substantial torsional stiffness.

For a fuselage structure, large ribs are spaced along pod length to provide radial strength and supports onto which a fiberglass shell could be mounted. Struts running the length of the pod connect the ribs and contribute to axial and torsional stiffness. The ribs could be waterjet out of large metal plates and serve as the mount location for other subsystems. This would be a good option if a pressurized “passenger” cabin is needed.

Finally, in a ladder frame two parallel rails or tubes run the length of the pod with cross supports for strength and stiffness. This concept is simple to manufacture, requires fewer jigs and welding than other concepts, and provides easy access to mounted components. However, this concept has a low torsional stiffness.

One of the major considerations when choosing a frame concept was whether or not the “passenger area” needed to be structurally protected. Due to our initial decision to have the aerodynamic and structural systems be separate, we decided to focus on minimizing weight and cost instead of creating a structurally sound passenger compartment.

We ultimately selected the ladder frame concept (see Fig. 5.1) due to the direction of our most significant loads (braking and acceleration), which occur in the axial (x) direction. As loads in other directions are small in comparison, this design allows us to place all our mass in the direction of critical loading. It is relatively easy to build and less expensive than other concepts, and provides a large number of easily accessible mounting points for other subsystem components along with easy access when the aerodynamic shell is removed. This design also allows for flexibility during testing. If we discover a component needs to be repositioned within the frame, we can fairly easily drill new holes in the aluminum, which would be substantially more difficult for a carbon fiber monocoque. Moreover, the simplicity of the frame design and manufacturing allows valuable resources to be spent on components that are more unique and important to the feasibility of a Hyperloop system.

5.1.2 Detailed Design

The frame was designed with two L-shaped rails on either side, directly above the levitation skis and suspension, as shown in Fig. 5.1. These rails are held together by two end plates, with the dummy chair mounted to the front and the Pod Receiver Interface (PRI) mounted to the rear. The end plates have a machined index to set the width of the frame and allow for proper alignment during assembly. All main frame components are $3/8$ ” stock and are bolted together with $3/8$ ” bolts. Four hoist points are included on the corners of the frame to ensure lifting above the center of gravity. Small L-brackets are positioned along the main rails for attaching the shell.

We laid out the electronic and brake components as shown in Fig. 5.2. Each component group indicated was mounted separately for several reasons. First, this modularity made it easy to test individual subsystems. Second, each subsystem required different mounting stiffness. For instance, the control PCBs were bolted to $1/8$ ” aluminum plates attached to soft rubber mounts, while the brake hydraulics were mounted to a much stiffer $5/4$ ” aluminum plate. Finally, not having one continuous platform maximized access to other subsystems located below these, such as the lateral control modules.

Fig. 5.3 shows the arrangement of the brakes, lateral control modules, and the low-speed system. The lateral control modules are attached using $1/4$ ” aluminum plates and L-brackets (hidden in the figure to not obscure the actual modules). The lateral control modules were positioned as far from the center of gravity in x -direction as possible, to give them a large enough moment. The center of gravity needs to be in the




Figure 5.1: Structural frame overview.




Figure 5.2: Electronics lay-out on frame.

center of the levitation skins to ensure there is a constant magnet gap height along the levitation skins. The brakes need to be positioned close to the center of gravity to minimize the effect of the pod pitching on the functionality of the brakes. The brakes are also the heaviest component and are therefore positioned quite close to the middle of the levitation skins.




Figure 5.3: Vehicle dynamics components relative to frame. Note L-brackets for lateral control module attachment are not shown.

5.2 Aerodynamic Design

Note that the majority of this section has been published in the conference proceedings of the 35th AIAA Applied Aerodynamics Conference.¹⁷

This section details the design approach used for the aerodynamic shell, and discusses any analysis tools used in the design. Furthermore, the main trade-offs for the aerodynamic design of the MIT Hyperloop pod are explained. We also discuss the final aerodynamic shape, its performance at competition speeds, and characterize its performance at transonic speeds.

5.2.1 Requirements

The main performance requirement for the aerodynamic shell is to keep the overall $C_D A$ for the pod below 0.5, to adhere to the overall pod run time requirement. Furthermore, the rules require that a dummy is carried in the pod “in a reasonable position” during the test. We chose to use a 3 ft dummy. Other self-imposed requirements are to keep the total weight of any aerodynamic covers below 10 kg, and to be able to access the inside of the pod within 2 minutes.

Of course, the aerodynamic shell has to cover the structural frame, while leaving enough room for the dummy, electronics, hydraulics, etc. inside the pod, as shown in Fig. 5.4. In this case, the other subsystems drive the size of the shell.

5.2.2 Design Methodology

Because there has been little research into the aerodynamic design of Hyperloop pods, let alone design methodologies for such an aerodynamic design, we first carefully considered the aerodynamic design




Figure 5.4: The aerodynamic shell covers the systems inside the pod.^a

methodology. We had to take into account an unconventional flow regime, and designed a tool to allow for straightforward parametric sweeps over the design space.

Flow Regime

The choice of analysis tools for the aerodynamics of this pod depends strongly on the flow regime the pod experiences. Even though the Hyperloop system consists of a large partially evacuated tube, the remaining air in the tube still necessitates a careful aerodynamic design of the outer shape of the pod. Therefore, the MIT Hyperloop pod has an aerodynamic shell to cover the internals of the pod. To reduce manufacturing risk, the aerodynamic shell is decoupled from the structural frame of the pod, motivated by the aggressive manufacturing timeline in this project.

Because a Hyperloop pod travels at high speed through a partially evacuated tube, it operates in an unconventional flow regime. For the SpaceX Hyperloop competition, the tube pressure is 860 Pa , and the pod will travel at 250 mph . Even though the pressure is quite low, the fluid can still be modeled as a continuum. The Knudsen number Kn is a dimensionless parameter that is typically used to describe the boundary of continuum flow, and is defined as the ratio between the molecular mean free path λ and a characteristic length scale in the flow.¹⁸ For this flow condition, the Knudsen number is $Kn = \mathcal{O}(10^{-6})$, which is still far below the continuum limit ($\mathcal{O}(10^{-1})$).¹⁹

The Reynolds number for this flow regime is around 60,000, which means that the flow will transition from laminar to turbulent flow somewhere on the pod. Capturing this transition accurately is crucial and therefore drives the selection of appropriate aerodynamic analysis tools. This Reynolds number is of the same order of magnitude as the full-scale Hyperloop pod in the Hyperloop Alpha paper which travels at a higher speed and is longer, but the pressure in the tube is ten times lower. Finally, the Mach number for the SpaceX Hyperloop competition is around 0.3, requiring compressibility of the fluid to be taken into account. For a full-scale Hyperloop at higher speeds, this Mach number is considerably higher.

As mentioned earlier, the Kantrowitz limit is not significant for this competition due to the lower speeds, and therefore no compressor is included in design analyses.

Flow Analysis

During the preliminary design phase, extensive flow analyses are carried out to characterize the design space. Relying on a 3D Navier-Stokes CFD simulation for each design tweak is simply too expensive.

^aThis render is kindly provided by Main & Partners.

Therefore, in the preliminary design phase we relied on an axisymmetric viscous/inviscid analysis code, MTFLOW.²⁰ Moreover, this viscous/inviscid analysis code allows for capturing transition of the boundary layer from laminar to turbulent accurately, which is critical to the design. MTFLOW is typically used to design axisymmetric bodies and axisymmetric flow passages. This throughflow code uses an integral boundary layer method to solve for the laminar or turbulent boundary layer and a streamline curvature formulation to solve for the inviscid outer flow. The coupling between the inviscid and viscous flow is achieved through the displacement body model. MTFLOW uses the e^n method to capture boundary layer transition.^{21,22} We used MTFLOW v2.12, which improves its accuracy of blunt trailing edges.

In order to analyze the final 3D shape of the pod we resorted to a 3D Navier Stokes solver. Unfortunately, these solvers in general are worse at capturing transition than a viscous/inviscid analysis code like MTFLOW.




Figure 5.5: Design parameters for shell geometry.

Geometry

To facilitate rapid design iterations, we parametrize the aerodynamic shell in terms of a few geometric parameters, which can be used for trade-off studies. The curvatures in the geometry are generated using Lamé curves. The geometry for the aerodynamic shell as shown in Fig. 5.5 is parameterized in terms of the overall length of the pod, nose length l_n , width of the pod w , height of the nose h_n , tail height h_t , and the Lamé parameters for the nose λ_n , the side of the nose λ_{sn} , the back λ_b , and the side of the back λ_{sb} .

The geometry shown in Fig. 5.5 is a three-dimensional geometry, whereas for preliminary design studies we rely on an axisymmetric solver. The conversion from such a three-dimensional geometry to an axisymmetric one is never perfect; for one the pod actually does not travel in the middle of the tube but travels towards the bottom of the tube. The tube is also not perfectly circular because of the concrete base. However, by using a variable axisymmetric tube radius, we can minimize the discrepancies between the axisymmetric and three-dimensional geometries. The variable axisymmetric tube radius is chosen such that the area ratio between the axisymmetric tube and axisymmetric pod is the same as the area ratio between the actual tube and three-dimensional pod. This process is illustrated in Fig. 5.6. For the axisymmetric shape we revolve the top centerline of the outer mold line on itself.




Figure 5.6: Illustration of recovering the axisymmetric geometry from the full three-dimensional geometry.

5.2.3 Detailed Design

Using the design methodology from Section 5.2.2, we first designed the axisymmetric shape of the shell and then used that to inform the three-dimensional shape of the pod.

Kantrowitz Limit Considerations

When a pod travels at transonic speeds through a tube, it could choke the flow around the pod. This happens when the Mach number of the flow around the pod is equal to 1. When the pod speeds up further, a large pressure increase in front of the pod results, because the mass flow that can go around the pod is limited. This sonic condition is known as the Kantrowitz limit.¹⁰ There are two main ways to avoid the Kantrowitz limit. The first option is to increase the ratio between tube cross-sectional area and pod cross-sectional area, thereby allowing relatively more air to pass around the pod at a lower velocity. The other option uses a compressor that sucks in air through the front of the pod and feeds the compressed air through a duct in the pod which subsequently exits through a nozzle at the back, which is the approach taken in the Hyperloop Alpha white paper.¹ Neither option is perfect, it either means limiting the cross-sectional area of the pod, hence decreasing payload or increasing tube construction costs, or it means adding an expensive, high-maintenance compressor to each pod. Furthermore, transonic compressors at such low Reynolds numbers would require a large research and development effort, because they are not in use in any aerospace application today.

As shown in Fig. 5.4, the pod does not have a compressor, because it is not worthwhile to use one




Figure 5.7: Area ratio (pod-to-tube) versus Mach number. M_{ext} is the maximum Mach number of the flow around the pod – if $M = 1$ the flow is exactly choked. The red line indicates the Kantrowitz limit.

for the SpaceX Hyperloop competition. Fig. 5.7 shows the variation of the external Mach number to the pod-to-tube ratio and the freestream Mach number. The external Mach number is the maximum Mach number of the flow around the pod, the flow is choked when that external Mach number equals 1.0. Fig. 5.7 clearly shows that for a reasonably sized pod without a compressor, the Kantrowitz limits the speed of the pod without an additional drag increase. At an area ratio of 0.3 and a Mach number of 0.3, the flow around the pod is not even close to choking, as seen in Fig. 5.7. Because compressors also have a large risk associated with them in terms of design, manufacturing, and cost, the MIT Hyperloop team decided not to use a compressor.

Axisymmetric design

First, an axisymmetric aerodynamic shell was designed using MTFLOW.²⁰ All of the results in this section are generated using MTFLOW. In the aerodynamic design we relied on sweeps over the design parameters in Fig. 5.5, rather than going for a purely numerical optimization method. The main reason for this was to gain more physical insight in this design problem with a large unexplored design space, and to use constraints that would be harder to capture in mathematical statements. Additionally, this allowed for a more aggressive design schedule.




Figure 5.8: Laminar flow separation on a Hyperloop pod at a Reynolds number $Re = 60,000$ and Mach number $M_{\infty} = 0.3$. The boundary layer and wake are indicated in gray.

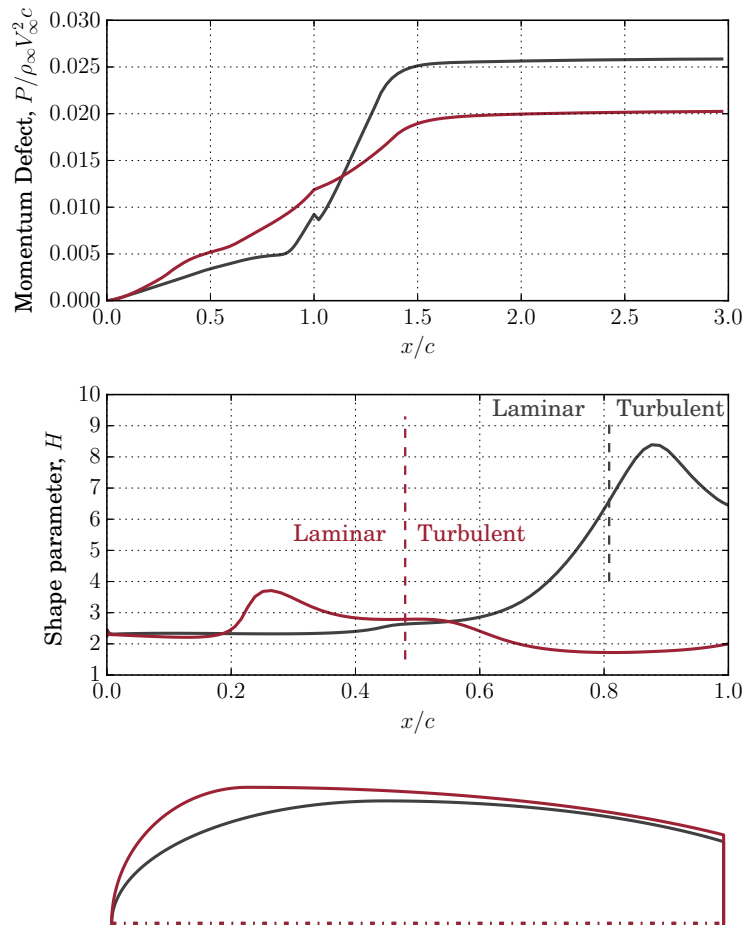


Figure 5.9: Shape parameter and momentum defect for different dummy positions. The shape in gray positions the dummy in the middle resulting in a flatter shape ($C_D A = 0.0553$), the shape in red positions the dummy in the nose ($C_D A = 0.0431$).

At these low Reynolds numbers ($Re \simeq 60,000$) the boundary layer separates very easily, and proper aerodynamic design is needed to reduce that as much as possible to keep the pressure drag to a minimum. As an example, consider the sleek axisymmetric shape in Fig. 5.8. The laminar boundary layer separates at a very low adverse pressure gradient – just after the inflection point on the geometry – and therefore this shape results in large amounts of pressure drag. Thus, even though the skin friction drag is low due to laminar flow existing on the surface, the pressure drag is high because of the large wake. This laminar separation has to be prevented in order to reduce drag. Turbulent boundary layers can handle larger pressure gradients, and a key part to this aerodynamic design is therefore to ensure that the boundary layer transitions to turbulent towards the nose of the pod.


The transition location depends on N_{crit} , which is a measure for the ambient disturbance level as well as some degree of receptivity.²³ Because we expect the environment in the tube to be fairly noisy – e.g. dust in the tube or vibrations due to track disturbances – we use $N_{crit} = 4$ throughout the design.

The dummy is the largest “component” that has to be fitted inside the shell, and its position is therefore an important consideration in the design of the aerodynamic shell. Here, we show the trade-off between two options. The first option is to put the dummy in the nose of the pod in an upright position, thereby leaving


more room for components towards the rear of the pod but resulting in a higher nose. The second option is to lay the dummy more flat over the components inside the pod, thereby reducing the height of the pod and allowing for a more gradual ramp-up to the highest point of the pod.

Fig. 5.9 shows the results for both of these shapes. For the shape with the dummy in the nose, the boundary layer transitions much closer to the nose, therefore delaying separation and reducing pressure drag. For the shape with the dummy laying flat, the laminar boundary layer separates close to the highest point on the pod, resulting in large pressure drag. Therefore, even though the shape with the dummy in the nose has a larger cross-sectional area, the drag is lower. The reason for this is the blunt nose which is known to promote early transition.²⁴

Several different sweeps over design parameters have been performed during the design stage, although only a few of them are discussed here. Fig. 5.10 shows the influence of the nose Lamé parameter and the tail height on the drag coefficient. When the nose is too shallow (e.g. $\lambda_n = 1.5$) transition will occur later on the pod and more pressure drag results. However, too blunt of a nose increases the curvature on the highest point of the nose, which also induces separation. For the tail height, the higher the tail the higher the pressure drag. However, too low of tail does not add any benefit because the flow separates anyway.



(a) Nose bluntness influence on $C_D A$ for tail height of 12 in.



(b) Tail height influence on $C_D A$ for nose Lamé parameter $\lambda_N = 2.0$.

Figure 5.10: Nose bluntness and tail height influence on drag for axisymmetric pod.

We also investigate the use of an aerodynamic tail section to reduce drag. The idea is to keep a straight section of most of the components on the pod to provide maximum payload capacity, and then add a lightweight aerodynamic tail section to keep drag to a minimum. We compare such a design to our final design in Fig. 5.11. The concept with an aerodynamic tail has a smaller cross-sectional diameter to keep the internal volume (excluding tail) similar. The momentum defect on the pod surface is much better for the pod with an aerodynamic tail because there is no adverse pressure gradient on the pod. However, the flow separates as soon as the aerodynamic tail is reached, dramatically increasing the momentum defect. The large increase in pressure drag therefore renders the aerodynamic tail useless.

The flow field around the final geometry is shown in Fig. 5.12. As mentioned several times, if we can

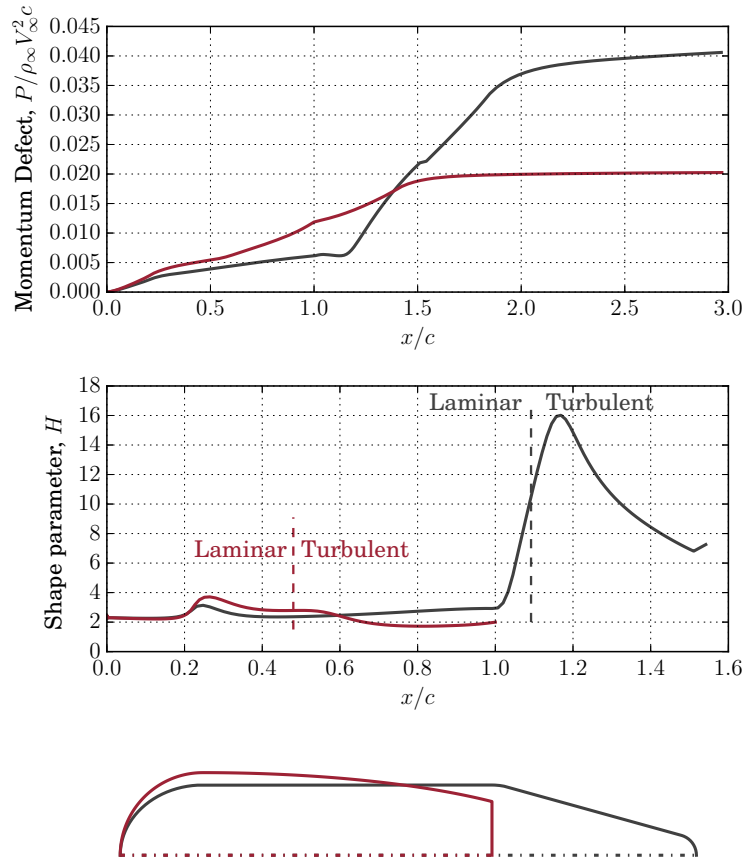


Figure 5.11: Comparison between designs *with* aerodynamic tail (gray – $C_D A = 0.0950$) and *without* aerodynamic tail (red – $C_D A = 0.0431$).

delay separation, we can dramatically reduce the pressure drag, because a turbulent boundary layer can handle adverse pressure gradients much better. Therefore, the final design has a blunt nose section which triggers transition slightly downstream of the tallest point of the pod. This results in no flow separation until the very back of the pod. The back of the pod is not tapered down further, because trying to close it completely (i.e. such that the cross-sectional area of the back is 0) is futile because the flow would separate anyway. Additionally, this shape results in clean separation off the back of the pod, which aids stability, although this is less of a problem here because the ratio of aerodynamic forces to inertial forces on the pod is quite low.

Lastly, we investigate the sensitivity of the performance of the final axisymmetric design to Reynolds number. The variation of drag as a function of Reynolds number is similar to the well-known variation of drag over a cylinder as a function of Reynolds number.²⁵ The results for low turbulence free stream are computed with $N_{\text{crit}} = 7.0$. We see that the roughness is quite important for the drag of the pod. Trying to trip the boundary layer near the nose could therefore be helpful. This could be achieved by roughening up the nose section. However, Re_θ may not be high enough to effectively use an physical trip on the nose section^b. In the end, we decided against tripping the boundary layer, because this would require detailed wind tunnel experiments to get right, and it was decided to spend valuable resources elsewhere instead.

^b Re_θ is the Reynolds number based on the momentum thickness of the boundary layer. Transition from laminar to turbulent flow is only possible when locally $Re_\theta \geq 150 \dots 250$.






Figure 5.12: Flow field around the final axisymmetric design at Reynolds number $Re = 60,000$ and Mach number $M_\infty = 0.3$. The boundary layer and wake are indicated in gray.



(a) Drag coefficient versus Reynolds number

(b) Drag versus tube pressure. Note that the tube pressure is proportional to Reynolds number, because the pod velocity and tube temperature are kept constant here.



(c) Momentum defect

(d) Shape parameter over body


Figure 5.13: Sensitivity of final axisymmetric shape to Reynolds number.

Final Three-Dimensional Shape


To generate the three-dimensional shape from the final axisymmetric shape, the axisymmetric shape is used as the centerline for the pod. The other design parameters from Fig. 5.5 are determined from packaging constraints with the other subsystems (e.g. structural frame, electronics, hydraulics, etc.). We study the performance of this design using STAR-CCM+. For these simulations we solve the laminar Navier-Stokes equations on a very fine grid (2.58 million cells). Adding a turbulence model would overestimate the transition location to be too far upstream and therefore underestimate separation to occur too far downstream. These simulations have to be unsteady because flow separation off the back is an inherently unsteady phenomenon. Total flow conditions (total temperature, total pressure) are set at the inlet to the domain and a pressure outlet is used at the outflow.

The results for the laminar Navier-Stokes simulations over the pod are shown in Fig. 5.14. We can see a small degree of vortex shedding off the back of the pod, but the overall influence on the drag coefficient is low, as shown in Fig. 5.14a. Vortex shedding is to be expected at these low Reynolds numbers.²⁶ The drag

coefficient is of course quite different from the axisymmetric case due to 3D interference effects. However, the trends that are deduced from the axisymmetric flow results are still valid. We see that the flow over the top of the pod stays attached until the back, agreeing with the axisymmetric results. The drag could be lowered by covering the skis with aerodynamic covers. However, this would reduce accessibility to critical components on the skis, and therefore no aerodynamic covers for the skis are used.



(a) Drag coefficient versus non-dimensional time \bar{t} .



(b) Velocity field contours at $\bar{t} = 48$.

Figure 5.14: Laminar Navier Stokes solution of the final aerodynamic design 110 m/s ($Re = 60,000$, $M = 0.32$).

5.2.4 Performance at Transonic Speed

Although the pod will not see any transonic speeds during the competition, we still investigate its performance at higher speeds here because it is such an important part of the Hyperloop concept. In this part of the work, we rely on turbulent 3D Navier-Stokes simulations. At these higher speeds – we increase the pod velocity, thereby increasing both Mach number and Reynolds number – turbulent flow is expected over a larger portion of the pod, and therefore using a turbulence model is more justified. We use the κ - ω SST turbulence model,^{27,28} which is known for handling separated flows and adverse pressure gradients well.²⁹

Any simulation near and above the Kantrowitz limit has to be unsteady. The choking of the flow around the pod and the subsequent pressure build-up is an inherently unsteady phenomenon, and prevent a steady simulation from converging. This also requires careful set-up of the inlet and outlet boundary conditions, i.e. specifying total pressure and temperature at the inlet, and a pressure outlet at the outflow boundary. Note that the meshes for each run are different to ensure that $y^+ \leq 1$ everywhere on the pod boundary.

