UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

THE DESIGN AND MANUFACTURE OF A VERTICAL TAKEOFF AND LANDING FIXED WING UNMANNED AERIAL SYSTEM FOR USE IN ATMOSPHERIC SAMPLING

A THESIS

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

MASTER OF SCIENCE

By

JORDAN JOSEPH LOGUE
Norman, Oklahoma
2019

THE DESIGN AND MANUFACTURE OF A VERTICAL TAKEOFF AND LANDING FIXED WING UNMANNED AERIAL SYSTEM FOR USE IN ATMOSPHERIC SAMPLING

A THESIS APPROVED FOR THE
SCHOOL OF AEROSPACE AND MECHANICAL ENGINEERING

BY THE COMMITTEE CONSISTING OF

Dr. THOMAS HAYS, Chair

Dr. PHILLIP CHILSON

Dr. WEI SUN

This work is dedicated to the application of technology as a force for good.

Acknowledgments

Just as it takes a village to raise a child, so too does the creation of an unmanned aerial system. My adviser, Dr. Thomas Hays, had by far the single largest impact on the success of the project outside of my own. Through the thick and the thin he consistently challenged me to do my best work. There would have been no project at all without Dr. Phillip Chilson and Dr. Elizabeth Pillar-Little of CASS. Their complete support served as a constant bulwark against the stresses and strains found throughout the project. CASS staff engineer Bill Doyle was absolutely instrumental in the implementation of the avionics system, one which supported the platform in the air from the very first flight. It was Tony Segales' Ardupilot programming and quadcopter piloting skills which ensured that all bases were covered on the aircraft control side of things during design, final implementation, and test flights. A big thanks to Billy Mays, Greg Williams, and the entire Aerospace and Mechanical Engineering machine shop for making the time in their busy schedules to produce some of the most critical structural components for the aircraft. I also wish to thank the Tom Love Innovation Hub's Fabrication Lab and its employees for allowing the use of their CNC router and laser cutters in support of this project. Many undergraduate students lent their energies towards the manufacturing of the aircraft; many thanks to Santiago Mazuera, Chris Hughes, Bryan Boone, Martha Nguyen, Meghana Fathepure, and Callen Koiner for their assistance. I am so grateful for the unflinching support of my loved ones despite my prolonged absence in the pursuit of this project. Finally, I would like to thank Dr. Thomas Hays, Dr. Phillip Chilson, and Dr. Wei Sun for taking the time to sit upon my master's committee.

Table of Contents

List of Tables

List of Figures

Abstract

The creation of a prototype vertical takeoff and landing fixed wing unmanned aerial system for use in sampling of trace gases, aerosols, and volatile organic compounds is described. A conceptual design framework is devised based upon desired performance characteristics and conclusions drawn from background research into platform layout, energy sources, and construction methods. Optimized designs are produced according to this framework, and the most appropriate option for the application serves as the foundation in producing a detailed design of the platform. The methods employed in manufacturing the aircraft's individual components as well as the assembly of the system as a whole are described. The testing involved in validating critical systems is presented, culminating in an accounts of the aircraft's flight tests. The lessons learned from this first implementation are highlighted and then applied to produce an improved conceptual design.

Chapter 1

Introduction

1.1 Concept of Operations

In the spring of 2018, the University of Oklahoma School of Aerospace and Mechanical Engineering (AME), in association with the Center for Autonomous Sensing and Sampling (CASS), identified the need for a high payload-capacity, long endurance unmanned aerial system (UAS) to facilitate study of the vertical and horizontal heterogeneity of important chemical species such as trace gases, aerosols, and volatile organic compounds within the first kilometer of the atmospheric boundary layer. With human activity mostly confined to the lower troposphere, the behavior of these species within this region of the atmosphere was of particular interest. Aerosols often manifested themselves as air pollution widely recognized to limit visibility and to induce deleterious health effects in humans. High ground level concentrations of certain trace gases such as ozone (O_3) also promoted health problems, while their presence in the upper troposphere served to contribute to greenhouse effects (Wang et al. 2017).

Frictional interactions between the air and the Earth's surface contributed to high degrees of wind shear within the atmospheric boundary layer, while diurnal surface heating/cooling cycles contributed to temporal thermodynamic gradients. These effects combined to produce highly complex atmospheric behavior within the region of the atmosphere with the most immediate impact on human activities. The nonlinear vertical and horizontal distributions of the chemical species of interest required low altitude, high-resolution in situ measurements in order to form continuous profiles capable of validating surface air quality forecasts based upon ground or satellite measurements (Brady et al. 2016). Existing measurement platforms proved to be of limited usefulness in fulfilling these requirements due to comparatively poor sampling density and altitude restrictions (Li et al. 2014).

Given CASS' pilot and operational expertise with remotely-controlled systems, a solution of this type would allow the fullest utilization of existing resources. An unmanned aerial system capable of carrying a suite of chemical sensors could fill this sampling gap as lower flight speeds would allow for high horizontal measurement resolution while lack of on-board human operators and smaller size would lower the minimum safe operating altitude to beneath that already possible for sampling by stationary methods such as towers. The ability to operate from remote settings in comparatively high winds would widen the range of surface-atmosphere interactions possible for investigation, while sufficient operational endurance supporting various payloads would lend sufficient opportunity for development of continuous data trends for tuning and validation of air-quality prediction models. In summary, a platform satisfying the performance requirements listed in Table 1.1 was projected to fill the existing platform-driven data gap:

Table 1.1: Platform Performance Requirements

Operational Endurance:	1 hour
Maximum Payload Capacity:	5 pounds
Maximum Wind Resistance:	30 knots

1.2 Literature Review

1.2.1 Fixed Wing versus Rotorcraft

The question of whether to use a fixed wing type aircraft or a rotorcraft was one that had to be revisited in the preliminary investigations of any new aerospace design process as its answer surely depended upon the specific requirements and constraints of the given application. As the eventual site of operations, Kessler Atmospheric and Ecological Field Station, was capable of supporting remotely-piloted airplanes and helicopters, the core requirements and constraints affecting the answer to the question of fixed wing vs. rotorcraft were: sensor packaging, endurance, payload, and maximum wind resistance. To ensure

accurate sampling of the chemistry of the atmosphere, it was of paramount importance to package the sensor inlets to allow for accurate representation of the craft's surroundings. To this end, the sensor placement options were seen to be more numerous on a fixed wing aircraft of the application's size scale than on a similarly sized rotorcraft. This was partially due to the latter's rotor placements, displacing air downward over the entire craft's densely-packed airframe, whereas fixed wing craft would have greater freedom over their propeller(s)'s placement. Conventional packaging of a fixed wing UAS included a wing and fuselage, large portions of which would be out of the flow stream of disturbed air from the propeller(s).

Even more concrete than these packaging concerns were the payload-specific endurance values for the two aircraft types. The maximum reported values of each were plotted together in Figure 1.1 for a number of craft that run the gamut of UAS size scales and configurations[1]:

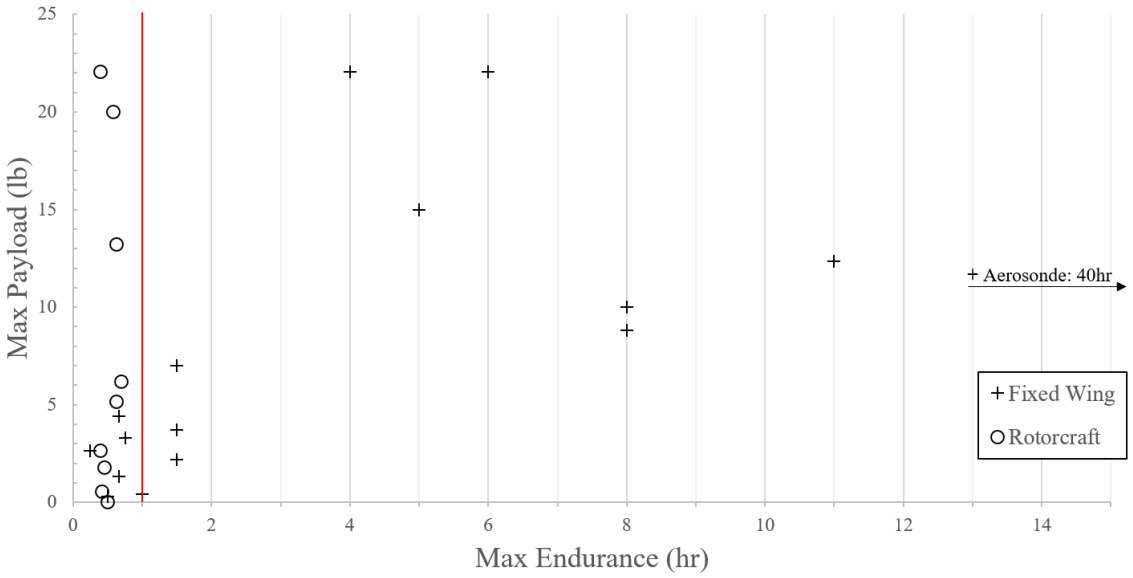

Figure 1.1: Maximum Possible Payload vs. Maximum Possible Endurance for Fixed Wing and Rotorcraft

[1] Source data available in Table A.1

One of the most striking conclusions drawn from Figure 1.1 was that none of the surveyed rotorcraft appeared to surpass one hour of endurance, CASS' desired requirement for endurance. This upper bound seemed to persist for all payload classes of rotorcraft and even across examples of quad, hex, and octocopter configurations. Further investigation into the possibility of increasing copter endurance through augmentation of battery capacity for representatives of each rotor configuration produced Figure 1.2[2]:

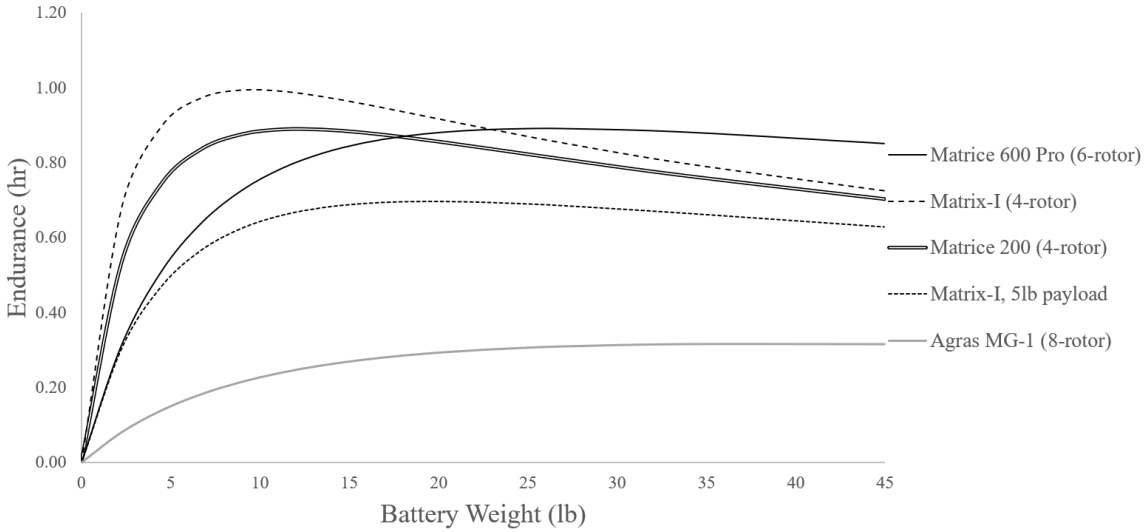

Figure 1.2: Endurance vs. Battery Weight for Representative Rotorcraft

Note that each craft's default curve represented its hovering endurance in no wind while supporting no payload. Under these allowances, the Matrix-I achieved a maximum possible endurance at the desired requirement for CASS. However, with the addition of the five lb payload requirement, the Matrix-I's performance envelope shifted well below this threshold. Qualitatively, Figure 1.2 showed that each rotorcraft achieved a maximum endurance at a finite battery weight and that any further addition of batteries only served to decrease endurance. By modifying the platforms for use with larger propellers to maintain equivalent disk loading, it might have been possible to push one or more of the commercial rotorcraft beyond the one hour mark. However, such allowances would require work beyond the scope of this comparative analysis between commercially-available platforms.

[2]See Footnote 1

This overall endurance vs. battery weight trend was also seen using fixed wing data such as MARIA's[3]:

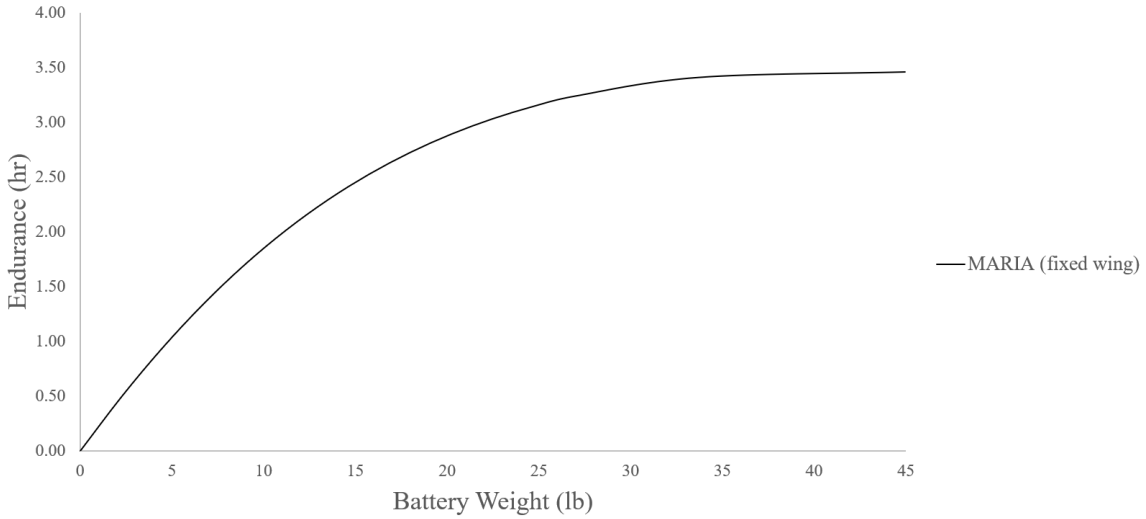

Figure 1.3: Endurance vs. Battery Weight for a Representative Fixed Wing Aircraft

for which CASS' requirements easily fit within the craft's performance envelope.

The last performance requirement, wind resistance, came about due to CASS' desire to operate within a wider range of weather conditions. The desired threshold of thirty knots of wind, together with the five lb payload requirement, were outlined by the shaded region in Figure 1.4 along with data points corresponding to a number of commercially-available fixed wing and rotorcraft[4]:

[3]See Footnote 1
[4]See Footnote 1

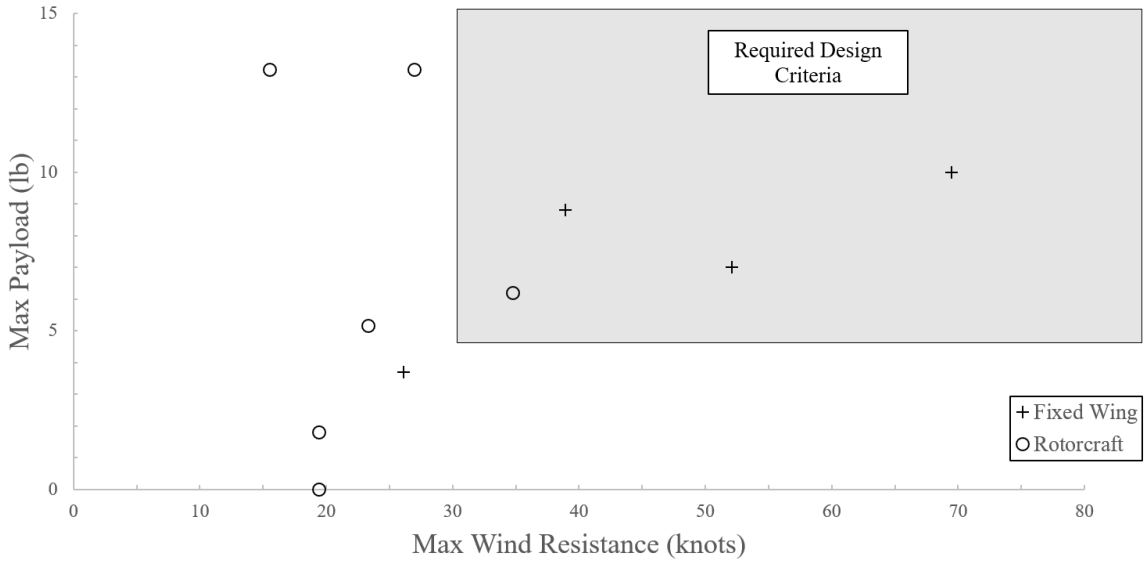

Figure 1.4: Maximum Possible Payload vs. Maximum Wind Resistance for Fixed Wing and Rotorcraft

Figure 1.4 exhibited similar trends as Figure 1.1 in that the lion's share of rotorcraft data points, across all payload ratings, fell short of the wind resistance requirement; the fixed wing data points appeared to trend towards increasing wind resistance at higher payload scales. The exception to this rule was the Drelio single-rotor copter UAS which was designed for use in imaging windy coastal regions but whose maximum endurance was unreported (Delacourt et al. 2009).

Based upon the requirements for packaging, endurance, payload, and wind resistance and their respective trends observed in the collated performance of commercially available platforms, a platform utilizing fixed wing lift to carry out the bulk of its mission was deemed to be necessary. Given the requirement for field operability, alternative launching and landing capabilities would be required. Bungee or rail-launching, coupled with belly landings, were the traditional options, but the possibility of VTOL was perhaps a more elegant and exciting route.

1.2.2 Platform Layout

The reduced operational velocities and exclusion of an on-board pilot allowed for wider exploration of design possibilities in terms of the placement of the platform's wing(s), control surfaces, payload, power plant, and even fuselage compared to human-scale aircraft. A high endurance UAS intended for atmospheric measurement, and one which exhibited evidence of this added design freedom, was the Aerosonde aircraft. This design, shown in Figure 1.5, featured a single fuselage that housed its payload and power plant at the fore and aft, respectively:

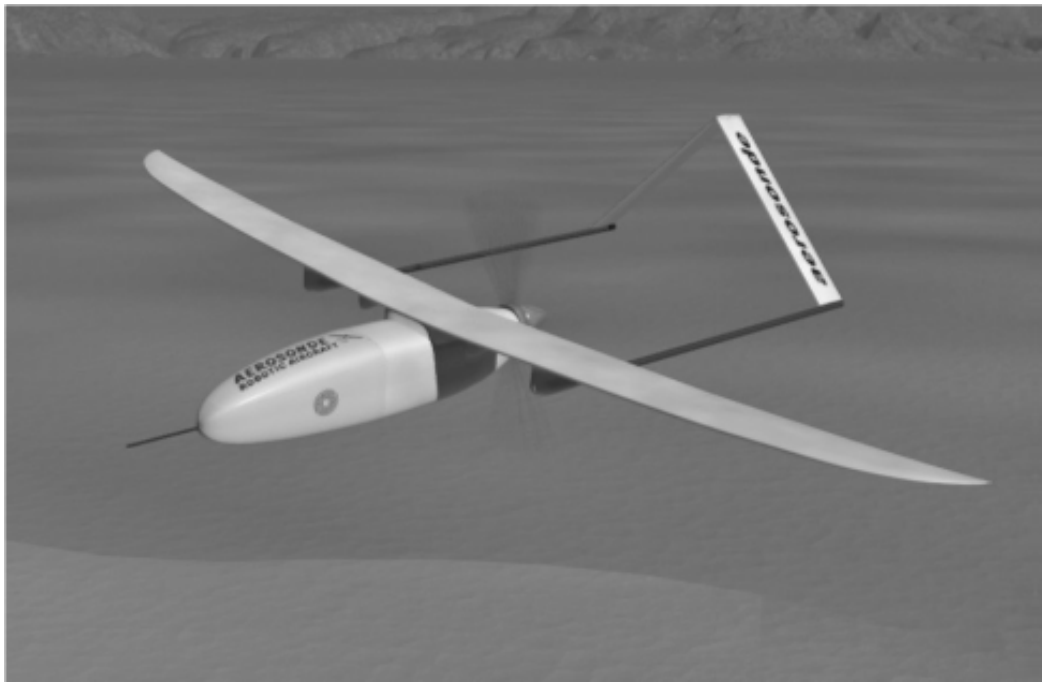

Figure 1.5: Render of the Aerosonde Mark 2 (Holland et al. 2001)

The placement of the power plant downwind of the payload's sensors insured against measurement contamination by the prop wash and the internal combustion engine's heat and exhaust. The choice of a boom-mounted inverted V-tail was noted, though no specific justification of this design choice was forthcoming. While the decision was most likely borne out of the packaging concerns associated with the rear-mounted propeller, the choice of an inverted V-tail also avoided what Raymer referred to as "adverse roll-yaw coupling"

in standard V-tails because its flipped geometry produced a rolling moment in the same direction as that desired while executing a banking turn. In addition, the tail's overhead boom-mounting alleviated the ground clearance concerns of fuselage-mounted inverted V's (Raymer 2012).

One example of the standard V-tail in use by an atmospheric-sampling UAS was that of the Cruiser UAS:

Figure 1.6: Image of the Cruiser UAS (Schrod et al. 2017)

This tail configuration necessitated additional control and mechanical complexity due to the adverse roll-yaw coupling inherent to any deflection of a rear control surface, though the aircraft retained the luxury of rear propeller placement, thereby allowing the packaging of sensors throughout the fuselage (Raymer 2012). The placement of the wing over the fuselage allowed for the location of a sensor intake at the bottom of the fuselage. This placement likely explained the tall, robust, and permanently deployed landing gear (Schrod et al. 2017).

BAE Systems' Manta UAS followed the trend of "boomed," V-type empennages, though with a main wing situated midway-up its flat fuselage:

Figure 1.7: Image of BAE Systems' Manta UAS (Raytheon 2018)

The Manta's flat but wide fuselage may have allowed for superior packaging of two-dimensional payloads such as circuit board-mounted sensors. The rear-mounted propeller and permanently deployed landing gear would offer similar advantages and disadvantages as those described on the Aerosonde and the Cruiser UASs. The Manta's boom-mounted rear empennage ensured a wide static margin, allowing for flexibility in payload placement even after its manufacture.

The Manta's low-profile fuselage and smooth wing-body blending bears greater similarity with flying wing configurations such as that found in the Micro Air Vehicle (MAV) described by Allred (2007):

Figure 1.8: Image of the Micro Air Vehicle (Allred et al. 2007)

This simple platform was based upon the commercially available Stryker RC aircraft (Elston et al. 2015). Pitch and roll control was provided for by the two long elevons, with the two vertical surfaces aiding in yaw and roll stability. The inclusion of strakes at the leading edge of the wing raised the stall angle of attack (AoA) of the aircraft, a useful addition to a lightweight aircraft prone to pitch disturbance by wind gusts due to its light foam construction (Raymer 2012). The rear placement of the propeller freed the entire fore of the craft for placement of sensors.

Hardly more conventional than the previous few aircraft was Ma's (2004) Miniature Robotic Plane Meteorological Sounding System (RPMSS), shown in Figure 1.9:

Figure 1.9: Image of the RPMSS UAS (Ma et al. 2004)

This unique configuration paired an H-tail and high-mounted polyhedral wing with a canard along with a pusher-type power plant configuration. This so called "three-surface configuration" had the potential to produce less induced drag versus traditional canard-wing designs as its rear empennage allowed for the creation of pitching moments without affecting the lift produced by the wing, thereby allowing for the usage of better optimized wing lift distributions. However, this efficiency might have been tempered by the negative aerodynamic effects experienced on the wing as a result of its operation in the canard's wake (Scholz 2017). The polyhedral shape of the main lifting surface produced stable roll behavior as perturbations around this axis leading to side-slip triggered a naturally stabilizing behavior akin to that of a falling, oscillating leaf, acting to right the aircraft's orientation. In the case of a perturbation in yaw, this wing shape could lead to adverse roll coupling, introducing either additional complexity to the autonomous control system or to the human operator when maintaining a desired flight direction. The twin verticals

of the H-tail operated on either side of the fuselage's wake, thereby avoiding the issue of blanketing experienced by single vertical surfaces at high angles of attack. The endplate effect of the twin verticals on the horizontal improved its lift distribution and might have lowered the span required. However, the reinforcement required for this tail still tended to be heavier compared to conventional empennages (Scholz 2017).

More traditional in layout was UASUSA's Tempest, shown in Figure 1.10:

Figure 1.10: Render of the Tempest UAS (UASUSA 2018)

The Tempest combined a tractor-type propeller configuration with a high aspect ratio wing vertically placed approximately halfway down its slender fuselage and a T-type tail. Elston (2015) described how the craft's design was based upon competition sailplanes in the pursuit of similarly high levels of aerodynamic and structural efficiency (Elston et al. 2015). The T-tail required additional weight in the form of structural reinforcement within the vertical tail surface to support the horizontal's loads, but similar to the case of the RPMSS' H-tail, the endplate effect of the horizontal on the vertical resulted in a smaller required size for the vertical, thus reducing weight. In addition, by lifting the horizontal out of the

path of the main wing's wash, the surface's performance was increased, thereby possibly lowering weight due to reduced area requirements (Raymer 2012).

More on the cargo side of the UAS design spectrum was Avery's (2013) Mesocyclone Analysis Research and Investigation Aircraft (MARIA):

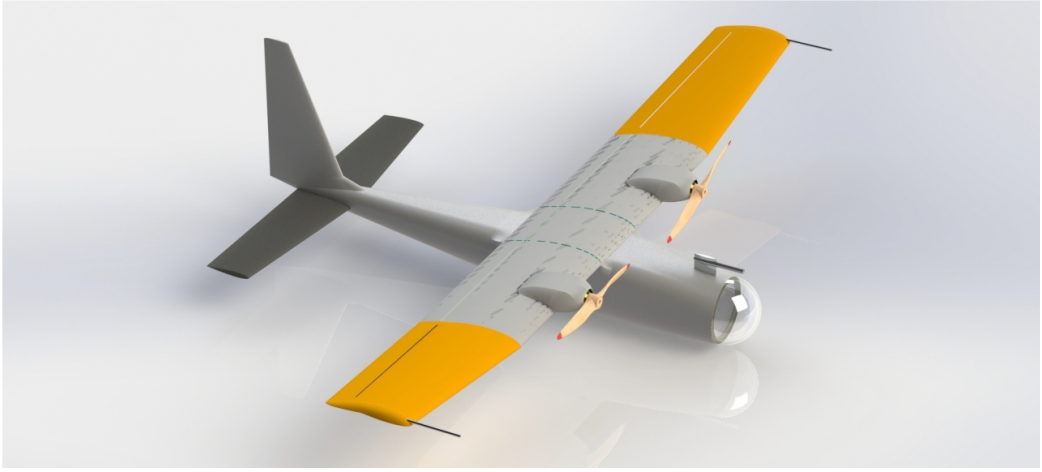

Figure 1.11: Render of MARIA (Avery 2013)

More so than its conventional tail layout, MARIA's dual-tractor propeller configuration stood out among the other aircraft here examined, most of which fell into the small and lightweight category of UAS. As MARIA was designed for transporting a payload of multiple weather sondes into extreme weather, the possibility of an engine-out scenario was a serious consideration during its conceptual design. As a result, the propulsive redundancy of a two-engine design played favorably into the final decision of MARIA's configuration (Avery 2013). In addition, the placement of the propellers on the craft's wings provided the same clean air for nose sensors as pusher-type configurations while maintaining the propellers themselves in clean air. The use of these wing-mounted powerplants, coupled with the aircraft's belly-landing mission requirement, likely contributed to the vertical placement of the main wing above the fuselage.