The change in C_{DA} as a function of Mach number is shown in Fig. 5.15. The small drag increase from $M_\infty = 0.5$ to $M_\infty = 0.65$ can be explained by the flow turning supersonic over the highest point of the

pod (as shown in Fig. 5.17a), which results in a small shock with associated wave drag. Although part of the flow over the pod is supersonic at those freestream Mach numbers, the flow has not choked yet, because the supersonic region does not reach all the way to the tube boundaries. The flow around the pod chokes around $M_\infty = 0.675$, resulting in a large drag increase for larger Mach numbers. Due to the fact that not all flow can continue past the pod, a pressure build-up in front of the pod results. Fig. 5.16 shows the pressure coefficient along the tube, 1 m above the pod, which clearly shows a large pressure increase in front of the pod for freestream Mach numbers higher than 0.675. The drag build-up due to exceeding the Kantrowitz limit is significant, the drag coefficient is three times as high for $M_\infty = 0.8$ compared to $M_\infty = 0.65$. Note that part of the drag increase is also the result of the wave drag increase due to the associated strong normal shock, as shown in Fig. 5.17c. Furthermore, the shock-induced boundary layer separation for $M_\infty > 0.70$ also increases the pressure drag.

The drag increase due to the exceeding the Kantrowitz limit is substantial: a three-fold increase in $C_D A$ between $M_\infty = 0.65$ and $M_\infty = 0.80$. That additional drag increase results in a power loss of 31 kW. However, it is questionable that a compressor used to avoid exceeding the Kantrowitz limit could compress the air to feed it through the pod for less power. For example, the Hyperloop Alpha concept's first stage compressor has a power requirement of 276 kW,¹ and Chin et al. found that for their configuration the compressor power requirement exceeded 300 kW.⁵ For our configuration, if we assume an isentropic efficiency for the compressor of 80% and a duct^c area of 0.05 m², the power requirement for the compressor is 204 kW at $M_\infty = 0.80$. Of course, if the power requirement for a compressor to avoid the Kantrowitz limit is higher than the power loss due to the additional drag from exceeding the Kantrowitz limit, adding a compressor would be futile.




Figure 5.15: $C_D A$ as a function of Mach number for the final design. The large drag build-up starts around $M_\infty = 0.675$. Note that the freestream density is kept constant, therefore the Reynolds number increases proportionally with Mach number.






Figure 5.16: Pressure coefficient at 1 m above the base of the shell along tube for different Mach numbers.

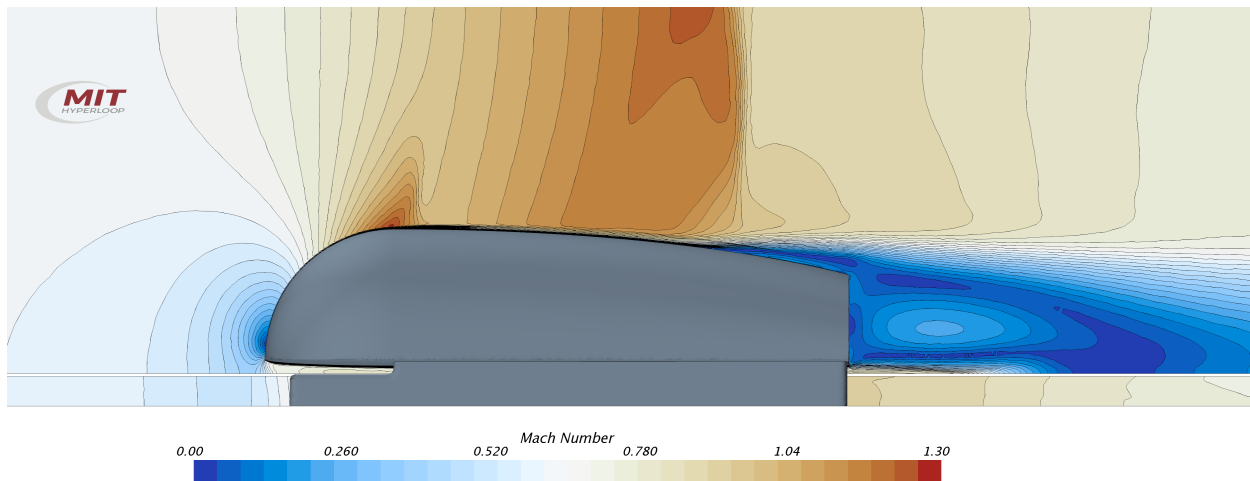
^cThe duct feeds the compressed air from the compressor outlet to the nozzle at the back of the pod.



(a) $M_\infty = 0.65$, below Kantrowitz limit.



(b) $M_\infty = 0.675$, below Kantrowitz limit, a small shock develops.



(c) $M_\infty = 0.70$, above Kantrowitz limit.

Figure 5.17: Contours of local Mach number around the pod for different freestream Mach numbers.

5.3 Shell Structural Design

The aerodynamic shell needs to be lightweight, because it has barely any forces acting on it, favoring either a fiberglass or carbon fiber shell. Furthermore, the intricate shape of the aerodynamic shell favors such materials. It was decided to use carbon fiber as the material for the shell.

To allow for easy shell removal and access to the dummy, the shell was designed into two separate parts, with the parting line located at the intersection of the curved nose and the flat bottom at the front. This separation also greatly simplifies manufacturing, as the sharp corners of that intersection and enclosure of the bottom front section of the shell would require a deep mold that would complicate layup part removal. In addition, the weight penalty from having the piece be made of two separate parts (overlapped laminates and fasteners) was insignificant compared to the advantages and fulfillment of requirements mentioned above.

5.3.1 Materials

The structure of the shell was designed through iterations in mechanical analyses. Key requirements included maintaining the Outer Mold Line (OML) geometry as previously defined in aerodynamic analyses, withstand maximum aerodynamic loads and handling/transport loads, and stay within weight budgets. Filamentary composite materials were selected to be able to meet these considerations. In particular, aerospace grade Carbon Fiber Reinforced Plastics (CFRP) were selected given their high mass-specific strengths and stiffness properties, which result in high performing lightweight structures as is increasingly being pursued in the aerospace industry (e.g. Boeing 787, Airbus A350). In addition, filamentary composites offer the manufacturing advantage of creating a thin structure that can follow the defined complex OML curves.

For this shell, the Hexcel AS4 plain weave fiber system was selected given its availability, high strength fibers (4.6 GPa), and wealth of characterizations/literature. The Pro-set INF 210 resin system was selected for its low viscosity, desirable modulus, and flexibility in offering a room temperature resin infusion or wet layup in fabrication, which has proven to be advantageous during fabrication.

5.3.2 Ply Layup Schedule

Multiple laminate designs were considered for the creation of a lightweight stiff shell structure. Two primary competing ideas include a monolithic laminate structure with both longitudinal and transverse CFRP stiffeners, or a sandwich structure laminate with a lower density core. However, the heterogenous stiffness of stiffener-reinforced monolithic laminates combined with the additional manufacturing complexity of fabricating/bonding stiffeners (e.g. intersecting hollow top hat stiffener interfaces) rendered the sandwich structure advantageous. Moreover, a sandwich structure offers a significant increase in the stiffness by placing the high modulus carbon fibers away from the neutral plane to carry extensional stresses, while the lightweight foam core served to carry shear. The uniform increase in the shell stiffness also allows for easier handling, transport, and mounting to the vehicle, thus outweighing the small weight penalty of a low-density core inclusion. A $1/4$ " 5 pound-per-cubic foot (pcf) scored Divinycell H80 PVC foam core was selected for its availability, good shear properties (27 MPa shear modulus and 1 GPa shear strength), drapability into the complex curved geometries, and compatibility with the epoxy resin system chosen.

The sandwich structure was first designed to be balanced and symmetric – with equal number of plies stacked at angles symmetric to the foam core to minimize any warping upon part removal. A Quasi-Isotropic (QI) sequence was selected for a uniform reinforcement in all directions and to minimize any bending-twisting coupling. 0° and 90° fibers were selected to be on the outside of the laminate (further away from the neutral axis) to carry more of the hoop and longitudinal loads on the shell, while $\pm 45^\circ$ fibers

were alternated in between to increase in-plane shear stiffness desirable for load cases such as twisting of the shell from asymmetric handling. As a result, a minimum QI sandwich structure that satisfies the considerations above would entail a 2-core-2 $[0/90 \pm 45 \text{ core}]_S$ layup schedule. In terms of manufacturing, a 2-core-2 laminate is also more amenable to infusion over a 1-core-1 due to the availability of open flow channels along the fabric crimps, and thus allows better permeability for resin flow to minimize void formation and dry spots.

For the shell areas that require fastening to the frame, a monolithic laminate is more attractive for the higher compressive stiffness and resistance to hole edge delaminations. Therefore a 45° core pandown was designed two inches from the bottom of the shell to form a reinforced lip region, with the 2-core-2 laminate transitioning to a 6-ply QI laminate. The outer and inner two plies of this lip region are continuous with facesheets of the sandwich laminate to minimize delamination initiation on the surface of the layup transition. Two $0^\circ/90^\circ$ degree plies were inserted in the middle between the facesheets to form the symmetric and balanced monolithic laminate throughout all bottom edges of the shell.

5.3.3 Structural Analysis

Micromechanical models were used to predict the carbon fiber/epoxy ply-level properties as inputs in a finite element analysis of the shell structure under maximum load cases, such as ambient tube pressures in the event of a preliminary test run. Microscale properties of the carbon fibers were first obtained through manufacturing testing data, with Hexcel AS4 fibers having mean strengths of 4.6 GPa and stiffness of 231 GPa , and epoxy having a stiffness of 3.1 GPa . An estimate of 40% fiber volume fraction was used, which is typical of woven aerospace laminates. By knowing the density of the fibers, a 197 gsm Hexcel AS4 fabric was estimated to have longitudinal and transverse stiffness of 48 GPa (plain weave with same warp and weft fibers).

The shell's structural properties are analyzed using Abaqus. The CFD results are used to get the correct loads on the structure. The shell is analyzed for the structural design case (full speed – 250 mph at $1/10$ th of an atmosphere – 10133 Pa) and for a handling case, where the shell is pulled up from its six suspension points.

The success criteria for the structural design was defined here to limit deformation to 1 mm as is consistent with the anticipated clearance between the inner lip of the shell and the bottom closeout, and to ensure no failure at the bolt location. For the $1/10$ atmosphere structural design case results in concentrated stresses close to the first holes towards the front of the vehicle with a peak stress of 14 MPa . However, these values are significantly lower than the anticipated failure strength, even with a large factor of safety of 2. If we assume locally only the half of the fibers are engaged in the direction for failure, then an allowable stress of $1/2 \cdot 986.2$ (unidirectional strength) = 493 MPa for the CFRP can be estimated. The Margin of Safety (MS) is then,

$$MS = \frac{\sigma_{\text{allowable}}}{\sigma_{\text{limit}} \cdot SF} - 1 = \frac{493}{2 \cdot 14} - 1 = 17.$$

It is also important to note that even if a B-basis allowable (generally in the range of 80% of the mean strength value depending on scatter) was used here, the MS would still be 13, and thus we do not anticipate failure for the shell under $1/10$ th atmospheric loading at top speed.

The maximum deflection of 0.5 mm at the nose of the shell is also well within the 1 mm limit as outlined before, and thus the structural design fulfills the intended bounding operating load case.

5.4 Shell Manufacturing

Manufacturing the shell also posed some key challenges, most notably the two month build time constraint and a lack of access to an autoclave system. Therefore, an out-of-autoclave process was utilized. To

yield a high quality part with low void fractions and high mechanical properties (consistent with material allowables used in design) we employed a close mold composite manufacturing technique here called Vacuum Assisted Resin Infusion (VARI). In contrast to a typical wet lay-up, VARI was implemented as follows

1. dry woven carbon fabrics were draped onto a female mold,
2. peel ply and porous distribution flow media were placed on top of the fabrics,
3. the whole assembly was bagged and sealed,
4. vacuum was pulled onto the part, and resin was pulled through the part.

5.4.1 Mold Design

MDF sheets were chosen as the mold material, to be cut using a CNC router. A direct female composite mold was chosen to minimize intermediate steps, and to create the tool side on the outer mold line of the part for higher surface quality. It is important to note, our shell design is deeply coupled with manufacturability, and the selection of MDF sheets also sized the length of the shell to be under 8 *ft* since the stock sizes of MDF sheets are 0.75" \times 4' \times 8'. To offer flexibility in part release, a total of three separate molds were designed. One for the top shell, the bottom closeout, and the rear cover.

It was decided that the creation of the rear cover as a separate piece to be glued in place as advantageous to allow easier access from the rear of the mold to pop and slide the part off without damaging the part. In addition, having a separate piece allows for a 0° draft angle (to conserve longitudinal space) and sharper corners at the rear edge of the shell (a minimum corner radius on the order of 0.5" would be needed to allow for fabrics to drape properly into the female mold without encountering fiber bridging).




Figure 5.18: The shell mold is built up using 34 MDF sheets ($\frac{3}{4}$ " \times 4" \times 8").

The shell mold is shown in Fig. 5.18. Our sponsor Whitecap Composites graciously allowed us to use their CNC machine to cut the mold, and also allowed us to conduct the layup at their facility. To drastically reduce cut time, the strategy is to perform 2D outline cuts on one MDF ply at a time with alignment features (holes) such that only the contour corresponding to each $\frac{3}{4}$ " thick slice is cut through the sheet and the center island removed. After each sheet was cut, they were stacked, glued and aligned with a steel dowel. The mold surfaces were then sanded with 40 and then 80 grit sand paper until no ridges could be seen.

5.4.2 Layup and Infusion

The layup and infusion consisted of three distinct steps. For the first step, two layers of carbon fiber are placed in the mold (first layer $0^\circ/90^\circ$, second layer $\pm 45^\circ$). After this, peel ply is placed in the mold, as well as green mesh and spiral tubing to ensure the resin can easily flow through. After vacuum bagging the mold, the part was pulled vacuum at 25 inHg , before starting the infusion process. After the infusion process was finished, the part was kept under vacuum for 24 hours to cure.

For the second step, slices of foam core sections are trimmed and fit into the interior. A mix of epoxy and silica particles were then applied to the faces to glue to the cure. This whole setup is vacuum bagged yet again, and held under vacuum overnight to allow for curing.

The final step of the infusion process is the same as the first, by placing the $\pm 45^\circ$ first and the $0^\circ/90^\circ$ second, and then infusion the part again following the same procedure as before.

Before sending the shells to the paint shop, we still postcured the shell. Since we are using a room temperature curing epoxy resin, its inherent glass transition temperature is still quite low ($\sim 60^\circ\text{C}$). To safeguard the shell from warping in the competition environment (anticipated that the sun may result in large temperature swings), we need to postcure the part inside of an oven for an 8 hour time period at 80°C .

CHAPTER 6

ELECTRONICS

This chapter describes the design of the electronics system for the pod. First, the requirements are discussed, then the sensing and communication systems are discussed in Section 6.1, followed by the power systems in Section 6.2. Finally, Section 6.3 focuses on the wiring harness. The electronics systems are placed on top of the frame for easy access, as shown in Fig. 6.1.

6.1 Sensing & Communication

Our pod has over 35 sensors and two actuator systems (Braking and Low Speed). Reading from these sensors, making decisions, and controlling the actuators is the function of the sensing and communications electronics.




Figure 6.1: Electronics packaging on pod.^a

^aThis render is kindly provided by Main & Partners.

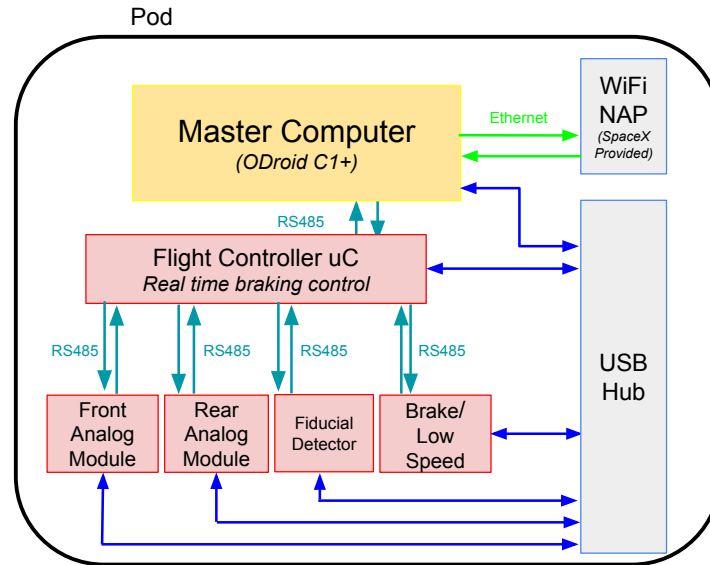


Figure 6.2: Structure of sensing & communications electronics.

The basic structure of the sensing and communications electronics are summarised in Fig. 6.2. This diagram shows seven important blocks:

1. Master Computer^b
2. Flight Controller
3. Front Analog Module
4. Rear Analog Module
5. Fiducial Detector
6. Brake/Low Speed Controller
7. Network Access Panel (provided by SpaceX)

Custom PCBs were made for each of the first six blocks.

Using one microcontroller to read from all of these sensors was not possible in our case. Firstly, there are too many sensors, and therefore there are not enough analog, SPI, UART ports on a given microcontroller for all of our sensors. Secondly, this provided additional flexibility during the design process. Moreover, state estimation can require a substantial computing effort, therefore it is better to have a dedicated microcontroller for this. Finally, before using the fiducial sensor, we were unsure how complicated reading from it would be. To be safe, we decided to have a dedicated microcontroller. Summarizing, we have five microcontrollers that are each responsible for different groups of sensors and actuators.

6.1.1 Computer Selection

The Master Computer's needs to have at minimum support for the Ethernet communication protocol, three USB ports, and needs to run a dedicated OS (Linux or Android). An ODROID-C1+ was selected as the master computer. It is a single-board computer – RAM, SSD/HD, microprocessor, graphics card and ports are all on the same circuit board.

^bThe Master Computer is also sometimes referred to as "Flight Computer".



To simplify the design and reduce the number of PCBs, we decided to use the same microcontroller for each of the remaining modules. The requirements of these modules informed us of what kind of microcontroller we needed. The requirements for the microcontrollers are shown in Table 6.1.

Table 6.1: Requirements for microcontrollers.

| Module | Requirement |
|----------------------------|---|
| Flight Controller | Convenient to program state estimation: <ul style="list-style-type: none"> • Floating point arithmetic • Good math libraries that are easy to use • Can connect to 2 IMUs |
| Front Analog Module | <ul style="list-style-type: none"> • Sampling rate of at least 60Hz. The expected dynamics on the pod are at a maximum of 30Hz – this comes from the natural frequency of the pod’s magnetic-spring mass suspension. • I2C comms to read 36V battery data |
| Rear Analog Module | <ul style="list-style-type: none"> • Sampling rate of at least 60Hz – any microcontroller can do this • At least 10 analog input ports • I2C comms to read 24V battery data |
| Fiducial Detector | <ul style="list-style-type: none"> • 10 Digital I/O pins for talking to 2 of our chosen fiducial sensors. • 2xRS232 ports for the 2 fiducial sensors |
| Brake/Low Speed Controller | <ul style="list-style-type: none"> • 2xRS232 ports to talk to the motor controllers • PWM outputs to control pump speed and proportional valve • 8 digital I/O pins for setting various relays on the brake interface PCB |

Thousands of microcontrollers satisfy the above requirements. The deciding factor was ease of use, minimizing the burden on the software engineers. The Arduino is larger than a bare microcontroller integrated circuit. Relative to our pod, the Arduino is still quite small so this was not necessarily a problem here. The Due model was chosen, because it satisfies each of the module requirements listed in Table 6.1, team members had previous experience with this microcontroller, and there is more computation power available for state estimation.

6.1.2 Module Description

This section describes each of the modules in Fig. 6.2. Table 6.2 holds the legend to the schematics on the next few pages.

1. Master Computer

This module – see Fig. 6.3 – is the bridge between the remote laptop and the on-board microcontrollers. The data flows from the pilot laptop (i.e. the laptop we use to monitor and control the pod remotely) to the Odroid in the following path: Laptop → WiFi → SpaceX NAP → Ethernet → Odroid.

Table 6.2: Legend for Figures 6.4 to 6.6, 6.8 and 6.18.

| Abbreviation | Description | Abbreviation | Description |
|--------------|---------------------------|--------------|------------------|
| GH | Gap Height laser | A1 | 1D Accelerometer |
| IMU | Inertial Measurement Unit | C | Camera |
| T | Temperature | OF | Optical Fiducial |
| P | Potentiometer | ME | Motor Encoder |
| V | “Vacuum” pressure sensor | | |

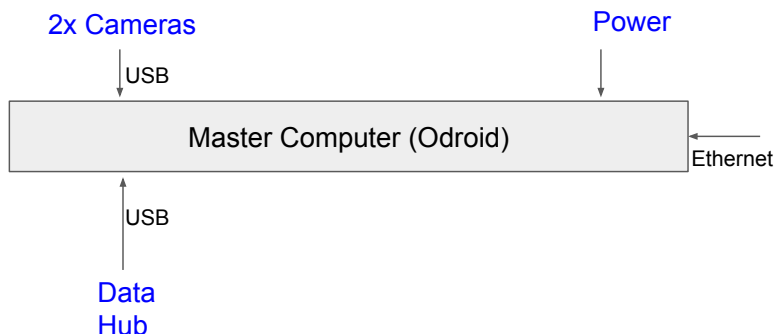


Figure 6.3: Schematic of Master Computer module.

The Odroid is also connected to the USB port of each Arduino via a USB Hub. This means that the Odroid can flash the Arduino with code updates. The USB port can also be used to stream data. The front and rear cameras use two USB ports on the Odroid.

2. Flight Controller

This module – see Fig. 6.4 – is responsible for the most important decision during a flight - when and how to brake. Its sole job is to perform state estimation so that the pod knows where it is in the tube.




Figure 6.4: Schematic of Flight Controller module. Legend is shown in Table 6.2.

3./4. Front & Rear Analog Module

The Front and Rear Analog Modules – see Fig. 6.5 – read data from non-critical sensors (ie. not involved in the decision to start braking). These are used for monitoring purposes.

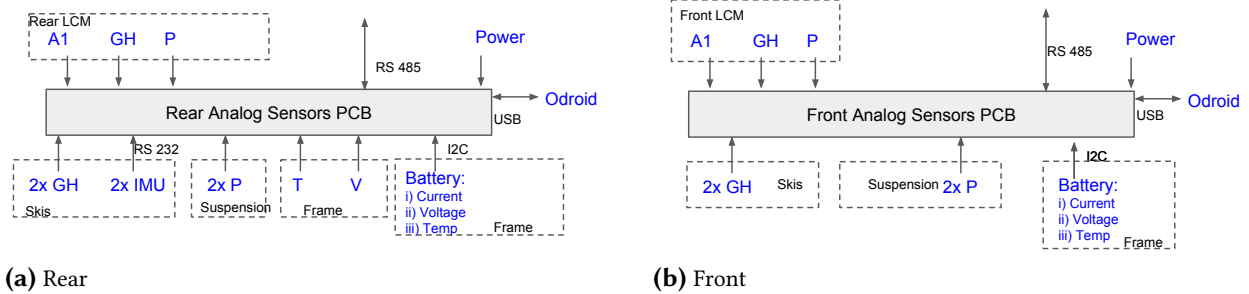


Figure 6.5: Schematic of Analog modules. Legend is shown in Table 6.2.

5. Fiducial Sensor

The Fiducial Detector – schematic shown in Fig. 6.6 – is used to measure how far the pod has traversed down the tube. The inner walls of the tube are lined with 4" retroreflective orange tape, spaced every 100 ft. By counting the strips we can tell how far the pod has traveled, more information about the state estimation can be found in Chapter 7. The sensor shines light against the walls and reads the reflected signal, as illustrated in Fig. 6.7.

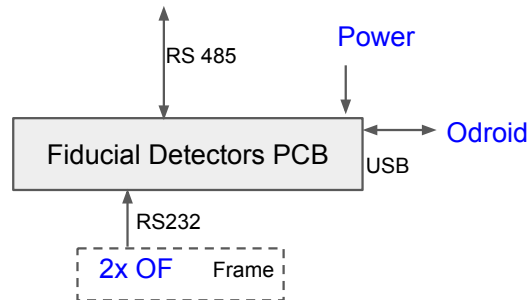


Figure 6.6: Schematic of Fiducial Sensor module. Legend is shown in Table 6.2.




Figure 6.7: Illustration of the state estimation using the fiducial sensor (not to scale).

6. Brake/Low Speed Controller

This module – see Fig. 6.8 – is responsible for commanding all actuators on the pod. It also reads from sensors related to the braking system.

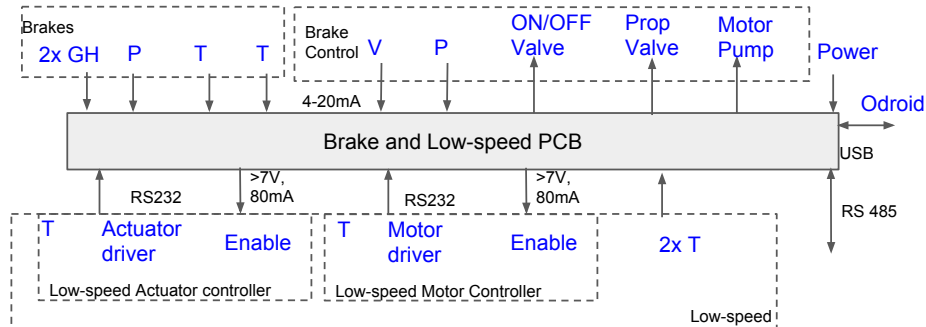


Figure 6.8: Schematic of Brake/Low Speed Controller module. Legend is shown in Table 6.2.

The braking system has three hydraulic actuators: an on/off valve, a proportional valve, and a pump. The other actuators are electromechanical and relate to the low speed system: the low speed actuator engages the wheel with the rail, and the low speed motor rotates the wheel in either direction.

6.1.3 PCB Design

The 6 modules mentioned in the previous section all use the same Arduino Due microcontroller (except the Master computer Odroid). To simplify design and save a lot of time, we decided to design one PCB to be used for 5 different modules. We refer to this as the “Main” or “Generic” PCB. The PCB schematic for the Main PCB is included in Appendix B.

Notable features of the Main PCB:

- 2x Isolated RS232
- 2x Isolated RS485
- Non-Isolated I2C
- 2x LEDs
- 7x 4-20mA receivers
- 4x 0-3V3 receivers
- Analog power supplies: 1.6V, 3.3V, $\pm 5V$, 24V (Isolated 40W 24V converter)
- Digital power supplies: 3.3V, 5V, 9V
- Scaling for 24V digital output of the Fiducial
- Scaling Arduino outputs to 5V for the brake PCB interface
- Differential to single ended conditioning for the 1D accelerometer

When constructing the ground planes on the PCB, we tried to keep analog and digital signals flowing in separate planes - in order to reduce noise to the analog signals. This was most likely overly conservative, since we sample the analog signals at ~ 100 Hz (very low frequency). Therefore we would not have noticed high frequency noise anyway.

One issue that came up during the October testing weekend required a fix to these PCBs. During test weekend, a 24V wire accidentally shorted to the shield inside a connector. I2C isolation was added to

prevent ground loops and keep digital and power ground completely separate, the schematic of which is shown in Fig. 6.9.

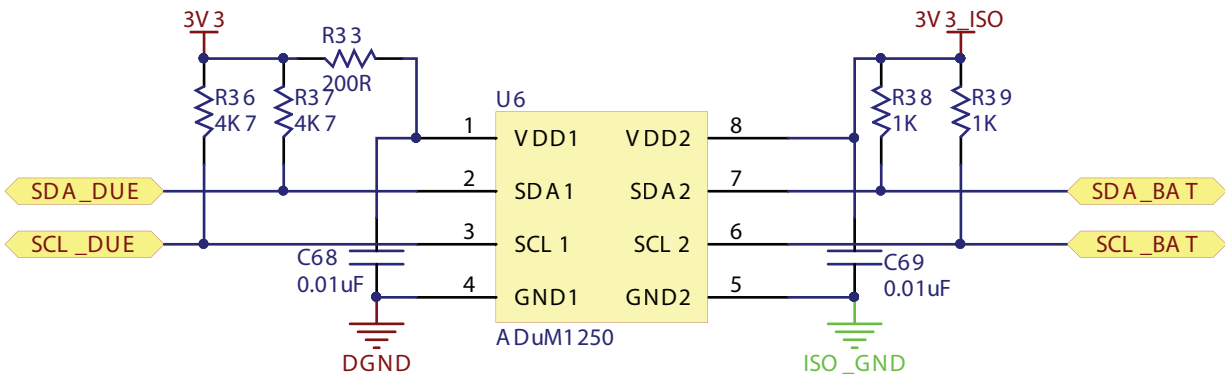


Figure 6.9: PCB schematic for I2C.

Finally, the ODROID PCB is quite simple. It contains a 24/5V converter to power the ODROID and the USB hub, isolated RS485 communication, and mounting holes for the ODROID standoffs. The PCB schematic for the ODROID PCB is included in Appendix B.

6.1.4 Shielding & Grounding

There are two isolated grounds on the pod:

- PGND (Power/Chassis ground) - where the 24V and 36V power supplies are grounded at one point (to prevent ground loops).
- DGND (Digital ground) - all PCBs are connected to a single digital ground via the USB.

The enclosures rest on the chassis of the pod, which is connected to PGND. We decided to leverage this when grounding the shields of the cables.

The shielding is connected to the connector heads. If the enclosure conducts, then that provides a path from the shield to the chassis. All aluminum plates on which the enclosures rest are also electrically connected to the rest of the chassis (even if they are mounted on dampers).

Silver and copper conductive spray were sprayed on the outside of the enclosures. Continuity checks showed that the shields were indeed connected to PGND via the sprayed enclosures.

Shielding is particularly important for wiring that carries motor controller currents, as well as signal wiring *near* motor wiring.

6.1.5 Sensor Selection

The pod relies on an array of sensors necessary for not only for telemetric sensing and data collection, but also for active control of the pod. More than 35 sensors collect information about pod position, velocity, acceleration, roll, pitch, yaw, pressure, temperature, power and ride height. All sensors interface with five PCBs that accommodate multiple analog and digital output sensors, and send data to the master computer (the ODROID) for health monitoring and piloting purposes.

Most of our analog sensors use the 4-20mA output standard. This is an industrial standard used since the 1950's. There are two types, 2-wire and 3-wire. 2-wire has 2 pins: Power and Analog Output. 3-wire has an additional pin, which is a ground connection. The advantages for this type of sensor are,

- Robust to resistance of wire (i.e. length of wire),
- Robust to supply voltage change,
- Much simpler than digital protocols like I2C, SPI, UART,
- Fault detection. If 0 A are flowing, something is wrong, either the loop is broken, sensor is connected incorrectly, et cetera,

while the disadvantages are,

- Cannot be read directly by ADC – sense resistor is required,
- Maximum of 1 sensing variable per loop,
- Need to be careful to avoid ground loops when using more than one of these sensors.

Here, we could work around those disadvantages without much difficulty, and therefore selected the 4-20mA output standard.

In the following, we describe all the sensors used on the pod, starting with the sensors for telemetry.

Sensors for Telemetry

Gap height sensor – Wenglor OPT2001

All eight gap height sensors located on the pod are used for telemetry. Two are situated on each ski and on the braking system and one on each LCM (front and back). The gap height sensor had to be able to measure gaps of 20-50 mm with resolution of 100 μm . In accordance with these requirements, the Wenglor OPT2001, a diffuse laser distance measurement sensor that has a sensing range of 30-80 mm , $< 8 \mu m$ resolution and measurement frequency of 1500 Hz , is used. The sensors operates under 18-30 VDC while outputting 4-20 mA to the front and rear analog PCBs (three sensors each), and the brake and low-speed PCB (two sensors).

Linear Variable Inductive Transducer – Omega LDI-119-050-A020A

One type of Linear Variable Inductive Transducer (LVIT) can be found with the skis (4) and the LCM (2). These sensors are also sometimes referred to as potentiometers.

Pressure Sensor – Omega PX459

An Omega PX459 vacuum range pressure transducer was selected to track the pod's environmental pressure changes. This sensor can detect pressures between 0 to -15 psi while outputting 0 to 5 VDC and taking in 10-30 VDC.

1D Accelerometer – Dytran 7504A2

Each LCM has one 7504A2 1D accelerometer. These sensors have a 5 G range with 800mV/ G sensitivity and operate with 5 to 28 VDC input. Each sensor sends data to rear and front analog PCBs.

Sensors used for Active Control

Linear Variable Inductive Transducer – Omega LDI-119-150-A020A

Two of these LVITs are involved in the feedback control of the brakes. The LDI-119-150-A020A takes in 18-30 VDC, outputs 4 – 20 mA , and has 12.5 μm resolution. Its update rate is 300 Hz , and its stroke length is 150 mm .



Inertial Measurement Unit – VectorNav VN-100 Rugged

These are integral to the pod’s state estimator for active control. Two of them, located on the front and back electrical panel, deliver data necessary for piloting the pod. The IMUs located on each ski are strictly for telemetry. The telemetry IMUs are connected to the rear analog sensor PCB and the remaining IMUs are connected to the Flight Controller.

Optical Fiducial – Micro-Epsilon OT-3-LD-500-23

The Micro-Epsilon OT-3-LD-500-23 fiducial sensor tracks fiducial markers at sampling rates greater than 4 kHz at 250 mph. This is used as an absolute position reference during the flight.

6.2 Power Systems

A full diagram of the power system of the pod is shown in Fig. 6.10. In this section, we address the design decisions and considerations on practical implementation on the following four parts:

- Battery selection and testing
- Brake system related electronics and connections
- Low speed system related electronics and connections
- Power system protection and connection (including grounding, relays, switches, etc)

6.2.1 Battery Selection and Testing

The battery selection process was focused on finding a battery that functions safely in vacuum, provides enough power and has the capacity for 2-3 runs at the minimum, is lightweight and compact, affordable, and easy to interface with.

Specifically for our pod, by using a passive levitation scheme, we eliminate the use of bulky and power-hungry coils, thus largely reducing our power needs from the batteries. The high power system is mainly used by the brakes and low speed system. The high power system runs on 36V, such that components draw a lower current than at 24V, and more importantly there are many available off-the-shelf motors at this voltage.

The low power system is used for onboard computers, Arduinos and sensors. 24V is a suitable level because many sensors are rated at this voltage. The requirements on the power and the capacity are shown in Table 6.3.

Table 6.3: Power and energy requirements for 24V and 36V systems.

| Parameter | Computers & Sensors | Proportional valve | On/Off valve | Low speed actuator | Low speed motor | Brake actuator |
|---------------------|---------------------|---------------------|-------------------|---|-------------------|------------------|
| Voltage | 24 V | 24 V | 24 V | 36 V (PWM to 24 V) | 36 V | 36 V |
| Peak Power | 50 W (2.1 A) | 14.4 W (0.6 A) | 48 W (2 A) | 300 W (8.3 A) | 974 W (27 A) | 180 W (5 A) |
| Average power | 50 W (60 mins) | 14.4 W (5 mins) | 48 W (5 mins) | 300 W (1 min) | 500 W (20 min) | 60 W (30 sec) |
| Energy capacity | 50 Wh (2.1 Ah) | 1.2 Wh (0.05 Ah) | 4 Wh (0.17 Ah) | 5 Wh | 167 Wh | 0.5 Wh |
| Battery selection | 24 V, 10 Ah, 750 W | | | 24 V, 10 Ah, 1620 W | | |
| Total peak power | 112.4 W (4.7 A) | | | 974 W peak (27 A, only one of three on at once) | | |
| Total average power | 50 W (2.1 A) | | | 500 W (14 A consider low speed motor only) | | |

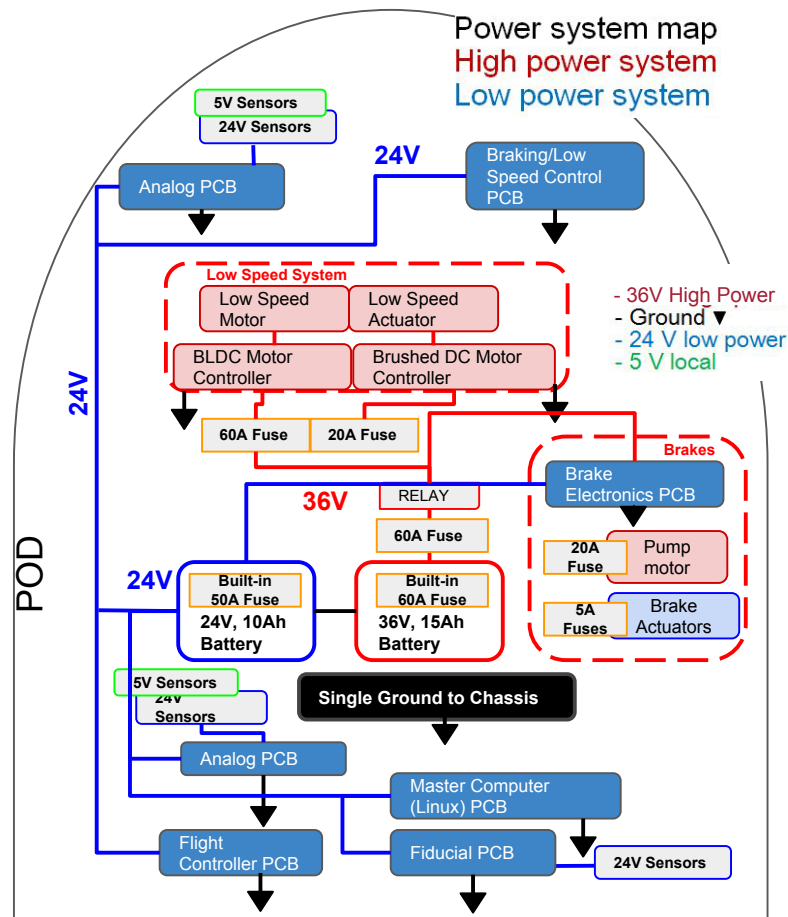


Figure 6.10: Power system diagram for entire pod.

When sizing cells for the high power system, the power rating decided the cell type, rather than the cell capacity. Thus we prefer high-discharge rate batteries. For the low power system, we would like to keep the computer running as long as possible, thus we prefer low-discharge rate high capacity batteries. Considering the functionality in vacuum, the requirement on size and the price, we rule out lead-acid batteries and special batteries designed for space/military applications. This narrowed our selection to off-the-shelf lithium-ion or lithium-polymer batteries.

Two battery types (3 different cells) were tested: lithium-ion manganese and lithium-polymer. Polymer cells get bloated in the vacuum but the lithium-ion cells function well. We tested two lithium-ion manganese cells in both vacuum (3000 Pa) and ambient sea level conditions.

They are both IHR 18650 BN cells rated at 2.5 Ah. Each cell was discharged at 9 A (expected 6.5 A at our peak power) and their temperature was recorded, the results are shown in Fig. 6.11. As expected, the temperature increase is sharper for the cells in vacuum.

The actual discharge time at high current during the run is very short, thus the battery cell is considered thermally safe. More tests were done at the pack level later to ensure the whole pack is thermally safe.

The high power system uses a 36V 15 Ah pack, while the low power system uses a 24V 10 Ah battery. In the actual run, there were no major issue with batteries.




Figure 6.11: Temperature of cells during discharge in vacuum chamber.

6.2.2 Brakes Electronics

The electrical system of the brakes consists of five parts: 24V proportional valve, 24V on/off valve, a 24V pump motor, a 24V pressure sensor and two 24V linear potentiometers, as shown in Fig. 6.12. For more information on the sensors, please refer to Section 6.1.5.

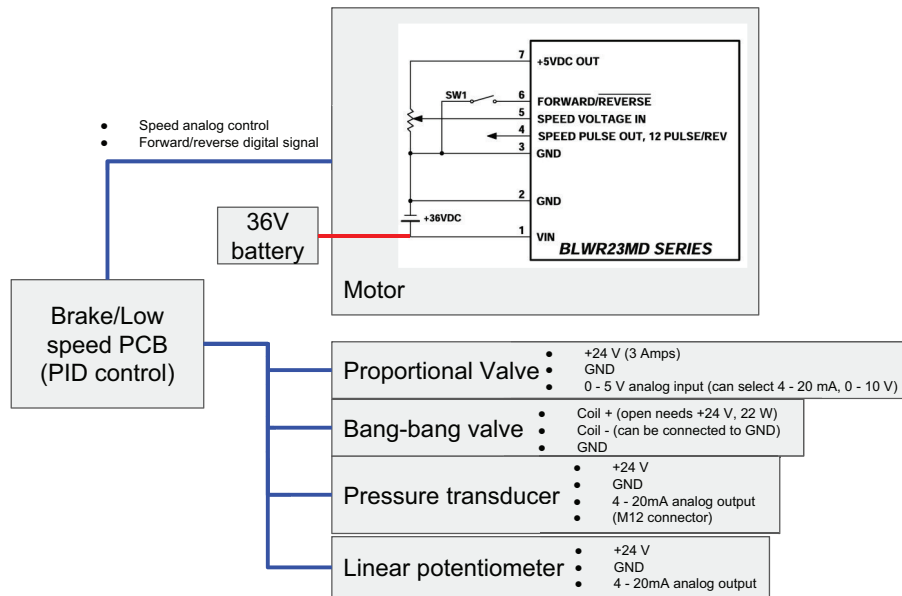


Figure 6.12: Brakes electronics system overview.

The valves and motors all have built in controllers, thus we care mainly about powering them and interfacing them with the Brake/Lowspeed PCB. This is realized by the brake interface PCB. One should note that the interfaces of the enable signals for the low speed motors and actuators are also realized on this PCB.

The brake interface PCB provides ground isolation between the signal ground and the power ground to ensure that if a high power ground is cut, it does not go through the low power ground and burn the computer systems, and to reduce noise.

Thus all speed signals go through isolators - if the speed signal is analog, it is converted to PWM before going through an isolator. All enable signals are used to enable other relays to realize the ground isolation. Fig. 6.13 is a good example of how different signals are processed. The pump motor is enabled by a relay CMX60D20. It also needs an analog speed signal, which is sent out by the Arduino as a PWM signal. Then it passes through an isolator and a PWM to ADC converter LTC2644. The motor sends out a pulse signal as the actual speed, which goes through the isolator and is received by the Arduino.

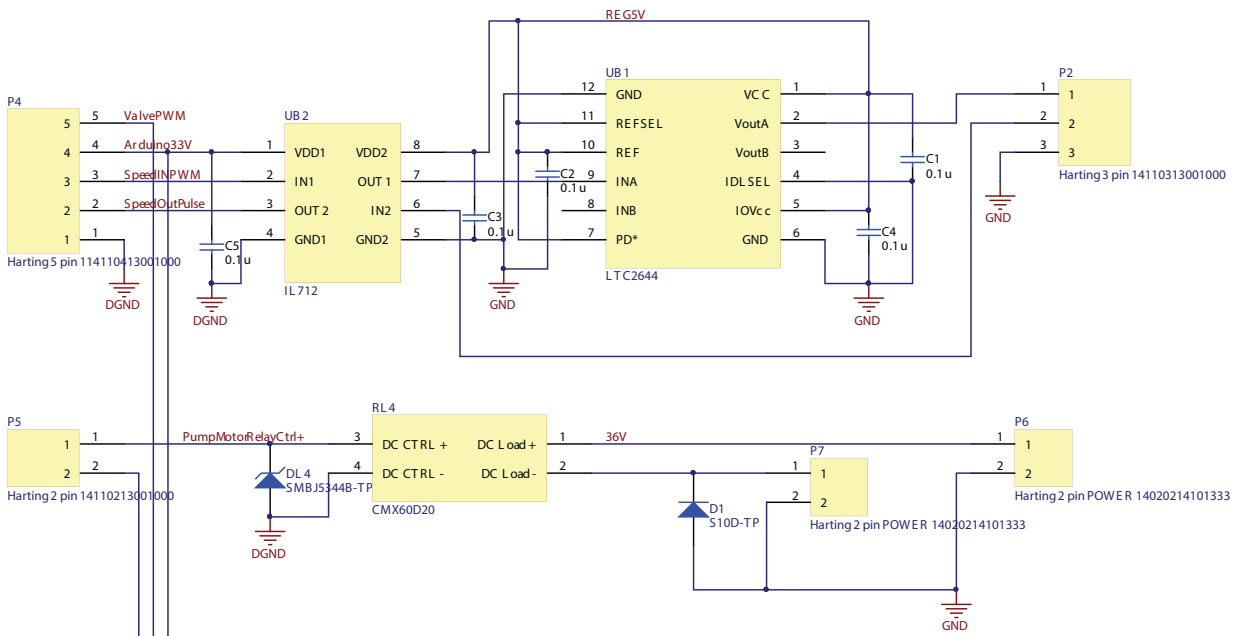


Figure 6.13: Schematics of pump motor interface circuits on the brake interface PCB.

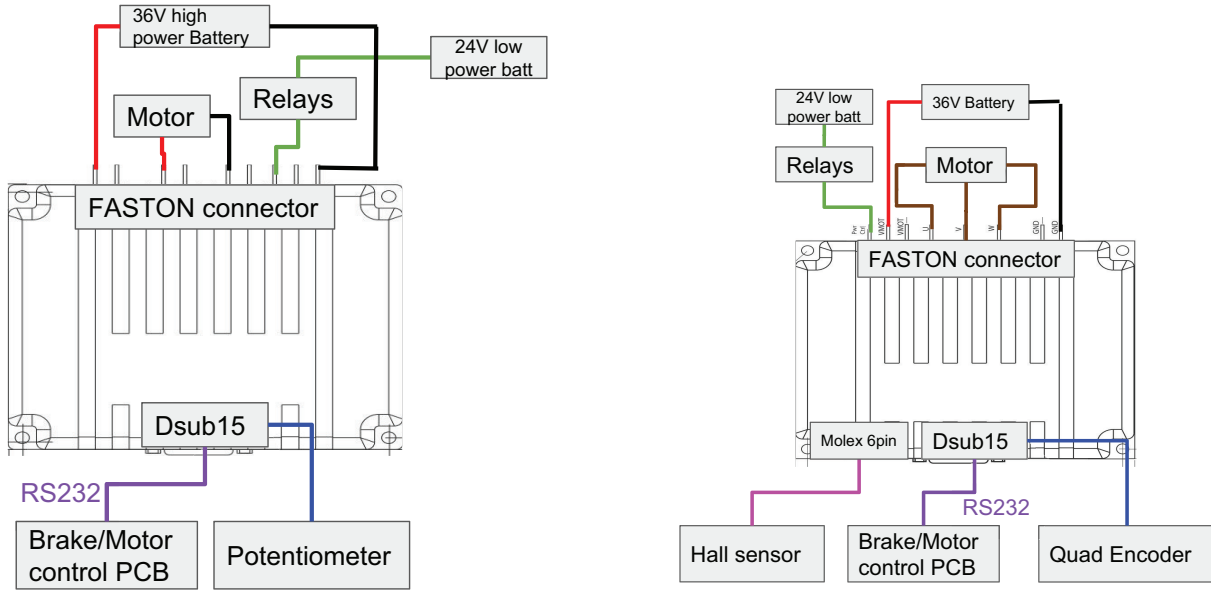
One important feature of the braking system is that only the pump motor is powered from the 36V bus, the proportional valve and the on/off valve (also referred as digital valve or bang-bang valve) are powered from the 24V line. This is a fail-safe measurement – even if we lose 36V battery, as long as the 24V battery is alive, the two valves can hold the brake closed. This requires an isolated 24V to 24V power supply, which is the most bulky component on the whole PCB.

6.2.3 Low Speed System Electronics

The low speed system consists of a 36V low speed motor (connecting to a gear and a wheel to drive the pod at low speed) and a 24V actuator (to deploy the low speed motor and wheel). The electronics system for them is shown in Fig. 6.14. Both are powered from the 36V battery. Enable signals are from the Brake/Lowspeed PCB (one of the 5 Main PCBs) and interfaced on the Brake Interface PCB. The local motor control is done by controllers from Roboteq, and they are directly interfaced with the Brake/Lowspeed PCB through RS232.

6.2.4 Power System Protection and Connections

Here we discuss the power system connections and safety measures, such as switches and e-stop, fuses, and the grounding scheme. The e-stop and switch diagram is shown in Fig. 6.15. We have one hand switch



(a) Low speed actuator connection diagram.

(b) Low speed motor connection diagram.

Figure 6.14: Low speed system electronics system overview.

for 36V, one hand switch for 24V, another e-stop and relay combination to shut down 36V using 24V. All of these are for safety during testing.

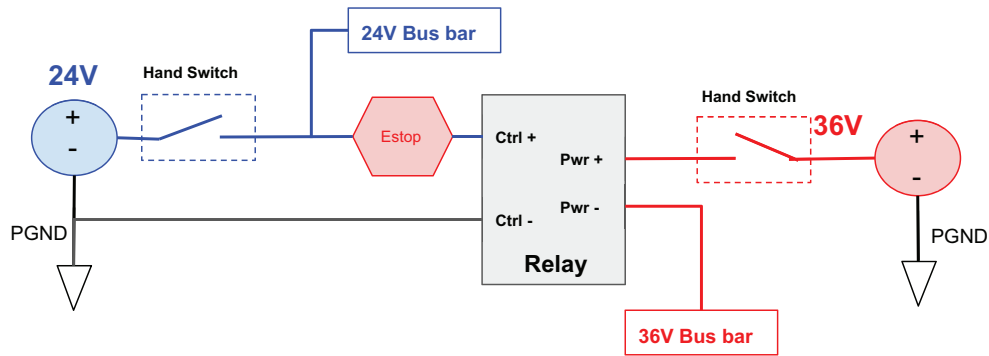


Figure 6.15: Diagram for E-stop and battery switches.