Insitu took a decidedly minimalist approach to its ScanEagle UAS' layout, as shown in Figure 1.12:

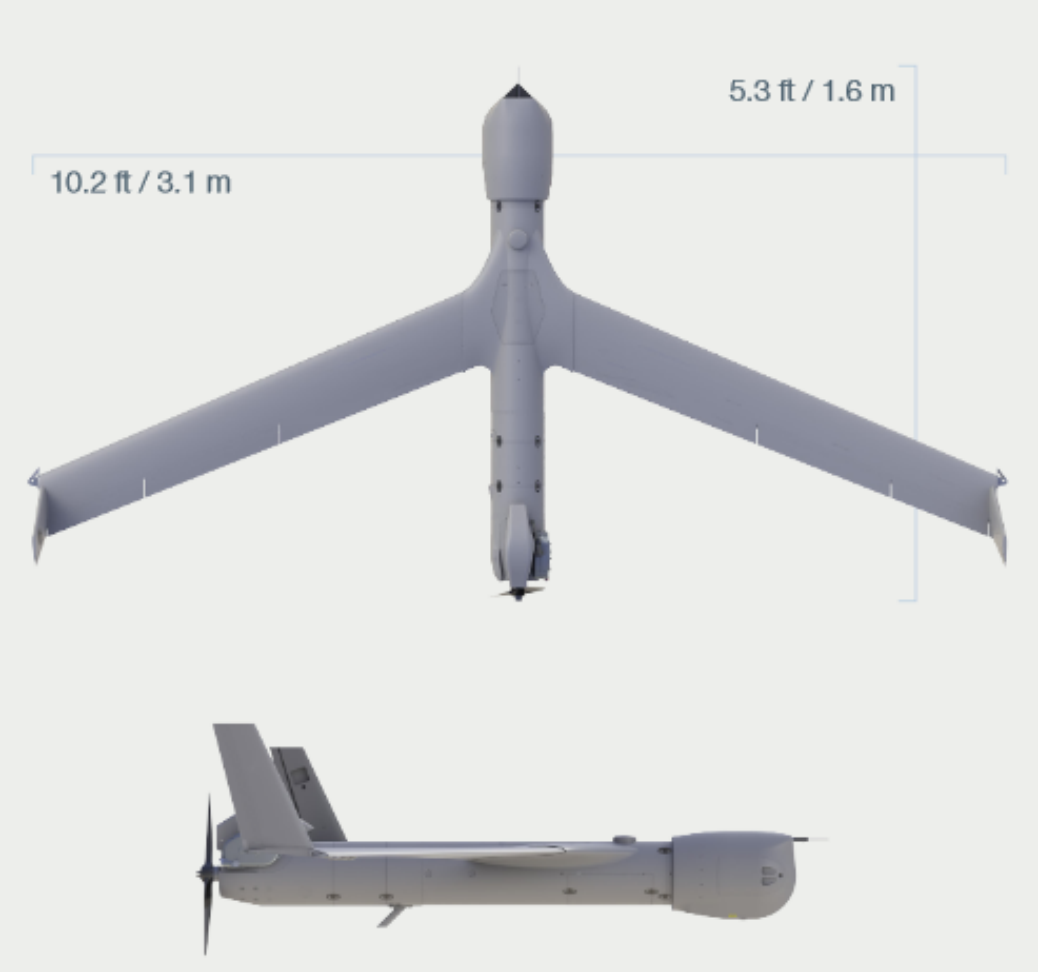

Figure 1.12: 2-View of Insitu's ScanEagle UAS (Insitu 2017)

The ScanEagle's lack of rear aerodynamic elements opened up the entire rear of the fuselage for the propulsion system, leaving the entire rest of the body for the payload. This was not without consequence, given that the air flowing into the propeller had likely been disturbed by the craft's fuselage and/or wing and thereby lowering propulsive efficiency (Raymer 2012). In addition, the combined effects of a rear-mounted propulsion system and no rear control surfaces led to an aircraft with higher sensitivity to center of gravity (CG) placement, limiting the flexibility of payload placement locations.

The application at hand required an aircraft with the ability to support a payload ranging from approximately zero to five pounds. This wide range of possible payload weights demanded a wide and flexible static margin in order to maintain static stability. The boom

mounted tails of the Aerosonde and Manta made the most sense to meet this need, though a more traditional empennage configuration like that of the RPMSS or MARIA would do so with less mechanical complexity. High lift production was crucial to support the aircraft under max loading, as well as to allow for a low enough stall speed for sufficient data sampling rates during measurements. In the case of VTOL capability, a high lift wing would do much to support the additional weight of the associated hardware during forward flight. Therefore, a traditional wing with a high aspect ratio was chosen to be utilized for the platform.

1.2.3 Energy Sources

Hays (2009) categorized the possible energy systems for UASs between three and twenty pounds into the following categories: direct energy storage systems, conversion based energy systems, and harvesting energy devices. Direct energy storage referred to those methods, such as batteries and capacitors, which directly stored electrical energy for use. The conversion based energy systems category was comprised by those options which required additional hardware for conversion of the energy source into a useable form, as in the case of internal combustion engines and fuel cells. Harvesting energy devices featured a time-dependent component to their energy density, typically necessitating long duration missions for viability; an example of this energy source was solar cells (Hays 2009).

According to Hays (2009), direct energy systems powered all manner of aircraft from AeroVironment's Wasp:

Figure 1.13: Image of AeroVironment's Wasp UAS (Griffis et al. 2009)

to Hobbico's NexSTAR EP Select:

Figure 1.14: Image of the SMARTSonde NexSTAR Aircraft (Chilson et al. 2009)

upon which Chilson's (2009) SMARTSonde UAS was based.

Most direct energy systems were comprised of batteries which stored energy electrochemically. As was illustrated in Figure 1.15, this was accomplished by way of an anode, the end of the battery hosting a build up of electrons, a cathode, the end of the battery featuring a dearth of electrons, and an electrolyte bath, which hosted the chemical reactions

allowing the build up of electrons on the anode while also preventing these same electrons from flowing back to the cathode (Forbus 2018):

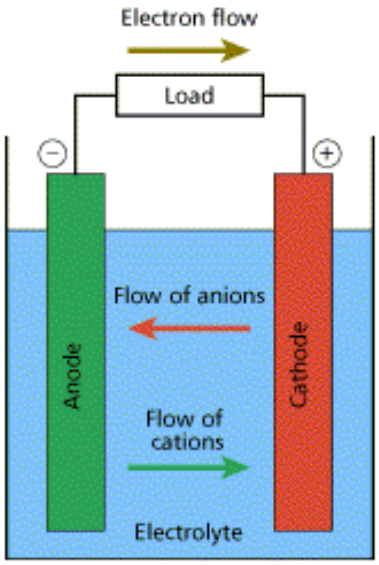

Figure 1.15: Schematic of a Simple Battery (Griffis et al. 2009)

As shown, the resulting flow of electrons could be routed through a load to power it.

In chronological order of development, the common battery types for UASs were Nickel Cadmium (NiCd), Nickel Metal Hydride (NiMH), Lithium Ion/Polymer (LiPo), and Lithium Sulfur (LiSu). NiCd batteries had the worst energy density of the battery types at 25 Watt-hours per pound ($^{WH}/_{lb}$) and were comparatively problematic; as a result, these batteries were less and less commonly used. NiMH batteries were thought of as the less problematic cousins to NiCds, and with an energy density of 32 $^{WH}/_{lb}$, they were still in use for some lower energy applications. LiPos were the batteries of choice for performance-intensive UAS applications. These batteries could provide an energy density of 75 $^{WH}/_{lb}$, and research is on-going to further increase this. A growing option for UAS batteries, LiS featured energy densities as high as 160 $^{WH}/_{lb}$, but its two-stage reaction behavior during discharge necessitated the use of specialty microcontrollers to manage the battery's output and load (Hays 2009). All of these types were rechargeable and had at some point been

used to successfully power an unmanned aircraft. The other major form of direct energy system is capacitors. Such devices stored energy electrostatically using two oppositely charged plates in close proximity for a layout reminiscent of batteries. Unlike batteries, most capacitors only held enough charge for quick but high-intensity power applications. Their main advantage compared to batteries was their negligible degradation over millions of recharge-discharge cycles (Hays 2009).

The next most common energy source for UASs was the conversion based energy system. As previously stated, this source type required additional hardware in order to make use of its ultimate energy source. The main forms of this energy source in-use by unmanned aircraft were hydrogen fuel cells, reciprocating piston engines, rotary engines, microturbines, and rocket engines.

Of this first option, the most widely used type of hydrogen fuel cells for UASs was the proton exchange membrane fuel cell, and it shared several characteristics with batteries as shown in Figure 1.16:

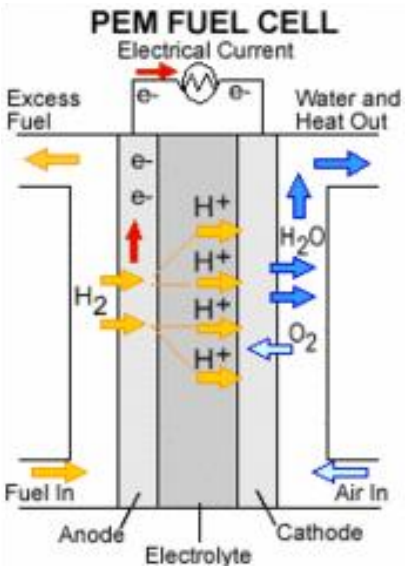

Figure 1.16: Schematic of Hydrogen Fuel Cell (Griffis et al. 2009)

Unlike the battery, the fuel cell was not a self-contained device; it required external sources of both hydrogen fuel and oxygen reactant, as well as provisions for exhausting excess

reactants and waste products. The process operated as follows: hydrogen fuel was introduced to a catalyst, typically made of platinum, that ionized the hydrogen molecules into their constituent electrons and protons. The proton ions flowed through the electrolytic membrane from this anode to the cathode while the electrons were prevented from doing so. The protons on the cathode reacted with the supplied oxygen in electrochemical oxidation, a process also requiring electrons. This impetus drove the flow of the electrons from the anode, through the load, to the cathode, before being exhausted away as waste (Griffis et al. 2009). The source of the hydrogen itself could be quite varied, with Hays (2009) listing cryogenic storage, glass microspheres, chemically-bound reformation, carbon nanotubes/graphene sheets, and metal hydride as options. Griffis (2009) reported a real-world UAS powered by a hydrogen fuel cell, the Protonex SpiderLion:

Figure 1.17: Image of the Protonex SpiderLion UAS (Griffis et al. 2009)

With a possible energy density of 454 $^{WH}/_{lb}$, the prospect of hydrogen fuel cells was very intriquing (Hays 2009). However, Griffis (2009) noted a substantial increase in system complexity accompanying the use of such fuel cells versus more conventional conversion-based energy sources.

The most common conversion based energy system by far was the reciprocating piston engine (RPE). Such engines powered UASs as varied as the Pioneer, utilizing a two-stroke RPE:

Figure 1.18: Image of the Israel Aerospace Industries Pioneer UAS (IAI 2002)

and the Predator XP:

Figure 1.19: Image of the General Atomics Predator XP UAS (Atomics 2018)

being powered by a four-stroke RPE. At their most simple, RPEs utilized a difference in pressure on opposing sides of a piston to induce linear motion which could then be directly harvested as mechanical energy or converted to electrical energy for storage and/or immediate use. The various types of RPEs varied in their sources for the creation of this pressure and in the geometries required for housing, capturing, and transferring the linear motion. Two and four-stroke RPEs typically made use of the combustion of a petroleum distillate fuel such as gasoline in order to create the necessary pressure differential for linear motion. A schematic of the four-stroke cycle was given by Figure 1.20:

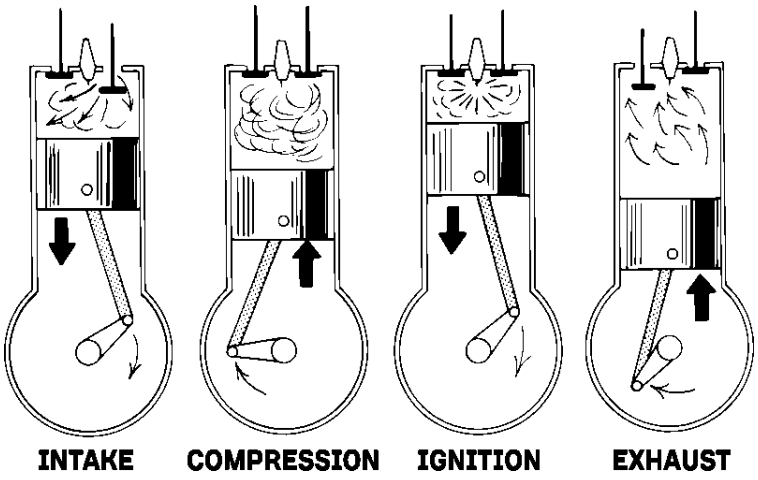

Figure 1.20: Outline of the Four-Stroke Combustion Cycle

The difference for a two-stroke cycle involved the combination of fuel intake with exhaustion, where the motion of the cylinder following combustion was utilized to both draw fuel and air into the combustion chamber and to force exhaust products out. A two-stroke engine produced energy once per revolution of the shaft to which the piston was connected (the crankshaft), whereas the four-stroke produced power once in every two revolutions. The two-stroke was simpler than the four-stroke engine, thus requiring fewer parts. The culmination of such differences resulted in a power increase associated with using a two-stroke over a four-stroke engine, being about 1.5 times for a given weight class (Hays 2009). However, the two-stroke's combination of intake and exhaust often resulted in inefficiencies due to incomplete combustion of reactants. In addition, a two-stroke engine accepted fuel directly into the chamber housing the crankshaft (the crankcase) prior to injection into the combustion chamber. As the crankshaft had to be lubricated to prevent friction losses, the two-stroke engine fuel was diluted with oil which lowered the system's overall energy density. A four-stroke engine accepted fuel directly to its combustion chamber, allowing the crankcase to be lubricated separately (Hays 2009). RPEs were numerous enough that examples existed which occupied every step of the energy density ladder. For example, the Rotax 914 turbo 4-stroke engine, a version of which powered the massive Predator XP in Figure 1.19, achieved an energy density of only 10 $^{WH}/_{lb}$ though it developed over 110

horsepower (hp) (BRP Rotax GmbH and Co. 2016). Meanwhile, the RCV single-cylinder 4-stroke engine, weighing only 4.2 lb, boasted an energy density of over 300 $^{WH}/_{lb}$ but also produced only 2.7 hp at full tilt (RCV Engines Ltd. 2016).

A close relative of two and four-stroke engines was the diesel engine, in which ignition was achieved not through the use of a spark plug but by way of compression alone. Such engines were known for their higher efficiencies as their compression ratio, the ratio between maximum and minimum instantaneous combustion chamber volume, was necessarily higher in order to reach the fuel's auto-ignition temperature, which in turn related directly to the amount of power generated per engine cycle (Tian 2013). Different parameters of the diesel cycle could be rearranged to prioritize various performance characteristics such as in the case of a homogeneous charge compression ignition (HCCI) engine which mixed the correct fuel-air ratio prior to injection into the combustion chamber (Hays 2009).

If the diesel engine was a close relative of two and four-stroke RPEs, then the Stirling engine was a distant cousin. Rather than intaking new fuel each cycle, this type of RPE utilized alternating heating/cooling of the same volume of working fluid in order to produce the device's linear motion. Afework (2018) illustrated the cycle as Figure 1.21:

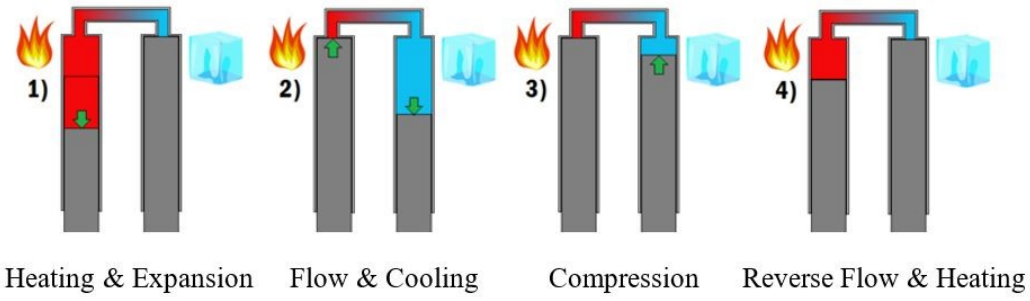

Figure 1.21: Outline of Stirling Engine Cycle (Afework et al. 2018)

An external heat source raised the temperature and pressure of the working fluid in **1)**, causing an expansion of the fluid that pressed the system's left piston downwards and generated work. The piston was then made to move up, pushing the fluid into the right cylinder as in **2)**, where it was cooled by an external source. The right piston then moved up to compress

the fluid, which required less energy than that produced in **1)** because the cool fluid was at a lower temperature in **3)**; any excess heat generated by this compression was removed by the cooling source. Finally in **4)**, the compressed fluid was introduced back into the left cylinder and the cycle repeated (Afework et al. 2018). The result was a highly efficient (30%) but low energy-to-weight device (11-43 $^{WH}/_{lb}$) (Hays 2009). As the heat/cooling sources were external to the engine, these had to be provided for separately, and their removal from direct contact with the working fluid resulted in reduced responsiveness compared to more popular forms of RPEs (Afework et al. 2018). However, this design element could serve as an advantage for high altitude UASs for which the existing low ambient temperature could serve to drastically increase cycle efficiency (Hays 2009).

The Wankel rotary engine shared the same basic operational structure with the four-stroke, but whereas RPEs, as their name suggested, utilized the linear motion of a piston to generate energy, the rotary engine featured a rounded triangular rotor that rotated about an eccentric shaft within an epitrochoidal housing (Scitech 2009). As the rotor rotated, each of its three sides formed a separate volume against the wall of the housing; at any given point, each of these sides could be experiencing fuel intake, compression, combustion, fuel exhaustion, or some transitory state as was highlighted by Figure 1.22:

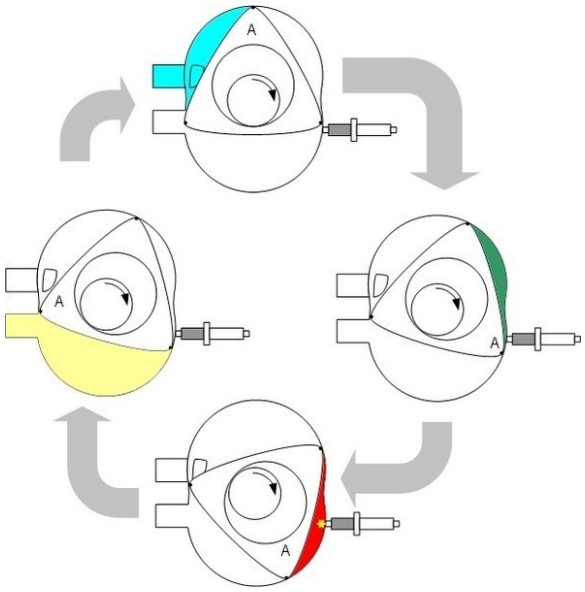

Figure 1.22: Outline of the Wankel Rotary Engine's Cycle (Scitech 2009)

where the path of the highlighted volume (representing the working fluid) began at the top with fuel intake. As shown, the geometry of the housing, together with the rotor's, combined to produce similar compression and expansion cycles to those found in RPEs. The Wankel has been used successfully in UAS applications, as in the case of the Shadow 200:

Figure 1.23: Image of the AAI RQ-7A Shadow 200 UAS

The engine's geometry made it smaller than RPEs of the same power class, leading to about twice the power output of conventional RPEs of similar internal volume. However, fuel economy in Wankels has traditionally been less than that possible in linear piston engines, and the relative lack of technical literature on rotary engines has often led to their being passed over in favor of better understood energy sources. Additionally, the same geometry which compacts the engines often lead to significant issues in maintaining proper sealing. Despite these drawbacks, commercial options existed such as LiquidPiston's (2018) X Mini 70cc engine, marketed as an "inverted Wankel engine" (LiquidPiston 2018). By utilizing a highly modified combustion cycle that combined aspects of those carried out by Diesel and RPEs, the X Mini reported an energy density as high as 344 $^{WH}/_{lb}$ when run off of gasoline, though it was capable of operating off of a wide variety of working fluids (Budapest University of Technology and Economics 2018). With a dry weight of 4 lb, though, the buy-in weight for this engine barred its usage in smaller UASs (LiquidPiston 2018).

Gas turbine engines operated on an axial layout, with intake at one end and exhaust at the other. Angularly symmetric fans rotated at high speed about centrally located shafts

and worked to either compress or expand air as it moved along the engine's length. Combustion occurred about midway along this path following the air's compression by fans. Petroleum-based fuels were sprayed into the flow and ignited to produce high pressures and temperatures. The energy of this reaction was directed rearwards along the engine's path; some of the energy was collected by fans to power the compression fans, but most was directed rear and out to produce thrust. The simplest version of a gas turbine engine, the turbojet, was given by Figure 1.24:

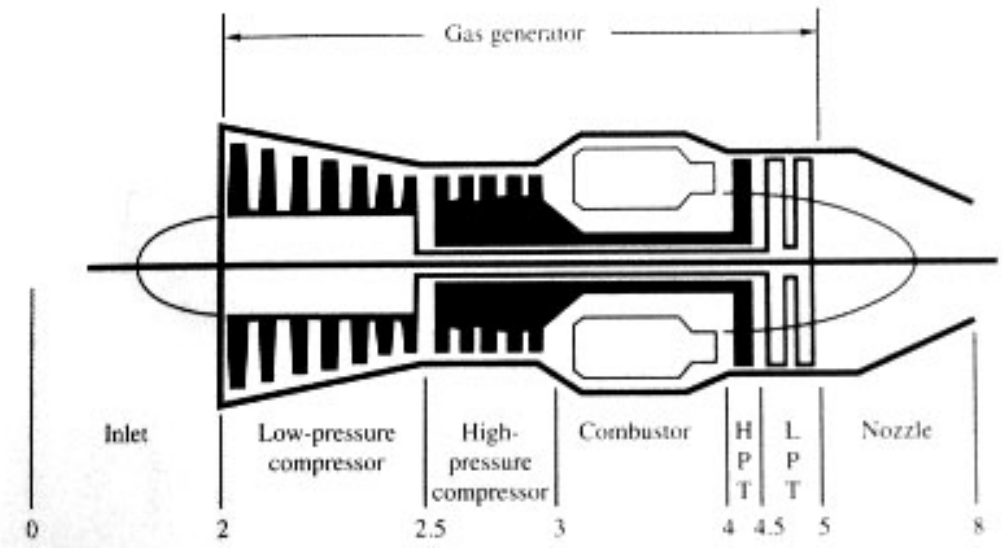

Figure 1.24: Schematic of a Turbojet Engine (Griffis et al. 2009)

The flow of air ran from **0** to **8**. Improvements to the turbojet included the addition of a larger diameter fan at the inlet to drive some flow around the core of the engine to raise efficiency. The turbofan engine housed this larger fan within an outer housing, while the turboprop left it exposed to the ambient flow. As the power produced by turbine engines was related to the fan diameters and rotational speed, the task of miniaturizing such engines for use on UASs was nontrivial and therefore only demonstrated in cases of massive UASs such as the 116.2 ft-wingspan Northrop Grumman Global Hawk (Griffis et al. 2009):

Figure 1.25: Image of the Northrop Grumman RQ-4A Global Hawk UAS (Northrop Grumman 2018)

Microturbines that could potentially power smaller UASs were slowly being developed, such as the PBS Aerospace EJ 20 Turbojet:

Figure 1.26: Image of the PBS Aerospace TJ20 Turbojet Engine (PBS Aerospace 2014)

Running off of kerosene, which exhibited an energy density of approximately 5400 $^{WH}/_{lb}$, the TJ 20's overall system energy density varied depending on flight speed. This has to due with the usage of system power efficiency, a function of forward velocity, in decrementing the overall energy density from that of the fuel's. At the CASS-specified sampling velocity of 52.49 $^{ft}/_{s}$, the engine achieved only 44 $^{WH}/_{lb}$, placing this option on the same energy density scale as batteries but with a substantial increase in system complexity. Such an energy source would have been better suited for much higher-velocity applications; indeed, the same TJ 20 would have achieved an evergy density of 567 $^{WH}/_{lb}$ at a flight speed of 400 knots.

The final type of energy source for UAS applications was harvesting energy devices. These produced usable energy through absorption from the surrounding environment. As such, energy harvesting devices could not be graded based upon energy density given that it was the devices' ability to convert otherwise useless energy that was the limiting factor in their performance and not the amount of energy available from their ultimate source

(often quite large, theoretically). These devices were also unique among the three major categories of energy sources as they featured a time-dependent component to the energy that they provided (Hays 2009).

The most popular form of energy harvesting device in use by UASs was photovoltaic solar cells (Hays 2009). Such cells utilized the photoelectric effect to convert the electromagnetic radiation from sunlight into usable electric current as was illustrated in Figure 1.27:

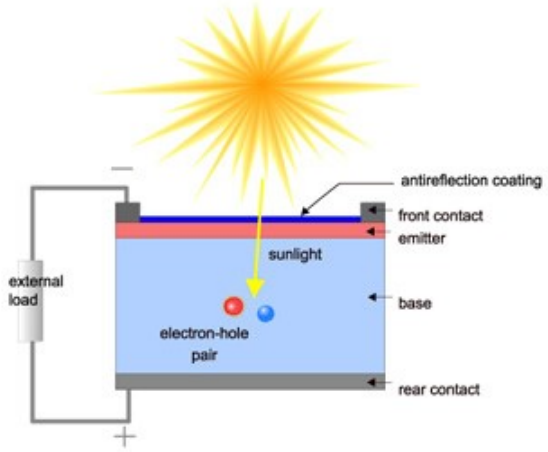

Figure 1.27: Schematic of a Photovoltaic Cell (Buonassisi 2013)

When particles of light, called photons, with sufficient energy impacted electrons in the solar cell's base metal, they dislodged electrons from the metal's atoms, thereby transferring some of their remaining energy and freeing the electrons to flow as current (Stierwalt 2015). The main drawbacks of photovoltaic cells revolved around the intermittency of sunlight in the context of UASs' continuous energy needs. An option for addressing this issue existed in the form of support batteries for "rainy-day" storage of energy, as utilized by the AC Propulsion SoLong UAS:

Figure 1.28: Image of AC Propulsion's SoLong UAS (Griffis et al. 2009)

which successfully demonstrated the efficacy of this energy configuration with flights of over 48 hours. As Figure 1.28 illustrated, UASs sourcing their energy from photovoltaics were limited to highly aerodynamically efficient designs that could operate entirely from the backup battery in times of poor sunlight. Additionally, adequate surface areas of solar cells had to be achieved to provide for the power requirements of the UAS, and special attention had to be paid to the climate in which the aircraft was to be flown in order to determine whether sufficient sunlight for flight was to be found there (Griffis et al. 2009).

Thermoelectric conversion made up the second possible harvesting energy source for UASs. Thermoelectric harvesting consisted of the conversion of heat energy into electrical energy as dictated by the Seebeck effect (Hays 2009). The Seebeck effect stated that if a temperature gradient was to be maintained across a thermoelectric material, a material with either free positive or negative particles, then the combination of particle diffusion due to

the temperature gradient and electrostatic repulsion would lead to a buildup of charged particles (and thus an electric potential) at one end of the material (Northwestern 2018):

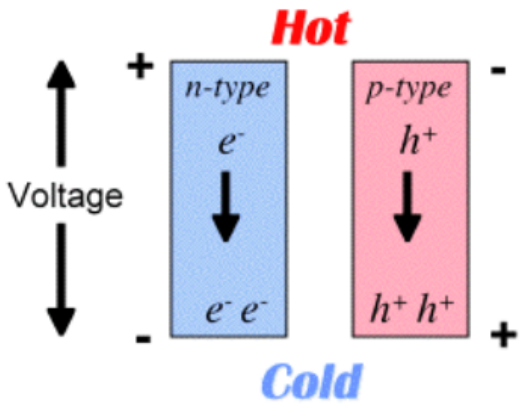

Figure 1.29: Schematic of the Seebeck Effect (Northwestern 2018)

As shown, if the thermoelectric material contained free positive particles, a p-type, then a positive charge developed on the cold side; likewise for an n-type material, a negative charge developed on the cold side. Had these oppositely charged materials' sides been connected via a load, the possibility for usable power would have existed. Use of thermoelectric devices in UASs included Fleming's (2004) use of them as efficient replacements for batteries in aircraft powered primarily by internal combustion engines. Devices were developed that achieved an overall energy density of 95.3 $^{WH}/_{lb}$ and which purportedly outperformed lithium batteries for endurances greater than twenty minutes (Fleming et al. 2004). This successful use also illustrated the main limitation of thermoelectric energy sources; namely, that they themselves required some source of thermal energy and as such were unable to power a UAS alone.