During the run, there are relays on the Brake Interface PCB controlled by the Brake/Lowspeed PCB to shut down high power components, there is no remote control to shut down 24V system. There are many reliable, low weight/small, high current battery switches for systems below 24V thanks to the automotive industry. For the 36V switch, we use a manual battery switch rated at 58 V and 300 A. There are limited options for switches for more than 36 V and 50 A, especially if the switch is to be lightweight and small.

Finally, the grounding scheme is shown in Fig. 6.16. As mentioned before, we isolate the 24V and 36V battery grounds except at one single point, to ensure that if a high power ground is cut, it does not go through the low power ground and burn the computer systems, and to reduce noise caused by high current motors.

The single connection between two ground is realized by a ground wire connecting between two busbars. Another single ground wire connects the 36V ground to the chassis.

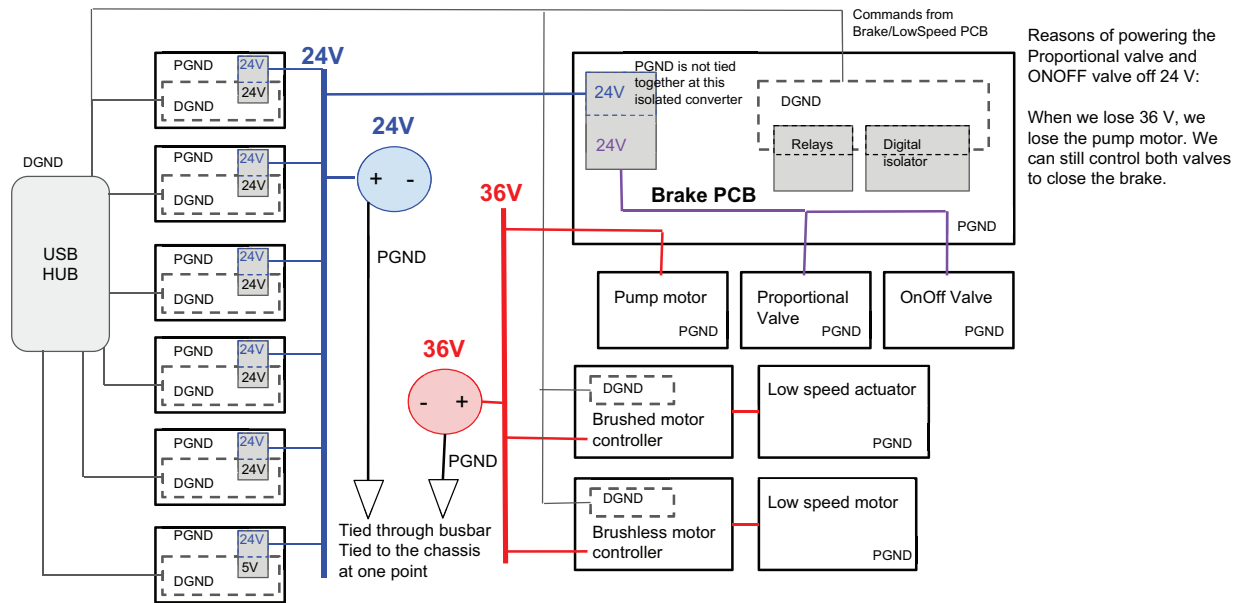


Figure 6.16: Grounding diagram for power electronics system.

6.3 Wiring Harness

Given the modular, distributed nature of the onboard computers, and the large number of sensors, we decided to have a separate cable to connect each component to each microcontroller. This enabled us to build the wiring harness rapidly, while being flexible enough to accommodate changes in components and their locations on the pod.

6.3.1 PCB Enclosures

Each microcontroller is mounted on a version of our Main PCB. This PCB interfaces the signal wires securely to the pins of the Arduino. These PCB assemblies are in polycarbonate enclosures which allow them to be physically secure and to prevent their contacts from being shorted accidentally. An example of a PCB enclosure assembly is shown in Fig. 6.17.

The placement of the panel feed-through connectors embedded in the walls of the enclosure is largely independent of sensor and enclosure location. This makes each PCB enclosure unit more compact, as we can decide the most convenient location for a given panel feed-through connector. The connectors are also arranged to keep our wiring neat and prevent unnecessary crossover of wires. The cutouts for these panel feed-through connectors were machined out using a CNC machine.

6.3.2 Sensor Placement

The location of sensors on the pod is a function of the sensor requirement. A simplified sensor map is shown in Fig. 6.18.

The high-power system (36V battery, motor controllers and actuators) are intentionally situated together, as shown in Fig. 6.19. This way, high current-carrying wires can be concentrated together making it easier to route signal cables away from these.

The Arduinos in the PCB enclosures are placed to minimise the average distance between a given




Figure 6.17: PCB enclosure assembly.

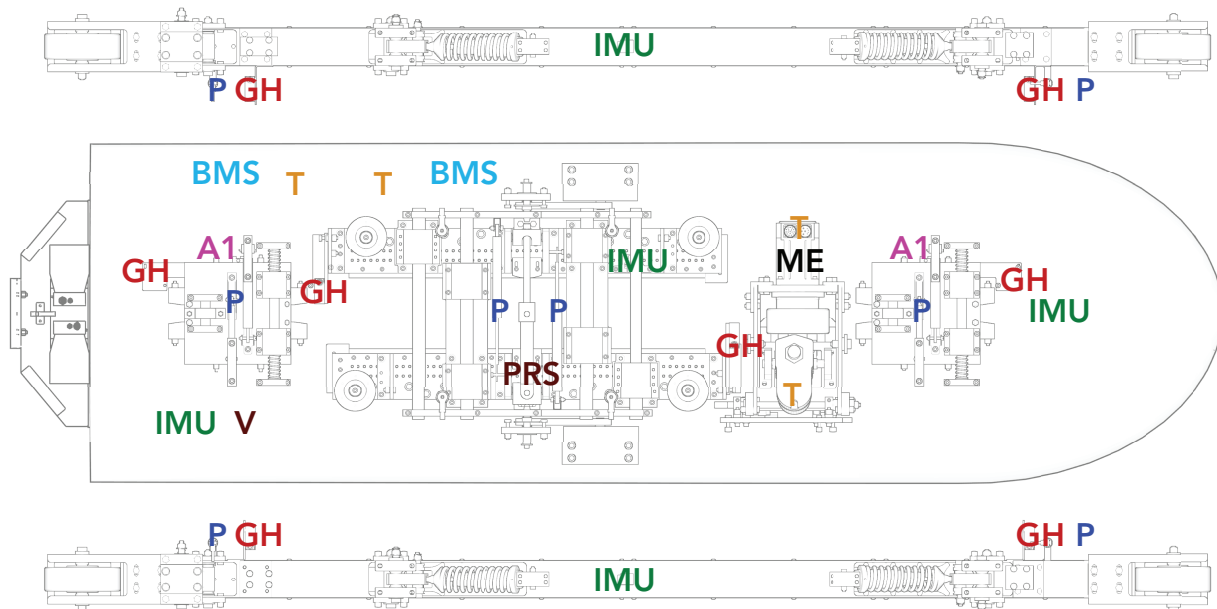


Figure 6.18: Map of sensor locations. Legend is shown in Table 6.2.

enclosure and all of its sensors and actuators. Since analog sensors are distributed throughout the pod, the analog sensor module is split to use two Arduinos: one in the front and one at the rear.

6.3.3 Cable and Connector Selection

Most sensors are analog sensors. Some have integral cables, while for others, a shielded 4-conductor PVC insulated cable is used.

PVC can be susceptible to outgassing in a vacuum; it undergoes 25% mass loss when heated to 125°C in a vacuum for 24 hours. We tested this effect by exposing our cables to the pressure conditions and time durations we expected to encounter at testing and competition (10-30 minutes). No outgassing was detected, making PVC cables a suitable choice.




Figure 6.19: Overview of electronics systems on the pod.

Highly flexible silicone insulated cables carry the large currents of the high power cables. Their wire gauge size is decided based on an ampacity rule (700 circular mils per amp) used for power transmission in a bundle, rather than a single cable. This conservative estimate is appropriate because the wires will see no convective cooling in a vacuum.

Shielded circular M12 connectors are chosen for the analog sensors and for the 24V transmission that run the Main PCBs. These provide both shielding and locking and are hand-assembled on one end of the cable.

The digital sensors and the RS485 bus both use shielded D-sub (DB9) connectors as they are standard and locking. USB connects all Arduinos to the main onboard computer (ODroid) and the USB cables have special locking panel feed-through connectors. A custom cover for the USB hub holds the cables in place (since a locking USB is not compatible here).

The PCB-mounted connectors themselves are two-piece connectors (har-flexicon from Harting) which snap together for easy attachment/removal. One piece was soldered to the PCB contact, the other piece had a spring and blade which held the wire conductor.

The power from the batteries comes through Anderson connectors, which prevent shorting when disconnected, and do not allow reverse-polarity connection. Ring terminals and blade terminals are used for other high power connections between the busbars, motor controllers and actuators. The low-speed drive motor requires its own circular M23 connectors.

6.3.4 Mounting of Electrical Components

Each of the sensors is mounted either on the frame or suspension component as required by its attachment points. A laser rangefinder was initially planned to find the distance to the end of the tube during the braking, but it was not mounted since we could not ensure it would always point at the end of the tube due to the pitching of the pod.

For mounting the other electronics components (PCB enclosures, batteries, busbars, motor controllers, USB hub, high power relay and emergency switch, NAP), there are four plates on the ladder-frame structure. All components are bolted down on these. For the PCB enclosures, batteries and NAP, which require frequent removal and replacement, bonded nut plates are used to facilitate easy mounting.

The plates and the entire frame act as the electrical ground. Two of the plates – the ones with the PCB



enclosures – are mounted on rubber bushings to isolate them from vibrations. This also isolates them electrically, so a separate cable bridges the mounting points and connects them to the common ground. The plates also act as heat sinks for the motor controllers.

For the batteries, it was found that the temperature rise was low for the expected duration of operation in vacuum: $1^{\circ}C$ per 7 minutes. This was determined using experiments with a single cell and then simulating the battery-pack. A single cell's parameters were determined to be - electrical resistance: $8\ m\Omega$, heat capacity: $15\ J/K$. The batteries are also covered by custom carbon-fiber enclosures to protect them.

Since the NAP is mounted on one of the electrical plates, i.e. within the pod's electrically conductive carbon fiber shell, we use the alternative mounting strategy for the wireless communication antennae. These are stuck on the PRI plate at the rear, facing upwards and at a right angle to each other.

6.3.5 Cable Routing

The cables from sensors underneath the electrical plates come up to the PCB enclosures through cut-sections in the plates located on the four corners of the frame. The side rails of the frame are mostly used for cables that need to run longitudinally on the pod (left rail for most signal cables, right rail for most power cables). The high power cables are short and have inline fuses.

All the cables are tied down to the frame using cable tie mounts with adhesive backings. These were vibration tested along with all the other electronics and did not come apart.

The need to splice cables was minimized. Some of the signal cable paths are longer than commercially available cables, or their end connectors are incompatible (e.g. locking USB-A to USB-B) necessitating splicing of cables.

While testing it was found that the signal to the brake hydraulics pump motor was noisy. This was because the power cable and signal cable were routed together (as they went from the same PCB to the same connector on the motor). The power converter in the motor operated at a switching frequency of $20\ kHz$, and EMI at this frequency was emitted by the power cable. The skin depth for copper at this frequency is on the order of $1\ mm$, and the signal cable shielding was much thinner than this. Thus, the solution was to route the two cables so that they were far apart.

The final state of the pod with all components mounted, cables attached and routed is shown in Figures C.1 and C.3. Note the carbon fiber battery enclosures, dummy seat and NAP are not shown in these photographs.

SOFTWARE

We employ a modular design in our pod’s control system, motivated by an underlying need to maintain simplicity and maximize safety while minimizing development time. Low-level or real-time-constrained sensing and control tasks are delegated to one of multiple low-level microcontrollers. Management of these boards, along with additional high-level, non-real-time, and complex tasks, are handled by a more powerful central computer. A user interface supporting data review and high-level flight commands is run on a laptop. Hierarchical controller design enables a simple organization of watchdogs and error detectors to ensure graceful failure under a variety of failure conditions, including loss of communications between any layer of this hierarchy.

All communications in the software system is handled through a single unified, broadcast-based message-passing system, namely the Lightweight Communications and Marshalling system (LCM^a) (originally designed and used by the MIT DARPA Urban Challenge and DARPA Robotics Challenge teams).³⁰ LCM unifies the process for defining and transmitting datagrams between any pairs of modules in the system, and additionally provides mature tools for data inspection, logging, and playback. By design, all relevant control and sensor state passes through the message passing system, ensuring that an LCM data log captures the complete state of our system over time.

Simulators are used to test the behavior of the high-level software systems. These again leverage the flexibility of our LCM-based message-passing system, as simulated modules are able to emit data in a manner identical to their real counterparts, removing the need for special-case code paths in the high-level controllers and making for more robust testing. This also enables hardware-in-the-loop testing in many cases.

Note that a subset of the code utilized by our team is available online^b.

7.1 Overview of System Modules

The MIT pod contains five low-level embedded microcontrollers (ARM Cortex M3 microcontrollers on an Arduino Due development board) and one onboard embedded flight computer (ODROID C1+), as already discussed in Section 6.1. The flight computer maintains communications to each five of the Arduinos via USB connections carrying 100kbaud serial data, and also maintains a network connection via wired ethernet to a LAN (which may include a wireless bridge) containing the pilot laptop. The ODROID runs high-level autonomy on the pod, communicates with high-bandwidth sensors (webcams), monitors network state, and relays information between the Arduinos and a UDP broadcast network bridging the ODROID and pilot laptop. (See Fig. 7.1.)

^aUnless otherwise specified, all uses of the acronym LCM will refer to Lightweight Communications and Marshalling in this chapter.

^bThe code is available at https://github.com/MITHyperloopTeam/software_core

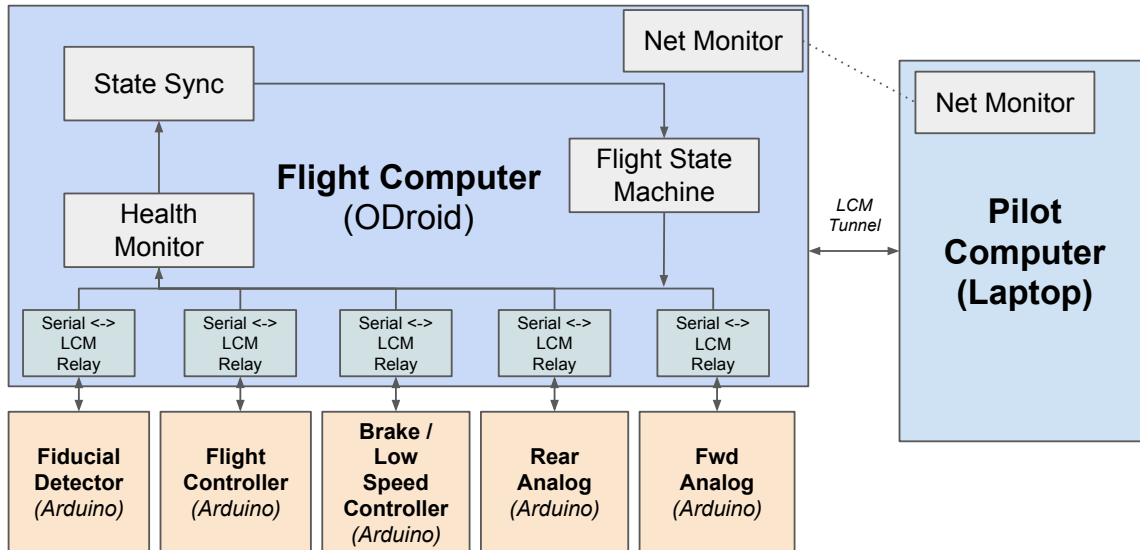


Figure 7.1: System-level diagram of the software system on the pod.

7.1.1 Embedded Modules

The five embedded modules handle tasks that require a real-time performance in close conjunction to sensors and actuators. In particular, these modules handle interactions with all sensors (except the cameras), all actuation commands to the low-speed and brake actuators, and perform the high-rate computation for sensor filtering, state estimation, and flight and braking control.

The mechanical design of our pod emphasizes passive stability where possible, which results in comparatively light specification requirements for the embedded computers. Simulation of the dynamics of the pod indicate that all control-relevant dynamics would be at or below 30hz, with a Nyquist rate of 60hz. We found that almost any modern embedded computer was capable of handling the simple actuation (a functionally single-DOF braking system) and state estimation (requiring only estimating position and velocity along a single axis) requirements of this system. Thus, to minimize toolchain development time, we chose a relatively-computationally-powerful Arduino variant with which team members had experience. The Arduino Due hosts a 32-bit 84Mhz SAM3X8E ARM Cortex M3 CPU, and exposes most of the digital and analog I/O capabilities of that board – and, crucially, is supported by the Arduino toolchain. This toolchain, which wraps the core SAM library for operating the SAM3X8E, was fairly easily extracted and incorporated as a direct dependency in our codebase. A direct result of this choice was that we had to write no lines of embedded hardware management code to run the SAM3X8E microcontroller, instead drawing on the collective work of the vast Arduino community and leveraging previous solutions and libraries for performing the tasks we needed. We believe this ultimately saved us hundreds of man-hours and enabled our relatively small software team to write a huge amount of embedded code.

At a high level, the five embedded boards managed:

- The Fiducial Detector (FD) module interfaces with an optical sensor for detecting optical marks (“fiducial” marks) on the inside of the tube, to enable much more precise absolute localization within the tube. The detection of these marks had the strictest timing requirements of any task: the 4" wide marks pass a laser sensor in $400 \mu s$ at a relative speed of $110 m/s$, which demands a $5 kHz$ minimum sampling rate. Out of an abundance of caution with respect to this critical timing requirement, and to provide flexibility to switch fiducial sensors down the road, we dedicated an entire embedded

module to this task.

- The Flight Controller (FC) module interfaces over RS232 with two navigation-dedicated IMUs on the structural frame (i.e. above the suspension) and tightly integrates the output of those IMUs to generate velocity and positions estimates of the pod. It communicates with the fiducial detector and filters fiducial marker detections in an Extended Kalman Filter (EKF) to produce a pod position and velocity estimate, and passes those estimates through a simple trajectory-following controller to produce brake commands for the pod.
- The Brake Controller, often referred to as the Low-Speed (LS) Module, manages all actuation of the pod. It communicates with motor controllers for the low-speed system, as well as with the hydraulic system controlling the brakes. It runs an inner-loop controller to control brake caliper position.
- The Front and Rear Analog (AF, AR) modules communicate with a range of sensors on the pod suspension, lateral control components, skis, and batteries.

7.1.2 Overall System Comments and Review

While our system has proved capable of meeting its requirements and did fly our pod at the first competition weekend, the experience of building and testing it has given us insight into the consequences of our design choices:

The **Arduino Due platform** has generally performed admirably, but poses two kinds of communication bottleneck.

- **Pros:**

- The Arduinos have no thermal issues under vacuum.
- The Arduino library and community have removed much labor from most of our endeavors. Mature cross-compilation tools for the architecture made embedded code development relatively painless.
- The Arduino Due is programmed through a USB programming port, which ultimately communicates with a satellite Atmel 16U2 that functions as a USB/Serial converter and programmer. Our communication with the Arduinos happens through this same channel. This was an intentional choice, as it allowed us great flexibility for reprogramming in the field during testing and development, and still allowed serial communication at 115200 bps^c.

- **Cons:**

- The 115200 bps serial line proves to be a bottleneck for high-rate data transfer on modules with many sensors to report at a high rate. While this bandwidth limit was considered in our original spec and was not an issue for competition weekend, it does preclude the addition of further high-rate sensors on almost any module. Many alternative communication methods exist and were considered – chief among them, CAN, Ethernet, and alternative serial transmission protocols like RS232 and RS485. We opted for TTL UART over USB for simplicity and because it met our specifications. RS232 and RS485 may face similar bandwidth restrictions on the decoding (i.e. ODROID) side, as might CAN depending on the decoding hardware and software. Ethernet is a more modern solution that has become popular for robotics sensors and systems,

^cWe had an early hope that our communication could go much faster, as the Arduino Due theoretically supports baud rates in the low millions of bps. However, while the Arduino Due might support this, the ODROID we use does not support these kinds of rates on virtual COM ports.



and is supported by dedicated hardware on most modern microcontrollers; it would be a good direction to explore on a future system.

- The Arduinos have only 1 hardware UART, and 3 more software UART. Many of the sensors we interface with require some form of serial. These sensors thus had to be distributed among modules carefully, in some cases complicating communication and data synchronization. For example, the IMU mounted on the brakes is connected to the Front Analog Module instead of the Brake / Low Speed Module. As a result, the IMU is not easily available for use in high-rate state estimation of the brake system. Testing has proven that this sensor is not necessary for accurate brake performance, but this is certainly an area for improvement.

The **ODROID** has similarly impressed us with its versatility, but similarly poses a key communication bottleneck, and has proven less thermally stable than the Arduinos.

- **Pros:**

- The ODROID runs a very similar operating system to our pilot laptop (both running Ubuntu 14.04, the only difference being the ODROID is an ARM Cortex A5 running ARMv7, while our pilot laptops are universally utilizing x86-64) and shares all of the same tools. Thus, processes can be moved between the ODROID and pilot laptops with no additional development work, making testing significantly easier and more flexible.
- The ODROID has never failed for thermal reasons, even after extensive use in vacuum. However, see below for discussion of some reservations about this result.

- **Cons:**

- Beyond the virtual COM port limitations mentioned above, the ODROID has a fairly limited core filesystem. The ODROID utilizes an SD card for all storage beyond core memory, meaning the entire filesystem is stored on that SD card. Thus, filesystem speed is heavily dependent on the SD card performance. With our SD cards fare better we had write rates of around 10MB/s, but this still must be considered when planning data logging rates.
- During our first test weekend, we saw ODROID temperatures peak around $90^{\circ}C$ after being left in the sun within its enclosure for around an hour (without the aerodynamic shell). The small passive heatsink on the ODROID is likely to blame here. This issue was tracked and did not cause an issue at competition weekend.

7.2 Communications

To relay information throughout the pod during operation, we rely on the LCM broadcast system. LCM allows us to specify datagrams, called “types”, in a simple C-struct-like syntax, which allows us to define a variety of message types consisting of integers, floats, and variable-length arrays. An example of such a type is pictured in Fig. 7.2.

These datagrams are broadcast throughout the entire network to all listening programs. Where messages must be passed from one network to another, we use relay programs to pass the messages through (see Figures 7.3 and 7.4).

One of the biggest advantages of using this system is that, because the message-sending is implemented both across UDP networks and also within Arduino serial lines (using the ZCM library^d, which extends

^dZCM: Zero Communications and Marshalling, library available at <http://zerocm.github.io/zcm>

```

package mithl;
struct bac_state_high_rate_t
{
    int64_t utime;

    float distance_1;
    float distance_2;
    float brake_pressure;
    float estimated_motor_rate;
    float estimated_distance;
    float estimated_distance_velocity;

    float gh_front;
    float gh_rear;

    float PumpMotorPWM; // [0, 1]
    boolean OnOffValve;
    float ValvePWM; // [0, 1]

    float setpoint;
    float error;
    float integrator;
}

```

| Channel | Type | Num Msgs | Hz | 1/Hz | Jitter | Bandwidth | Undecodable |
|-------------------------|----------------------------|----------|-------|-------------|-------------|-----------|-------------|
| _RADM_WHEEL_STATUS | adm_status_t | 44 | 4.00 | 250.07 ms | 243.27 ms | 0.39 KB/s | 0 |
| _RADM_WHEEL_CURCONF | adm_config_t | 44 | 4.00 | 250.07 ms | 243.28 ms | 0.29 KB/s | 0 |
| _RADM_CLAMP_STATUS | adm_status_t | 43 | 4.00 | 250.07 ms | 242.98 ms | 0.40 KB/s | 0 |
| _RADM_CLAMP_CURCONF | adm_config_t | 43 | 4.00 | 250.07 ms | 243.01 ms | 0.29 KB/s | 0 |
| _NSUM | net_health_t | 77 | 7.00 | 142.90 ms | 902.85 ms | 0.92 KB/s | 0 |
| _NREQ | net_diag_t | 154 | 14.00 | 71.45 ms | 383.54 ms | 0.30 KB/s | 0 |
| _NREP | net_diag_t | 924 | 83.98 | 11.91 ms | 385.49 ms | 1.80 KB/s | 0 |
| _LS_OL | battery_reader_low_ra... | 22 | 2.00 | 500.14 ms | 491.10 ms | 0.04 KB/s | 0 |
| _FIDUCIAL_DESIRED_TE... | fiducial_teach_table_t | 5 | 0.00 | Infinity ms | -10.00 ms | 0.00 KB/s | 0 |
| _FIDUCIAL_DESIRED_CO... | fiducial_config_t | 5 | 0.00 | Infinity ms | -10.00 ms | 0.00 KB/s | 0 |
| _FIDUCIAL_CUR_TEACH... | fiducial_teach_table_t | 5 | 0.00 | Infinity ms | -10.00 ms | 0.00 KB/s | 0 |
| _FIDUCIAL_CUR_CONFIG | fiducial_config_t | 5 | 0.00 | Infinity ms | -10.00 ms | 0.00 KB/s | 0 |
| _FIDUCIAL_ACK_TIMEOUT | string_t | 9 | 1.00 | 1000.27 ... | 1244.83 ... | 0.02 KB/s | 0 |
| _FD_OL | battery_reader_low_ra... | 22 | 2.00 | 500.14 ms | 490.81 ms | 0.04 KB/s | 0 |
| _FD_O | fiducial_t | 541 | 48.99 | 20.41 ms | 25.92 ms | 2.30 KB/s | 0 |
| _FC_SE_STATE | state_estimator_state_t | 109 | 10.00 | 100.03 ms | 91.95 ms | 0.34 KB/s | 0 |
| _FC_SE_CONFIG | state_estimator_config_t | 22 | 2.00 | 500.14 ms | 490.80 ms | 0.18 KB/s | 0 |
| _FC_SE | state_estimator_partic... | 326 | 29.99 | 33.34 ms | 24.61 ms | 2.40 KB/s | 0 |
| _FC_OL | flight_control_low_rate_t | 22 | 2.00 | 500.14 ms | 490.99 ms | 0.23 KB/s | 0 |
| _FC_OH | flight_control_high_rat... | 1028 | 92.97 | 10.76 ms | 4.28 ms | 5.90 KB/s | 0 |
| _ESTOP | image_sync_t | 55 | 5.00 | 200.05 ms | 192.80 ms | 0.08 KB/s | 0 |
| _BAC_TELEOP_CMD_ST... | bac_teleop_cmd_t | 44 | 4.00 | 250.07 ms | 243.80 ms | 0.12 KB/s | 0 |
| _BAC_STATE_L | bac_state_low_rate_t | 44 | 4.00 | 250.07 ms | 243.77 ms | 0.21 KB/s | 0 |
| _BAC_STATE_H | bac_state_high_rate_t | 330 | 29.99 | 33.34 ms | 24.33 ms | 2.02 KB/s | 0 |
| _BAC_MODE_STATE | bac_mode_t | 22 | 2.00 | 500.14 ms | 491.05 ms | 0.03 KB/s | 0 |
| _BAC_CONFIG_STATE | bac_config_t | 22 | 2.00 | 500.14 ms | 491.06 ms | 0.21 KB/s | 0 |
| _BAC_AUTO_CMD_STATE | bac_auto_cmd_t | 109 | 10.00 | 100.03 ms | 91.09 ms | 0.20 KB/s | 0 |
| _BAC_ARM_COUNTDOWN | vectorxf_t | 109 | 10.00 | 100.03 ms | 91.05 ms | 0.21 KB/s | 0 |
| _AR_OM | analog_rear_medium_r... | 330 | 29.99 | 33.34 ms | 25.24 ms | 1.29 KB/s | 0 |
| _AR_OL | analog_rear_low_rate_t | 22 | 2.00 | 500.14 ms | 490.59 ms | 0.05 KB/s | 0 |
| _AR_OH | analog_rear_high_rate_t | 1034 | 93.97 | 10.64 ms | 2.71 ms | 6.24 KB/s | 0 |
| _AR_BATTERY | battery_status_t | 22 | 2.00 | 500.14 ms | 490.58 ms | 0.20 KB/s | 0 |
| _AF_OM | analog_front_medium... | 330 | 29.99 | 33.34 ms | 24.06 ms | 1.17 KB/s | 0 |
| _AF_OL | analog_front_low_rate_t | 22 | 2.00 | 500.14 ms | 490.62 ms | 0.05 KB/s | 0 |
| _AF_OH | analog_front_high_rat | 1036 | 93.97 | 10.64 ms | 2.23 ms | 4.04 KB/s | 0 |

Figure 7.2: LCM type example (left) and LCM spy utility (right).

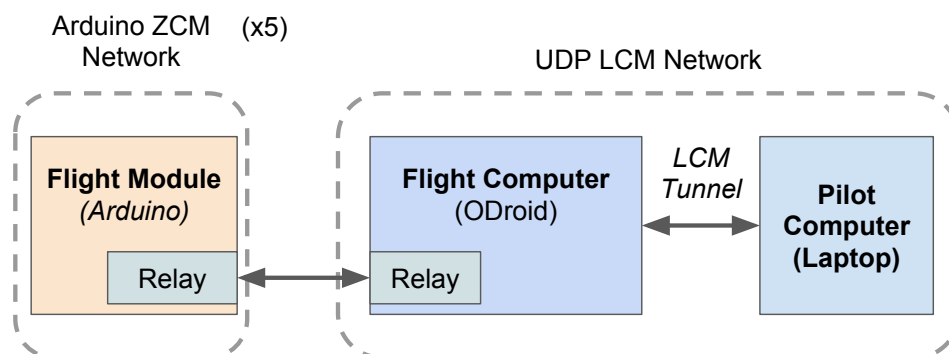


Figure 7.3: High-level depiction of the system-wide LCM broadcast network. Embedded systems are included in the overall network via individual relays that funnel messages from the shared UDP broadcast network down to each embedded computer. The pod and pilot computers are linked via a dedicated message tunnel process that rate-limits high-bandwidth channels.

vanilla LCM to other transport mechanisms, as well as embedded systems), we can transmit data from any module in our system to any other module flexibly and with little overhead. This interprocess communication enables monitoring of system status, sending of real-time control commands, and other functionality. It also makes our system easy to extend, with very few changes needed to bring a new line of communication into the system.

7.2.1 LCM Logging

The LCM software package includes the LCM Logger, which records all LCM traffic passing through a given network. These recordings can be played back later or converted directly into timeseries data for a given field.

Because all sensor information is being broadcast constantly, we can use the LCM Logger to perform after-the-fact diagnosis and data processing for any run of our system.

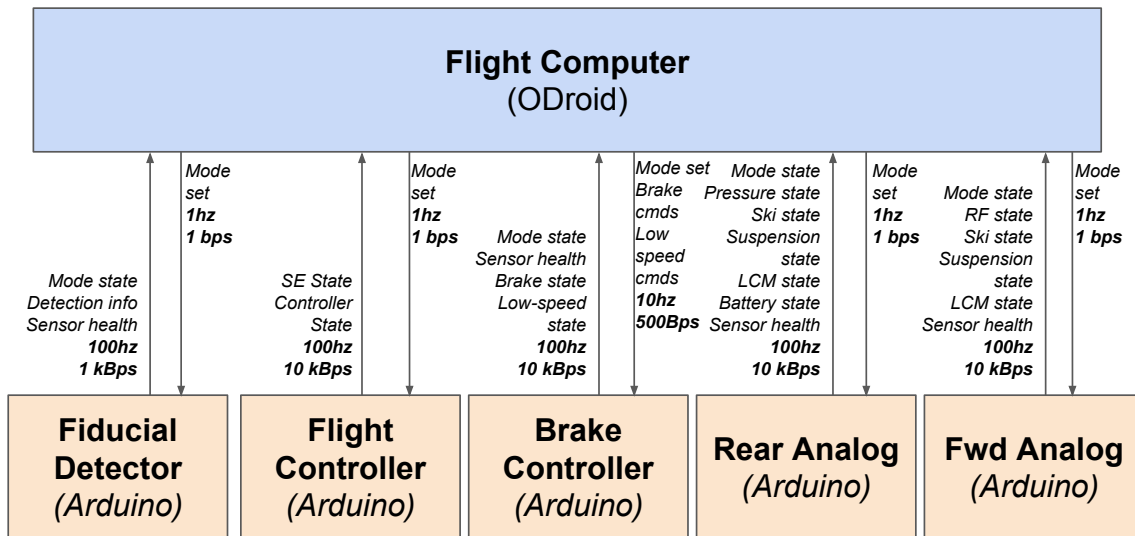


Figure 7.4: Communication architecture on the pod.

7.2.2 Process Management

We rely on the `libbot` library, which provides a suite of LCM-oriented tools originally motivated by robotics applications and utilized by the MIT DARPA Grand, Urban, and Robotics Challenge teams. For process management, we utilize the process manager or `procman` utility. This tool allows a single pilot machine to control spawned processes on multiple worker machines, making it easy to launch new processes where needed. Because these executions are spawned as independent processes, this tool will leverage multiple cores if they are available on a machine. These workers are kept separate from the UI which controls them, and thus they can continue running even if the UI is closed, as explained in Section 7.4.

Our system consists of two workers, or *deputies*: “pilot” and “odroid”. The “odroid” deputy runs on the onboard computer and launches all arduino and sensing utilities as well as the flight state machine. The “pilot” deputy runs on the pilot laptop and is responsible for launching UI, control, and simulation utilities.

7.2.3 LCM System Comments and Review

Using a broadcast-based communications structure greatly simplified system design and made it easy to create a modularized and thoroughly inspectable and log-able system. While the message broadcasting had the potential to waste bandwidth and CPU cycles, our bandwidth was sufficiently low that this was only a difficulty on the network link that connected the ODROID and pilot laptop. To address this, we selectively throttled the highest-bandwidth channels (e.g. the webcam data) as they crossed the ODROID-pilot laptop network bridge, in order to keep the bandwidth below the limit.

LCM in particular proved an excellent choice for managing our core communications. The team’s prior experience with the library was a great help during development, and its origin in robotics meant that it was well-suited to a semi-autonomous system project like this one. ZCM’s support for LCM message parsing on embedded systems allowed us to bring the embedded systems into this shared broadcast network, which vastly simplified communication to those modules (as each LCM/ZCM message has a defined and auto-generated serialization procedure).

7.3 Guidance, Navigation, and Control

The pod's control system provides autonomy for the duration of each flight. While the vertical and lateral levitation dynamics are entirely passive, the brake actuators and braking trajectory must be actively controlled to ensure a successful flight.

7.3.1 Flight State Machine

The high-level autonomy and decision making in our system is encoded in the Flight State Machine (FSM) that is always active on the pod. The FSM encodes our expectations of the process of readying, flying, and driving the pod. The FSM operates in conjunction with the many drivers, estimators, and controllers in our system: the FSM subscribes to sensor, state estimator, controller, and UI state to make informed transitions, and the controllers and drivers listen to FSM state in order to exact changes and emergency stops (ESTOPs). This system allows for shared autonomy: the pilot may request transitions between certain pairs of states (including from any state to ESTOP, in order to rapidly make the pod safe). The FSM may also automatically transition between certain states during flight. See Fig. 7.5 for a block diagram of the FSM.

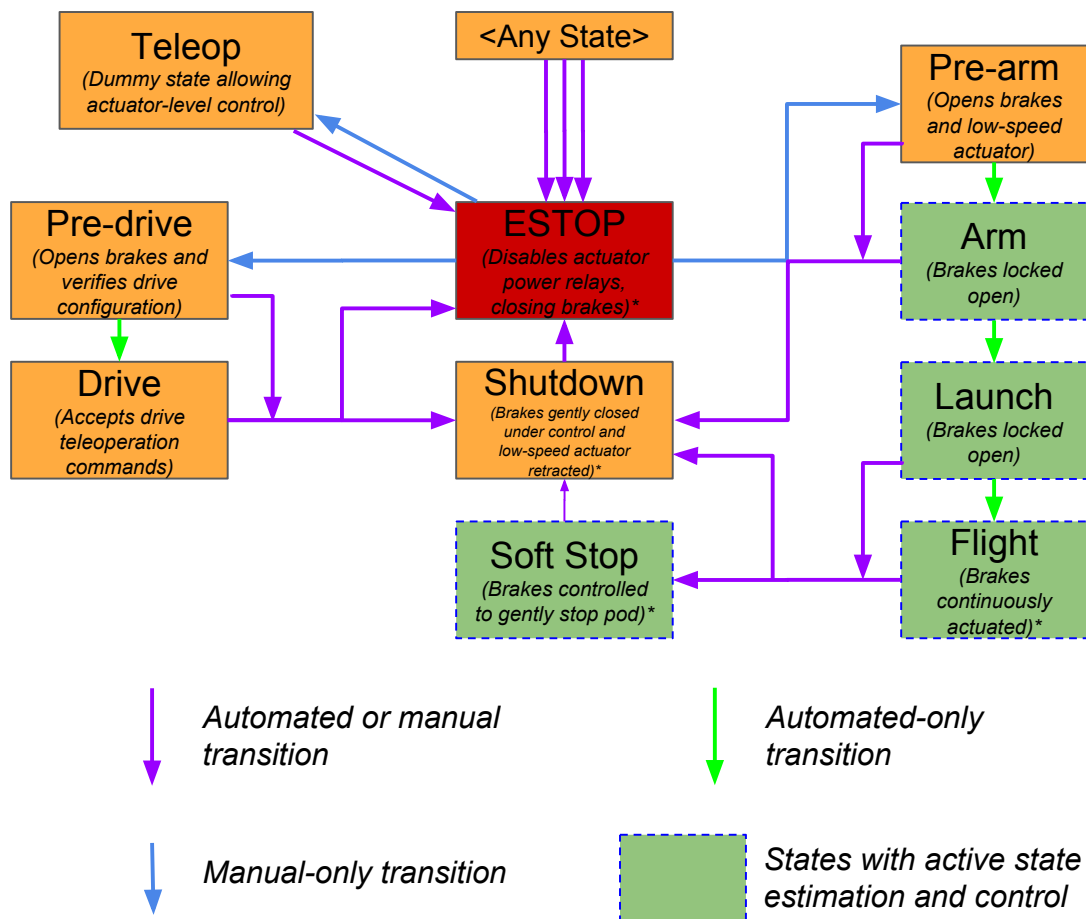


Figure 7.5: State diagram of the Flight State Machine (FSM). Note that the Arm Timeout is implemented at the embedded level, and overrides commands from the state machine when active.



ESTOP

The ESTOP state is the startup state of the system, and cuts off power to the actuators by disabling the actuator power relays. This results in immediate hydraulic pressure loss and rapid closing of the brakes. Any state can transition to the ESTOP state automatically, or by the pilot issuing an ESTOP command.

Automatic ESTOP Conditions include:

- Abnormal sensor readings (unexpected tube conditions, evidence of pod instability, sensor or CPU overtemperature, low battery voltage, safety-critical sensor failure).
- State estimation failure.
- State estimate predicts crash risk.
- Brake control failure (e.g. significant oscillation detected).
- Safety-critical communication watchdog timeout triggered.
- Total flight duration timeout.

Arming Timeout

A critical requirement of SpaceX during the first pod competition was that the pod would never engage its brakes while the pusher was actively pushing the pod. To prevent closing the brakes during launch, an Arm Timeout lockout is implemented on the actuator interface hardware in the LSM (see Section 7.1.1) to prevent brake closure until sufficient time has passed after the most recent Arm state was seen. This lockout supercedes all other commands to the LSM. The launch checklist ensures no launch will occur until our system is verified to be in the Arm state, so this ensures that no transition, software failure, or external command to the LSM can cause the brakes to close on the pusher.

Drive States

A set of states allow convenient management of the low-speed drive system:

- **Pre-drive** manages deployment of the low-speed system. It enables power to the low speed system actuator controllers, ensures that their configurations are nominal, verifies that the clamp is open, and finally opens the brakes and locks.
- **Drive** keeps the brakes open. In this state, the FSM acts as a simplistic driving controller, translating high-level drive commands from the interface into low-speed motor and clamp commands.
- **Shutdown** safely shuts down the system by opening the low-speed clamp (if it is active) and then gently closing the brakes.

Enabling Drive Mode requires the system to be in ESTOP, and transitions to Pre-drive. Pre-drive automatically transitions to Drive if successful (but can fail or be canceled to Shutdown). Drive transitions to Shutdown, which automatically closes out to ESTOP.

Flight States

A set of states manage the autonomous flight procedure:

- **Pre-arm** prepares the system for flight by opening the brakes (and opening and then disabling the low-speed actuator, if it is active), and initializing the state estimator and controller,
- **Arm** keeps the brakes open and waits for IMU information suggesting the pod is being pushed,

- **Launch** keeps the brakes open while waiting for IMU information suggesting the pod is no longer being pushed,
- **Flight** performs continuous brake adjustment to keep the pod on a safe, pre-planned trajectory,
- **Soft Stop** controls the brakes to an intermediate braking position to bring the pod to a quick stop (but gentler than ESTOP),
- **Shutdown** safely shuts down the system by opening the low-speed clamp (if it is active) and then gently closing the brakes.

Enabling Flight mode requires the system to be in ESTOP, and transitions to Pre-arm. Pre-arm automatically transitions to Arm if successful (but can fail or be canceled to Shutdown). Transitions from there occur automatically but can be automatically or manually transitioned to appropriate stop conditions.

7.3.2 Flight and Braking Control

For flight and braking control, two questions are critically important: how to brake, and when to brake. The brake hydraulics control, and the state estimation are discussed in this section.

Brake Hydraulics Control

The brake actuator system, as presented in Section 4.3.3, contains three actuators:


- A variable-speed, fixed-displacement hydraulic pump, which increases pressure in the brake cylinder and hence opens the brakes,
- A proportional valve, which releases pressure from the brake cylinder and hence closes the brakes,
- An on-off valve, which releases pressure from the brake cylinder and hence closes the brakes.

The brake caliper position is detected via a pair of linear potentiometers, which are lightly filtered to produce smoothed brake position estimates. Given a desired brake openness setpoint, our brake controller operates in either a *closing* or *opening* mode. These distinct modes were employed due to strong nonlinearities in the brake caliper dynamics when both the motor *and* proportional valve were being modulated. In the *opening* mode, all valves are closed and the motor is modulated using a slightly-overdamped PD controller to open the brakes to the desired position. In the *closing* mode, the motor is stopped, the on-off valve is opened, and the proportional valve is modulated using a separate slightly-overdamped PD controller to close the brakes to the desired position. Switching between these control modes was limited to 4 Hz , and in practice occurred very rarely.


Fig. 7.6 presents a characterization of this brake hydraulic controller. The pump motor saturated quickly, leading to a linear ramp of response times for opening (i.e. positive step) commands, with an initial hump near zero due to nonlinearities in the underlying system (including the dynamics of the fluid pressure, and deadband). The closing mode (i.e. negative step) is more complex, due to more dramatic nonlinearities in the proportional-valve-to-brake-velocity mapping. Response times were quite consistent, with little overshoot.

State Estimation in the Tube

The state estimation problem for an autonomous flight, in our case, requires estimating only a handful of parameters: the position, velocity, and acceleration of the pod along the length of the tube, relative to a desired fixed stopping location. Our state estimator fuses information from the two navigation IMUs



(a) Response times



(b) Overshoots

Figure 7.6: Step response behavior of the brake hydraulic controller, across a spectrum of requested brake position step sizes. Each point indicates a random brake position setpoint (evenly distributed between fully open and fully closed) issued to the brake controller, plotted by the change in position setpoint from the last command. Step commands are in inches, with positive numbers indicating commands to open, and negative numbers indicating commands to close. Response time indicates the time it took the controller to first get to within 0.1 inches of the setpoint. Overshoots indicates the farthest past the setpoint the brakes went during the control window.

(VectorNav VN-100s mounted on the pod body, above the suspension) with fiducial strip^e detections from a high-rate optical color detector. The IMUs are integrated to continuously produce estimated velocities and positions, while the strips are leveraged to periodically correct accumulated drift.

Specifically, we consider the state to be $x_k = \{q_k, \dot{q}_k, \ddot{q}_k\}$, for position along the tube q_k at time step k , and a covariance estimate of the state Σ_k , summarized in a Gaussian distribution $\gamma_k^{\text{state}} = \{x_k, \Sigma_k\}$. Between fiducial strip detections, we iteratively update the state estimate with a standard EKF process update.³¹ We use a first-order double-integrator process model

$$x_k = x_{k-1} + \delta t[\dot{q}_{k-1}, \ddot{q}_{k-1}, 0]^T.$$

IMU measurements, alongside a pod stopping detector, provide direct observations of \ddot{q} and \dot{q} respectively.

^eSpaceX furnished the tube with a 4" wide retroreflective orange strip every 100' on the inside of the tube; see Section 6.1.2 for details.

Fiducial detection status is reported from the fiducial module as a *time since last strip* and a paired *velocity estimation from last strip*. A fiducial strip detection event is detected by an observed decrease in the *time since last strip* back to near zero. Passing this around as a *time since last strip* resolves the message-dropping issue one might experience if detection events were broadcast on their own. Each strip detection must be treated carefully, for two reasons. First, each strip detection might be a false-positive. Because each strip detection must be weighted fairly heavily by the estimator to successfully counteract drift, outliers must be pruned carefully lest they corrupt the state estimate. Second, and more subtle, is that the position information from a strip detection is ambiguous in a discrete way: updating based on a strip detection inherently involves solving a correspondence between the positive detection and the set of strips available to be observed. Classification as a false-positive can conveniently be seen as an extra choice in the correspondence problem.

To consider those correspondences, when a detection event is detected, we consider each of a set of distinct hypotheses based on an investigation of the strip detection event in the context of our current state distribution γ^{state} . One hypothesis is that this detection was a false positive: this hypothesis is assigned a fixed false-positive-rate probability, which is estimated beforehand from experimental data. For the i^{th} candidate strip position (and its uncertainty) $\gamma_i^{\text{strip}} = \left(q_i^{\text{strip}}, \Sigma_{\text{strip}_i} \right)^{\text{f}}$, we create a true-positive detection hypothesis with posterior likelihood over pod state $\gamma_{tp=i}^{\text{state}} = \gamma^{\text{state}} \cdot \gamma_i^{\text{strip}}$. This product of Gaussians encodes the combination of the prior state likelihood with the i^{th} hypothesized strip position. We continue our state estimator with whichever of $\gamma_{fp}^{\text{state}}, \gamma_{tp=1}^{\text{state}}, \gamma_{tp=2}^{\text{state}}, \dots$ has highest total likelihood.

We note, and have prototyped, the relatively simple extension of this to a particle filter. The particular measurement model of the fiducial strip detection is well-suited to a dual-proposal particle filter.³²

7.3.3 GNC Comments and Review

Our brake hydraulic controller proved adequate for this test, and its consistency and robustness are positive factors when considering it an element in before-the-fact trajectory optimization. However, its relatively slow response times, even for small movements, would limit the ability of our pod to follow more complex braking trajectories, and to reject large disturbances during braking.

Our state estimator performed admirably during our test flight: though the tube lacks a ground-truth positioning system, videos from onboard the pod during flight indicate that we stopped a handful of meters in front of a strip just over 500 *m* down the tube. Our estimator agreed precisely, with a final position of 516 *m* several meters before a strip. Our first test flight focused on the sensing and levitation of our pod; implementing brake trajectory generator and controllers around this state estimator is the immediate next step, relying on dynamics data recovered during our test flight to form a predictive pod model for trajectory optimization and control.

As our final flight did not utilize the entire tube or require precise braking, we did not utilize any nontrivial brake trajectory generator or controller in our final system. Braking during our final run was executed from simple rules including maximum run time and farthest distance allowed. The above components – careful system identification of actuators, alongside a reliable and accurate state estimator – are direct prerequisites. Our experimental run indicates that our system would provide an excellent foundation on which to extend to much longer runs with active brake control.

^fThese candidate strips are the set of strips that we believe we may have detected during this detection event. In our system, we choose them with a simple heuristic: the N nearest strips to q , with Σ_{strip} populated from prior knowledge of strip localization precision. In our system, we use $N = 1$, which can be demonstrated to always be sufficient to detect the most-likely correspondence if state and strip position distributions are parameterized with Gaussians.

7.4 UI

Our flight control UI is written in PyQt, allowing for easy modification and augmentation of each interface. Our UI is organized as a set of panels, each focusing on exposing some modular part of our system functionality[§]. The hierarchical and modular organization allowed by PyQt widgets means that these individual panels, usually independently developed during subsystem testing, can be easily tied together into a larger unified interface window (see Fig. 7.7). Subpanels allow control and visualization of every system and diagnostic of network and sensor health.