The two main factors affecting energy source selection were sensor measurement influence and desired endurance. In addition to system complexity, most conversion-type energy sources produced waste that, if not exhausted carefully, would have hampered accurate measurement by many of the chemical sensors on-board. Even the sterling engine, which featured no exhaust gases, would have had to be packaged so that the payload was

shielded against the engine's external heating source, which could also be detrimental to the sensors. Direct and harvesting energy sources did not produce waste products, and any heating due to operation could be more easily cordoned off from the payload. However, as the design endurance was on the order of one hour, the time-dependent component of the harvesting energy source's operation would have likely prevented successful operation over such a short period of time, thereby limiting this option's efficacy. Lack of readily-available hardware, at least in comparison to direct energy sources such as batteries, and increased complexity were major downsides to the use of more exotic energy sources such as photovoltaic cells and hydrogen fuel cells. Finally, the same minimum endurance requirement that barred implementation of harvesting type energy sources also placed the application beyond the realm of capacitive direct energy sources due to their energetic but relatively short operation. This left batteries as the most reasonable energy source for use with the given application.

1.2.4 Construction Methods

Weight has always been an enemy of aerospace performance, whether it be measured in top speed, endurance, or maximum possible forward acceleration, and this remained applicable on the scale of UASs. In fact, the difficulty in miniaturization of many of the energy sources available to full-scale aircraft served to exacerbate this effect. Consequently, the methods most often utilized for UAS construction were chiefly concerned with meeting the structural requirements of the aircraft as efficiently in weight as possible. This was tempered by the related effects of greater ease of manufacturability brought about by the crafts' smaller scales and cost. The unmanned aspect of UASs allowed for their use by hobbyists while simultaneously removing the consideration of an on-board pilot's safety.

Light but stiff, petroleum-based foams were often the material of choice for cost-efficient, small-scale UAS airframes. They were essentially matrices of hollow, interconnected plastic shells. Such materials achieved their shape through extrusion, which could

only produce simple linear shapes, or heated-expansion of large numbers of plastic beads that filled with gas or air, allowing the plastic to form to the shape of complex molds. In either case, foam could easily be cut to shape using heated cutting implements, sanded to fit, or repaired with glue, making the material one of the easiest to work with and maintain. The oldest type of foam used in UAS applications was polystyrene. Manufactured through extrusion and expansion, the latter variety was preferred for UASs for its comparably higher strength to weight (Reyes 2018). Expanded polystyrene (EPS) was most commonly found in home-built UASs or prototypes as its low hardness left the material prone to indentation during crashes or landings (Glover 2013). The more common foams found in modern UASs were expanded polyethylene (EPE) and expanded polypropylene (EPP). Each foam had its own advantages and disadvantages relative to the other, but both of these materials exhibited superior damage resistance versus EPS. According to Reyes (2018), EPE had a smoother surface finish than EPP and was easier to mold accurately, but EPP exhibited higher thermal resistance and material stiffness. In addition, comparative hardness testing of JSP Type C expanded bead foams revealed significantly higher hardness values for EPP vs EPE for any given stiffness between about 20 and 40 grams per liter (ARPRO 2018). An example of EPP's use was in the TuffWing UAS Mapper:

Figure 1.30: Image of the TuffWing UAS Mapper (TuffWing 2018)

Though durable, neither EPP nor EPE were stiff enough to entirely support the expected flight loads on most UASs' wings. In the case of the TuffWing, carbon fiber spars added the necessary stiffness. There were also numerous proprietary foam blends that sought to

optimize specific material characteristics; some of the names seen in UASs included EPO, Aerocell, Elapor, and Arcel (Reyes 2018).

Foam-cored manufacturing took many of the advantages of purely foam construction and complemented them with the increased strength and stiffness of composites. The concept was as follows: the desired three dimensional shape of a UAS's component was produced by way of computer numerical controlled (CNC) machining or hot-wire cutting of foam stock. This geometrically correct but structurally unsound component was then reinforced by a flexible fabric composite such as fiberglass or carbon fiber infused with a liquid epoxy matrix that dried and held the fabric to the shape of the foam core. The part could then be surface finished as desired. In this way, many of each materials' advantages were combined without the need for separate mold manufacturing or much of the intermediate assembly labor required of female-molded composite parts. Complex shapes could be produced using this technique, as shown by the Archytas Vertical Takeoff and Landing (VTOL) UAS:

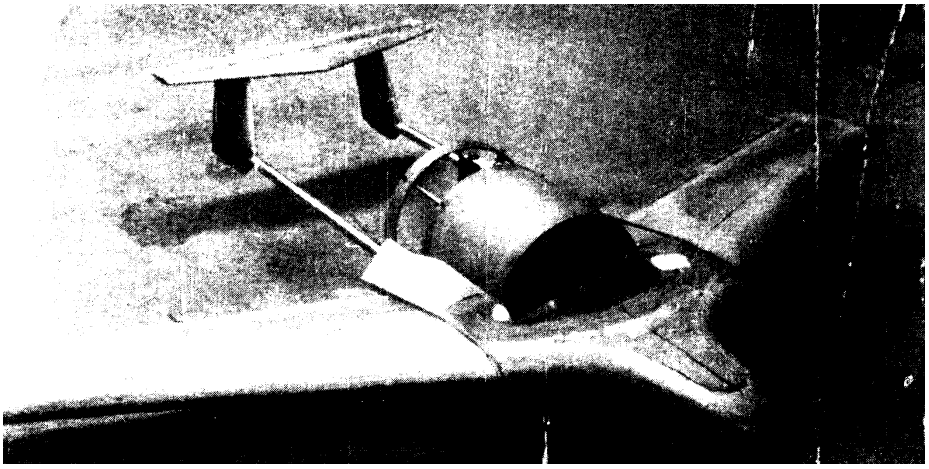

Figure 1.31: Image of Archytas VTOL UAS (Ellwood 1990)

The main drawbacks to using a foam-cored construction vs female molded composite construction were comparatively higher weight and lower surface quality. The bulk of the foam core, though durable, was excessive when compared to the bulkheads and air that filled the same volume in monocoque constructions. The surface finish of a foam-cored

component was inferior to that of a molded part due to the absence of prolonged pressure against a smooth mold surface during skin curing.

The traditional remote-controlled (RC) aircraft hobbyist's construction method of choice, balsa and doped fabric had a long and proven history of use in UASs. Balsa's comparatively high strength to weight and ease of manufacturability have long lent its use to the construction of nearly any structural component in small aircraft. Modern manufacturing tools such as precision laser cutters made the material's use even easier, with designs such as a Jamara Pitts S2B being possible:

Figure 1.32: Balsa Construction of Model Pitts S2B (Carpenter 2018)

Individual balsa parts were typically notched in order to aid full-component assembly, while white wood glues, aliphatic resins, and cyanoacrylate glues were common choices for adhering the parts permanently. The aircraft's skin could be formed by lightweight fabric such as silk or tissue paper that was doped, a process wherein fabric was soaked with a specialty liquid that adhered and shrunk the fabric to the balsa skeleton as it dried, or

by thin plastic sheets designed to shrink-to-fit when exposed to a heat source such as an iron or heat gun (Carpenter 2018). The result was an extremely lightweight construction; unsurprising, given that the majority of volume taken up by the construction was actually filled by air. The method was scalable up to large size scales, such as the 5.7 ft-wingspan Hobbico NexSTAR EP Select shown in Figure 1.14. Aircraft constructed using balsa were generally much lighter than similarly sized foam aircraft. There was, however, a considerable rise in labor and manufacturing time associated with this method, not least due to the numerous gluing and doping stages during which components had to be left alone to dry for extended periods of time. Balsa and fabric construction would also be particularly prone to damage as a consequence of normal field operations. These factors made foam and even foam-cored composite manufacturing more appealing for trainer platforms or prototypes.

The final construction method commonly found in UAS manufacturing involved the use of load-bearing composite skins in conjunction with an internal supporting skeleton, a construction method known as monocoque. As the skin was load-bearing, the supporting skeleton could be much less extensive than that seen in solely balsa and fabric construction as in the case of MARIA:

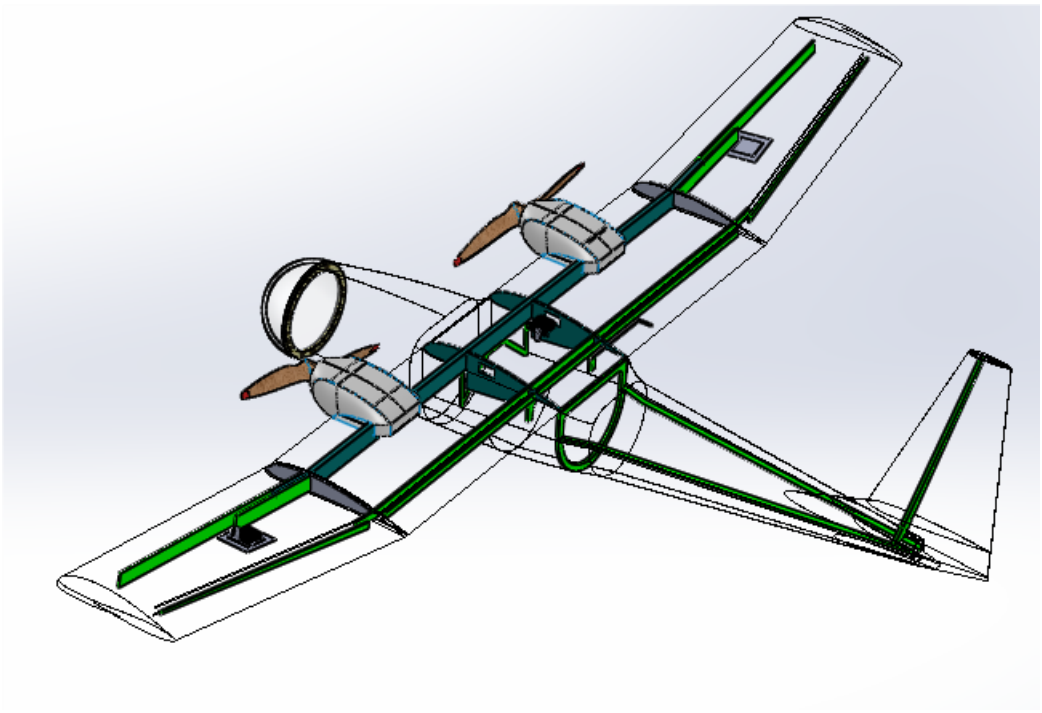

Figure 1.33: Structural Skeleton of the Composite MARIA (Avery 2013)

This method represented some of the state of the art for higher end UAS manufacturing, and it was the style of choice in applications where performance could be prioritized over cost. This point was crucial as the additional time and material cost associated with female-molded composites made it the costliest option of the four UAS construction methods. What was gained for this investment was greater freedom when it came to possible aircraft shapes and the best durability for the associated weight. Modern composites such as fiberglass, carbon fiber, and aramid fiber were flexible prior to infusion and curing of epoxy, allowing them to form to nearly any molded surface. In addition to greater design flexibility, the combination of sparse skeletal components, often manufactured similarly to those utilized in balsa and fabric construction, and an otherwise air-filled internal volume resulted in high weight efficiency for the right applications. Composite fibers' high strength and stiffness-to-weight ratios made them ideal for UAS size scales from medium up, but the initial weight buy-in for such materials degraded the weight benefits for smaller UASs. However, for UASs such as MARIA, monocoque construction allowed for a flight-worthy

aircraft capable of withstanding the unusually high flight loads resulting from operations within severe weather.

A durable construction capable of withstanding airborne loads from inclement weather and normal operating conditions in the field was highly desirable to CASS. This requirement alone eliminated the possibility of balsa and fabric construction due to this method's susceptibility to skin puncture and tear. Higher quality foams, such as EPP, would have fit the bill for certain cases, but past experience operating aircraft manufactured from such materials led to an aversion of their usage by the stakeholder. Examples of why included inadequate strength and stiffness in resisting the flight loads encountered in the field and the comparative lack of durability against handling and operational mishaps. It was also noted that the lion's share of commercially available UASs moved away from foam construction at the payload scales desired by CASS. Instead, most utilized some form of composite construction, and this route also made the most sense for the given application. A fully monocoque airframe was the most weight-efficient option, but the inflexibility of molded construction was recognized as a non-trivial drawback given the prototype aspect of the design. Foam-core components could have been more quickly manufactured than monocoques and thus allow for higher flexibility in design changes even into the manufacturing stage. However, the lack of readily available internal hardpoints inherent to foam-core components would have made the attachment of control surfaces and payloads difficult or impossible depending on the loads involved. Therefore, a molded wing-fuselage paired with flat, balsa-core tail surfaces were selected as the manufacturing method of choice in order to expedite structural design of the main body while mitigating manufacturing time required by the remaining airframe. A simple payload pod bolt pattern would allow for a high degree of flexibility in the CASS' applications without any need for changes to the main airframe.

Chapter 2

Design

2.1 Conceptual Design

The conceptual design stage consisted of the high-level decision-making processes involved in transferring CASS' requirements into a design framework. By combining this framework with the conclusions described in Section 1.2, a preliminary computer-aided design (CAD) of the aircraft could be drawn. The fundamental parameters defining this simplified geometry in turn served as the decision variables optimized through multidisciplinary, multidimensional optimization according to a parametric simulation of the proposed mission. In this way, the vast number of design permutations was trimmed down to comparatively few locally optimal configurations.

2.1.1 Preliminary CAD

In order to expedite the design of the aircraft, a number of simplifications were established early on. This included the use of a single airfoil for the main wing, being the SD 7062 (Selig et al. 1989):

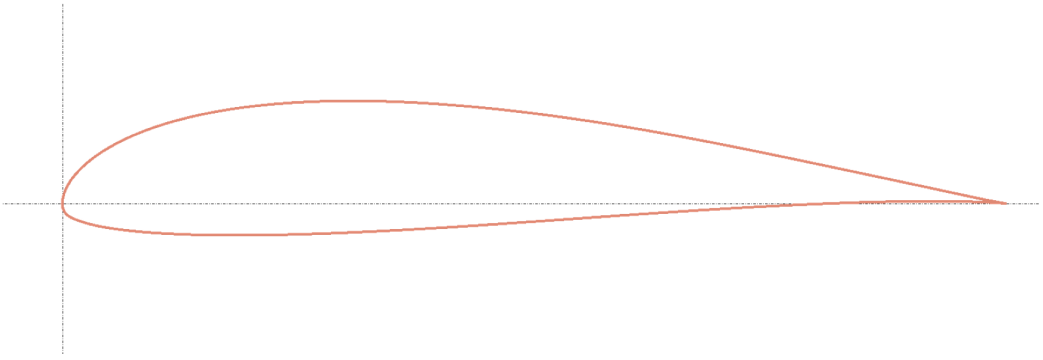

Figure 2.1: Outline of Selig-Donovan 7062 Airfoil

which was selected for its manufacturability and high lift characteristics; note that Figure 2.1 is scaled according to a chord of unity and is dimensionless. Additionally, a "box" wing planform with no twist or chord changes was selected. Rudder(s), elevator, and ailerons would serve as the craft's control surfaces. The ability to takeoff and land vertically, together with the variable-payload and endurance requirements, informed much of the layout decisions for the aircraft. Given that prop wash would negatively affect the airborne measurements of centrally-located sensors, centrally-located propellers would have required greater packaging complexity than simply utilizing dual-propellers on either side of the fuselage. This eliminated the possibility of a single-rotor helicopter configuration for VTOL but not multirotor options. Two additional rotors, so placed to balance the force from the two existing rotors about the aircraft CG, would allow for utilization of existing quad-rotor control software. This configuration also allowed for rotation of the front two powerplants into a dual tractor configuration for forward flight. Though this configuration left the rear two motor and props as "dead weight," it was decided that the comparatively high efficiency of fixed-wing flight versus solely copter-powered flight would purposefully be leveraged to achieve the desired endurance while retaining the stakeholder's requirement for field operations. Redundancy in structural weight was eliminated through placement of the aircraft's four powerplants on the existing tail struts, which connected an H-tail to the main wing. An H-tail easily accommodated flat horizontal and vertical tail components for ease of manufacturability, while the wide static margin possibilities supported the variable payload weights expected. The payload pod would be a simple prismatic shape for ease of manufacturing and would be placed underneath the main wing. Examples of some configuration concepts were shown in Figures 2.2 and 2.3:

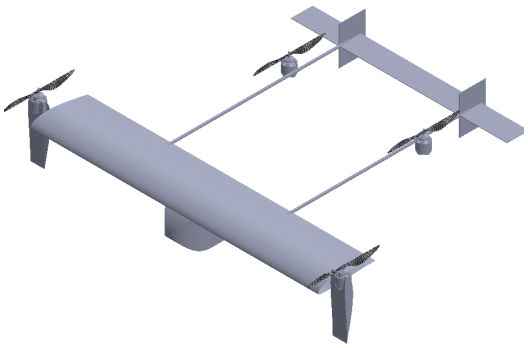

Figure 2.2: Preliminary CAD of Wingtip-Rotating Concept

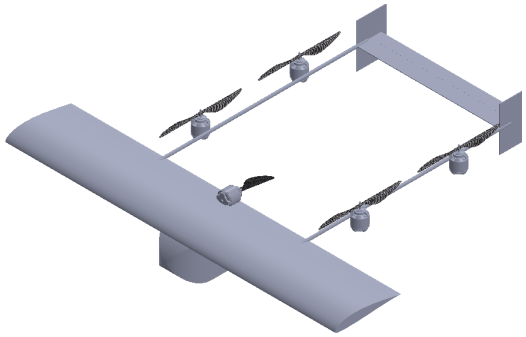

Figure 2.3: Preliminary CAD of 5-Motor Concept

Generation of such concepts was particularly helpful in determining the platform's intrinsic dimensions for optimization..

2.1.2 Mission Simulation

Software leveraging optimization, mission simulation, and modeling of aircraft performance was crucial in providing a solid foundation for the platform's design, especially given its prototype nature. The basic framework of the aircraft's mission was depicted by Figure 2.4:

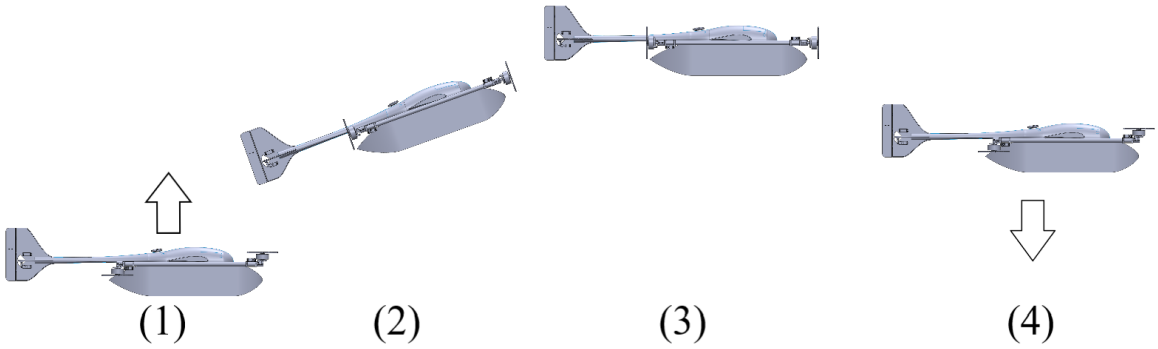

Figure 2.4: Outline of Mission Legs

with mission legs: vertical takeoff (1), climb to altitude (2), operation and descent to landing-altitude (3), and vertical landing (4). Each leg required a certain amount of power,

and these were calculated according to equations derived from several found in Raymer (2012). For (1), the power required to vertically ascend at a velocity V_{to} was given by:

$$Power_{to} = \eta_{tot}(W\sqrt{\frac{W}{2(\rho)S_{disk}}} + \frac{1}{2}(W)V_{to}) \tag{2.1a}$$

$$\eta_{tot} = \eta_{elec}(\eta_{ctrl}) \tag{2.1b}$$

where the subscript *to* abbreviated "take-off," η_{tot} denoted the total power efficiency factor as defined by Equation 2.1b, W stood for the total aircraft weight and whose value depended upon the sum of all the individual models' predicted weights, ρ represented atmospheric density and was taken to be 0.002377 $^{slugs}/_{ft^3}$, and S_{disk} was the sum of all rotor disk areas and was therefore a function of propeller diameter. η_{elec} represented the presumed electrical power losses and was taken to be 2, and η_{ctrl} book-kept the presumed power losses due to control inputs and was taken to be 1.2.

Following transition, the aircraft would climb to altitude under fixed wing lift (2); the power required to accomplish this was calculated according to:

$$AR = \frac{b^2}{S_{wing}} \tag{2.2a}$$

$$K = \frac{1}{\pi(e)AR} \tag{2.2b}$$

$$D = \frac{1}{2}(\rho)V_{cl}^2(S_{wing})(C_{Do} + KC_L^2) = Q(S_{wing})(C_{Do} + KC_L^2) \tag{2.2c}$$

$$Thrust_{cl} = D + W(\sin\gamma) \tag{2.2d}$$

$$Power_{cl} = \eta_{tot}(Thrust_{cl}V_{cl}) \tag{2.2e}$$

where AR stood for wing aspect ratio as defined by Equation 2.2a, b was the wingspan, S_{wing} represented wing planform area, K was the induced drag coefficient as defined by Equation 2.2b, e was the lift efficiency factor and was taken to be 0.85, D stood for drag force, the cl subscript denoted "climb," C_{Do} was the parasite drag coefficient, C_L represented the lift coefficient, and $Thrust_{cl}$ was the thrust force required to climb under fixed wing lift at a climb angle γ. Note that Q represented the dynamic pressure and that its velocity corresponded to the given mission leg.

In the (3) leg, the assumption was made that the aircraft would be operating at the minimum power requirement to maintain steady level unaccelerated flight (SLUF) in order to represent the aircraft's best endurance. This power requirement at a velocity V_{sluf} was dictated by Equation 2.3:

$$Power_{sluf} = \eta_{tot}(Q(S_{wing})C_{Do} + \frac{K(W^2)}{Q(S_{wing})}V_{sluf}) \tag{2.3}$$

Note that the use of total aircraft weight in Equation 2.3 stemmed from the equation of this weight with lift force, thereby minimizing the drag force due to lift production in the pursuit of maximum endurance.

In vertically descending during leg (4), the power required to do so mimicked that defined by Equation 2.1 and was taken to be:

$$Power_{la} = \eta_{tot}(W)\sqrt{\frac{W}{2(\rho)S_{disk}}} \tag{2.4}$$

with the difference being the removal of the power requirement due to gravity and with la denoting "landing."

In order to provide a starting point for design optimization, the aircraft's overall design configuration was condensed down into the following parameters:

Table 2.1: Aircraft Decision Variables

Parameter	Definition	Range	Units
S	Wing Area	[1.076, 53.82]	ft^2
b	Wingspan	[0.3281, 11.98]	ft
cap	Total Battery Capacity	[1, 10000]	WH
Vto	VTOL Takeoff Velocity	[0.1640, 16.40]	ft/s
Vla	VTOL Landing Velocity	[0.06562, 6.562]	ft/s
Vclim	FW Climb Velocity	[0.9843, 98.43]	ft/s
Climbang	FW Climb Angle	[1, 60]	deg
propdia	Propeller Diameter	[0.1, 0.5]strutL	ft

where WH denoted Watt*hours. It was from these intrinsic variables that all other aircraft values were derived and by which the aircraft design space was explored during the optimization. Note that "strutL" refers to the given configuration's tail boom length, itself a derived variable dependent upon wingspan, so that the the range of possible propeller diameters ranged from one tenth to one half of the strut length to capture some measure of packaging constraint.

The optimization technique utilized by the mission simulation software was the direct search algorithm known as Leapfrog. This algorithm was developed by Rhinehart (2019) as a robust alternative to other routines utilizing function gradients or Hessians which could fail for certain complicated applications involving data saddles, inflections, ridges, discontinuities, or stochastic behavior (Rhinehart 2019). The technique began with initialization of decision variables; this required specification of the possible ranges of values for each design parameter and the desired quantity of "players," being the individual design possibilities. For the given investigation, the decision variable ranges were as given by the "**Range**" column in Table 2.1 and the number of players was specified as being ten times the number of decision variables as recommended by Rhinehart (2019). The actual initialization of each player involved random assignment of each parameter according to:

$$Val = Rand(maxVal - minVal) + minVal \tag{2.5}$$

in which *Rand* represented a function whose output was a random value between and including zero and one, *maxVal* was the maximum value of the given decision variable, and *minVal* was the minimum value. The only deviation to this was propeller diameter, whose maximum and minimum values were only set for a given player following random assignment of its wingspan.

With the decision space fully populated with a randomized set of players, the next step involved evaluation of each combination's objective function value in order to rank the players' relative fitness. For the given investigation, this involved determination of a player's endurance according to Equations 2.1 - 2.4, which itself required detailed modeling of each configuration's weight and aerodynamics, as well as determination of electrical components to match. In the case of the weight model, the general order of operations consisted of determination of component geometry according to decision variable values before determination of component weight according to a known pseudo-density quantity, though some components such as the Pixhawk 2.1 flight controller were simply constants. Each component's weight relation was as defined in Table 2.2, where bold quantities denoted decision variables:

Table 2.2: Weight Model

Component	Geometric Relation	Weight Relation	Density
Battery	n/a	$\rho_{batt}(\mathbf{cap})$	$\rho_{batt} = 0.01429\ ^{lb}/_{Watt}$
Wiring	$L_{wire} = 1.5(\mathbf{b} + \mathbf{strutL})$	$\rho_{wire}(L_{wire})$	$\rho_{wire} = 0.05000\ ^{lb}/_{ft}$
Wing Skin	$S_{wet,wing} = 2(\mathbf{S})$	$\rho_{wing}(S_{wet,wing})$	$\rho_{wing} = 0.3000\ ^{lb}/_{ft^2}$
Tail Skin	$S_{wet,tail} = 0.19279(\mathbf{S})$	$\rho_{tail}(S_{wet,tail})$	$\rho_{tail} = 0.2000\ ^{lb}/_{ft^2}$
Tail Struts	n/a	$2(\rho_{strut})\mathbf{strutL}$	$\rho_{strut} = 0.08922\ ^{lb}/_{ft}$
Payload	n/a	22.241 N	n/a
Pixhawk	n/a	0.4805 N	n/a
Motor	n/a	4(motorsize)	n/a
Hardware	n/a	$0.15(W_{strut}) + 0.1(W_{payload} + W_{motor} + W_{wing}) + 0.0(5W_{tail})$	n/a
Additional Structure	n/a	$2(0.05W_{wing}) + 2(0.03W_{tail})$	n/a
Servos	n/a	7(servosize)	n/a

Note that "motorsize" and "servosize" were additional functions specifically intended to output motor and control servo weights appropriate to the given player's VTOL takeoff power and control surface torque requirements, respectively, as chosen from databases of commercially available components. The resulting overall weight value was then utilized

in the power requirement equations as part of determining the given player's objective function value.

The next model utilized was that for the configuration's aerodynamic properties during fixed wing operations. The bulk of the calculations were defined by a drag buildup as the coefficient of lift, C_L, was defined assuming an equivalency of weight and lift as dictated by SLUF (Alonso 2009). A quadratic drag model was chosen to represent the breakdown between parasitic and lift-induced drag where:

$$C_D = C_{Do} + \frac{C_L^2}{\pi(AR)e} \tag{2.6}$$

as used previously in the power-required expressions. C_{Do} itself was calculated according to:

$$C_{Do} = \frac{1.09 \sum C_{f,i}(k_i) S_{wet,i}}{\mathbf{S}} \tag{2.7}$$

where 1.09 served to book-keep a presumed 9% overall increase in parasite drag due to skin roughness, $C_{f,i}$ denoted a given component's von Karman skin friction coefficient, k_i represented a given component's form factor, and $S_{wet,i}$ was a given component's wetted area (Alonso 2009). 9% was chosen as a simplified though conservative estimate of skin roughness given that detailed determination would have been impractical given the large numbers of configurations to be tested. The von Karman skin friction approximation was based upon experimental data for flat plates with varying degrees of roughness. For fully turbulent plates, the data could be represented by:

$$C_f = \frac{0.455}{(\log Re)^{2.58}} \tag{2.8}$$

where Re was the standard Reynolds number as defined by component length. This base coefficient value was then altered by the form factor k in order to account for the increase in drag for "thick" components over commensurate flat plates because of increased surface

velocities and pressure drag (Alonso 2009). This factor depended upon a given component's "fineness ratio," taken to be the ratio of thickness or height versus flow-wise chord length. In the case of wings, this relationship was more explicit; for any SD 7062, its t/c was 14%, yielding a k of 1.34. For all other components, the form factor was determined based upon a trendline produced from Alonso's (2009) K versus Fineness Ratio figure for bodies of revolution:

$$k = 0.005313(\tfrac{t}{c})^2 - 0.1241(\tfrac{t}{c}) + 1.793 \tag{2.9}$$

Finally, all semi-flat components' S_{wet} were taken to be twice their planform areas and revolved components' were taken to be equal to all exposed surface area. The components whose contributions to parasite drag were book-kept were the wings, struts, horizontal stabilizer, vertical stabilizers, payload pod, and propellers. Note that the propellers were simply taken to be solid cylinders of equal diameter for the purpose of drag contribution. In this way, both coefficients of lift and drag were determined for use in a given player's objective function evaluation. Additionally, the optimal sampling velocity for minimum required power could then be determined according to:

$$V_{samp} = \sqrt[4]{\frac{4}{3}\frac{W}{\mathbf{S}}^2 \frac{K}{C_{Do}(\rho^2)}} \tag{2.10}$$

With all parameters defined, it only remained for the mission simulator itself to determine the max endurance possible for each player. This was accomplished by first calculating the power required to takeoff vertically as defined by application of the given player's weights and aerodynamic properties to Equation 2.1a. This done, the resulting battery capacity decrement could be calculated according to:

$$Cap_{used} = Power(t) \tag{2.11}$$

where *Power* was the power required to complete the given leg and t was the time required to do so in hours, so calculated according to the corresponding decision variable velocity value. In the case that a given leg exceeded the given configuration's remaining battery capacity, the failure to complete the mission resulted in a zero value for endurance so as to penalize that player. Assuming the configuration survived takeoff, the decrements corresponding to fixed wing climbing and then vertical landing were similarly calculated and subtracted, again with checks for non-zero battery capacities. The battery capacity that remained following these subtractions represented that left for operational endurance, calculated using Equation 2.11 solved for t and using sampling power-required.

Having determined each player's objective function value in their operational endurance, all configurations were ranked according to this value. The "leaping" in Leapfrog referred to the treatment of the worst-ranked player, whose decision variable values were altered according to:

$$Var_{new} = Var_{best} - Rand(Var_{best} - Var_{worst}) \tag{2.12}$$

where Var_{best} denoted the value of a given decision variable held by the best-ranked player, Var_{worst} denoted that variable's value held by the worst-ranked player, and Var_{new} represented the new value of this decision variable to be held by the previously worst-ranked player. In this way, the worst player "leapt" over the best player in the decision space by a random amount. This previously worst-ranked player's objective function was then re-evaluated, and the whole of the database of players re-ranked, thereby saving computational cost by only requiring evaluation of all players during their initialization. The algorithm then continued until all players converged to within a preset convergence criterion for their objective function values. The overall operation of the software was depicted in Figure 2.5:

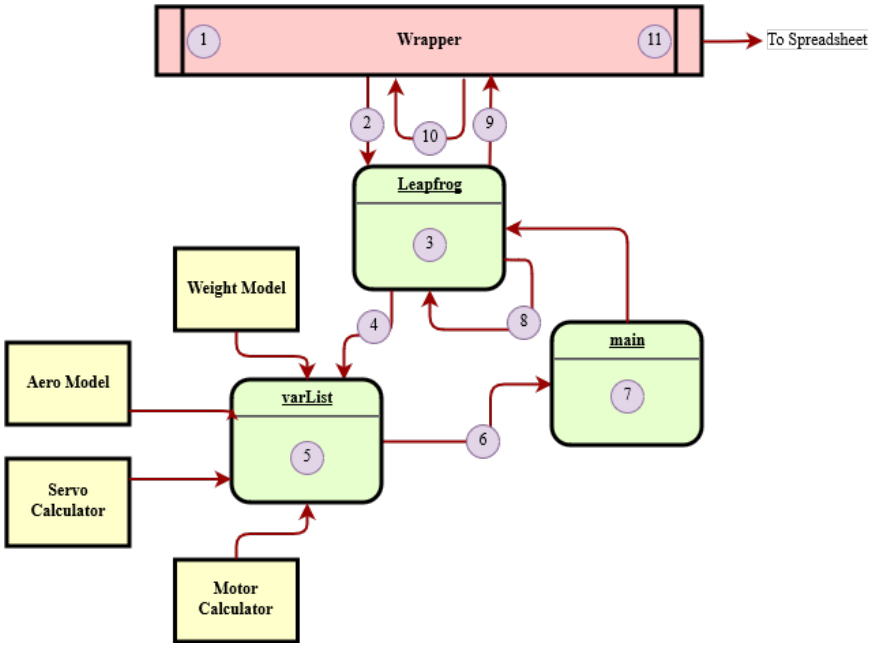

Figure 2.5: Graphical Overview of Mission Simulation Software

1. User defined amount of optimization trials to run
2. Wrapper made function call to Leapfrog optimizer
3. Leapfrog populated the decision space with 10(#DecisionVariables) random players
4. Each player's aircraft setup was determined through a function call to varList
5. varList utilized the various individual models to calculate all remaining aircraft parameters
6. Function call was made to main for each player
7. main calculated the physics of each player's performance, outputting max possible operational endurance
8. Players ranked according to endurance with worst performing being randomly redefined; process continued until all players converged
9. Endurance and setup of converged design outputted to wrapper
10. Wrapper ranked locally optimal designs by endurance
11. Ranked set is exported for post-processing

The algorithm was altered with regards to wing aspect ratio and platform weight. In the case of AR, a maximum value of 10 was imposed upon each player so as to prevent wing designs too skinny for manufacturing. As for weight, it was noticed that the optimizer tended to produce top designs with battery capacity values equal to the maximum allowed limit. To combat this, the original max battery capacity limit was raised to the very large value listed in Table 2.1 and different investigations were performed that implemented maximum total weight values of 15, 20, 25, 30, 35, and 40 pounds. Figure 2.6 presented a summary of the optimal endurance values as functions of maximum allowed weight:

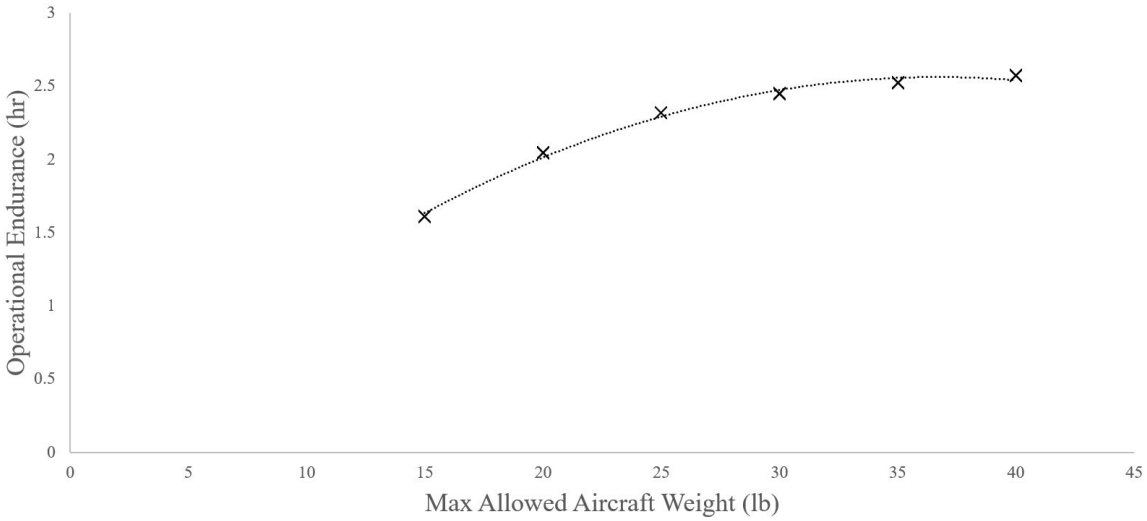

Figure 2.6: Top Design Operational Endurance vs. Maximum Allowed Total Weight

where all endurance values were decremented by half in order to account for the low fidelity nature of the analysis. It was interesting to note the similarity in the shape of Figure 2.6 with those of Figures 1.2 and 1.3, suggesting the oversize role of battery capacity in determining aircraft endurance despite the additional complexity of the other features found in the mission simulation.

2.1.3 Final Configuration

In the end, the configuration corresponding to a maximum allowed total weight of 20 pounds was chosen as the desired size for CASS, in terms of portability, that maintained

fulfillment of the performance requirements. Table 2.3 summarized the objective function values of the chosen design:

Table 2.3: Final Conceptual Design Configuration

Wing Area (ft^2)	**Wingspan** (ft)	**Battery Capacity** (WH)	**VTOL Takeoff Vel.** (ft/s)
6.848	8.274	642.7	8.373
VTOL Landing Vel. (ft/s)	**FW Climb Vel.** (ft/s)	**FW Climb Ang.** (deg)	**Prop Diameter** (ft)
4.452	18.57	24.34	1.803

This design configuration was predicted to achieve an operational endurance of 2.050 hours at a sampling velocity of 51.38 ft/s and was the top-ranked from 4,196,115 tested permutations. Figure 2.7 depicted a render produced according to the chosen configuration:

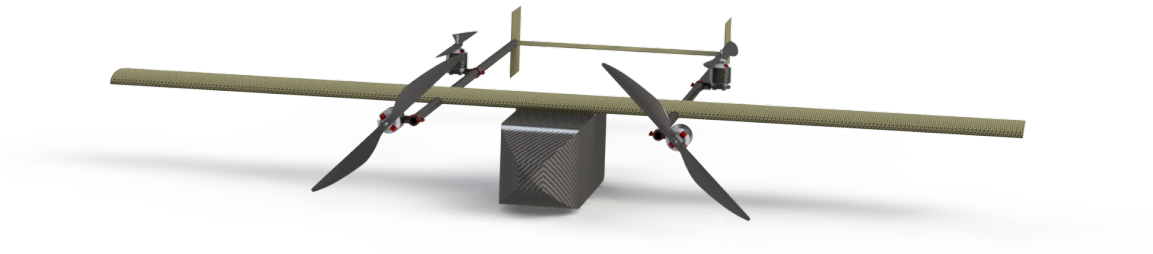

Figure 2.7: Render of Conceptual Design Aircraft

2.2 Detailed Design

The detailed design stage represented the portion of design devoted to detailed implementation of the relatively insular subsystems of the UAS subject to the design framework outlined in Section 2.1 as well as their eventual packaging into a complete platform. Doing so prior to the start of manufacturing was paramount in order to mitigate potentially costly delays due to lack of forethought in aircraft stability, propulsion, avionics, mechanical design.

2.2.1 Static Margin

Stability analysis was solely applied to the glide behavior of the UAS in its fixed wing configuration as its VTOL stability would be mostly incumbent upon proper tuning of the active control inputs from its four lift motors rather than the effects of airflow over its flight surfaces. There were two major categories of stability: pitch and yaw. Pitch stability was most greatly affected by relative size and placement of the craft's horizontal tail and by overall CG location. Yaw stability likewise depended upon CG location, though also vertical tail size and positioning (Drela et al. 2006). A well behaved aircraft would tend to maintain a given heading in the absence of control inputs as well as to level itself out if left to its own devices in the middle of a maneuver.

With respect to placement of the longitudinal CG, there existed a hypothetical location known as the Neutral Point (NP) at which a craft would exhibit neutrally stable pitch characteristics. Placement aft of this location would result in an aircraft whose tendency would be to drift in the direction of a pitch disturbance, resulting in pilot fatigue or loss of the platform. However, too far forward of the NP and the aircraft would exhibit sluggish responses to pitch inputs. The relative degree of pitch stability was represented by a parameter known as Static Margin (SM), defined as (Drela et al. 2006):

$$SM = \frac{x_{np} - x_{cg}}{\mathbf{c}} \tag{2.13}$$

where x denoted the distance of the given feature relative to some forward reference point and $\mathbf{c}$ represented the chord of the main wing. As determination of x_{np} would require greater detail in the design than was available immediately following conceptual design, the following approximation was utilized (Drela et al. 2006):

$$V_h = \frac{S_h(l_h)}{\mathbf{S}(\mathbf{c})} \tag{2.14a}$$

$$\frac{x_{np}}{\mathbf{c}} \simeq \frac{1}{4} + \frac{1 + 2/AR_{wing}}{1 + 2/AR_h}(1 - \frac{4}{AR_{wing}+2})V_h \tag{2.14b}$$

where the h subscript denoted horizontal, l_h referred to the distance between the CG and the horizontal tail quarter chord, and V_h represented a parameter known as the horizontal tail volume coefficient. Drela (2006) reported well-behaved aircraft V_h values of 0.3 - 0.6 for typical static margin values of 0.05 - 0.15. There was also an analogous parameter known as vertical tail volume coefficient, defined as:

$$V_v = \frac{S_v(l_v)}{\mathbf{S(b)}} \tag{2.15}$$

which provided the measure of yaw damping effectiveness of the vertical tail; typical values were given to be 0.02 - 0.05 (Drela et al. 2006).

For the purpose of achieving improved stability, the dimensions of the conceptual design's tail had to be altered to produce a reasonable static margin. Assuming a desired V_h of 0.4000 and lowering l_h from 5.249 to 3.937 feet, a new S_h of 0.7737 ft^2 was determined from Equation 2.14a. Assuming a constant AR_h of 7.260, Equation 2.14b predicted a x_{np} of 0.4573 feet. Using the most updated CAD of the design at the time, the CG location was approximated to be at 0.3182 feet for a static margin of 0.1524. Though this value lay just above the upper side of well-behaved SM values, slightly sluggish pitch behavior was deemed acceptable for a platform intended for steady, non-acrobatic flight. Assuming an l_v equal to l_h, as well as an unaltered S_v, V_v was calculated according to Equation 2.15 to be 0.03340, well within the acceptable range.

2.2.2 Stability Analysis

With an appropriate theoretical static margin achieved, the next step in ensuring the aircraft's stability was performed using an analysis tool known as XFLR5 (XFLR5 2014). This software expanded upon Mark Drela and Harold Youngren's original XFoil airfoil analysis program to include full stability analysis capability for model sailplane configurations.

The process began with individual analyses of the 2-D airfoils utilized in the full 3-D aerodynamic surfaces. These consisted of the SD 7062 airfoil (shown in Figure 2.1) in the main wing and a NACA 0003 to stand in for the flat elements found in the tail:

Table 2.4: XFLR5 Airfoil Geometry

Name	**Thickness (%)**	**at (% chord)**	**Camber (%)**	**at (% chord)**	**Points**
NACA 0003	3.000	29.35	-0.1300	0.08000	100
SD 7062	14.00	27.25	3.970	38.37	150

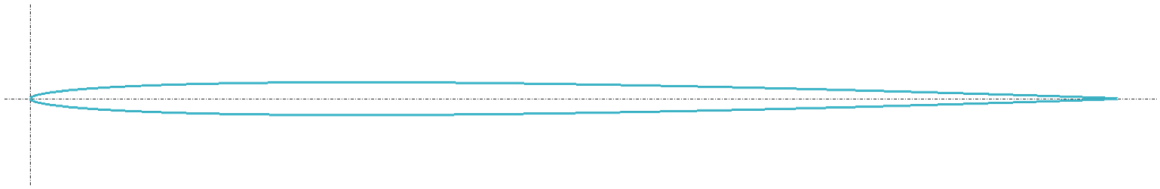

Figure 2.8: Outline of NACA 0003 Airfoil

where **Points** referred to the number of coordinate points used to defined the airfoil outline of unit chord length. Each airfoil's theoretical 2-D performance was evaluated for an angle attack range of -10 - 15 degrees and for Reynolds numbers from 10,000 - 1,010,000, producing figures such as Figure 2.9:

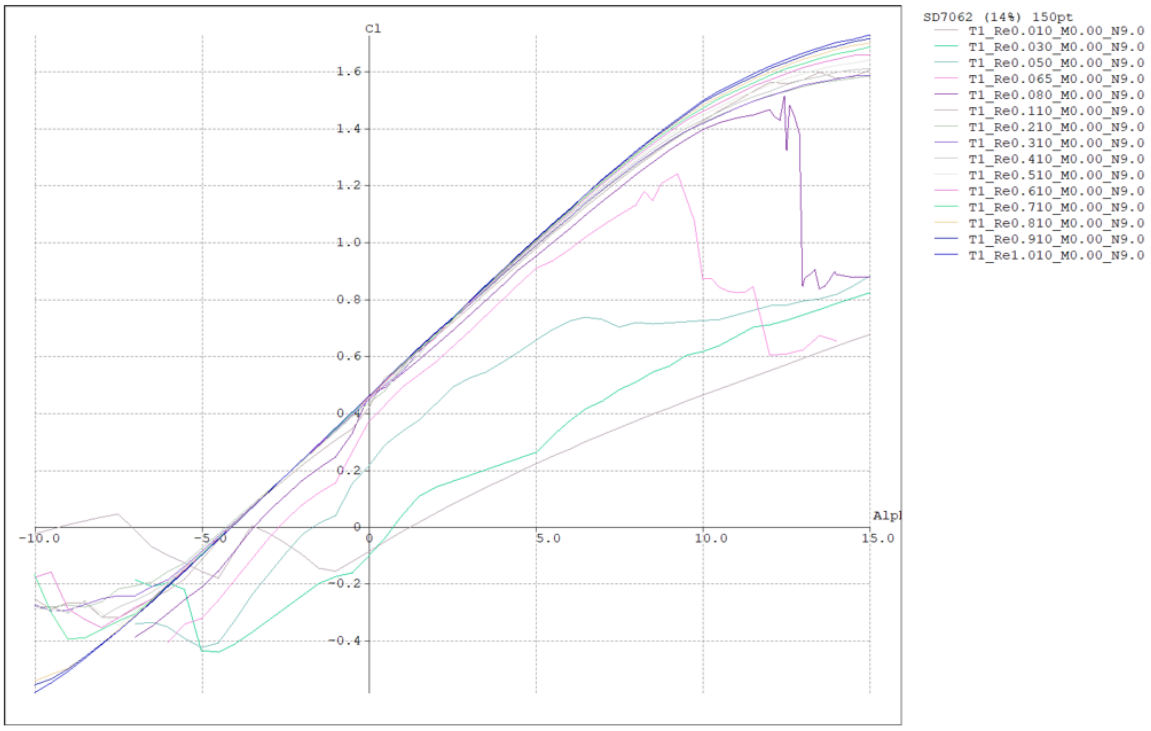

Figure 2.9: SD 7062 C_l vs. α

With the behavior of the individual airfoils established across a broad range of orientations and scenarios, the next step involved modeling the entire aircraft. This consisted of the specification of dimensions such as wingspan as well as relative placement of the aerodynamic surfaces. More intensive was individual mass placements for the purpose of approximate CG calculation; the exact components for which these were accounted for were listed in Table 2.5:

Table 2.5: Individual Component CG Locations

Component	**Weight** (lb)	**CG Location** (ft)
Wing	1.761	(0.3215, 0.0, 0.02624)
Horizontal Tail	0.1089	(3.832, 0.0, 0.003281)
Vertical Tail (2x)	0.05869	(3.881, +/– 1.066, 0.003281)
Payload Pod	5.000	(0.2297, 0.0, -0.3543)
Pixhawk	0.1080	(0.2067, 0.0, 0.08858)
Batteries	9.180	(0.2067, 0.0, 0.08858)
Front Tilt Servos (2x)	0.01984	(-1.102, +/– 1.260, 0.0)
Aileron Servos (2x)	0.01984	(0.5774, +/– 3.022, 0.0)
Rudder Servos (2x)	0.01984	(3.937, +/– 1.079, 0.0)
Elevator Servo	0.01984	(3.848, 0.0, 0.003281)
Front Motors (2x)	0.1720	(-1.102, +/– 1.260, 0.0)
Rear Motors (2x)	0.1720	(1.808, +/– 1.260, 0.1214)
Struts (2x)	0.3803	(1.404, +/– 1.073, 0.0)
Wiring Stand-in	0.9403	n/a
Hardware Stand-in	0.7617	n/a
Additional Structure Stand-in	0.1896	n/a

where **CG Location** was taken with respect to the center of the main wing's leading edge. The entries labeled as "Stand-in" represented components whose CG locations were indeterminate. Having fleshed out the aforementioned details, the following model was produced:

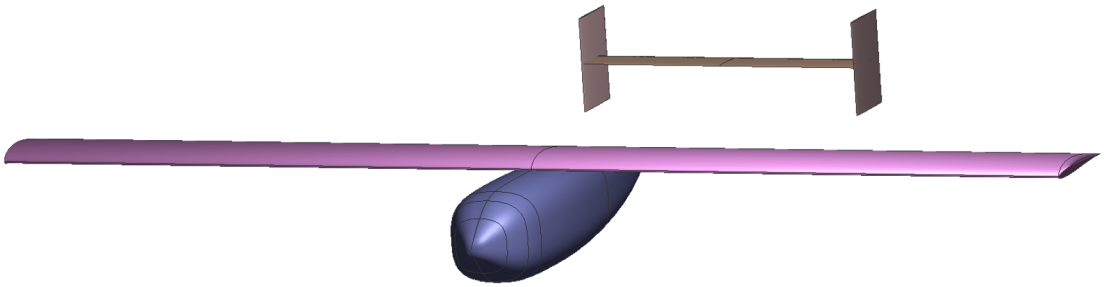

Figure 2.10: XFLR5 3-D Aircraft Model

Note that the payload pod shown in Figure 2.10 was used solely for CG placement and aesthetic purposes as recommended by the tool's guidelines. The overall configuration was tested over an angle of attack range of -2.5 - 13 degrees using the Vortex Lattice Method

which modeled the perturbation of the flow by a lifting surface through a summation of individual vortices distributed across the surface's length.

One of the first investigations performed on the design involved confirmation of a stable pitching moment with respect to angle of attack; the result of this analysis was depicted in Figure 2.11:

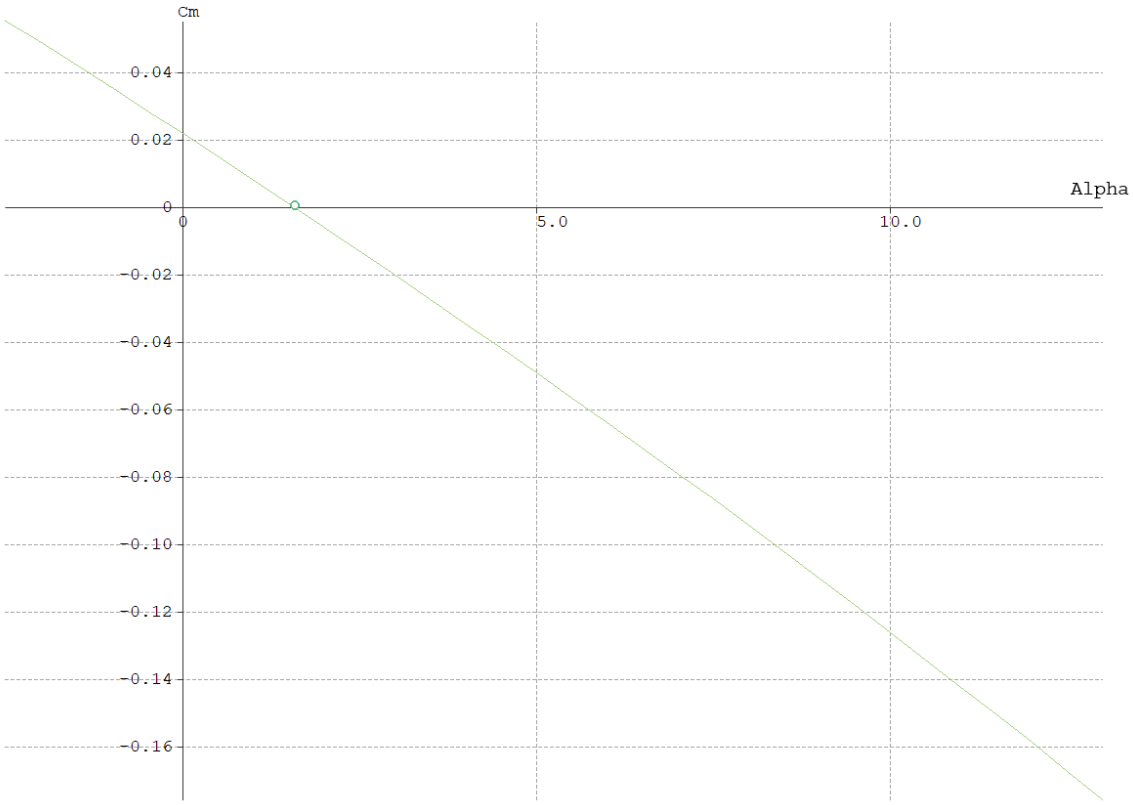

Figure 2.11: XFLR5 Coefficient of Pitching Moment vs. Angle of Attack

where C_m stood for coefficienct of pitching moment and the data mark represented the location of steady state. The negative slope of the curve was crucial as it ensured that any perturbation in pitch would be met by a pitching moment acting in the opposite direction to the perturbation. Also important to check was the zero-pitching moment lift coefficient, visible within the left curve of Figure 2.12:

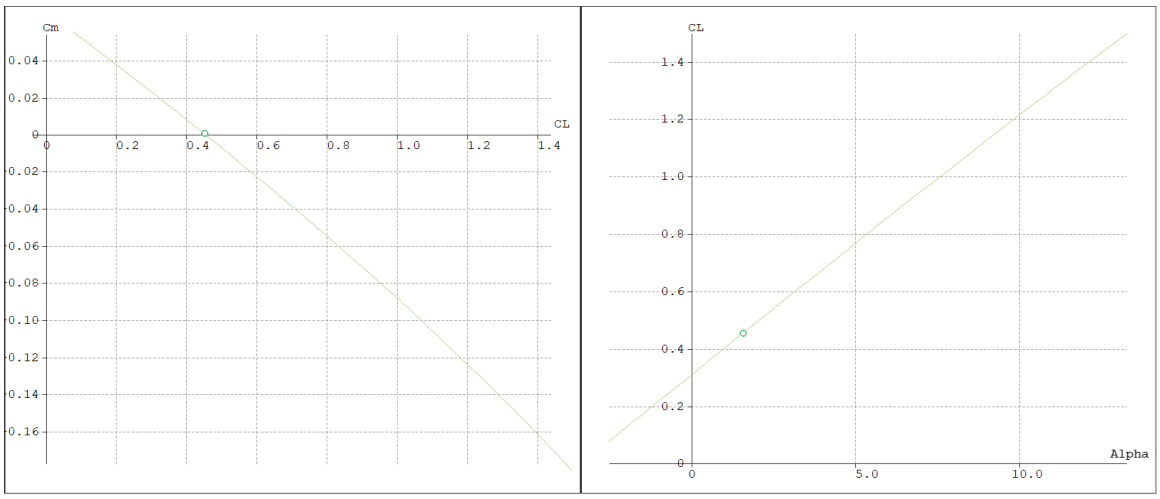

Figure 2.12: XFLR5 Coefficient of Pitching Moment vs. Coefficient of Lift and Coefficient of Lift vs. Angle of Attack

Under no control inputs, it was assumed that the aircraft would experience a zero-value of pitching moment; therefore, the marked point on the left curve of Figure 2.12 represented the overall steady state lift coefficient of the aircraft. Referencing the right curve of Figure 2.12, it was shown that this lift coefficient, as marked, fell well within the range of possible values.