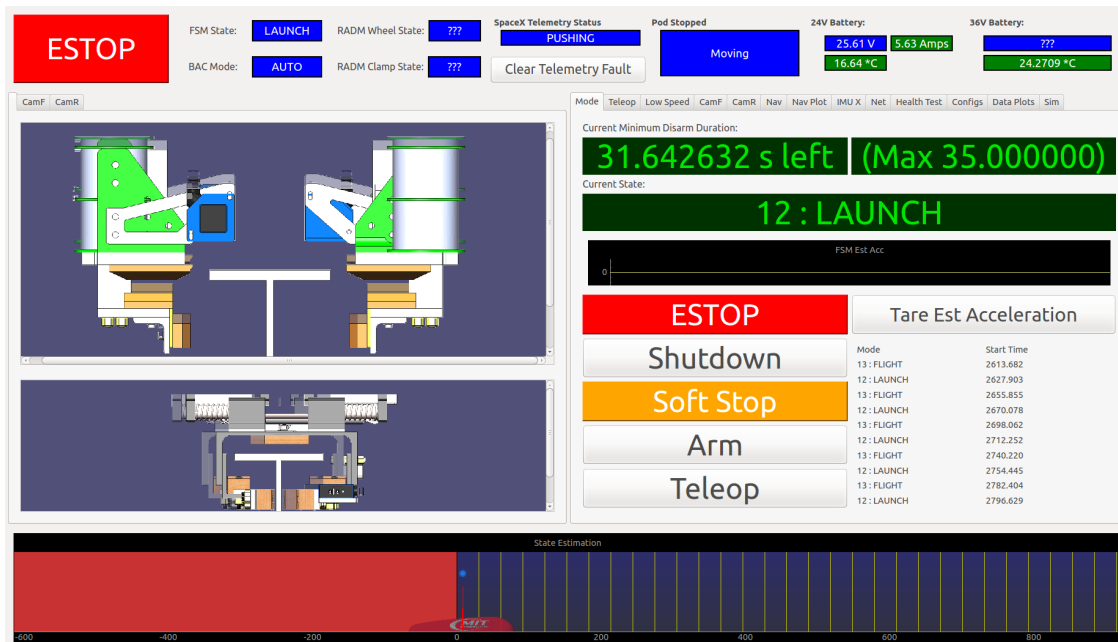


Figure 7.7: A view of our primary pod pilot interface.

The procman task launcher allows us to start all of the modules in our system, whether for flashing code, running a simulation, or running the full system during an actual test.

7.5 Data Logging

Data logging is straightforward and consists of recording all communications traffic passing through our system. This is done using the LCM Logger tool, which is a high-performance tool capable of streaming a large volume of distinct LCM messages to disk for later playback.

Our total network bandwidth peaks around 500 *kB/s*. Experiments on the ODROID show that we can consistently attain >6 MBps writing to our 64GB memory card. Hence, this simple logging system comfortably meets our requirement of logging multiple runs in a single day. However, if we needed to log much-higher-rate acceleration data, or higher-quality video, we would likely need both better hardware for a higher write speed, and likely a non-LCM based onboard logging system to handle this additional data without violating network traffic limits.

[§]A video of the UI during our competition run is available at <https://youtu.be/U12rGwQfNb4>.

7.6 Simulation

We developed a simulator alongside the rest of our system in order to test control logic and allow vertical integration testing of many system components. Our software's modular design based on universal LCM-based communication infrastructure made it possible to run almost the entirety of our real system against simulated sensor values to get significant coverage of the same code paths to be used during real runs. We used a simulator with published simulation IMU readings and both analog and digital signals that were piped (via a simulation-only intermediary layer) to simulations of the appropriate embedded modules.

The pod simulator models the pod's physical location in the tube, and models one-dimensional acceleration effects from the pusher, the brakes, and the low-speed actuators. Accelerometers and other physical sensors are simulated so that the control systems can be tested. The simulator is decoupled from the control state machine, which is run alongside it to ensure correct behavior in the face of raw noisy sensor data.

The ability to do both complete-, as well as hardware-in-the-loop-simulation proved invaluable for rapid development and testing of our brake actuator controller, state estimator, brake trajectory generator, and flight state machine. It made demonstration and certification of system components by SpaceX painless, as components under scrutiny (e.g. stopping detector, state estimator, and state machine) could be demonstrated quickly and easily in simulations and against data streaming live from the pod. It also made it possible to quickly add new features to the system without worrying about compromising the next run, as each change could be tested incrementally on the full-stack simulation to ensure that it did not produce undesired errors during vertical integration.

PART II

TESTING

LEVITATION TESTING CAMPAIGN

In order to validate the simulation results for each of the maglev systems (brakes, LCMs, and levitation skis), two test rigs were designed and tested. A record player style rotary test rig allowed the testing of the exact LCM geometry for speeds up to 50 m/s . It also allowed the testing of a small section of the brake array. A linear test rig was also constructed to test a representative lift ski to see gap height impact, wavelength, and array thickness up to 20 m/s . Test results matched well with simulation and provided confidence for large scale pod testing with the designed array geometries.

8.1 Rotary Test Rig

The motivation for building the rotary test rig was to measure the performance of both the lateral control modules and the eddy current brakes at speeds as close to those of the actual pod as possible. It was found that high enough speeds were impractical with the linear test rig (outlined below) so a spinning test rig was chosen instead.

8.1.1 Design

Two configurations were considered for the rotary test rig: a “record player” configuration where the array would be parallel to the face of the disc, and an “edge riding” configuration where the array would be curved over the edge of a disc.

The record player configuration is significantly simpler to manufacture since all the components are 2D extrusions, but has the downside of transverse (radial) variation in relative speed between the magnet array and the disc as well as variation in the angle between the disc velocity and the magnet orientation due to the fact that the magnet is only exactly tangent to the disc at one point. These variations can be mitigated by using a larger disc, but the rig is limited in size by what can be reasonable manufactured and safely spun up to speed.

While the edge riding configuration addresses the two speed issues above, it was dismissed because it would mean magnets would not be flush next to each other and it would also be difficult to manufacture suitable back iron. We had determined in simulation that even small gaps between adjacent magnets would result in significant changes to L/D . Alternatively, non-cuboid magnets could have been used but these would likely have been prohibitively expensive to custom manufacture. One method for creating this configuration is to use a large disc thicker than the width of the magnet arrays (~ 3 ”). This option, however, is dangerous and expensive and does not allow comparison of different track thicknesses. Another option is to mount a narrow ring to the outside of a central hub, but this would also be prohibitively expensive to manufacture and add significant complexity to the rig design. For these reasons a sufficiently large record player style rig was designed, an overview of which is shown in Fig. 8.1.




Figure 8.1: Overview of Rotary Test Rig.

Based on the typical lift/drag versus speed curves from 2D simulations (Fig. 8.2), it is found that a speed of 50 m/s allows for observing the entire drag peak with the wavelengths of magnets used for the LCM and brakes ($4''$ and $2''$ respectively). With the required tip speed defined, the diameter of the wheel was sized based on minimizing the transverse speed variation to less than 10%. Since the LCM magnet array is $2''$ wide, this set the minimum wheel radius at $20''$ for a wheel diameter of $40''$.

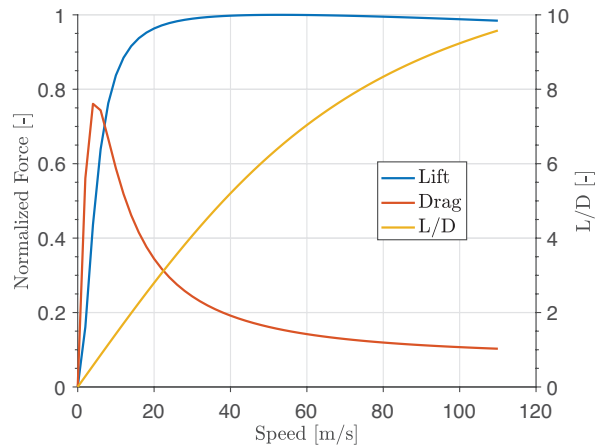


Figure 8.2: Lift and drag curves versus speed, obtained from 2D simulations.

Initially a single piece disc of constant thickness was considered for simplicity. Multiple discs would be made to test the LCM on a plate thickness of $0.312''$ (the thickness of the rail web) and to test the eddy current brakes on a plate thickness of $0.413''$ (the thickness of the rail flanges). It was found that the natural frequencies for plates of these thickness would be too low to allow for safe operation of the rig.³³ To alleviate this issue, a multi piece rotor was designed with a $1''$ thick central hub to add rigidity onto which thinner outer rings are bolted. This gives a $> 3\times$ safety factor between the lowest natural frequency of the




Figure 8.3: Outer disc of rotary test rig.

plate and the rotational speed of the system.

The diameter and the required speed results in a required rotational speed of approximately 100 *RPM*. The power requirements are dictated by the maximum drag produced by the LCM of 50 *N*. At the tip speed of 50 *m/s*, this results in a power dissipation of 2.5 *kW*. A 5 *hp* (3.5 *kW*) 1200 *RPM* motor was chosen to achieve these requirements.

Rig Construction

The rig was designed to be quick to manufacture and assemble. A large cast iron frame motor is used to provide a rigid base for the rotating components. An off-the-shelf sprocket is used as a drive hub as it already incorporates the necessary keyway and set screws for mounting on the motor shaft. The outer teeth are not used. The front face of the sprocket was faced on a lathe and had 6 holes tapped into it to allow for the mounting of the central hub.

The central hub is constructed from 1" thick MIC-6 aluminum sheet. This is a dimensionally stable and precision ground cast aluminum alloy with flatness to within 0.005". The rough outer diameter and all mounting holes are cut on a waterjet. The final diameter and central bore (the locating feature for aligning the hub and the motor shaft) were finished on a rotary table on a manual milling machine to within 0.002" concentricity. 38 press-in-studs were attached to the hub to provide a means of clamping on the outer discs and two reamed holes and shoulder bolts were used to align the two pieces.

The outer discs were cut on the waterjet and all holes deburred to provide a flat clamping interface between the hub and the discs. Balancing of the assembly was done by attaching extra washers and nuts to the light side of the disc until no distinct heavy spot could be found when mounted on the motor shaft. This effectively limits the imbalance to the drag from the motor bearings which is extremely small in comparison to the inertia of the wheel.

A 12" x 48" optical breadboard clamped to a heavy workbench served as a mounting base for the motor and a convenient means of mounting various test fixtures for measuring the lift and drag of the LCM and




Figure 8.4: LCM 2-axis load cell setup for simultaneously measuring lift and drag.

the brake module. For the LCM testing, two moment compensated (single-point) load cells were joined at right angles to one another to create a two axis force measurement device to simultaneously measure lift and drag, as shown in Fig. 8.4.

Rig Electronics and Control

The rig speed is controlled with a 3-phase variable frequency drive. Strain gauge amplifiers are used to amplify and condition the signal from the load cells and allow for it to be recorded by an analog to digital converter. A standard computer power supply provides a source of clean DC power for the amplifiers. A hall-effect sensor is used to measure wheel speed by counting the passing frequency of the mounting studs for the outer rings.

8.1.2 Data Analysis

The rotary test rig is used to measure the lift (lateral restoring) force and drag forces on the LCM. Forces are measured at all possible combinations of offset ranging 0 – 13 *mm* and speed ranging 0 – 50 *m/s*. Interpolated surface plots for the lift and drag data are shown in Fig. 8.5.

The rig shows excellent repeatability between runs and good agreement with numerical results. Comparison between four consecutive runs is shown Fig. 8.6.

The key parameter that we hoped to validate with the testing was the stiffness of the LCM as this was critical for the pod to maintain lateral stability and heading. The comparison of the LCM stiffness at 50 *m/s* compared with 3D simulations using ANSYS Maxwell are included in Fig. 8.7. The maximum error is less than 7%.

At lower speeds, the errors between measured and simulated lift force were significantly larger (15-20%). This is illustrated in the plot of lift and drag as a function of speed at a 10 *mm* offset. This error could result in larger pod movements than the LCM’s design targets at low speed.

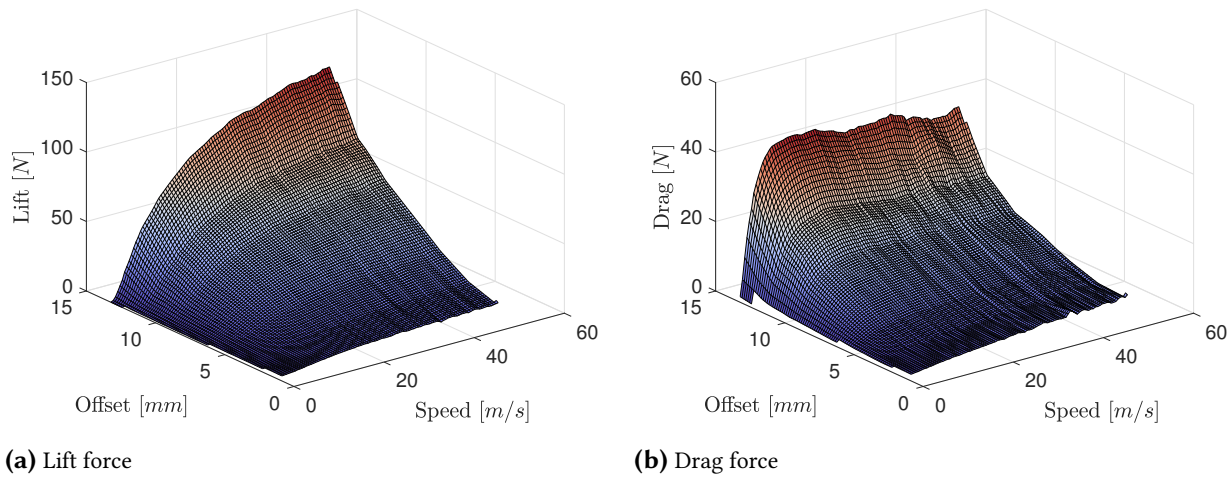


Figure 8.5: Experimental results from the rotary test rig for the LCM configuration.




Figure 8.6: Repeatability between runs of rotary test rig.

8.2 Linear Test Rig

Besides testing the LCM and brake arrays, it was desired to also test the levitation arrays. The primary goal was to design a rig that could be used to test an array that was identical in configuration and as close as possible in size to the full-scale array, such as to minimize any uncertainty about the performance of our levitation skis, while corroborating our simulation data.

8.2.1 Design

To keep the configuration as close as possible to our full ski array design we decided the nominal test array would have the same thickness ($\frac{3}{4}$ "), width (2"), gap height (10 mm), magnet grade (N42), and number of periods (2) as the final ski design. Furthermore, the aluminium plate on which the configuration was riding had to be the correct alloy (6101-T6) and have the same thickness ($\frac{1}{2}$ ") as the actual track. Note that at the initial time of design the thickness and width values of the array were slightly different ($\frac{1}{2}$ " and 4", respectively) but this did not dramatically affect the design of the rig.




Figure 8.7: Comparison between stiffness of LCM configuration as found from experimental and simulation results.




Figure 8.8: Comparison of experimental and simulation data for LCM configuration at 10 mm offset, as measured using the rotary test rig.

A rotary test rig was considered for the levitation array testing as well, again considering a “record player” and “edge riding” configuration. The edge riding configuration was dismissed for the same reasons as before. However, because the levitation arrays are larger than the LCM and brake arrays, also the record player configuration is problematic. For the levitation arrays the speed variation across the transverse direction of the array for a reasonably sized plate would be too high. Furthermore, it would be prohibitively expensive to obtain an Al 6101-T6 plate large enough to make such a disc out of.

A test rig where the array is stationary but the aluminum plate shoots under it, was also considered. For cost and safety reasons, this idea was not further pursued.

This left a linear test rig where the a magnet array moves relative to a stationary track. Given the short time frame, the linear test rig was designed with ease of manufacturing in mind. A description of the design of the linear test rig is provided below.

The track is shown in Fig. 8.9. As shown, only part of the track has Al 6101 I-beams on it, which kept the cost of the rig down because we only intended to measure at relatively high speeds. The total track is 75 ft long. A simple simulation of the cart dynamics showed that a 20 m long track would be sufficient to

achieve a high enough speed (20 m/s) at the start of the Al 6101 I-beams, to have a long enough (1 s) run time over the Al 6101 I-beam to collect enough data, and to have enough room for sufficient deceleration from the Al 6101 I-beams to reduce the speed before the end of the track. Due to standard sizes of I-beams in the US (25 ft), the eventual track ended up being 75 ft . A foam pit is installed at the end of the track to complete slow down the cart. This was the cheapest and safest way of slowing the cart down to a full stop. MIT Facilities was kind enough to let us use the basement in the MIT Stata Center during night hours for testing.




Figure 8.9: Linear Test Rig set-up.

The cart itself then is made almost entirely out of aluminum, for obvious reasons (i.e. magnets). An overview of the cart is shown in Fig. 8.10. The outer structure of the cart is designed for ease of manufacturing (all of the major pieces were cut using a waterjet). The gussets (two external brackets) give the cart rigidity and hold its shape. There are two sets of wheels (top and bottom) that were custom designed to be able to manage the target speeds and maximum loads. Ceramic bearings are used to avoid using ferritic material for the bearings, as they are placed close to the magnets. There are four wheels on the bottom but only two wheels on top to prevent over-constraining the cart as it rolls over unevenness in the track. It was important to put more wheels on the bottom because the upward force on the cart from the lift is far greater than the down force due to its weight. Flanged wheels are used to keep the cart aligned with the track. This is not the most efficient design because the flanges rub against the i-beam, but it allows for only using two sets of wheels and to just mount the wheels on bolts instead of needing pillow-block bearings (which would be needed for side wheels, such as on a rollercoaster).




Figure 8.10: Cart with magnets for Linear Test Rig.

The cart holds two load cells to measure both lift and drag. These were sized with a safety factor of two on the expected maximum loads. Single point load cells were used to avoid having to use configurations of multiple load cells for each direction, which would have required some type of summing circuitry that would have added complexity. The slider/track assembly in the cart decouples the lift and drag measurements, as it allows the drag load cell to only measure drag and the lift load cell to only measure lift. Finally, there are spacers between the lift load cell and the slider/top plate, which allow the load cell to deflect and therefore measure load, and to allow for changes to the magnet gap height.

A bungee was deemed to be the simplest and cheapest way of achieving the necessary velocities. Using a high speed winch was considered for a while, but the idea was rejected because developing a control system that would stop the winch or disconnect it from the cart before the end of the track was deemed too complicated. Also, no winch was found that could achieve the necessary speeds and loads. Speargun rubber is used for the bungee, for safety and ease of use. Speargun rubber is extremely tough and can be extended significantly before failure. It also has a small hollowed out core that allows for using certain adapters with constrictor knots. Here, special speargun fishing adapters are used to connect the speargun rubber to the cart and track.

An overarching requirement for this test rig was to demonstrate that the levitation arrays would be able to provide enough lift to lift off the pod. This meant we had to at least reach our liftoff speed, which we had estimated to be around 10 m/s . The maximum speed of the cart on this rig is therefore 20 m/s to be confident we could satisfy this requirement, and at the same time could try to characterize the drag peak.

8.2.2 Data Analysis

For the linear test rig testing, three different magnet arrays were tested at several gap heights to confirm the effect of wavelength, and validate our 3D simulations of stiffness and L/D. Below are the lift and drag profiles of the first two runs using the same geometry. The results are quite noisy due to the vibrations of the cart moving along the track but the data is still quite repeatable, as shown in Fig. 8.11.

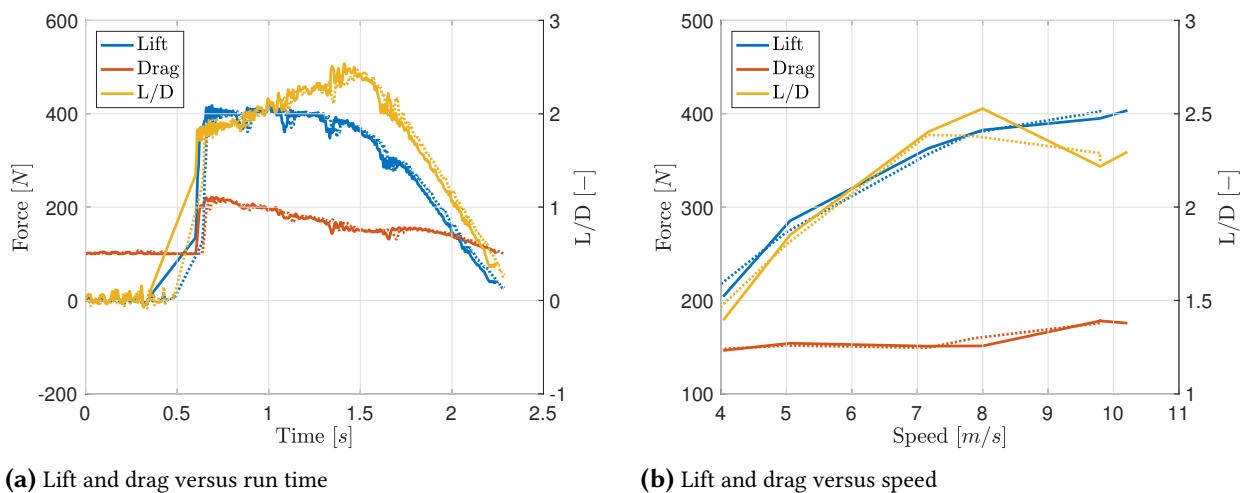


Figure 8.11: Experimental results for Case A from linear test rig, showing excellent repeatability.

The data is smoothed to remove the noise and better compare the lift and drag trends. The speed was measured at several points along the track as mentioned above, as shown in Fig. 8.12. Below is the data for several runs and the lift and drag vs speed. The data is corrected for inertial loads due to acceleration and deceleration of the cart.

An interesting note from this data are the dips in lift that occur through the run. These dips correspond to the connections between track plates where the eddy currents can not be continuous through the plate. At this junction, this is approximately a 5-10% decrease in lift as the eddy currents have to be generated again in the new plate. This effect will also exist on the full size track and is the primary vibration input to our suspension system. The frequency of this input is determined by the speed of the pod over the known track lengths.

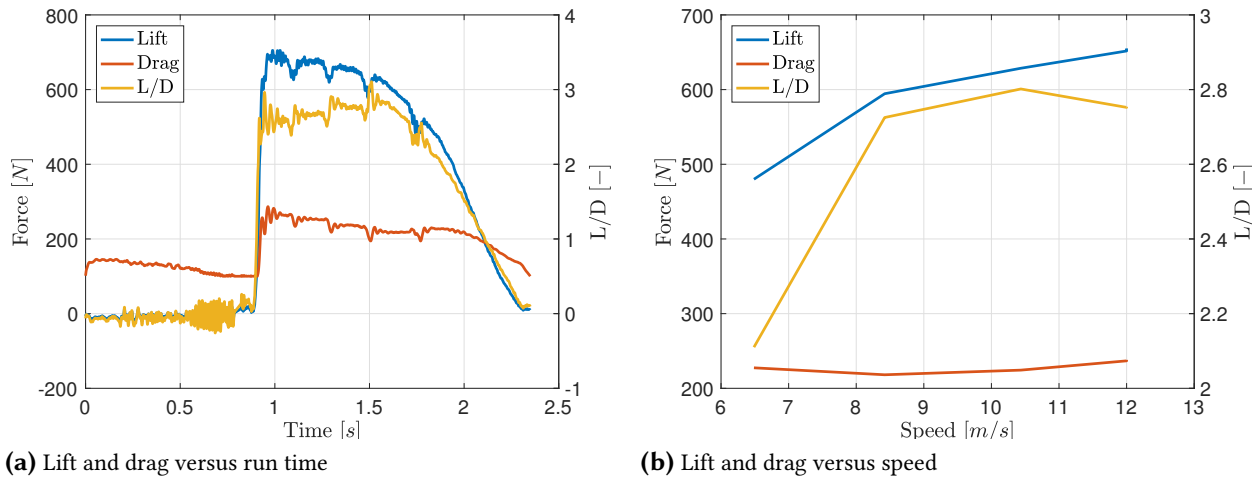


Figure 8.12: Experimental results for Case E from linear test rig.

We show a comparison between numerical and experimental results for several runs in Table 8.1 and Fig. 8.13. The numerical results are quite close to the experimental results, giving additional confidence to the levitation design which was based off of the numerical results.