Having confirmed static stability of the design, it then remained to investigate its various dynamic stability modes. These consisted of longitudinal modes, being short period and long period, as well as lateral modes, being roll damping, spiral, and dutch roll. Longitudinal modes occurred about the pitch plane's normal while lateral modes corresponded to behavior of the aircraft outside of this axis. Stability of the platform with respect to each of these modes was represented using a root locus chart featuring modal eigenvalues plotted on a complex plane, and Figures 2.13 and 2.14 depicted these for the two modal categories:

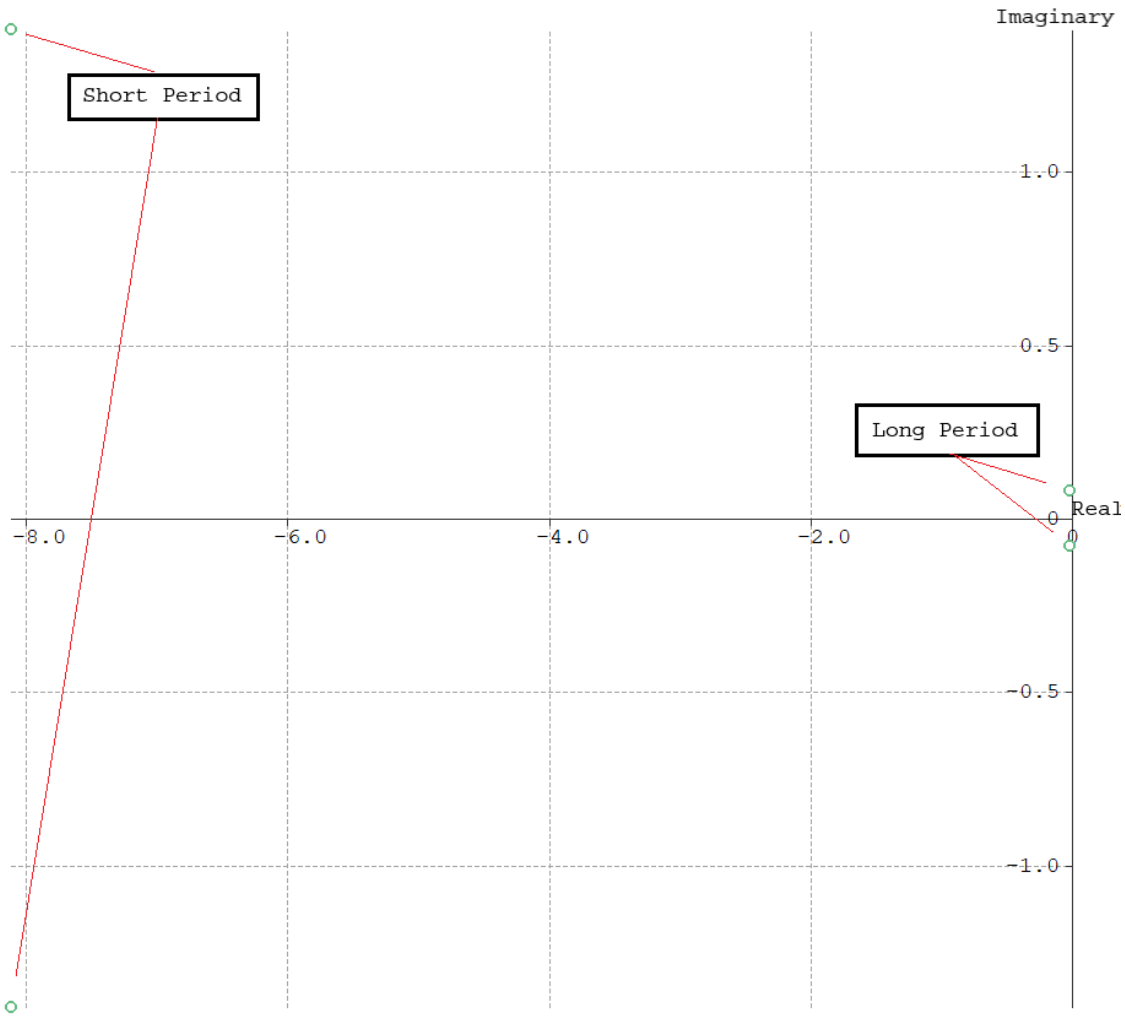

Figure 2.13: XFLR5 Longitudinal Mode Eigenvalues

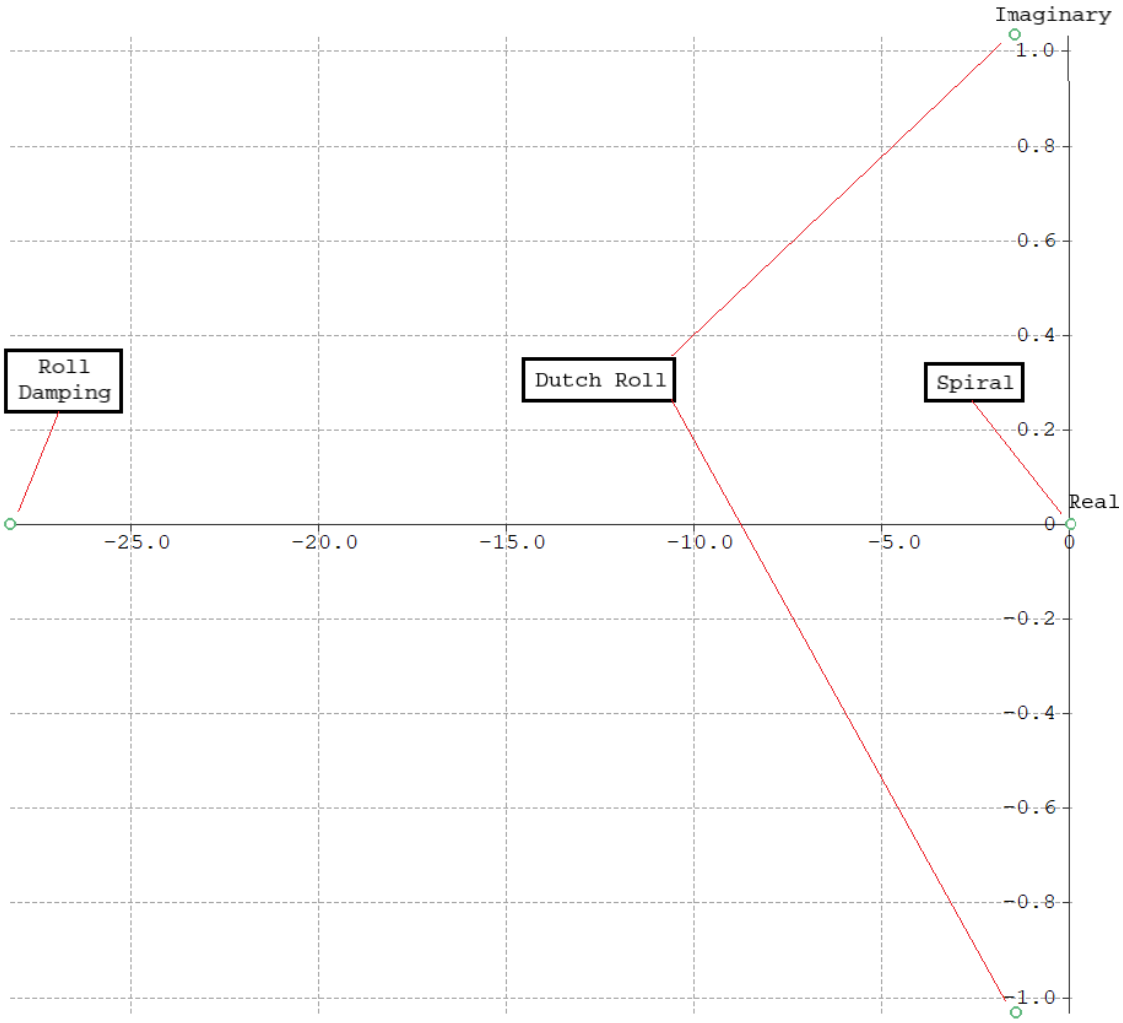

Figure 2.14: XFLR5 Lateral Mode Eigenvalues

where the imaginary components of each data point were plotted on the vertical axes and the real components on the horizontal axes. Theoretical stability required the eigenvalues of a given mode to feature negative real components. The roots of the short period mode, located on the left side of Figure 2.13, were well within the left hand side of the real axis. Those of the long period were also stable, though much closer to the break-even point for instability. As long period oscillations would occur over long periods of time, even some slight instability in this mode could easily be counteracted by human or autopilots. The far left data point in 2.14 corresponded to that of roll damping and was clearly stable. The pair

of points to the right side denoted the eigenvalues of dutch roll and were also stable. The spiral mode's eigenvalue was actually slightly unstable. However, similar to the case of long period behavior, the amount of time required for spiral to lead to divergent behavior would be quite large, and many successful aircraft designs featured slightly unstable spiral modes, so this was no cause for concern. In this way, the theoretical static and dynamic stability of the aircraft was confirmed.

2.2.3 Propulsion Design

The selection of a proper propeller and motor combination was of paramount importance to the performance of the aircraft. To be able to fly over a range of velocities, the motor-propeller system would have to spin at different speeds in order to match the thrust produced to the drag of the aircraft at the desired airspeed or the weight and desired VTOL velocity. The physical geometry of a propeller could be broadly defined by its overall diameter, the distance from propeller wingtip to wingtip, and its pitch, the idealized forward distance which the propeller would travel in space following one revolution. The differences in these values, along with spanwise twist and airfoil shape, lead to different ranges of efficient operation with respect to aircraft forward velocity, with propeller efficiency (η_p) defined as:

$$\eta_p = \frac{C_t(J)}{C_p} \tag{2.16}$$

where C_t stood for coefficient of propeller thrust, C_p represented the coefficient of propeller power-required, and J stood for a non-dimensionalized quantity known as advance ratio which book-kept both motion of the aircraft and of the propeller itself. Note that propeller

performance data typically consisted of tables listing η_p, C_t, and C_p as functions of J, with C_t, C_p, and J being defined according to:

$$C_t = \frac{T}{\rho(n^2)D^4} \tag{2.17a}$$

$$C_p = \frac{Power_p}{\rho(n^3)D^5} \tag{2.17b}$$

$$J = \frac{V}{n(D)} \tag{2.17c}$$

in which T denoted propeller thrust force, n represented propeller angular velocity, D was overall propeller diameter, V stood for aircraft forward velocity, and $Power_p$ stood for the power required for the prop to be spun under the given conditions. Also crucial was the efficiency of the motor, given as:

$$\eta_m = \frac{Power_{sh}}{Power_{el}} \tag{2.18}$$

where $Power_{sh}$ referred to the shaft power actually delivered by the motor to the propeller and $Power_{el}$ was the electrical power originally delivered to the motor. These values were themselves defined by:

$$Power_{sh} = (I - I_o)(Volt - IR_{tot}) \tag{2.19a}$$

$$Power_{el} = Volt(I) \tag{2.19b}$$

where I was the current delivered to the motor, I_o was the motor's zero-load current, $Volt$ was battery voltage, and R_{tot} represented the total resistance of the electrical system from the battery to the motor. The goal of propeller-motor matching was to identify combinations satisfying the basic power requirements of the aircraft while doing so at the highest possible aerodynamic and electrical efficiencies.

In order to begin an investigation into the optimal propulsion system, the relevant data for numerous propellers and motors were required. The performance data for all investigated propellers were sourced from the University of Illinois Urbana-Champaign's Propeller Database (Brandt et al. 2015) or APC Propellers' Performance Data (APC Propellers 2017), though significant data processing was required to transform the resulting individual data files into compatible repositories. The company T-Motor provided all information necessary for an analysis using any of their motors (T-Motor 2017); T-Motor was selected as the motor manufacturer due to this information availability as well as in-house familiarity. Derived from the maximum power-required values from the mission simulation results and initial investigations of motor-propeller combinations, a minimum continuous power of 2,000 Watts served as a preliminary metric in eventually selecting the 19 possible motors, whose relevant parameters were collated in Table A.2.

The foundation for the propulsion analysis tool was provided by a code developed for use in electric propulsion instruction; additional infrastructure allowed for use of the aforementioned propeller and motor repositories. Performance ranking unique to the given application provided the top combinations subject to the thrust requirements of the VTOL and forward flight configurations. Figure 2.15 illustrated the workflow of the code, where the arrows denoted inputs and outputs and each concentric ring represented calculation loops that were completed over the specified range to satisfy a single calculation of the next most outward ring:

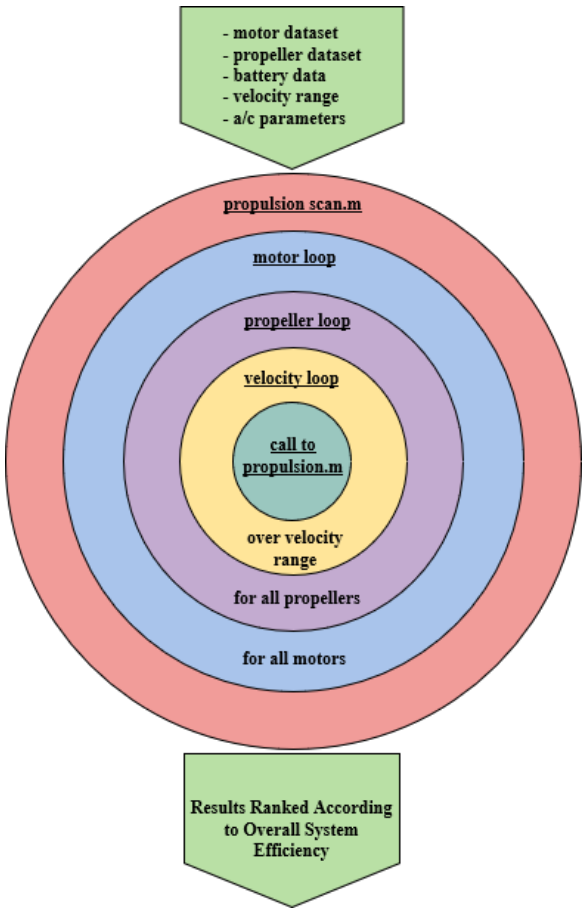

Figure 2.15: Graphical Overview of Propulsion Analysis Tool

The code's inputs besides the propeller and motor repositories included battery information as defined by the mission simulation, the desired velocity range as defined around the sampling velocity of 52.49 ft/s, and necessary aircraft parameters such as weight, drag coefficient, maximum coefficient of lift, and wing reference area. The remaining infrastructure simply looped for every combination so that each propeller was tested with each motor over the entire aircraft velocity range. The actual calculation was defined within a secondary function called propulsion.m whose inputs were those values being looped through in the outer wrapper. Within this calculator, the rotational velocity of the propeller and motor was incrementally increased from zero until the power delivered to the propeller (Equation 2.19a) matched the power required by the propeller to spin at the given aircraft

forward velocity (solving Equation 2.17b for $Power_p$). The state determined by this matching also defined the electrical current draw of the system. The method for actually altering the desired rotational velocity made use of battery voltage throttling wherein the throttle value simply served as a multiplier between zero and one for total battery voltage, and this throttled voltage defined the voltage value found in Equations 2.19a and 2.19b.

Following analysis of all possible combinations, a selection for the optimal system had to be made. Though the overall results were ranked according to overall propulsive efficiency:

$$\eta_{p,tot} = \eta_p(\eta_m) \tag{2.20}$$

a slightly lower efficiency system could prove superior in practice for a number of reasons, e.g. wider throttle and thrust ranges within the systems' respective motor maximum-power and maximum-current constraints, and this ended up being the case for the given application. The selected motor was the MN705-S KV260:

Figure 2.16: Image of T-Motor MN705-S (T-Motor 2017)

and the prop selected was the APC 20x13E:

Figure 2.17: Image of APC 20x13E Propeller (APC Propellers 2017)

The performance figures for this combination were plotted as Figures 2.18 and 2.19:

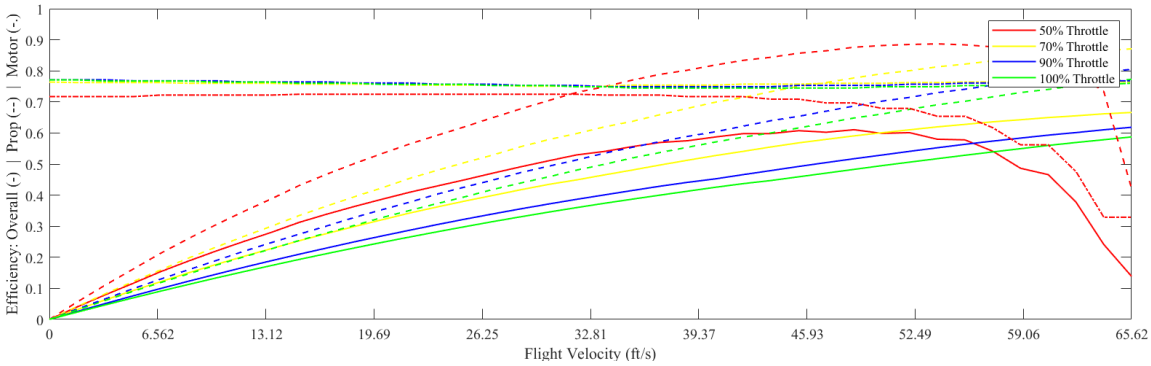

Figure 2.18: Propulsion Efficiency Curves

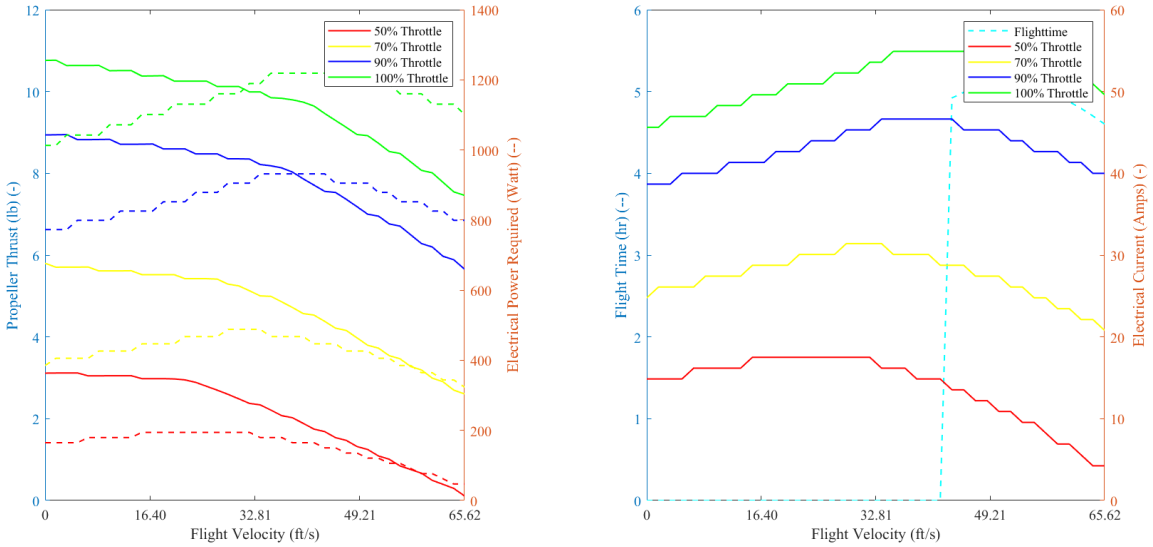

Figure 2.19: Propulsion Performance Curves

where each individual curve was identified by a dash pattern specified on the vertical axis and a color specified in the legend. Figure 2.18 clearly showed the effect of optimizing for a sampling velocity of 52.49 ft/s, with the overall efficiency curve for 50% peaking quite close. Maximum power draw never approached the motor's max allowable value of 2000 Watts, as shown in the left window of Figure 2.19, and the minimum requirement for vertical takeoff was determined to occur somewhere between 50% and 70% throttle. This would leave the propulsion system with ample additional power for maneuvering in the VTOL configuration. Figure 2.19's right-side window showed the effect of sampling velocity optimization, this time in terms of peak theoretical endurance. Additionally, even 100% throttle fell well beneath the max limit of 80 Amps across the entire forward velocity range.

2.2.4 Avionics System

The avionics system consisted of the various components and their connections involved in the storage, distribution, and usage of electrical power for aircraft control and propulsion. The flow of electrical power, denoted by arrows, was outlined for the system in Figure 2.20 with Table 2.6 serving as the key:

Table 2.6: Avionics System Components

Label	Component
1-5	Tattu 6,000 mAh Batteries
6,7	Mauch HS-200-LV Current Sensors
8	Mauch Sensor Hub X2
9	Mauch Power Cube 3 v3
10	Pixhawk 2.1
11	RFD900+ Modem
12	Holybro Airspeed Sensor
13	Here+ GPS Module
14	FrSky X4R-SB Receiver
15-17	Castle Creations BEC Pros
18-25	Hitec HS5645MG Servos
26,27,30,31	T-Motor 100A LV ESCs
28,29,32,33	T-Motor MN705-S KV260 Motors
34,35	Wing LEDs

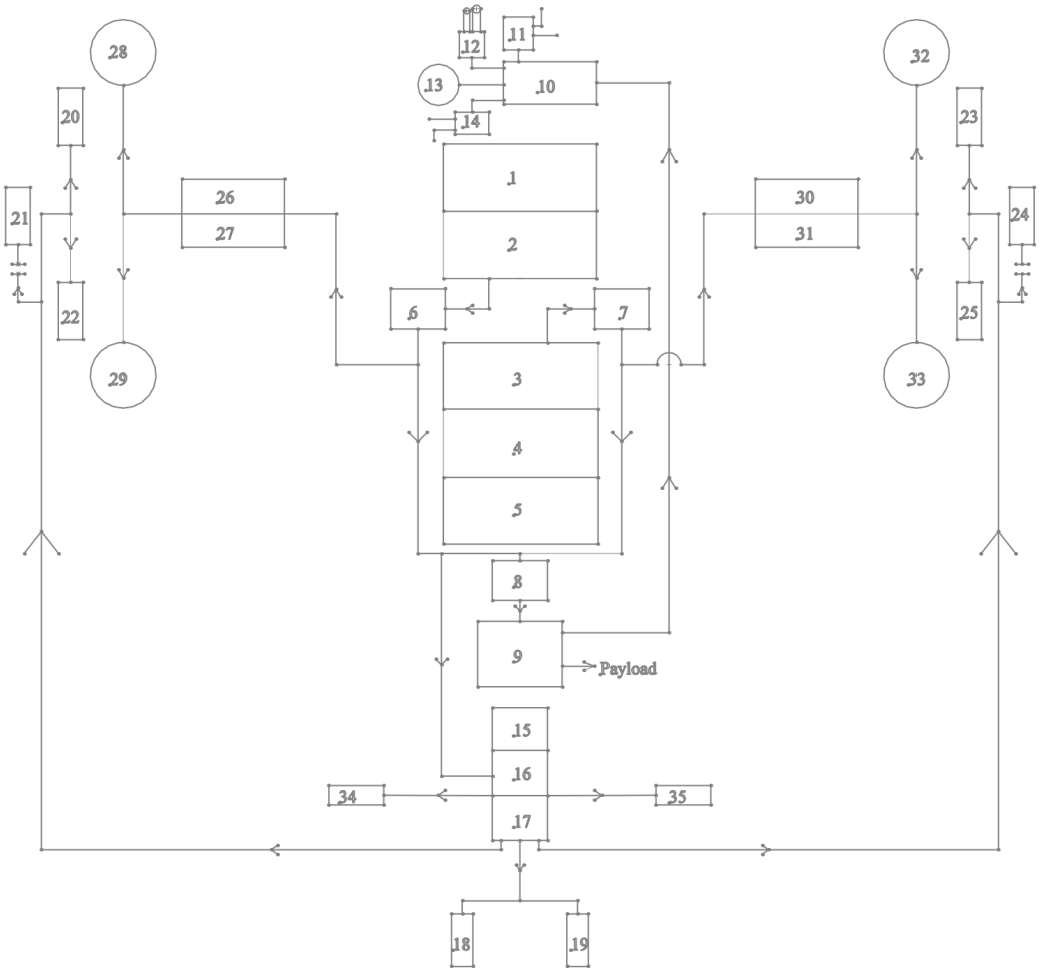

Figure 2.20: Avionics Wiring Diagram

The Pixhawk 2.1 flight controller made use of the Ardupilot autopilot software. Its main purpose was to run the QuadPlane autopilot code that would control the motor throttles, tilt servos, and control surface servos in order to follow predefined missions. The Pixhawk provided control signal connections for these components as well as interfaces for the GPS module and airspeed sensor, which provided inputs to Pixhawk on its location and velocity, and the FRSky receiver and RFD900+ modem, which received and delivered communications to the ground station and transmitter. The path of power began with the batteries, arranged in banks of two to the fore of the aircraft body and three aft. Batteries of the same bank shared a common power rail that led to one of the two 200 Amp current sensors. These sensors fed into the sensor hub which monitored and summarized the current and voltage from both battery banks, passing this information on to the Pixhawk through the power supply. The Mauch Power Cube 3 v3 replaced the standard Pixhawk power supply board due to the high power rating of the aircraft. The power supply board provided dual redundant 5.3 Volt power to the Pixhawk, passed on the current and voltage monitoring summaries from the sensor hub, and would provide 12 Volt power to the payload. The Pixhawk itself then provided for the small power requirements of the airspeed sensory, FrSky receiver, and RFD900+ modem. Also stemming from the two current sensors was the main power rail, to which were connected the four electronic speed controllers (ESCs) and the three battery elimination circuits (BECs). One to each motor, the 100-Amp ESCs converted the direct current (DC) from the batteries into alternating current (AC) for use by the motors and interpreted the signals from the Pixhawk dictating motor throttle. The BECs, each capable of supporting up to 20 Amps, stepped down voltage for use by the four control servos, the four tilt servos, and the two wing LEDs. Finally, information communication for the sensors and the payload was provided for through the Pixhawk's I^2C bus.

2.2.5 Mechanical Design

Mechanical design consisted of the detailed implementation of the structures that would physically support all of the subsystems heretofore described. For example, the whole of the aircraft's skin was reinforced throughout with a skeleton of 1/8" thick aircraft lite ply internals. The ribs, bulkheads, and stringers comprising this skeleton supported the monocoque skin against buckling and provided mounting locations for many of the avionics components. Closely tied to this skeleton were the main wing's load-bearing spars, which would ultimately support much of the UAS's weight during forward flight. Nesting carbon fiber tubes were chosen to serve this role while allowing for the outer halves of both wings to be removed for transport, as illustrated by Figure 2.21:

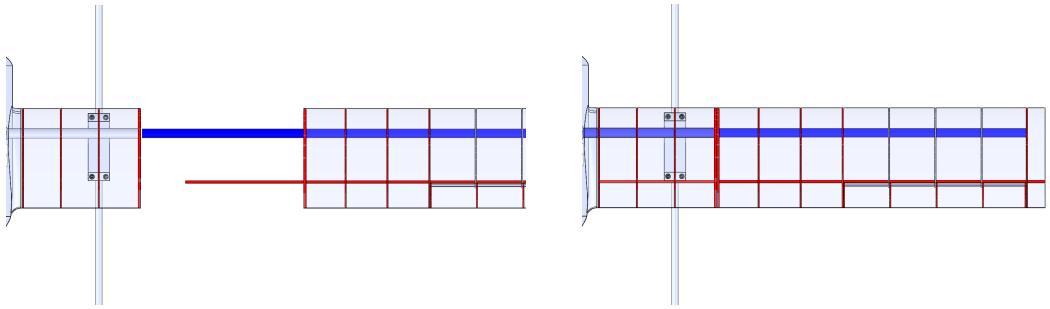

(a): CAD of Wing Conversion for Transport (b): CAD of Wing Conversion for Flight

Figure 2.21: CAD of Wing Conversion

where the inner spar, shown in blue, slid through the outer spar of the inner wing during operations. The trailing edge spar, shown in red, fit through corresponding holes placed in the inner wing's ribs. The trailing edge spar and the two wing-break ribs were designed as 1/4" thick components. The inner spar was secured in place during flight through a bolted connection between interfacing metal inserts from both wing spars as was illustrated by Figure 2.22:

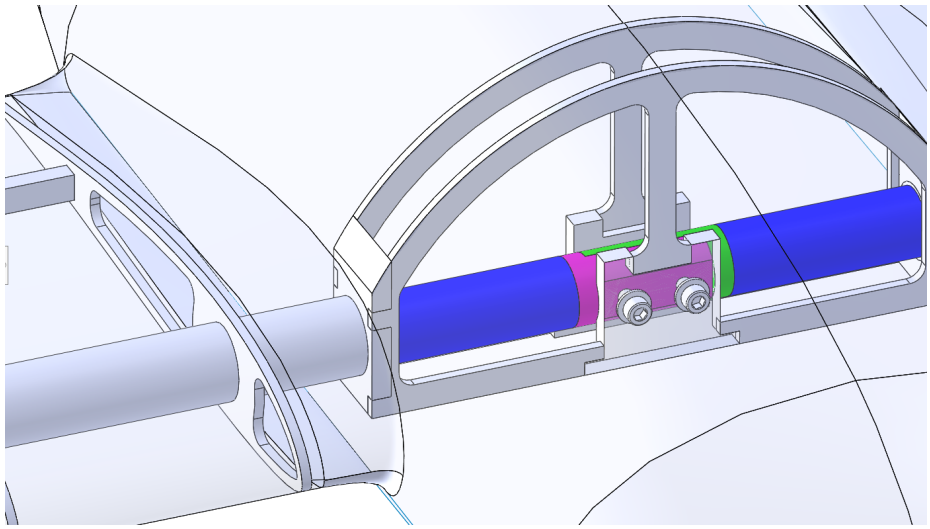

Figure 2.22: Inner Spar Anchor CAD

with the two spar inserts depicted in purple and green and the two spars in blue. Titanium 10-32 bolts fit through fore and aft 1/4" thick interface plates, threading into nylock nuts at their ends. In this way, the spars were prevented from displacement during flight while allowing for removal of the outer wing during transportation.