Table 8.1: Comparison of experimental results with numerical results for linear test rig

| Array length [in] | Magnet configuration | | | | Simulation data at 10 m/s | | | Simulation data at peak speed | | | |
|-------------------|----------------------|----------------------|---------------|-----------------|---------------------------|----------|---------|-------------------------------|----------|---------|-----|
| | Array width [in] | Array thickness [in] | # periods [-] | Gap height [mm] | Lift [N] | Drag [N] | L/D [-] | Lift [N] | Drag [N] | L/D [-] | |
| A | 16 | 2 | 1/2 | 2 | 12 | 426 | 243 | 1.8 | 400 | 210 | 1.9 |
| B | 16 | 2 | 3/4 | 2 | 12 | n/a | n/a | n/a | 680 | 290 | 2.3 |
| C | 16 | 2 | 3/4 | 2 | 9 | n/a | n/a | n/a | 850 | 310 | 2.7 |
| D | 16 | 2 | 3/4 | 2 | 6 | 1242 | 604 | 2.1 | 1100 | 390 | 2.8 |
| E | 16 | 2 | 3/4 | 1 | 6 | 700 | 310 | 2.3 | 720 | 260 | 2.8 |
| F | 16 | 2 | 3/4 | 1 | 9 | 600 | 250 | 2.4 | 550 | 250 | 2.2 |
| G | 16 | 2 | 3/4 | 1 | 12 | 500 | 200 | 2.5 | 450 | 220 | 2.0 |




Figure 8.13: Comparison of numerical and experimental results for linear test rig.

VIBRATION TESTING CAMPAIGN

Several components were vibration tested at MITRE, one of the team’s sponsors. This was done to test that these critical components could withstand an environment with potentially a lot of vibrations, such as disturbances from the track before lift-off, or disturbances introduced by the SpaceX pusher.

During these tests, the plates with electronics, the low speed system, and the brakes were tested. All systems were vibration tested in all three directions with shock tests, sine sweeps, and random vibrations. Afterwards, the components were mechanically inspected, and the pod was run with the electronics system, after which no failure was observed. During testing also no natural frequencies of the system were excited, as shown in Figures 9.1 to 9.3. The frequencies the components were tested at were considerably higher than what would be expected during a run, so there was additional confidence that no natural frequency would be excited in the system during a run.

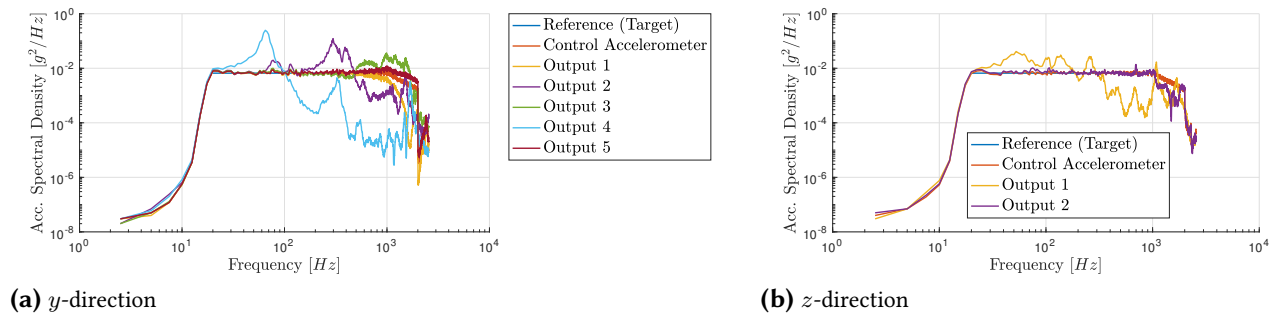


Figure 9.1: Vibration testing results for the electronics system for a random vibration test.

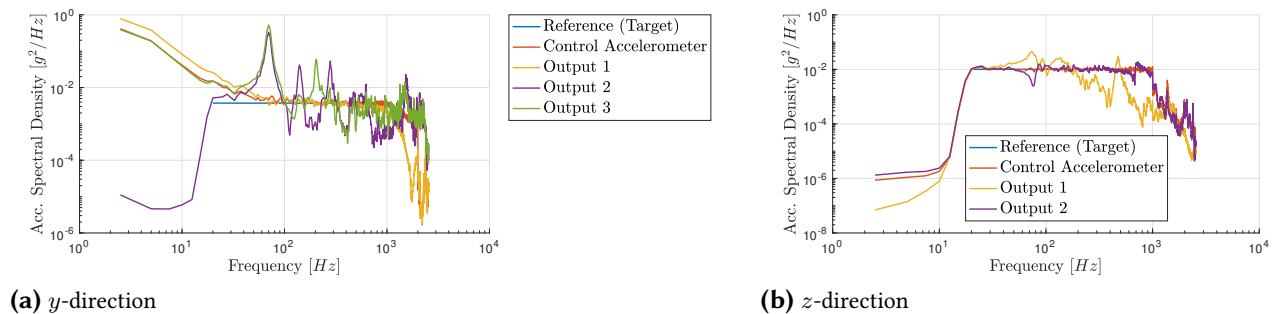


Figure 9.2: Vibration testing results for the brakes for a random vibration test.




Figure 9.3: Vibration testing results for the low speed system for a random vibration test in y -direction.

CHAPTER 10

COMPETITION LOG

This chapter gives a brief overview of the last phase of the competition in Hawthorne, CA. The official test week before competition ran from Sunday January 22nd through Thursday January 26th. The competition itself ran from Friday January 27th through Sunday January 30th. During the test week and the start of the competition, teams had to pass several tests before being allowed to do a vacuum Hyperloop run on the last day of competition. These tests consisted of showing critical functionality, such as smooth operation under vacuum conditions, checking fault-tolerance of system such that the pod does not brake while being pushed, demonstrating levitation on an small external subtrack, et cetera. In the end, only three teams – including the MIT Hyperloop Team – were allowed to do a vacuum Hyperloop run on the last day of competition.

Friday January 20th, 2017 – Saturday January 21st, 2017

The test crew arrived in LA on Friday. After the first test weekend in October 2016, the pod had been stored in the garage of a few team alumni and here the pod was being prepared for the first tests. Friday consisted on upgrading some parts which were improved after test weekend. Furthermore, the electronics had to be reinstalled, which had been taken back to Cambridge to continue trouble shooting. On Saturday, the team performed several tests with the pod on a cart to demonstrate the state machine was working properly.

Sunday January 22nd, 2017

Unfortunately, the first day of testing was canceled due to a severe storm, causing the test site to be under water. The team did attend the team-specific safety briefing at SpaceX's headquarters.

Monday January 23rd, 2017

Monday was spent mostly on setting up everything for the rest of the week. The truck was unloaded, and the tent area was set up. We also passed the structural inspection, part of the functional test, and got approval for all our testing procedures, navigation test and state diagram test (apart from Arm Timeout, which has to be done with Steve Davis). Unfortunately, the BMS from both batteries started playing up, not allowing for battery temperature read-outs.

Tuesday January 24th, 2017

On Tuesday, a lot of time was spent to solve the BMS issue. In the morning, the 24V battery communications got fixed. In the afternoon, it was decided to switch to just using the motor controller for the voltage of the 36V battery, and use a temperature sensor on the case of the battery. This allowed us to pass all functional tests.

Wednesday January 25th, 2017

The vacuum chamber was on the schedule first. There was a small problem during the vacuum chamber test that was solved by reflashing code, other than that the pod worked perfect. The rest of the day was spent passing the state diagram test, improving the electronics reliability and waiting for a slot on the external subtrack. Here the compatibility with the pusher was tested, as well as brake actuation at low speed, which all worked as expected. Note that this test was performed using dummy skis (i.e. skis without magnets in them), allowing us to focus on the operation of the brakes. Otherwise, the pod would slow down so quickly due to the magnet drag that the brakes would not have a chance to engage. At night, we continued to work on the electronics and also installed the magnet skis overnight.

Thursday January 26th, 2017

We were on the external subtrack early in the morning, now with magnet skis. Due to the large amount of magnet drag at low speed the SpaceX pusher had troubles getting us up to speed and could only get us up to *7 mph*, which was not enough to levitate. We found out that the track tolerances were worse than expected and had to increase our magnet gap height to prevent the magnet covers from scraping the track. Later that day, the pod went into the large Hyperloop tube to test our low speed system and the tolerances on that track. We had to play around with the actuation force of the low speed actuator, after which it was able to propel the pod at low speed. The magnet covers were still scraping due to worse tolerances and therefore at night the magnet gap height was increased even further. The fiducial sensor was also tested in the tube, working great.

Friday January 27th, 2017

The goal for Friday was to do a push test in the Hyperloop tube in ambient conditions. Unfortunately, there were several teams who wanted to do the same and the pusher was also having issues. In the end we went in the tube at 3pm. The pod was supposed to get up to *25 mph* but the pusher kept shutting down after *7 mph*. It turned out that our pod's antennas were interfering with the pusher's communication system. Although this was a quick fix, we were pulled off the track and had to try again on Saturday. At night, a mount for the antennas on top of the shell was 3D printed.

Saturday January 28th, 2017

Due to several teams wanting to do a push test, we had to wait for most of the day. We installed the 3D printed antenna mount in the morning. We also did our design review in front of several SpaceX judges. We found out that it was taking teams too long to load their pod on and off the track, therefore we spent the afternoon improving our operational procedures to ensure we could load our pod quickly. We were finally allowed to go into the tube at 9pm, and were out 55 minutes later. The test went very smooth, demonstrating we were ready for a vacuum Hyperloop run. After the run, a quick structural inspection was performed. Later that night, we got confirmation that we were indeed allowed to run on Sunday.

Sunday January 29th, 2017

We were running the first vacuum Hyperloop run of the day – the first ever vacuum Hyperloop run – which meant the pod needed to be ready by 9.30AM. The early morning was spent trying to get the pod ready for flight, and trying to figure out what we had to demonstrate during the run. Unfortunately, the SpaceX pusher could only push all pods up to *55 mph* instead of our design speed (*250 mph*). This meant that our



pod would be operating around the drag peak from the magnets, resulting in a large pod deceleration once the pusher releases the pod. It was therefore anticipated that we would not reach the end of the tube.


Our whole run lasted 35 seconds (30 second push, 5 second coast), the data of which is shown in Chapter 11. During this run, we successfully demonstrated stable magnetic levitation^a.

During the award ceremony later that night, the team won the **Safety & Reliability Award** and **third place in Design & Construction**.


^aFootage from the competition run is available at <https://youtu.be/0JEZkczITFk> and <https://youtu.be/U12rGwQfNb4>.

COMPETITION DATA ANALYSIS

The plethora of sensors on the pod recorded our runs in great detail, to be used for the design of future Hyperloop pods. In this chapter we show the data from our final competition run on Sunday January 30th, 2017.



(a) Acceleration and velocity throughout competition run.




(b) Deceleration versus velocity during the deceleration phase. The drag peak is clearly visible around ~ 6 m/s.

Figure 11.1: Acceleration and velocity data from the competition run. Acceleration data is filtered from IMU recordings, while the velocity is estimated from integrated IMU data and the fiducial markings.


The acceleration and velocity profile of that run is shown in Fig. 11.1. The pod is accelerated around

1 m/s^2 by the SpaceX pusher after which the pod is released after $\sim 35 \text{ s}$. Because the maximum speed attained by the pod is so low (due to the pusher limitations), the drag from the magnets is quite high, which causes the pod to decelerate quite fast. That deceleration is plotted as a function of the pod velocity in Fig. 11.1b. That plot clearly shows a drag peak around $\sim 6 \text{ m/s}$, which agrees with the simulation data that also showed a drag peak around that speed (see Fig. 8.2).

The data from the ski gap height sensors is shown in Fig. 11.2. These sensors measure the distance from the ski to the ground and are used to track the magnet gap height throughout the run. Fig. 11.2a clearly shows that levitation is achieved throughout the run. The gap heights as a function of pod velocity are shown in Fig. 11.2b, which shows that the pod fully achieves levitation around $\sim 7 \text{ m/s}$ and the increase in gap height with pod velocity is approximately linear.



(a) Ski gap heights versus run time.



(b) Ski gap heights versus pod velocity.

Figure 11.2: Data from ski gap height sensors showing magnetic gap height variation during run. Note that the front gap height sensors stopped recording half-way through the run due to a loose cable.

To show the suspension travel during the run, the data from the ski potentiometers is shown in Fig. 11.3. The suspension is compressed more in front than the in the back. This is because the pusher pushes the pod above its center of gravity, causing the pod to pitch forward. Then during the last part of the deceleration phase the pod again pitches forward.

Finally, we investigate the lateral movement of the pod by showing the data from the gap height sensors – measuring the lateral distance to the rail – and the potentiometers – measuring the suspension




Figure 11.3: Data from ski potentiometers showing the suspension travel throughout the run. Note that the front potentiometers stopped recording half-way through the run due to a loose cable.

movement in the lateral control module – inside the lateral control module in Fig. 11.4. For most of the run (especially when the pusher is pushing the pod) the lateral control modules seem to be maxed out against the hard stops. This is also the reason why the footage of the run is quite loud^a. Once the pod is coasting the lateral variation is substantial. During a subsequent run the magnet gap in the lateral control module could have been reduced to increase the restoring force in the lateral control modules.




Figure 11.4: Data from lateral control module sensors showing the lateral movement throughout the run. Note that the front sensors stopped recording half-way through the run due to a loose cable.

^aFootage from the competition run is available at <https://youtu.be/0JEZkczITFk> and <https://youtu.be/U12rGwQfNb4>.

BIBLIOGRAPHY

- ¹ SpaceX, “Hyperloop Alpha,” *SpaceX*.(Online Article). http://www.spacex.com/sites/spacex/files/hyperloop_alpha.pdf, 2013. Cited on pp. 15, 16, 19, 23, 57, and 64.
- ² KPMG, “Pre-feasibility study Stockhold–Helsinki using Hyperloop One technology – Short summary,” <https://home.kpmg.com/content/dam/kpmg/pdf/2016/07/fs-links-pre-feasibility-study-summary.pdf>, July 2016. Cited on p. 15.
- ³ Schafer, A. and Victor, D. G., “The future mobility of the world population,” *Transportation Research Part A: Policy and Practice*, Vol. 34, No. 3, 2000, pp. 171–205. Cited on p. 15.
- ⁴ Zimmerman, L., “MIT students win first round of SpaceX Hyperloop contest,” *MIT News*, February 2016, Accessed 10/23/2016 from <http://news.mit.edu/2016/mit-students-win-first-round-spacex-hyperloop-contest-0201>. Cited on p. 15.
- ⁵ Chin, J. C., Gray, J. S., Jones, S. M., and Berton, J. J., “Open-Source Conceptual Sizing Models for the Hyperloop Passenger Pod,” *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, Kissimmee, Florida, January 5–9 2015. Cited on pp. 15, 19, and 64.
- ⁶ Decker, K., Chin, J., Peng, A., Summers, C., Nguyen, G., Oberlander, A., Sakib, G., Sharifrazi, N., Heath, C., Gray, J. S., and Falck, R., “Conceptual Sizing and Feasibility Study for a Magnetic Plane Concept,” *55th AIAA Aerospace Sciences Meeting*, Grapevine, Texas, January 9–13 2017. Cited on p. 15.
- ⁷ Goddard, R. H., “The Limit of Rapid Transit,” *Scientific American*, Vol. 101, No. 21, 1909, pp. 366. Cited on p. 16.
- ⁸ Bachelet, E., “Levitating Transmitting Apparatus,” March 1912, US Patent 1,020,942. Filed April 2, 1910, *United States of America*. Cited on p. 16.
- ⁹ Salter, R. M., “The very high speed transit system,” Tech. Rep. RAND-P-4874, RAND Corporation, Santa Monica, California, 1972. Cited on p. 16.
- ¹⁰ Kantrowitz, A. and Donaldson, C., “Preliminary Investigation of Supersonic Diffusers,” Tech. Rep. Advance Confidential Report L5D20, National Advisory Committee on Aeronautics, Langley Memorial Aeronautical Laboratory, Langley Field, Virginia, 1945. Cited on pp. 16 and 57.
- ¹¹ Torenbeek, E., *Synthesis of subsonic airplane design : an introduction to the preliminary design, of subsonic general aviation and transport aircraft, with emphasis on layout, aerodynamic design, propulsion, and performance*, Delft University Press Nijhoff Sold and distributed in the U.S. and Canada by Kluwer Boston, Delft The Hague Hingham, MA, 1982. Cited on p. 19.
- ¹² Raymer, D. P., *Aircraft Design: A Conceptual Approach and Rds-student, Software for Aircraft Design, Sizing, and Performance Set (AIAA Education)*, AIAA (American Institute of Aeronautics & Ast, 2006. Cited on p. 19.

- ¹³ Drela, M., “N3 Aircraft Concept Designs and Trade Studies – Appendix,” Tech. Rep. NASA CR-2010-216794/VOL2, 2010. Cited on p. 19.
- ¹⁴ Lukaczyk, T., Wendorff, A. D., Botero, E., MacDonald, T., Momose, T., Variyar, A., Vegh, J. M., Colonno, M., Economon, T. D., Alonso, J. J., et al., “SUAVE: An Open-Source Environment for Multi-Fidelity Conceptual Vehicle Design,” *16th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 2015, p. 3087. Cited on p. 19.
- ¹⁵ Khatait, J. P., “Design and development of orifice-type aerostatic thrust bearing,” *SIMTech Technical Reports*, Vol. 6, No. 1, 2005, pp. 7–12. Cited on p. 20.
- ¹⁶ Meeker, D., “Finite Element Method Magnetics,” *FEMM*, Vol. 4, 2010, pp. 32. Cited on pp. 25, 27, and 29.
- ¹⁷ Opgenoord, M. M. J. and Caplan, P. C., “On the Aerodynamic Design of the Hyperloop Concept,” *35th AIAA Applied Aerodynamics Conference, Denver, Colorado, Paper No. AIAA-2017-3740*, June 5 – 9 2017. Cited on p. 54.
- ¹⁸ Knudsen, M., “Die Gesetze der Molekularströmung und der inneren Reibungsströmung der Gase durch Röhren,” *Annalen der Physik*, Vol. 333, No. 1, 1909, pp. 75–130. Cited on p. 55.
- ¹⁹ Chen, G., *Nanoscale Energy Transport and Conversion: a Parallel Treatment of Electrons, Molecules, Phonons, and Photons*, Oxford University Press, New York, 2005. Cited on p. 55.
- ²⁰ Drela, M., “A User’s Guide to MTFLOW 2.01,” Tech. rep., Massachusetts Institute of Technology, Cambridge, MA, 2010. Cited on pp. 56 and 58.
- ²¹ Van Ingen, J. L., “A Suggested Semi-Empirical Method for the Calculation of the Boundary Layer Transition Region,” Tech. Rep. Rapport VTH-74, Technische Hogeschool Delft, Vliegtuigbouwkunde, Delft, The Netherlands, 1956. Cited on p. 56.
- ²² Smith, A. M. O. and Gamberoni, N., “Transition, Pressure Gradient, and Stability Theory,” Tech. Rep. ES 26388, Douglas Aircraft Co., El Segundo, CA, 1956. Cited on p. 56.
- ²³ Drela, M., *Flight Vehicle Aerodynamics*, MIT Press, 2014. Cited on p. 59.
- ²⁴ Dodbele, S. S., “Effects of Forebody Geometry on Subsonic Boundary-Layer Stability,” Tech. Rep. CR NAS1-17919, Vigyan Research Associates, Inc., prepared for the National Aeronautics & Space Administration, Hampton, Virginia, 1990. Cited on p. 60.
- ²⁵ Achenbach, E., “Experiments on the Flow Past Spheres at Very High Reynolds Numbers,” *J. Fluid Mech.*, Vol. 54, No. 3, 1972, pp. 565–575. Cited on p. 61.
- ²⁶ Bearman, P., “On Vortex Shedding from a Circular Cylinder in the Critical Reynolds Number Regime,” *Journal of Fluid Mechanics*, Vol. 37, No. 03, 1969, pp. 577–585. Cited on p. 62.
- ²⁷ Menter, F. R., “Zonal Two Equation $\kappa-\omega$ Turbulence Models for Aerodynamic Flows,” *23rd Fluid Dynamics Conference*, July 6–9 1993. Cited on p. 63.
- ²⁸ Menter, F. R., “Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications,” *AIAA journal*, Vol. 32, No. 8, 1994, pp. 1598–1605. Cited on p. 63.
- ²⁹ Menter, F., Kuntz, M., and Langtry, R., “Ten Years of Industrial Experience with the SST Turbulence Model,” *Proceedings of the Fourth International Symposium on Turbulence, Heat and Mass Transfer*, edited by K. Hanjalić, Y. Nagano, and M. J. Tummers, Begell House, 12-17 October 2003. Cited on p. 63.



- ³⁰ Huang, A. S., Olson, E., and Moore, D. C., “LCM: Lightweight communications and marshalling,” *Intelligent robots and systems (IROS), 2010 IEEE/RSJ international conference on*, IEEE, 2010, pp. 4057–4062. Cited on p. 87.
- ³¹ Gelb, A., *Applied Optimal Estimation*, MIT Press, Cambridge, Massachusetts, 1974. Cited on p. 96.
- ³² Thrun, S., Fox, D., Burgard, W., and Dellaert, F., “Robust Monte Carlo localization for mobile robots,” *Artificial intelligence*, Vol. 128, No. 1-2, 2001, pp. 99–141. Cited on p. 97.
- ³³ Leissa, A. W., “Vibration of Plates,” Tech. Rep. NASA SP-160, Scientific and Technical Information Division, National Aeronautics and Space Administration, Washington, D.C., March 1969. Cited on p. 102.