One of the most important mechanical requirements was a motor mount capable of transitioning between vertical and forward flight. No commercially available options existed for the loads expected from the propulsion system, so a custom design was pursued. Figure 2.23 depicts the initial concept CAD for such a mount:

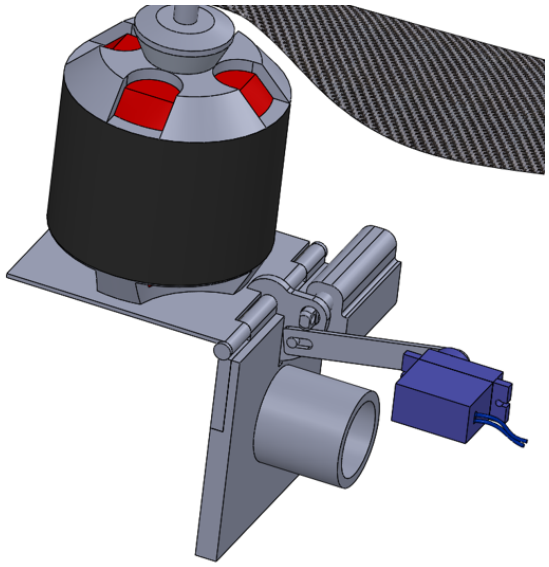

Figure 2.23: CAD of Initial Hinge Concept

The motor bolted directly to the raised platform, which was attached to the vertical plate via a hinge. This vertical plate fit over the front end of one of the tail struts. A linear actuator's arm provided a mechanical lock on the rotation of the motor, fitting into one of two holes in the central plate. This plate was attached to the motor platform and would rotate along with the motor platform, presenting its other hole to the linear actuator for re-locking once complete transition was achieved. The tilt servo, shown in blue, would be mounted on a platform attached to the strut. Some of the issues with this design included the displacement of the lift force vector away from the hinge in the VTOL configuration and the strong possibility of a miss-match between the linear actuator's arm and the central plate's holes during transition.

To address these and other issues, the motor mount design was iterated to a different concept. This concept moved the hinge in-plane with the strut center while leaving the tilt servo in-place, thereby directing the propeller force vector through the hinge at all orientations. Besides re-positioning of the hinge, this new concept implemented a dual-arm linkage for connecting the servo's output gear to the pedestal upon which the motor would be bolted. The shorter of these two arms would in fact be the servo's native metal

horn, albeit with a custom arm length; the longer arm would have to be a custom aluminum part. Though the majority of loading from the propeller would be directed through the motor pedestal's pivot outside of transition, some force through the linkage arm had to be anticipated. The design would direct this force through the center of the servo's output gear in both its fixed wing and VTOL settings rather than at some distance away, which would have required the servo to constantly resist a torque.

The correct layout and dimensions for this concept were determined using simple sketches in the form of Figure 2.24:

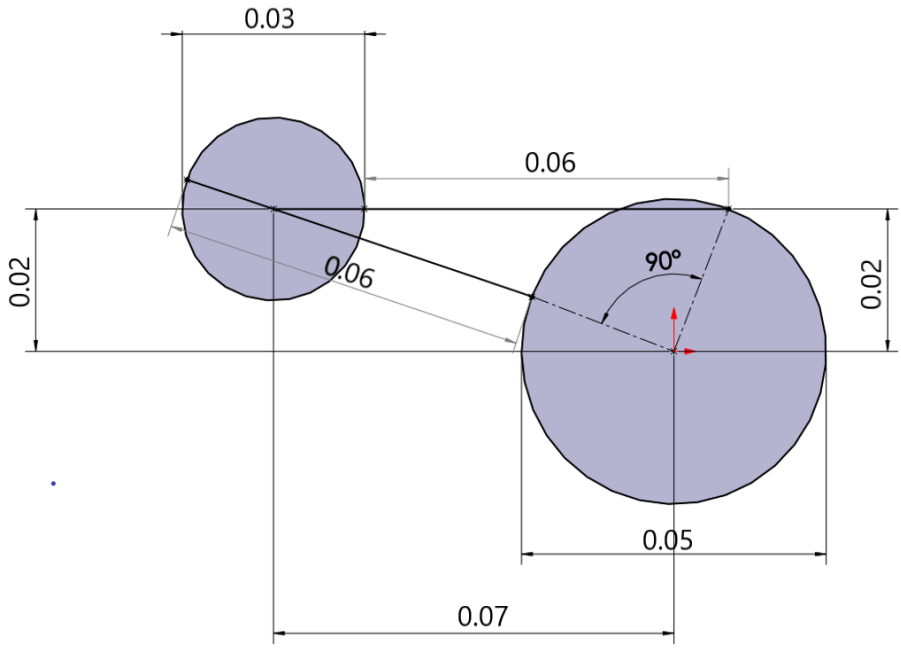

Figure 2.24: Motor Mount Linkage Sketch

in which the radius of the smaller circle represented the custom length of the servo horn, the radius of the larger circle represented the distance from the pedestal pivot to the pedestal's linkage hardpoint, and the two lines connecting the circles took the place of the linkage arm. The center of the bigger circle, being the pivot location, was fixed. The relative location of the servo's output gear was fixed according to packaging of the actual components, and the placement of the pedestal's hardpoint in the fixed wing setting was set according to the

pedestal's geometry. Knowing that the VTOL setting would angularly displace this hard-point by 90 degrees, this location was also defined. Given that the linkage arm necessarily had to remain the same length, the servo horn's custom length was iteratively adjusted until the two linkage arm lengths matched, fully defining the intrinsic dimensions of the design.

Even with this established methodology, it remained for physical implementation to confirm the efficacy of the design; this was performed using simplified geometry which nevertheless matched the intrinsic dimensions defined in the corresponding sketch. Figures 2.25(a) - 2.25(c) depicted three of the most major design iterations in the form of 3D-printed prototypes:

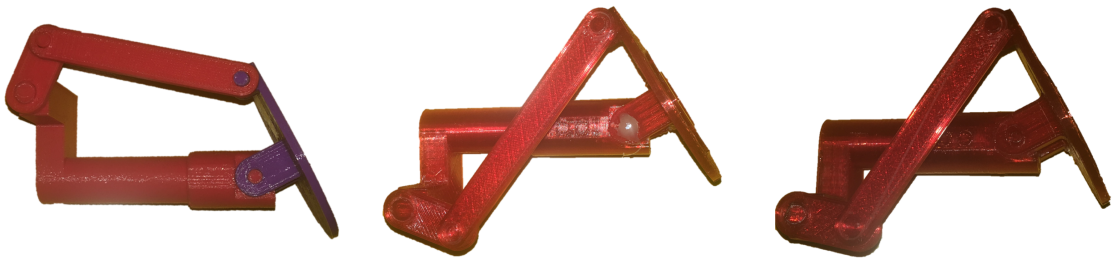

(a) (b) (c)

Figure 2.25: 3D-Printed Motor Mount Prototypes

where the designs progressed chronologically from (a) to (c). The issue encountered in prototype (a) surrounded binding of the linkage in the VTOL setting. This was alleviated by moving the the stand-in servo location beneath the strut, as in prototype (b). This alteration greatly lengthened the linkage arm, raising the possibility of miss-alignment during operation; therefore, prototype (c) shifted the servo as far forward as packaging would allow, compacting the assembly to a reasonable level. Figure 2.26 showed the CAD implemented from prototype (c):

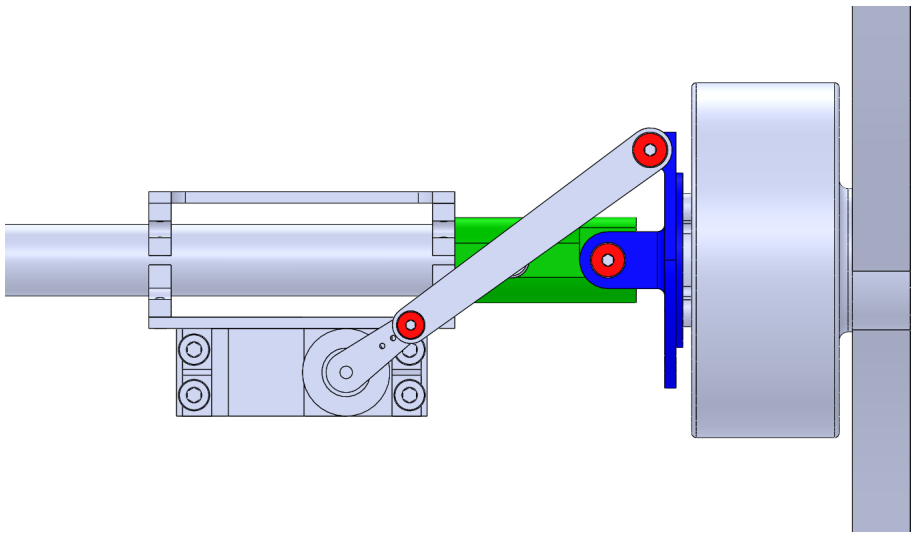

Figure 2.26: Motor Mount CAD

The three pivots consisted of steel shoulder bolts, shown in red, which provided smooth and precise hinge surfaces. The motor pedestal, depicted in blue, featured a raised bed at its interface with the motor in order to accommodate the MN705-S' exposed bottom spindle. This pedestal was connected to the strut via a sleeve, shown in green, which fit over the end of the strut; a bolt passing through both the sleeve and the strut complimented the two's existing bonded connection. This sleeve featured shelves that closely matched the curvature of the pedestal's tabs, providing mechanical stops to aid in constraining the rotating assembly once in its VTOL or fixed wing setting. The linkage arm was a simple flat design of the length determined from prototyping and attached to the servo horn and pedestal hardpoint via shoulder bolts. The servo itself was bolted to a simple bracket designed to fit commercially available tube clamps. Finally, a landing strut was implemented into the final design of the motor pedestal, designed to rotate with the motor so as to be deployed in VTOL and stowed in fixed wing flight.

The connection of the strut tube to the underside of the wing also went through a number of designs. It was clear from the beginning that some sort of curved interface would be required in order to match the bottom wing face with the circular strut. 3D-printed or computer numerical control (CNC) routered wood blocks spanning the entire wing chord

were initially envisioned. Due to concerns over the structural fitness of 3D-printing, this option was abandoned early on. Eventually, the idea of a chord-long block was iterated down to two discrete metal clamps towards the front and rear of the wing, saving weight and frontal area versus the alternative. Figures 2.27 and 2.28 illustrated the difference between these two concepts:

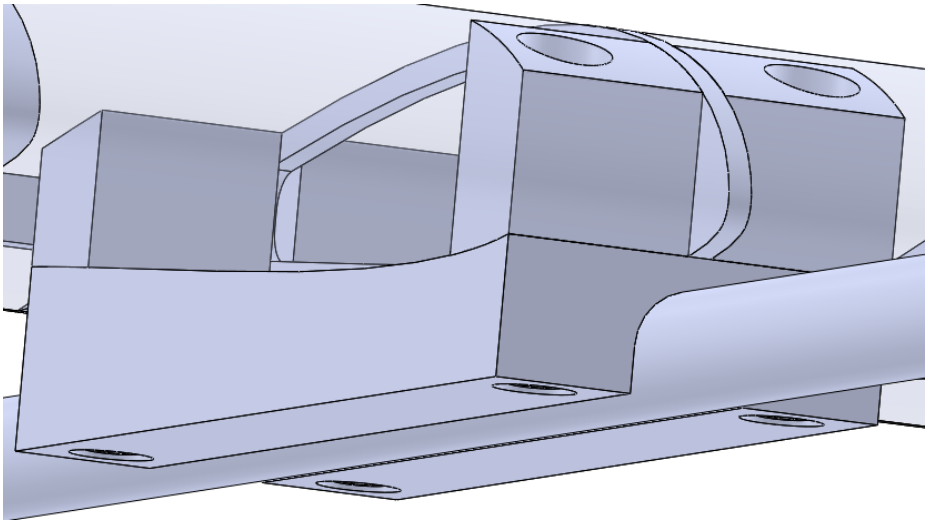

Figure 2.27: Initial Strut Interface CAD

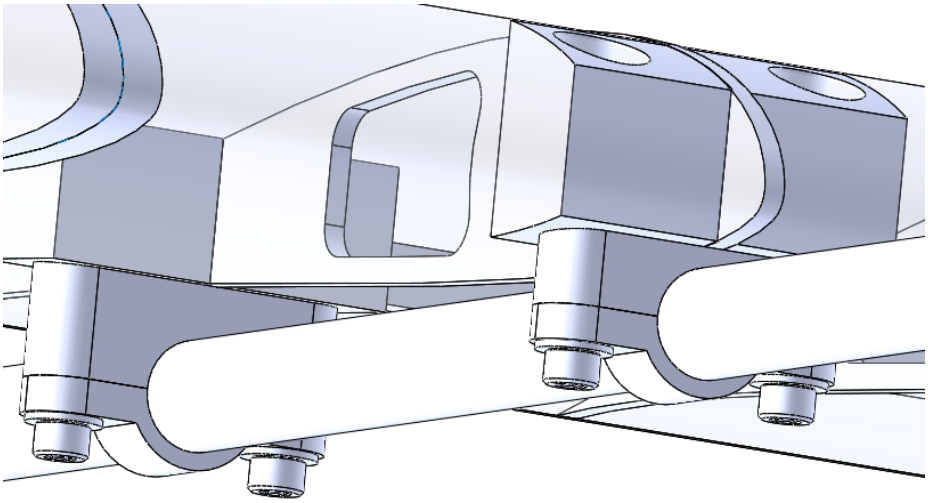

Figure 2.28: Final Strut Interface CAD

Also shown were the internal reinforcements to which both versions of the strut interface bolted. These were designed to be routered out of blocks of hardwood and bonded in

place. The bottom clamps shown in Figure 2.28 would be bonded to the strut, while the two interfaces would only be bonded to the underside of the wing, allowing for removal of the entire strut assembly if necessary. Attachment of the two halves was provided for by four 10-32 socket head bolts which threaded into matching nylock nuts bonded inside the internal reinforcement blocks.

The control surfaces required structures to support freedom of movement and to connect them with their respective control servos. In the case of the ailerons, some of the existing ribs were co-opted as horns for this purpose. The design initially located the servo on the suction side of the wing to allow for access after the plane was sealed up. However, this was abandoned in favor of fully internal mounting with a hatch to provide access. The linkage was also moved to pressure side of the wing to mitigate disruption of the flow over the top. Figures 2.29 and 2.30 depicted these two concepts:

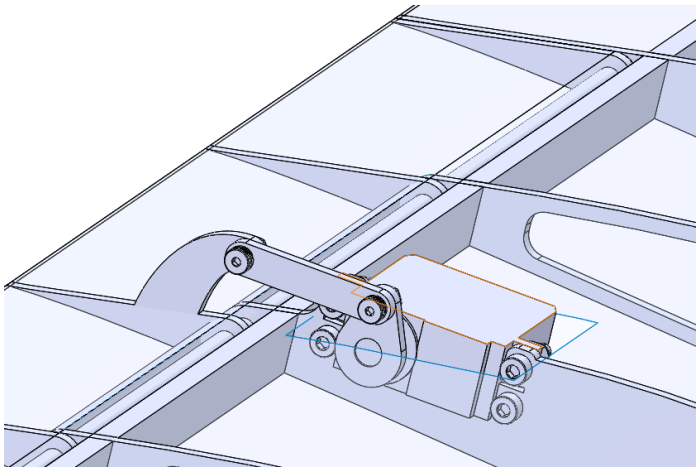

Figure 2.29: CAD of Initial Aileron Servo Linkage

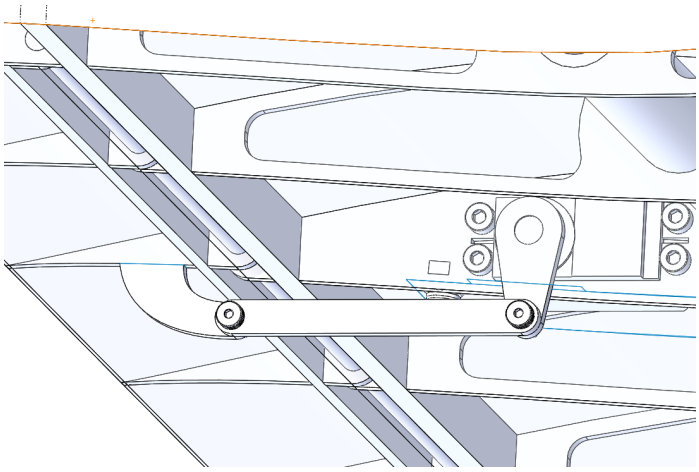

Figure 2.30: CAD of Final Aileron Linkage

Also visible was the hinge design, which implemented a solid carbon fiber rod as the hinging surface. This rod was itself supported by and bonded to the front of the aileron ribs and the rear of the associated wing ribs via holes included for this purpose.

The hinge design of the rear tail's surfaces had to be different as these surfaces were designed to be flat components with a solid balsa core rather than a skeleton. Instead, commercially available hinges that could be bonded inside of their respective surfaces were chosen, as shown by Figure 2.31:

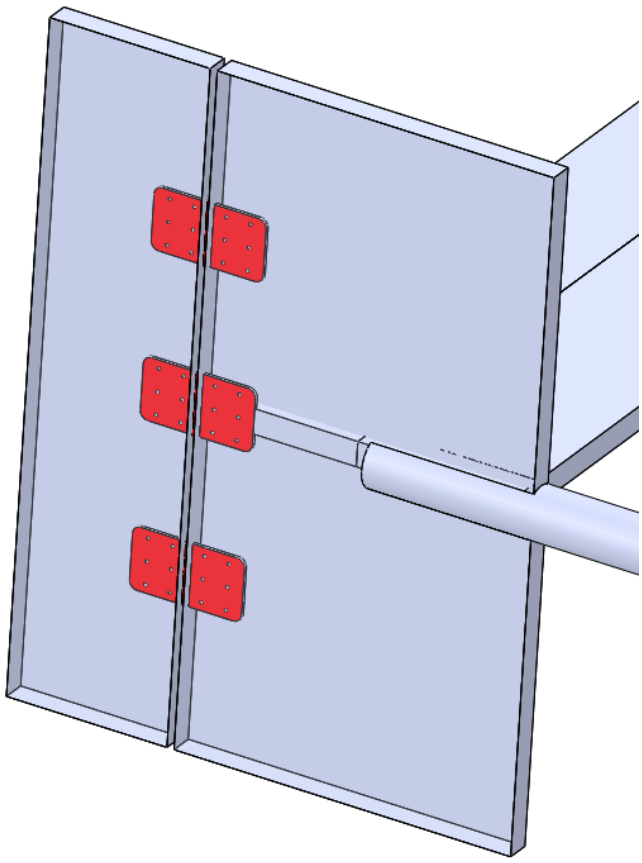

Figure 2.31: CAD of Vertical Tail Hinges

The red-colored hinges would be embedded into the surface's balsa core and bonded in place; a metal pin served as the actual hinge point.

It was at this time that a possible performance issue associated with the two non-transitioning rear motors was identified; namely, the natural tendency of the unpowered propellers to orient themselves in the flow during forward flight. The mounting for these propulsion assemblies resembled that of the tilt servos, being simple platforms bolted to tube clamps as depicted in Figure 2.32:

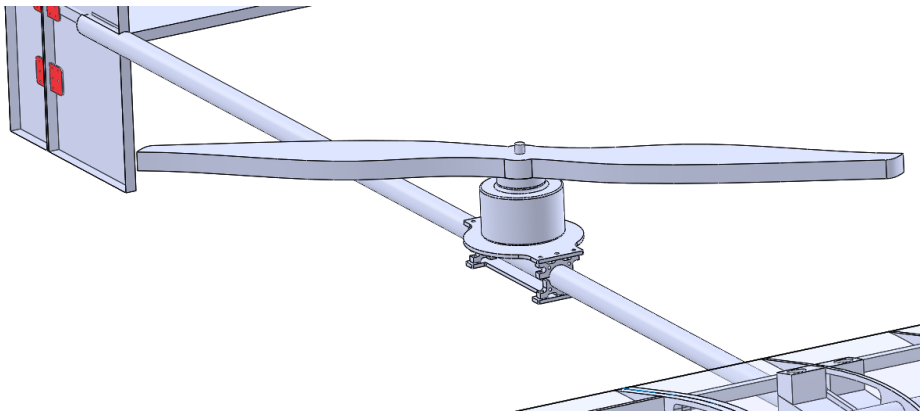

Figure 2.32: CAD of Original Rear Propulsion Mounting

Were the prop to become stuck sideways-on to the wind, the resulting drag increase would greatly diminish the maximum possible endurance of the aircraft. Therefore, empirical testing was performed to investigate this behavior. A propeller was mounted on a spindle, free to rotate. This setup was exposed to the wind outside of an automobile proceeding forward at the aircraft's sampling velocity of 52.49 ft/s. Unfortunately, the tendency of the propeller was to orient itself as feared, even at a range of different velocities above and below sampling. Though several options for addressing this behavior were considered, the only feasible route was to make use of the rear motor-prop combos as transitioning members along with their front-facing counterparts. In order to support this design change, the the tail had to be altered as the rear portions of the struts would impede transitioning of the rear propellers; a continuation of the existing body rearward to meet the tail was opted for as a replacement. Additionally, the change in tail mounting prompted a change in tail layout to bring the mass of the otherwise cantilevered vertical stabilizers inwards, lowering the need for structural reinforcement on the horizontal. Figures 2.33 and 2.34 illustrated the evolution in tail and rear motor mount layout:

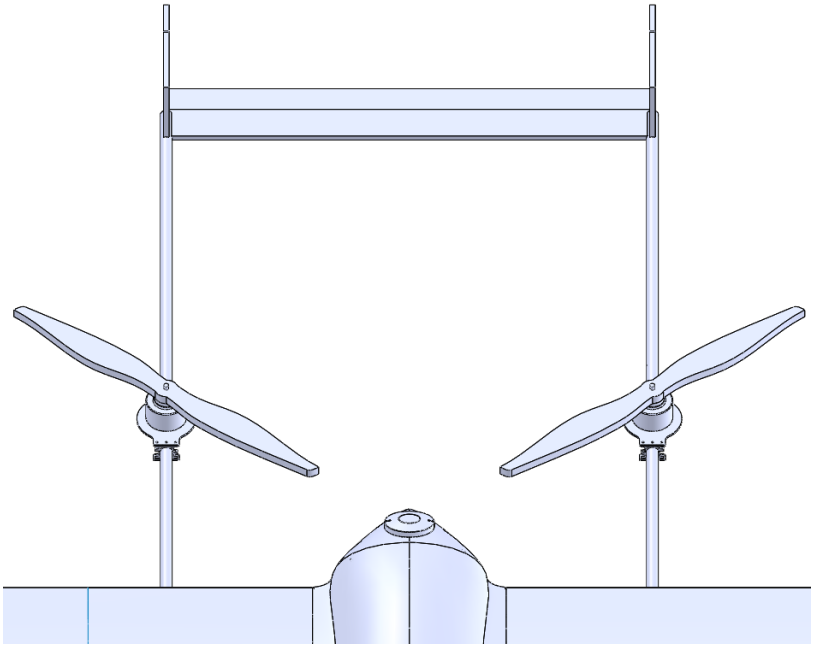

Figure 2.33: CAD of Original Tail Layout

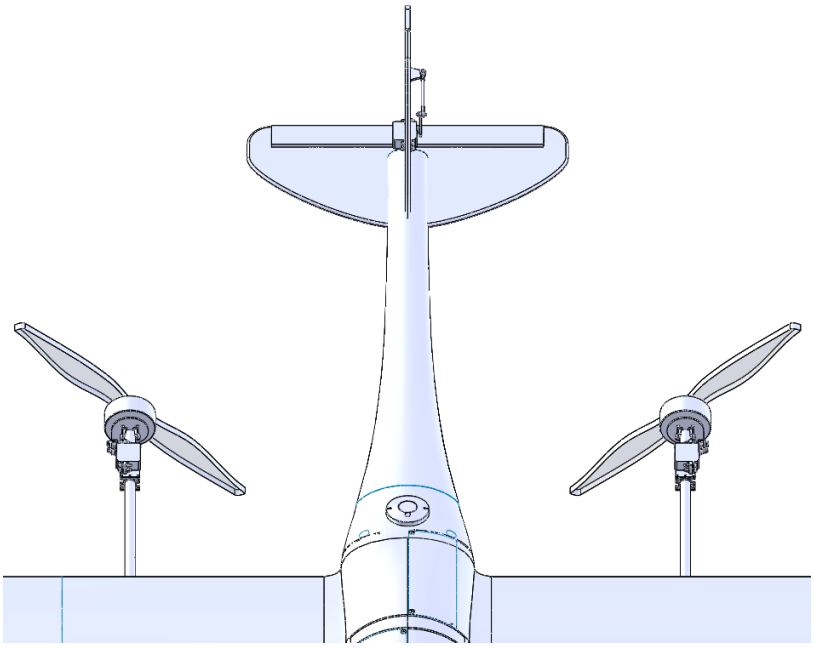

Figure 2.34: CAD of Final Tail Layout

where the tail volumes calculated for the H-tail were replicated in the subsequent conventional layout. The struts were shortened to allow for mounting of the rear motors on their

ends, which were mirrored from the front mounts in order to properly orient the resulting force vectors in VTOL and fixed wing flight. With the new tail design in place, the question of servo mounting was addressed as depicted in Figure 2.35:

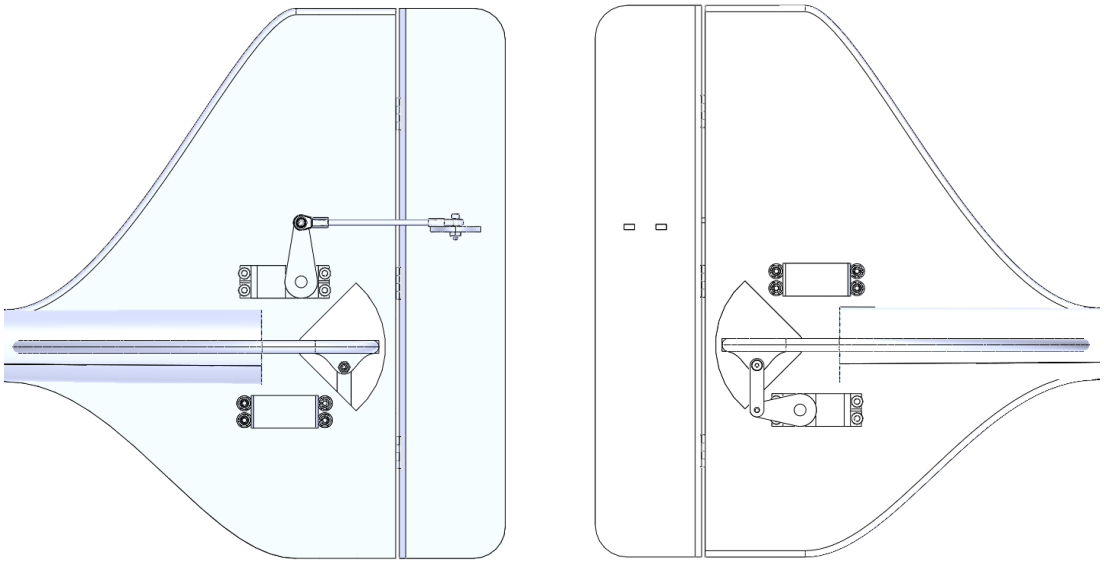

(a) (b)

Figure 2.35: Tail Servo Mounting and Linkages

where Figure 2.35(a) displayed the rudder servo mounting and linkage setup and (b) showed those of the elevator. In the rudder's case, a rod threaded into two spherical rod ends allowed for un-planar attachment between the servo horn and the rudder hardpoint. The elevator's linkage, being planar, was implemented similar to those of the ailerons. As a consequence of the tail's construction, the servos could be bolted directly to the vertical stabilizer via cutouts. A third cutout allowed for full rotation of the elevator through the vertical. The whole assembly would be bonded to the tail boom, fitting into slits cut into the end of the boom.

With the overall layout of the aircraft resolved, it only remained to package the aircraft's avionics components. In order to physically support the batteries, Pixhawk, modem, receiver, and airspeed sensor, 3D-printed "cribs" were conceived of which would feature cavities to house these components. In the case of the smaller components, mounting holes

were made use of for this purpose. The batteries, on the other hand, required some form of soft mounting. Velcro straps were one possible option, as was sketched in Figure 2.36 for the rear battery bank:

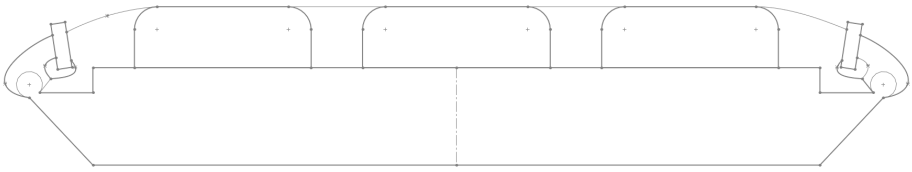

Figure 2.36: Sketch of Battery Crib Concept

Velcro straps with buckles at one end would loop through grooves in the crib back on themselves, securing the batteries in place during flight while allowing for their removal at mission's end. The implementation of this concept was successful, as evidenced by Figure 2.37:

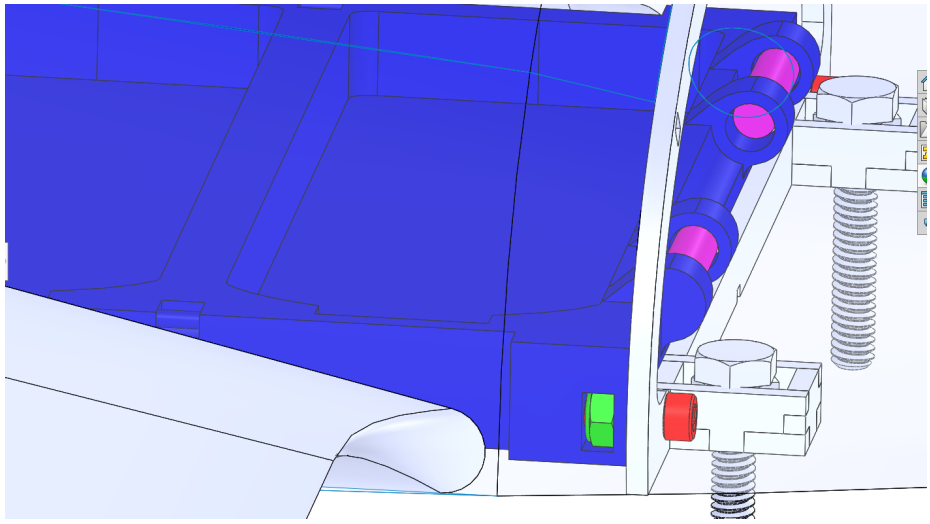

Figure 2.37: CAD of Battery Crib As-Packaged

where the crib itself was displayed in blue, the metal strap dowels in purple, the 8-32 crib mounting bolts in red, and the inset nylock nut in green. By insetting the nuts inside the crib, the whole assembly could be removed as one piece if desired. The dowels were to be bonded in place, providing additional stiffness and strength to support tightening of the

velcro straps. Also visible in Figure 2.37 were the rear two payload 1/4 - 20 mounting bolts, matched by two identical setups at the fore of the body.

Figure 2.38 featured all of the aircraft's centrally-located components in their final packaging:

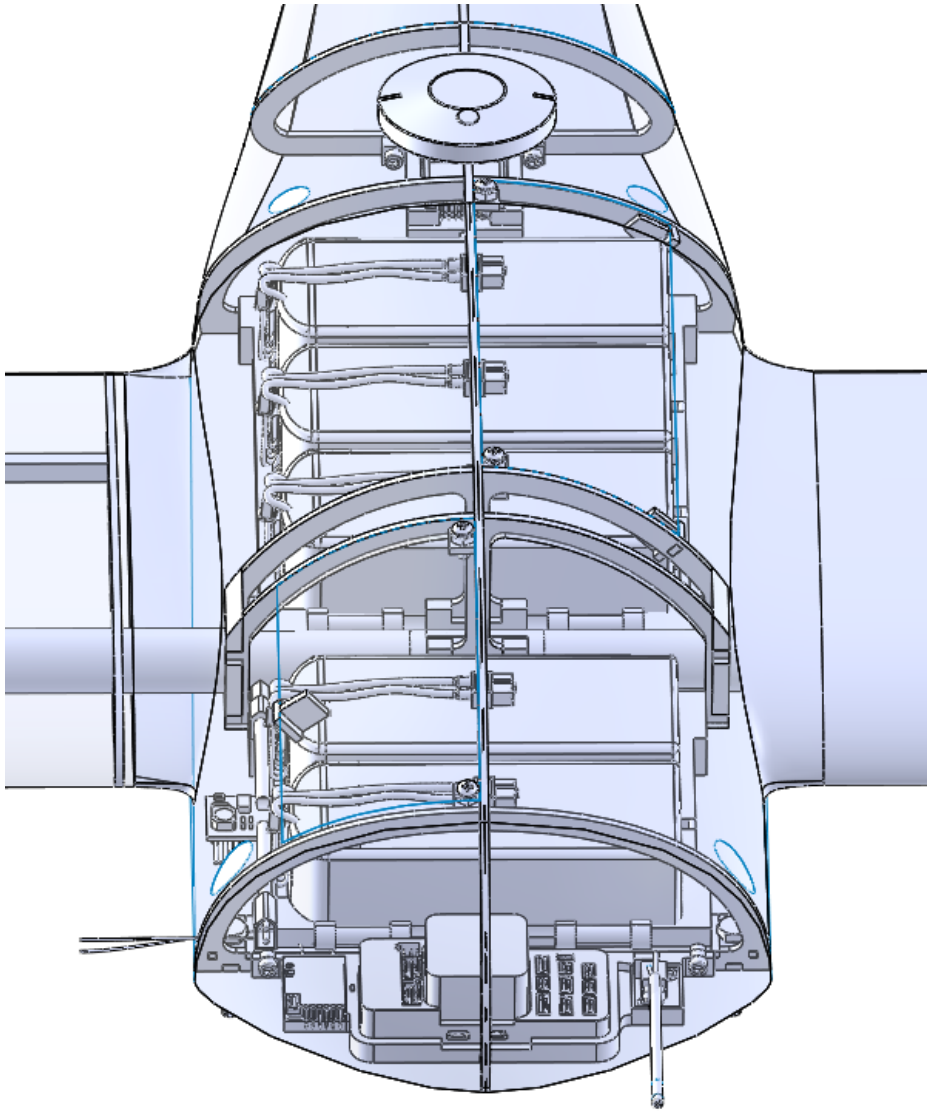

Figure 2.38: Detailed CAD of Body Packaging

depicting the fore and aft battery banks housed in their respective cribs around the wing spar and the Pixhawk, FrSky receiver, and airspeed sensor in their crib at the nose. Also

shown were the fore and aft body access panels, outlined in blue, including bolted hardpoints clustered around the top stringer and sliding softpoints opposite them on each hatch. The GPS module was placed externally towards the rear of the body, partially obscuring the RFD900+ modem in its own crib abut the rearmost visible bulkhead. With internal packaging implemented, the detailed design of the aircraft was complete. The final layout, including a placeholder payload pod, was displayed in Figure 2.39:

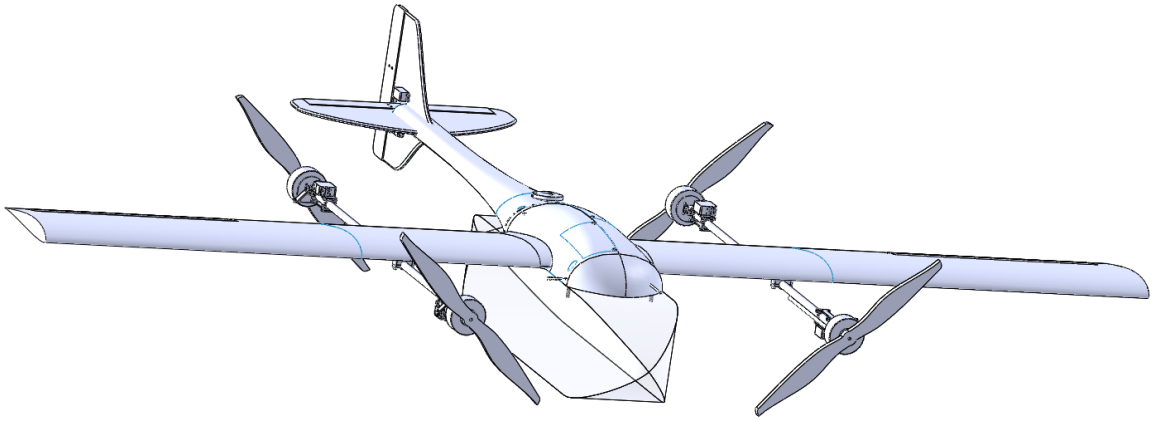

Figure 2.39: Overview CAD of Final Layout

Chapter 3

Manufacturing

3.1 Component Manufacturing

Component manufacturing encompassed the individual processes involved in the translation of every part's theoretical design into a real-world object. Barring commercial availability, every aspect of the aircraft's CAD would have to be laser cut, 3D-printed, CNC machined, laid up, or formed by some combination of these. The breakdown of manufactured components by method was listed by Table 3.1:

Table 3.1: Component-wise Manufacturing Methods

Laser Cutting	3D-Printing	CNC Machining	Laying Up
Ribs	Battery Cribs	*Motor Pedestals*	Outer Wing Skins
Bulkheads	Pixhawk Crib	*Strut Sleeves*	Fuselage Skins
Tail Surface Cores	RFD900+ Crib	*Tilt Servo Linkage Arms*	**Tail Surface Reinforcements**
Tilt Servo Platform Cores		*Strut Clamps*	**Tilt Servo Platform Reinforcements**
Stringers		*Front and Rear Strut Interfaces*	
Aileron Servo Linkage Arms		*Spar Inserts*	
Trailing Edge Wing Spars		Strut Mounting Blocks	
Aileron Spars		Outer Wing Top Mold	
Elevator Servo Linkage Arm		Outer Wing Bottom Mold	
Elevator Linkage Hardpoint		Top Fuselage Mold	
Rudder Linkage Hardpoint		Bottom Fuselage Mold	
Spar Interface Plates			
Access Panel Nut Shelves			
Access Panel Softpoints			

where the four bolded entries highlighted their mixed-method manufacturing and the italicized items denoted components machined by an external machine shop.

3.1.1 Laser Cutting

Laser cutting was the preferred manufacturing method for flat components due to its high degree of accuracy and speed. The method began by exporting a component's cross-section as a .dxf file in the CAD program. This type of file was readable by Adobe Illustrator,

allowing for arrangement of multiple files on a single window for cutting out of the same piece of material. The University of Oklahoma's Tom Love Innovation Hub had several Universal Laser Systems PLS6.75 laser cutters featuring 32" x 18" cutting beds. These laser cutters were only capable of cutting through certain materials; in the case of the given application, all parts manufactured using this method were either balsa, as in the case of the tail surface cores, or aircraft lite ply, being the case for the rest of the method's associated components listed in Table 3.1. In either case, the cut settings used corresponded to those for "General Medium Hardwood," with the thickness being set according to the given stock. Once all components for a given cut were arranged, the file was exported to "print" by the laser cutter; Figure 3.1 depicted some of the wing ribs and spars mid-cut:

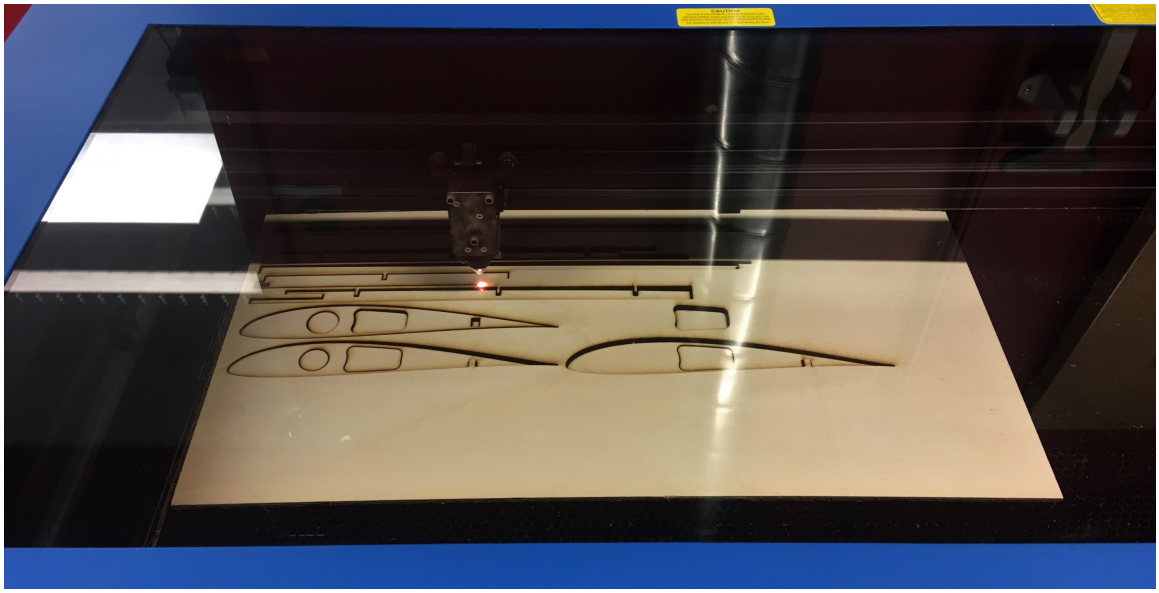

Figure 3.1: Image of Laser Cutting

Some of the longer components had to be separated into various segments in order to fit on the laser cutter's bed. These were notched in such a way as to allow for easy bonding using epoxy that was mixed with West Systems' 406 Colloidal Silica to add viscosity and bonding strength. Figure 3.2 depicted the assembled body skeleton and outer spars resting in the bottom body skin mold:

Figure 3.2: Image of Assembled Body Skeleton

Each skeleton was bonded together while resting in one of its respective molds to aid skin-skeleton agreement during assembly.

3.1.2 3D-Printing

3D-printing was the next most expedient manufacturing option. It allowed for the manufacturing of complex three dimensional geometries out of the realm of possibility for the laser cutters. However, it too was limited in terms of materials, in this case to various plastics such as PLA and ABS. Therefore, the components manufactured using the method could not be structural to the operation of the aircraft, making the cribs designed for holding the sundry avionics components in the body of the aircraft perfect candidates. The CAD geometry was exported as .stl files in order to be read by the CURA slicer used to prepare the

3D printer gcode, by discretizing the 3D geometry into numerous two dimensional slices in the vertical axis; Figure 3.3 came from the gcode preparation for the front battery crib and depicted one such slice:

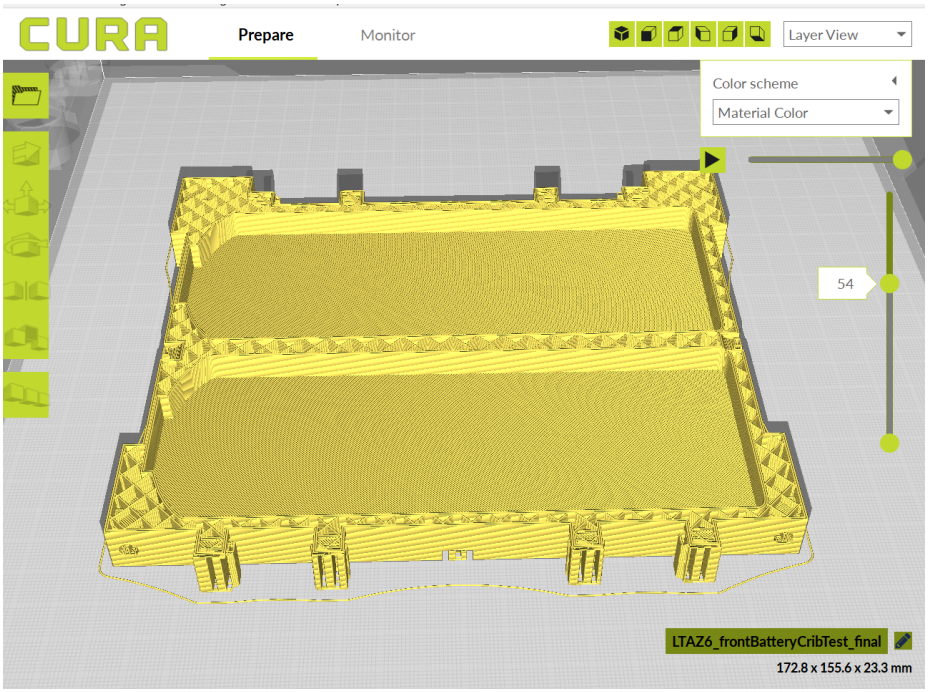

Figure 3.3: Mid-Print Slice of Front Battery Crib

The structure of the part was laid bare, revealing the small percentage of the part's internal volume taken up by material, a print value known as infill. The cribs were printed with a 20% infill to save weight while maintaining reasonable strength. Additionally visible in Figure 3.3 were the shells, those outer layers of solid material that provided much of the parts' stiffness; the cribs were printed with two of these layers. Additional print parameters such as temperature were set according to the CURA-recommended settings for the print material used: t-glase plastic. Multiple iterations of each part were printed to address packaging of each cribs' avionics components as well as of the cribs themselves. In the case of the battery cavities, some excess space was accorded to allow for swelling of the Li-Po batteries over time. Finally, the aluminum rods serving as the strap mounting points were bonded in place. The end result of the front battery crib was captured in Figure 3.4:

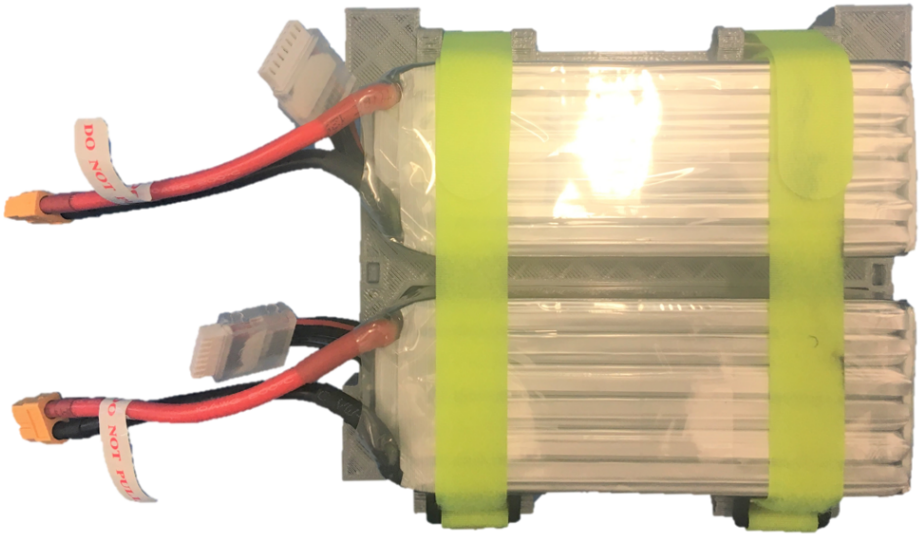

Figure 3.4: Image of Assembled Front Battery Crib

3.1.3 CNC Machining

Computer numerical controlled machining operated by way of precise removal of stock material until only the desired geometry remained and could encompass a slew of different individual machining methods such as milling, lathing, and drilling. As highlighted in Table 3.1, the two different categories of parts manufactured through this method consisted of the aluminum parts machined by an external machine shop and the molds routered in-house.

The complicated geometries of the CNC machined aluminum components necessitated their manufacture by The University of Oklahoma School of Aerospace and Mechanical Engineering machine shop. The shop was provided with part geometries in the form of .x_t parasolid files and paper part drawings highlighting important dimensions as in Figure 3.5:

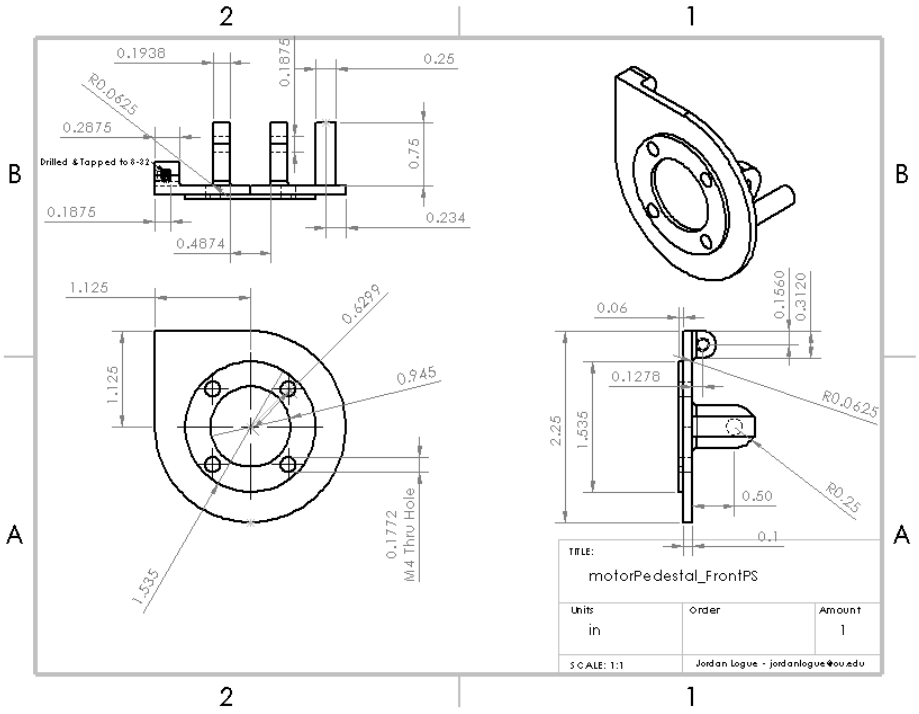

Figure 3.5: Front Right Motor Pedestal Drawing

as well as the 7075 aluminum stock to make them. The process for gcode preparation and machining was similar to that for the routered molds.

Routering was a subset of milling typically used for machining softer materials such as wood products and which typically utilized higher translational feeds and bit rotation speeds. The project utilized routering to manufacture the molds intended for shaping the aircraft's monocoque skins in a process known as direct female molding; "female" referred to molds shaped as the reverse of their associated part in the form of cavities as depicted in Figure 3.6:

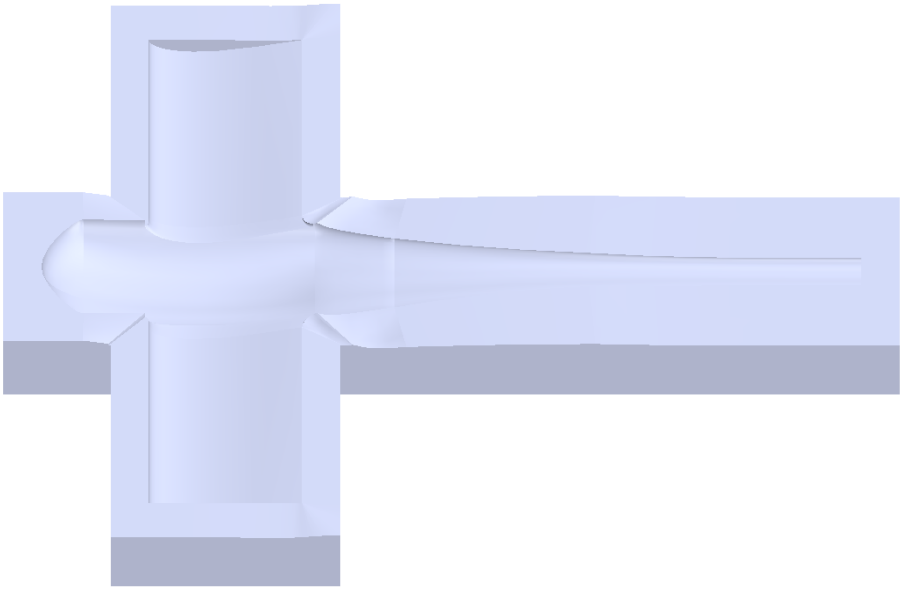

Figure 3.6: CAD of Top Fuselage Mold

whose shape was taken off of the top half of the aircraft's body and inner wings skin. The image also illustrated the necessity of splitting each portion of the monocoque skin into top and bottom halves in order to allow access for laying up each half and for bonding internal structures. Doing so also allowed for the implementation of flat, orthogonal flange surfaces at the parting line that would later aid in the layup process. The four original molds produced were the aforementioned top fuselage mold, the bottom fuselage mold, and the top and bottom molds of the outer wings which were designed for use by both sides to reduce machining time. In order to fit onto the router table available at The University of Oklahoma's Tom Love Innovation Hub, each fuselage mold was additionally split into front and rear halves. The material selected to form the molds was medium-density fiberboard (MDF) for its availability, machinability, and durability. The large boards were first cut into smaller sizes matching the outermost dimensions of the various molds before being bonded together using wood glue. These blocks of MDF formed the stock material used to create the router's toolpaths.