LIST OF ABBREVIATIONS

- CFD** Computational Fluid Dynamics. 55, 67
- CFRP** Carbon Fiber Reinforced Plastics. 66
- CNC** Computer Numerically Controlled. 30, 68
- DGND** Digital Ground. 76
- DOF** Degree of Freedom. 39, 88
- EDS** Electrodynamic Suspension. 20, 21, 23, 24
- EKF** Extended Kalman Filter. 89, 96
- EMS** Electromagnetic Suspension. 23, 24
- FEA** Finite Element Analysis. 37
- LCM** Lateral Control Modules. 34, 40, 41, 43, 44, 77, 101–106
- LCM** Lightweight Communications and Marshalling. 87, 90–92, 98, 99
- LVIT** Linear Variable Inductive Transducer. 77
- MDF** Medium Density Fiberboard. 68
- MS** Margin of Safety. 67
- NAP** Network Access Panel. 72, 86, 129
- OML** Outer Mold Line. 66
- PCB** Printed Circuit Board. 52, 71, 72, 75–78, 80–83, 85, 86, 125
- PGND** Power Ground. 76
- PRI** Pod Receiver Interface. 52, 86
- QI** Quasi-Isotropic. 66, 67
- VARI** Vacuum Assisted Resin Infusion. 68

APPENDICES

APPENDIX A

PROJECT SCHEDULE

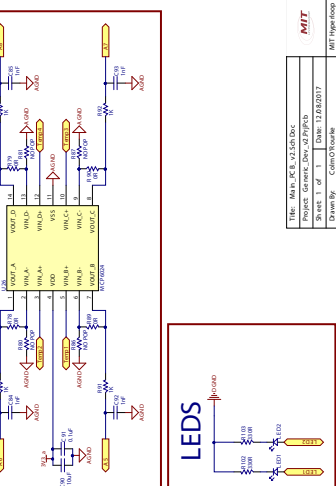
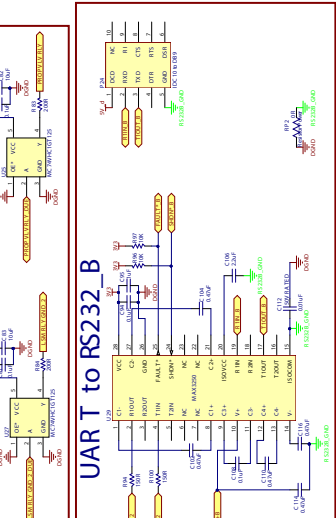
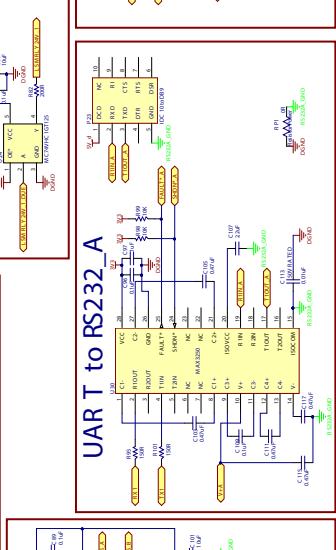
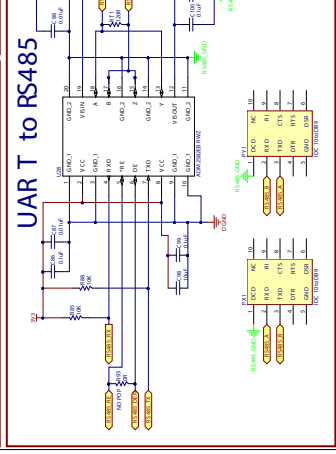

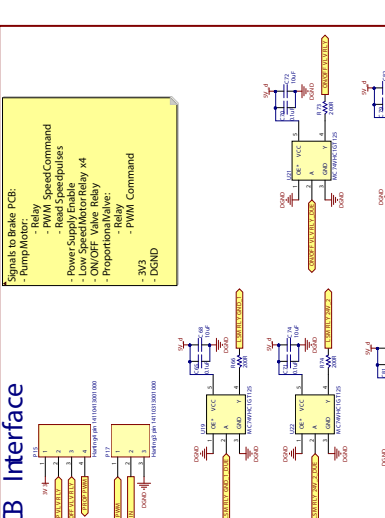
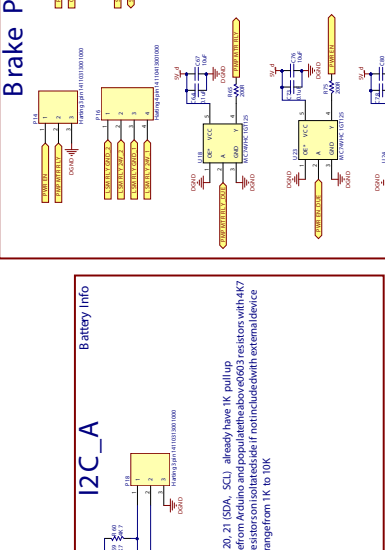
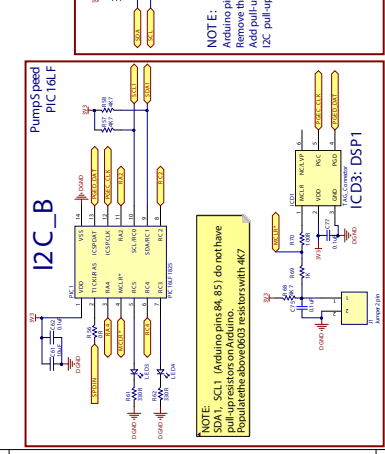
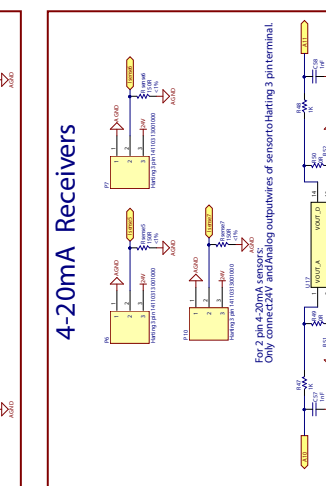
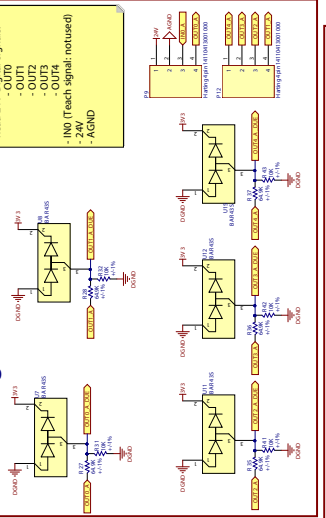
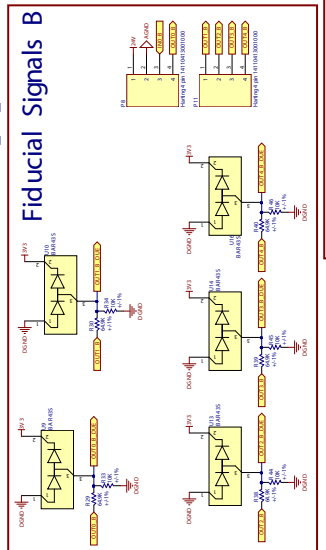
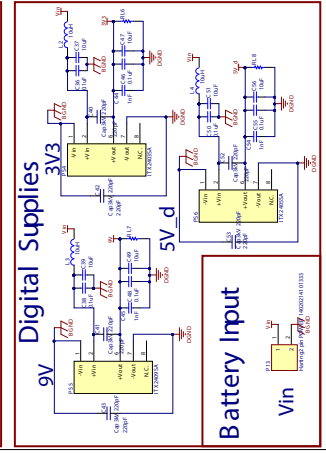
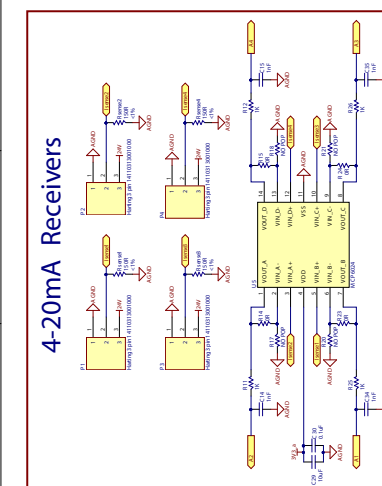
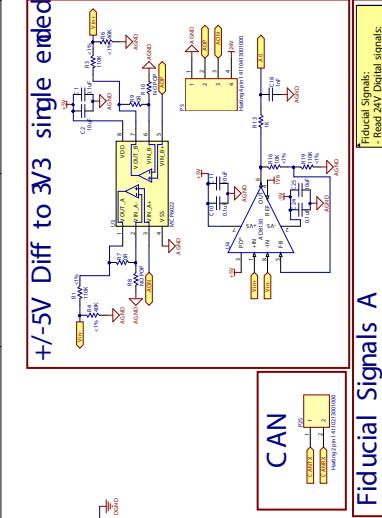
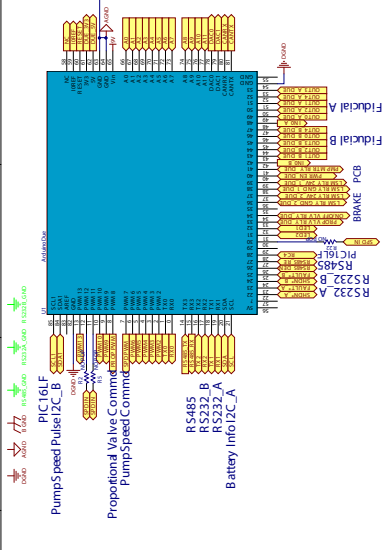
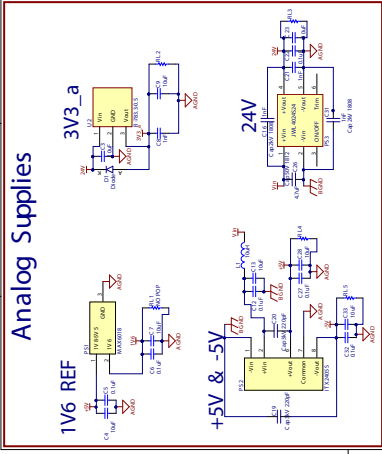
The overall schedule of the project is shown on the next page.

| | Summer 2015 | | | Fall 2015 | | | Spring 2016 | | | Summer 2016 | | | Fall 2016 | | | 2017 | | | | | | |
|---|-------------|-----|-----|-----------|-----|-----|-------------|-----|-----|-------------|-----|-----|-----------|-----|-----|------|-----|-----|-----|-----|--|--|
| | Jun | Jul | Aug | Sep | Oct | Nov | Dec | Jan | Feb | Mar | Apr | May | Jun | Jul | Aug | Sep | Oct | Nov | Dec | Jan | | |
| Organization | | | | | | | | | | | | | | | | | | | | | | |
| <i>Team founded</i> | ◆ | | | | | | | | | | | | | | | | | | | | | |
| Team setup | | | | | | | | | | | | | | | | | | | | | | |
| <i>Competition outline released by SpaceX</i> | | | ◆ | | | | | | | | | | | | | | | | | | | |
| Main recruitment | | | | | | | | | | | | | | | | | | | | | | |
| Pod unveil | | | | | | | | | | | | ◆ | | | | | | | | | | |
| Design Phase | | | | | | | | | | | | | | | | | | | | | | |
| Concept exploration | | | | | | | | | | | | | | | | | | | | | | |
| Conceptual Design | | | | | | | | | | | | | | | | | | | | | | |
| Preliminary Design | | | | | | | | | | | | | | | | | | | | | | |
| Detailed Design | | | | | | | | | | | | | | | | | | | | | | |
| <i>SpaceX rules update</i> | | | | | | | | | | | | | | | | | | | | | | |
| Aerodynamic shell redesign | | | | | | | | | | | | | | | | | | | | | | |
| Brakes redesign | | | | | | | | | | | | | | | | | | | | | | |
| Production Phase | | | | | | | | | | | | | | | | | | | | | | |
| Mold production | | | | | | | | | | | | | | | | | | | | | | |
| Metal component production | | | | | | | | | | | | | | | | | | | | | | |
| Carbon fiber lay-up | | | | | | | | | | | | | | | | | | | | | | |
| Assembly | | | | | | | | | | | | | | | | | | | | | | |
| Electronics mounting & wiring | | | | | | | | | | | | | | | | | | | | | | |
| Electronics troubleshooting | | | | | | | | | | | | | | | | | | | | | | |
| Testing Phase | | | | | | | | | | | | | | | | | | | | | | |
| Levitation testing campaign | | | | | | | | | | | | | | | | | | | | | | |
| Electronics component testing | | | | | | | | | | | | | | | | | | | | | | |
| Vibration testing campaign | | | | | | | | | | | | | | | | | | | | | | |
| Alignment checks | | | | | | | | | | | | | | | | | | | | | | |
| Low speed system test | | | | | | | | | | | | | | | | | | | | | | |
| SpaceX October test week | | | | | | | | | | | | | | | | | | | | | | |
| SpaceX January test week | | | | | | | | | | | | | | | | | | | | | | |
| Competition | | | | | | | | | | | | | | | | | | | | | | |
| <i>Preliminary Design due</i> | | | | | | | | | | | | | | | | | | | | | | |
| <i>Detailed Design due</i> | | | | | | | | | | | | | | | | | | | | | | |
| <i>Competition Design Weekend</i> | | | | | | | | | | | | | | | | | | | | | | |
| <i>SpaceX Build Approval</i> | | | | | | | | | | | | | | | | | | | | | | |
| <i>Test Week</i> | | | | | | | | | | | | | | | | | | | | | | |
| <i>Competition Weekend</i> | | | | | | | | | | | | | | | | | | | | | | |


APPENDIX B

PCB SCHEMATICS


The PCB schematic for the Main PCB is included on page 126, the PCB schematic for the ODROID PCB is included on page 127, and the PCB schematic for the Brake PCB is included on page 128.




5V Output to USB HUB




5V Output to Odroid




Digital Supplies





Components associated with the 70W 5V supply may not be needed. Footprints are only included just in case.



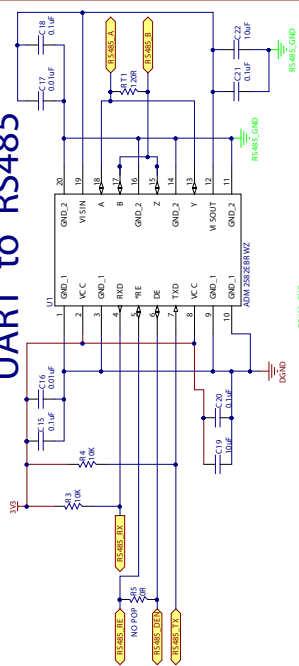
Battery Input




Odroid Signals




UART to RS485

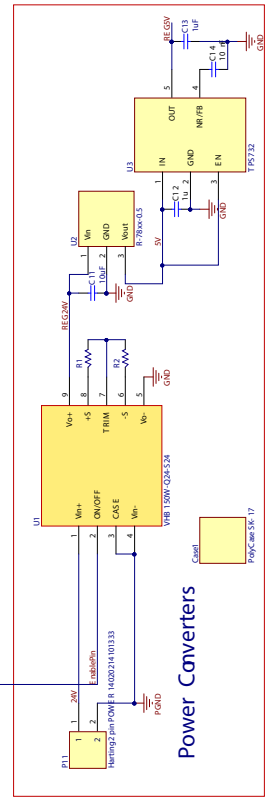

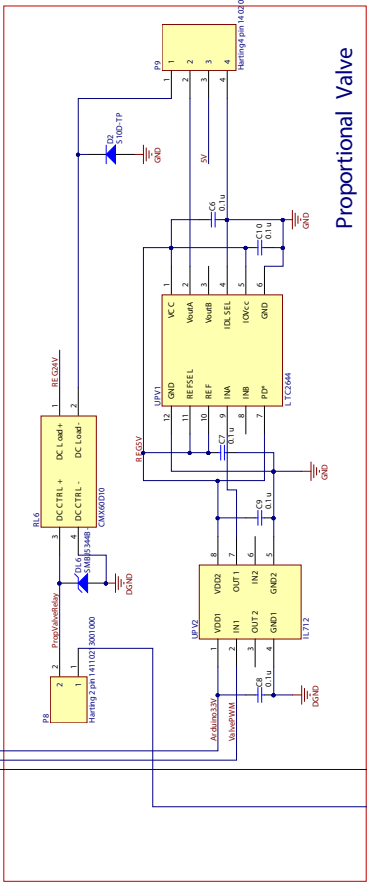
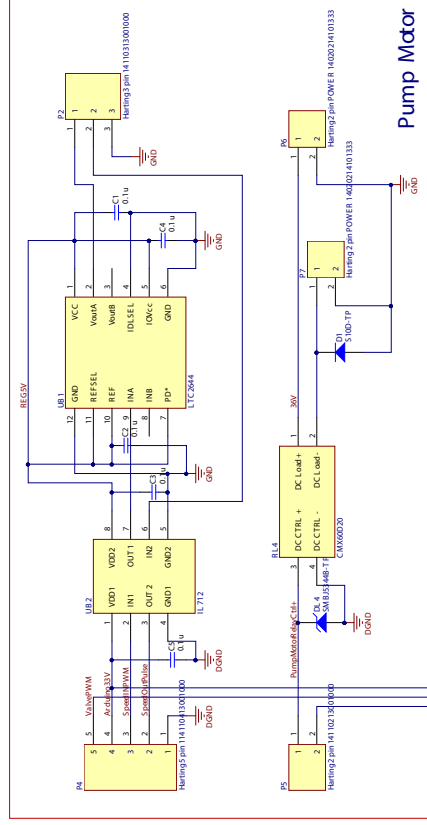
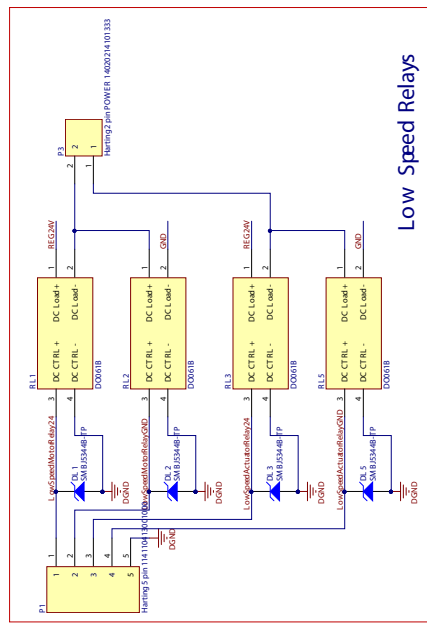
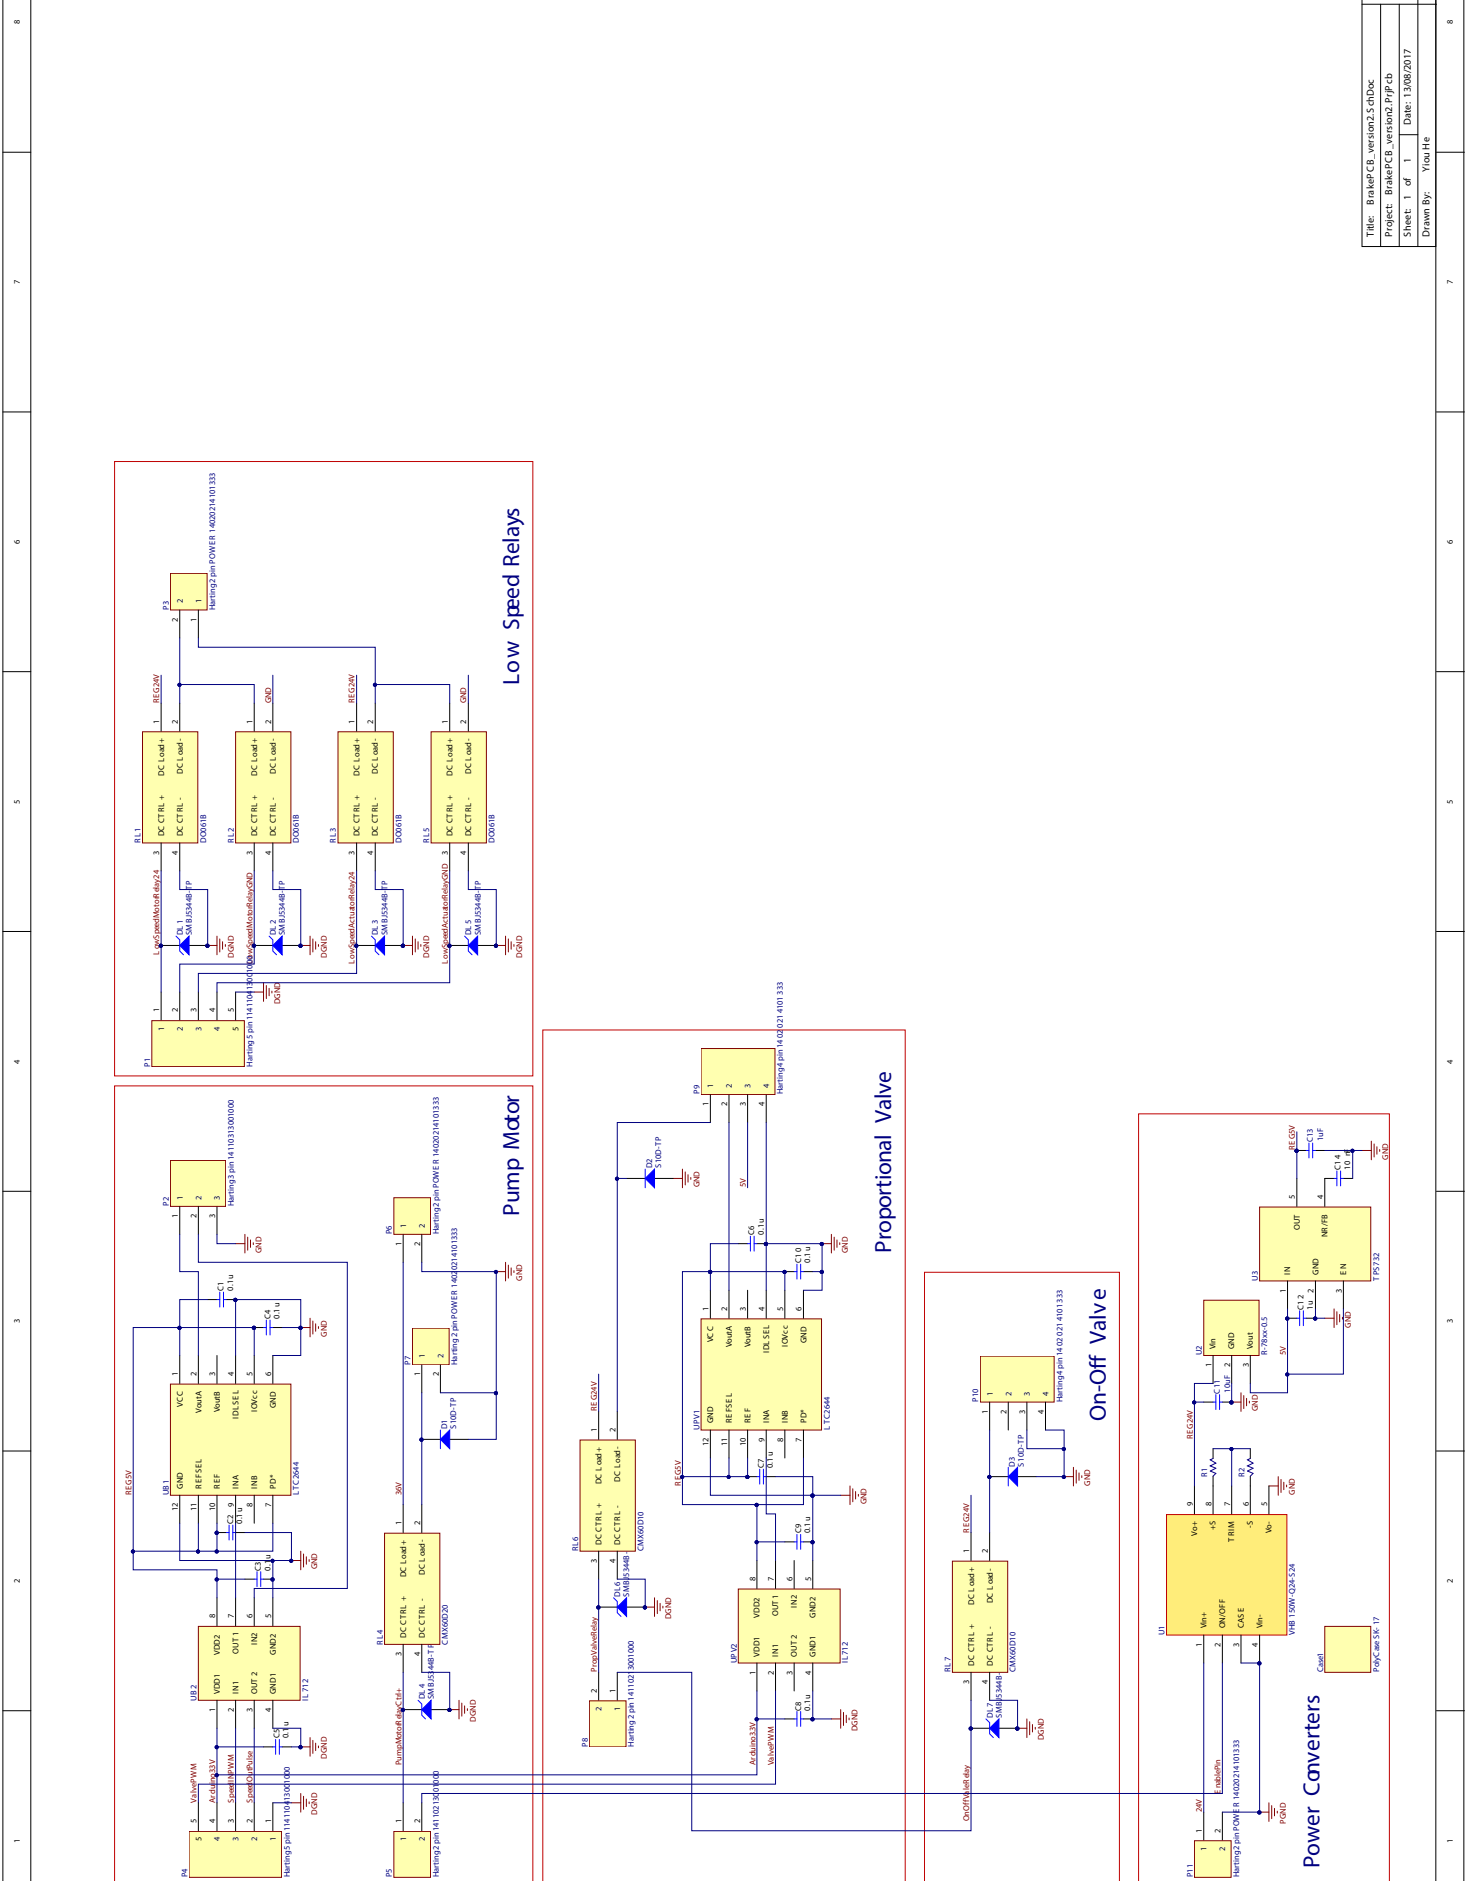


Odroid



LEDS





APPENDIX C

FINAL PROTOTYPE

This appendix shows some photographs from the final prototype that competed in the SpaceX Hyperloop competition in January 2017. Figures C.1 to C.4 were kindly shot by Fabrice Paget. Note the carbon fiber battery enclosures, dummy seat and NAP are not shown in Figures C.1 and C.3.




Figure C.1: Right side of final prototype.



Figure C.2: Right side of final prototype with shell.




Figure C.3: Top view of final prototype.





Figure C.4: View of LCM and brakes from the rear of the pod.



Figure C.5: Pod in the SpaceX Hyperloop tube in Hawthorne, CA before its competition run.



© 2017 MIT Hyperloop