The HSMWorks plugin for Solidworks was used as the computer aided manufacturing (CAM) software to generate the CNC toolpaths. These toolpaths, in the form of gcode,

dictated the feeds and speeds of the router, the tool or tools involved, and the positioning of the tool at any given point; visually, this appeared as in Figure 3.7:

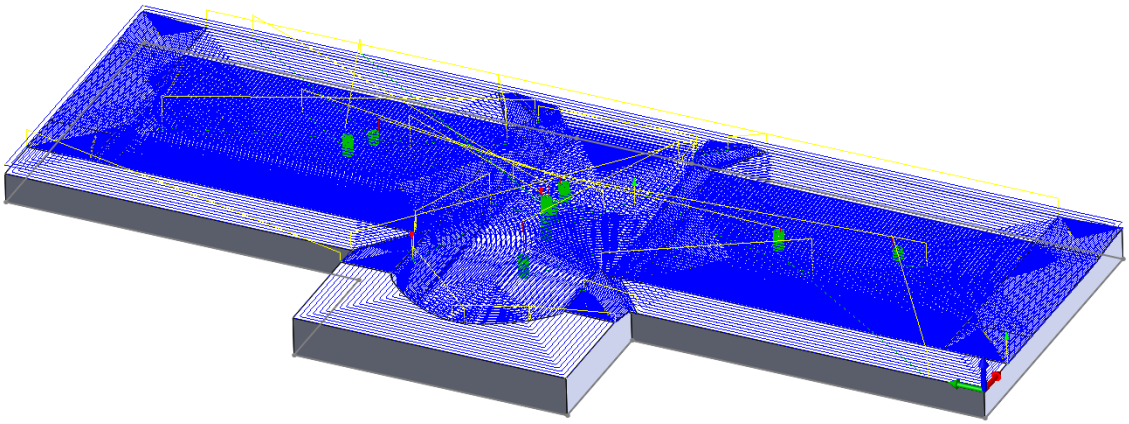

Figure 3.7: CAM of Front Half of Top Fuselage Mold

which depicted the toolpaths associated with the front half of the top fuselage mold, which itself was further split into its own top and bottom portions in order to fit within the router's range of vertical travel. The milling operations (ops) could broadly be categorized as those pertaining to voluming and surfacing. Voluming ops were concerned with removing large volumes of material from the stock and were accomplished using a $1/2$" flat end mill translating at 80 $^{in}/_{min}$ with a $1/2$" step depth and a $1/4$" step over. Surfacing ops finished the parts by cutting just enough into the part surfaces to produce the desired faces. As surfacing was performed using a $1/2$" ball mill, curved faces such as those on the wings could be captured much better than by the flat end mill; these ops were also run at 80 $^{in}/_{min}$ but with a 0.002" scallop height. By defining the stock dimensions of the CAM according to those of the real-world stock, the router could cut out the mold as desired:

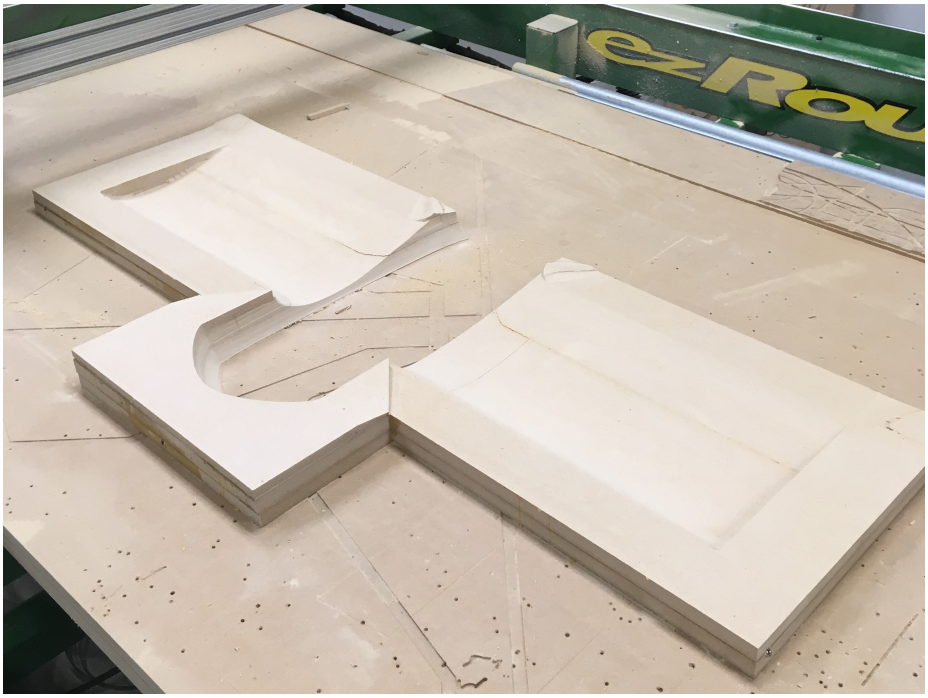

Figure 3.8: Topmost Portion of Top Fuselage Mold Front Half

The final step in mold assembly simply involved bonding the disparate pieces of the fuselage molds together, again with wood glue. The molds, though together, required additional surfacing in order to be ready for use. Sanding was of paramount importance to remove any remaining scallops in the mold surface from the surfacing op as well as to smooth out glue seams. All sanding was performed with a rigid sanding block if possible or at least a flexible foam pad, useful for confined regions. 180 grit sandpaper was used to remove the scallops, and 80 or 60 grit was utilized for the hard glue seams. Great care was taken to avoid sanding away the true mold surface. Once the surface was satisfactorily smooth, the next step involved coating the mold surface with graphite-infused epoxy. The epoxy system used for this was West Systems' 105 resin with either the 209 extra slow hardener or the 205 fast hardener depending on the requirements for work time. Graphite powder was added to the epoxy to harden the surface for use in lay ups and to provide color contrast between regions of sanded and unsanded mold surface; West Systems' 423 graphite powder was used for this purpose. The addition of this coating also introduced surface imperfections

which then had to be sanded out similarly to the scallops and seams. For divots or cavities that formed during this process, automotive body filler was applied to fill in these voids; its similar hardness to the surrounding MDF meant that neither would be sanded through prematurely. Three rounds of coating and sanding were typically performed for each mold, with spot-work as need. The last coating applied was a thin layer of Rust-Oleum 2-in-1 Filler & Sandable Primer spray to fill in the smallest divots:

Figure 3.9: Image of Top Fuselage Mold Coated with Primer

with this layer being otherwise sanded back off. Whereas previous coats of epoxy were sanded from 80 or 180 grit down to 500 grit, this last layer was incrementally sanded down to 2000 grit to achieve the desired smoothness for the skins.

3.1.4 Composite Lay Up

Laying up referred to the process of permanently molding composite materials to a desired shape. The specific technique used to produce the skins and tail surfaces was the wet layup method which involved manual infusion of the composite matrix as opposed to vacuum-assisted or prefabricated means. This technique was chosen due to in-house familiarity and lack of equipment required to perform the alternatives.

The method began with numerous preparatory measures. The first step involved placement of double sided tacky tape on the peripheries of the mold's flanges. This tape would eventually provide the airtight seal maintaining the necessary vacuum pressure. The mold surface itself then had to be sealed and protected against bonding. Sealing was achieved via application of wax to form a thin, durable, watertight layer between the actual mold surface and the part. Rexco Partall Paste #2 was used for this purpose. Microfiber cloths were used to apply the wax and then to buff off excess for a total of three layers. The second preparatory coating consisted of three layers of Rexco PARTALL FILM #10, a liquid polyvinyl alcohol-based mold release. This coating prevented bonding between the epoxy used in the lay up and the mold surface. These were applied in thin coats using foam brushes, with each layer requiring approximately fifteen minutes to dry.

With the mold itself ready, the actual constituent materials had to be procured. A schedule referred to the specification of a layup's number of individual fiber layers, their rotational orientations, and any core. The schedule for the given aircraft's main skins was four total layers of 2 oz fiberglass fabric with a $1/8$"-thick Divinycell foam core in the middle. The orientation of the outermost and innermost fiberglass layers was 45 degrees to the forward flight direction of the aircraft in fixed wing flight. This orientation provided better

resistance to shear loads on the skin, while the two inner layers' 0 degree orientation was more attuned to address axial loads. Mixing the layers' orientation as described lead to more isotropic behavior of the skins versus utilization of a single orientation. The individual fiberglass pieces were cut to fit within the boundary provided by the tacky tape. In order to fit the foam core into the empty skin found in between the skeleton components, each assembled skeleton had previously been used as a template to trace foam core patterns onto their respective molds, as in the case of the top outer wing mold shown in Figure 3.10:

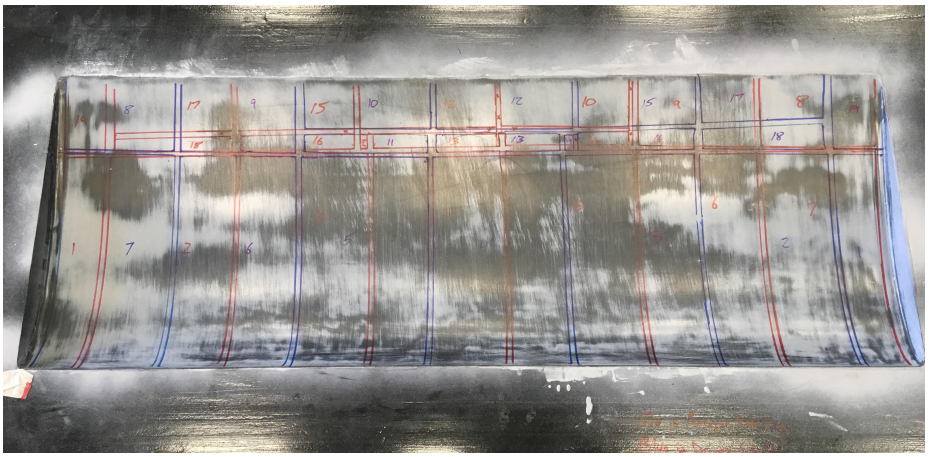

Figure 3.10: Image of Top Outer Wing Mold with Foam Core Outlines

These outlines were then used to form templates out of flexible plastic. Finally, the foam core pieces themselves were delineated and cut out according to these templates. Their edges were beveled so as to prevent voids between the top and bottom fiberglass layers at the foam core peripheries. The epoxy utilized for the composite matrix was also the West Systems resin system. The extra slow hardener was typically used to provide maximum working time. The amount used was geared towards a 1:1 weight ratio of of epoxy to fiberglass given the impreciseness of the wet layup technique employed; in practice, the amount of epoxy used was oftentimes greater than that dictated by this ratio in order to ensure full coverage of the fiberglass, though much of this excess would ultimately be removed from the part during the curing process.

The layup itself began with an initial coat of epoxy applied to all areas on the mold surface where fiberglass would come to rest. The outermost layer of fiberglass was then carefully laid down upon the mold surface, and any remaining dry spots were moisturized using epoxy-soaked foam brushes in a dabbing action. The next layer was then laid down, and the process was repeated. The foam core pieces were pre-coated with epoxy in an effort to prevent their moving around after placement. In practice, the foam pieces were apt to pop off of the mold surface and generally not remain in the spaces defined for them. Placement of the final two layers of fiberglass aided in restraining the core pieces.

Once all constituent composite materials were in place in the mold, the next layup layer to come was the perforated breather bagging. This bagging did not stick to the epoxy and thereby provided protection against unwanted bonding for the rest of the layup layers. Its perforations provided numerous small exit points for excess epoxy to escape into the next layup layer, being the breather cloth. This voluminous material was designed to intentionally soak up excess epoxy to save weight and to aid in debulking of the skins. The outer vacuum bagging served as the actual barrier supporting the layup vacuum pressure and was cut extra baggy in order to fit into the depths of the mold. The vacuum's path through the outer bagging was provided by two-part pucks which sealed around holes poked through the vacuum bagging, with the outer half of each puck featuring a connector for the vacuum pump's hose. With the inner pucks in place under the vacuum bagging, the bagging's peripheries were carefully pressed against the two sided tape at the mold's edges to seal in the mold, and the outer puck halves were afixed around their respective inners to begin the vacuuming process. Leaks were addressed with additional tacky tape until the mold held sufficient pressure to hold the uncured skin against the mold surface as was captured in Figure 3.11 for the top fuselage skin:

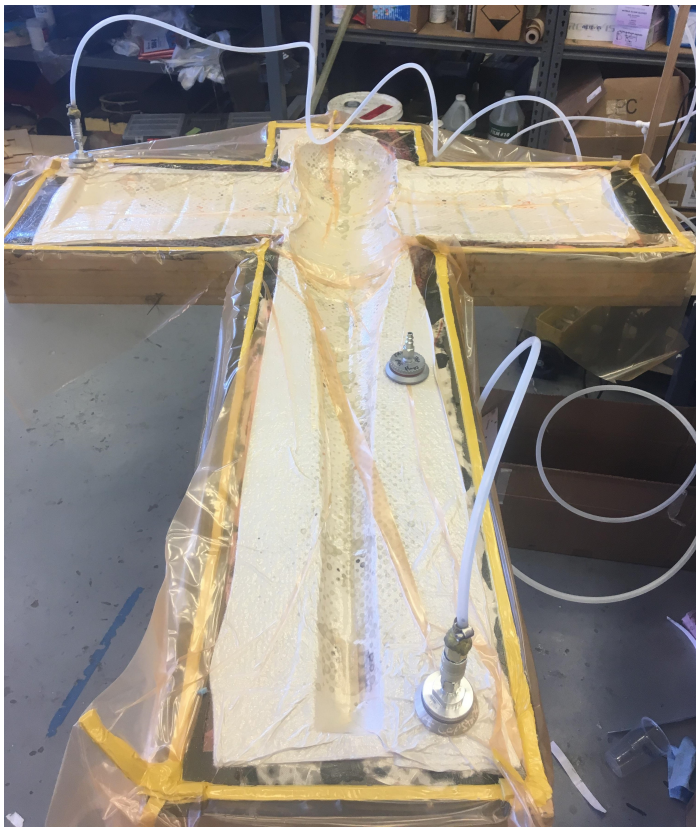

Figure 3.11: Image of Top Fuselage Skin Under Vacuum

The yellow striped material at the mold's edges shown in Figure 3.11 was the tacky tape, the white material on the mold's interior was the breather cloth, and the red plastic sheeting was the vacuum bagging. The amount of time spent under vacuum depended upon the type of hardener used, with fast cure layups being left for at least eight hours and extra slow cure layups for at least twenty four hours. The part produced by the layup shown in Figure 3.11 was shown in Figure 3.12:

Figure 3.12: Image of Cured Top Fuselage Skin

The excess fiberglass flanges at the skin's perimeter in Figure 3.12 were intentionally left for use as clamping surfaces during assembly.

3.2 Assembly

With all individual components manufactured as described in Section 3.1, the second stage of manufacturing shifted focus to the assembly of these disparate parts into a cohesive system. This included implementing previously designed packaging techniques as well as addressing unexpected fitment issues as they arose.

3.2.1 Structural Assembly

Despite bonding the lite ply skeletons together while resting inside their respective molds, their interface agreement with the skins was often poor. This was due in large part to the difficulty in restraining foam core pieces during their layup as described in Subsection 3.1.4. In places where core impinged upon the skeleton's interface, a Dremel tool with an abrasive attachment was utilized to remove this extraneous material. Though this disrupted the closed cell structure of the skin, epoxy added during sandwiching of the skins and their skeletons would eventually replace this lost material. Prior to their bonding with skins, the skeletons were also used as templates to mark the outlines of the ailerons and

access hatches, similarly to how they were used to outline the foam core pieces. Razor blades were utilized for trimming these out in order to reduce the resulting panel gap. For bonding, colloidal silica-infused epoxy was again used to attach each skin to its respective bottom skin; the bottom skins were chosen rather than the tops to allow for packaging of the avionics components and their cribs. In the case of the cribs, additional iterations were 3D-printed to address imperfect fitment with the body skeleton.

Aluminum-carbon fiber connections were found at the wing spars and spar inserts, at the struts and strut sleeves, and at the struts and strut clamps. Corrosion of such connections was known to occur over time as a result of galvanic coupling of the carbon in the tubes with the aluminum alloy (Banis et al. 2015). To prevent this reaction from occurring, an additive known as a bond line controller was used in the epoxy at these connections. The 0.014" diameter beads comprising this controller would serve as a physical barrier between the aluminum and carbon and thereby prevent occurrence of the reaction; the aluminum parts were designed to accommodate this barrier. The bonding agent used for these connections was Loctite 0151 Hysol epoxy.

Though the tail surfaces were laid up similarly to the rest of the composite skins, their assembly was different in that a solid balsa core had replaced the discrete skeleton of the monocoque body. As this core had been previously laser cut to include cavities where trimming would be required, such as tail servo mounting and assembly slots, assembly preparation was as simple as using a razor blade to trim off the core-less portions of fiberglass. In the case of the control surfaces, .dxf templates were printed off from CAD and used as guides for separating these pieces. Some gap was included between the front and rear faces of the control surfaces and main tail surfaces, respectively, to provide room for the nylon hinges. Each hinge half was inserted into slits cut into the corresponding tail surface for this purpose and bonded in place with cyanoacrylate adhesive. Each hinge was completed with the insertion of a brass pin into both hinge halves to serve as the actual

pivot. To fit onto the tail boom of the main body, slots had to be cut in the skin. A guide, shown by Figure 3.13:

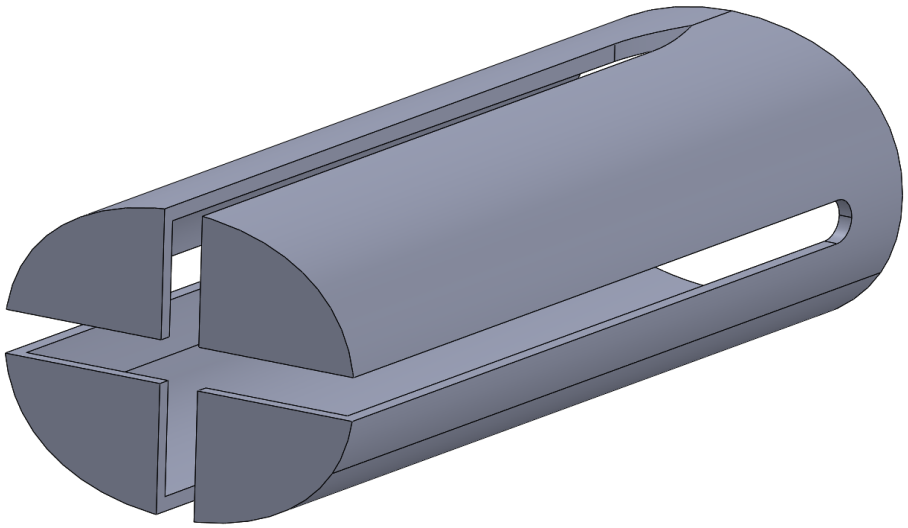

Figure 3.13: CAD of Tail Slot Guide

was designed to fit over the end of the tail boom and define the locations of these slots, which were cut out of the skins using the Dremel with a disk cutter attachment. The tail was fit into the slots formed in the bottom skin and skeleton and bonded into place using the epoxy-silica mixture.

Following completion of the avionics assembly described in Subsection 3.2.2, the aircraft was ready for sandwiching. This process involved sealing the plane's seams where the top and bottom skins met as well as bonding the skeleton to the top skin. The skeleton-skin bonding was achieved by placing a layer of very viscous silica-epoxy on the top faces of every part of the skeleton. Filling the skin seams made use of West Systems' 407 Low-Density Fairing Filler epoxy additive. Addition of the filler produced seams that were easier to sand than if the colloidal silica had been used. Figure 3.14 depicted what this setup looked like directly proceeding sandwiching of one of the outer wings:

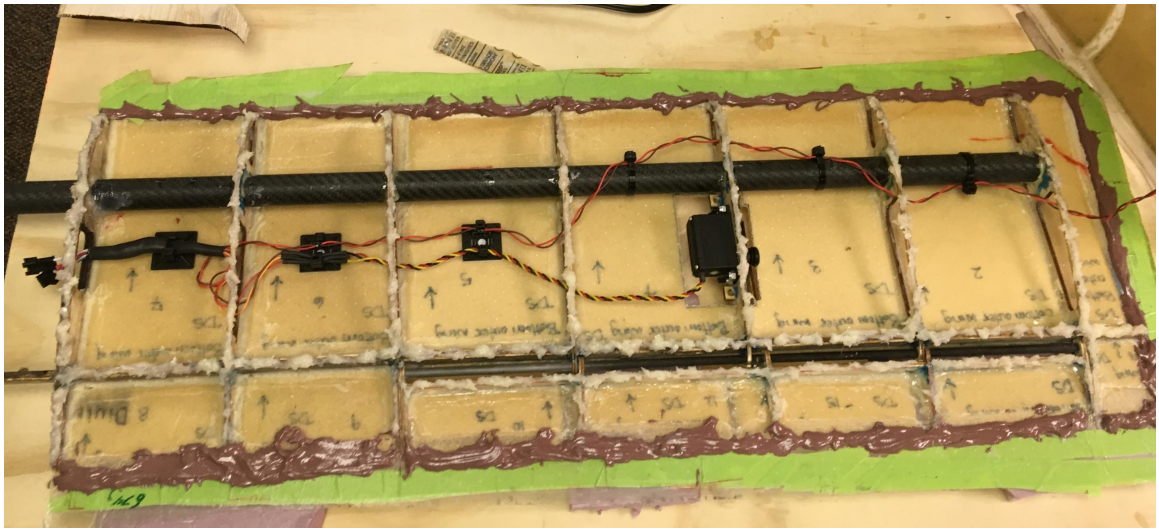

Figure 3.14: Image of Outer Wing Prior to Sandwiching

where the brown substance was the fairing filler-epoxy mixture and the off-white material atop the skeletal components was the silica-epoxy. The top skin could then be put in place. To ensure that as much of the top skin contacted the skeleton as possible, weights were placed across its surface. Clamps were used to do the same for the seams, as Figure 3.15 showed for an outer wing:

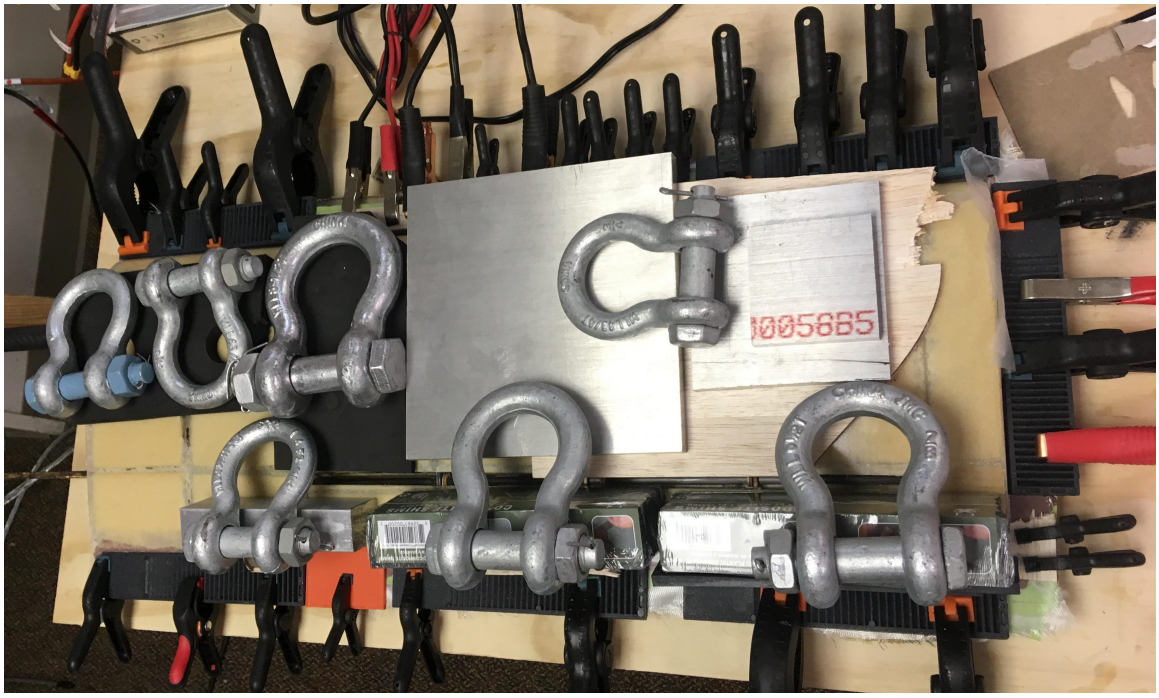

Figure 3.15: Image of Outer Wing Following Sandwiching

The shims visible between the clamps and the fiberglass flange helped to distribute the clamping force out across a wider area. As extra slow hardener had to be used in the bonding to provide maximum work time, the parts were left for over twenty four hours to cure. The flanges were then able to be trimmed off using the Dremel and cutting disk attachment. Care was taken to protect the part surfaces immediately abut the flanges with painter's tape. This tape provided a clear cut line while also protecting against unwanted abrasion. Any remaining flange was sanded down using 150 grit sandpaper and a long sanding block to preserve the straightness of the seams, the result of which was depicted in Figure 3.16 for one of the outer wings:

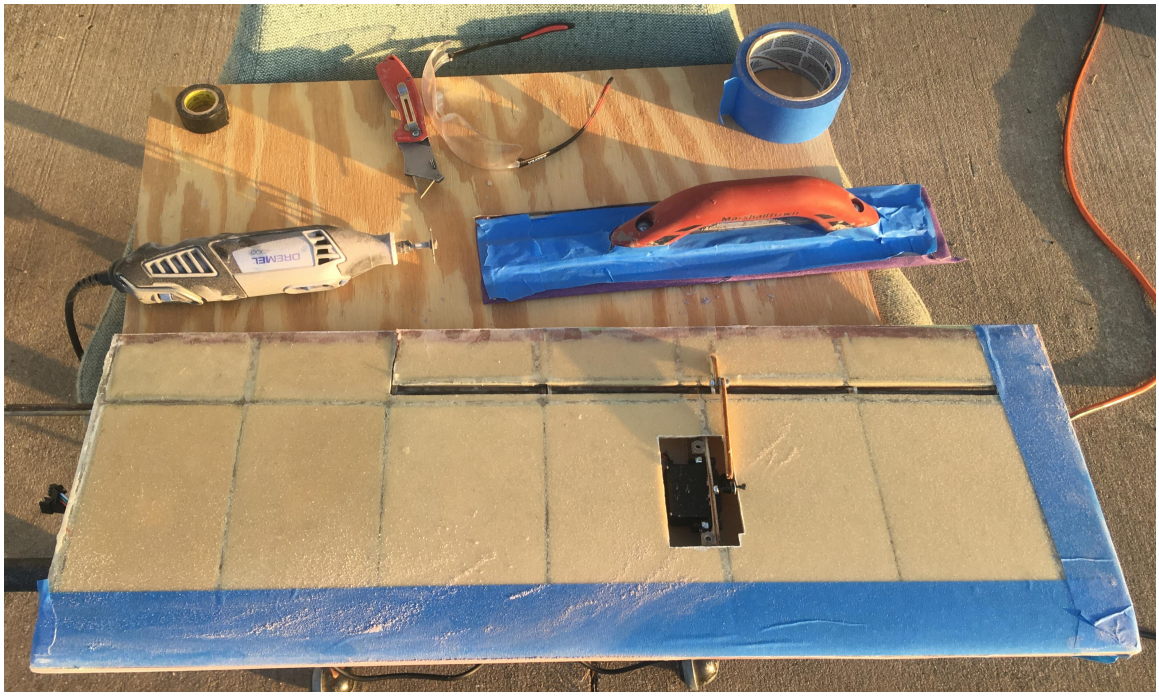

Figure 3.16: Image of Outer Wing Following Flange Trimming

The seams, though flush, were not as structurally sound as they would need to be; reinforcement was provided through the addition of fiberglass strips, cut to a 45 degree orientation to resist shear:

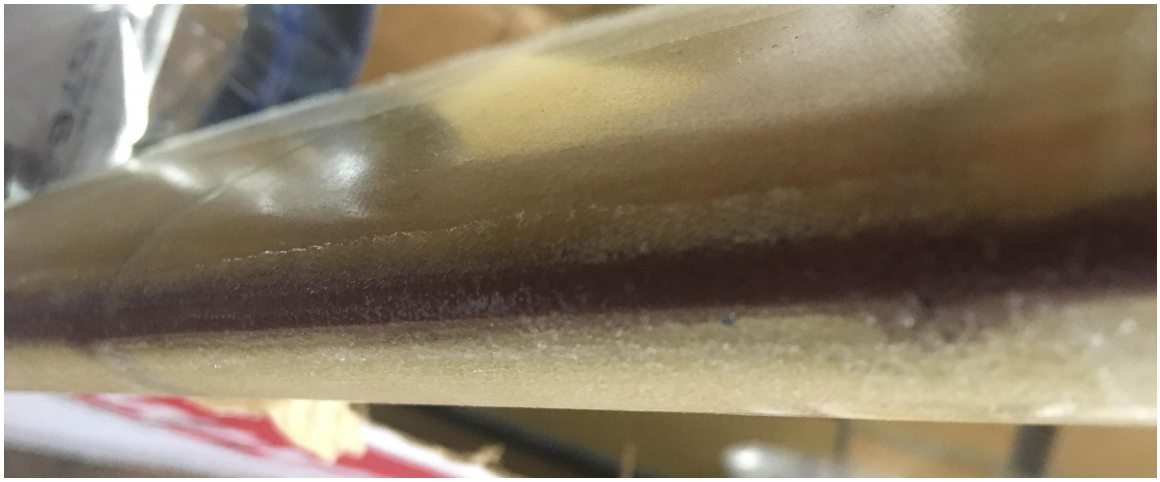

Figure 3.17: Image of Seam Reinforcement

The finishing step of each skin involved light sanding of this reinforcement to smooth the transition between it and the rest of the skin surface to avoid prematurely tripping the boundary layer flowing over it during forward flight.

This process of sandwiching, trimming, sanding, and reinforcement was followed for both outer wings and the fuselage, whose post-sanding state was depicted in Figure 3.18:

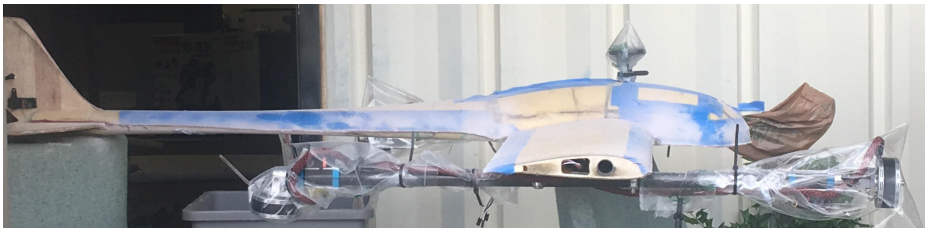

Figure 3.18: Image of Sandwiched Fuselage

where the plastic wrapping was put in place to protect the sensitive electronic components from sanding dust. Other than minor fitment issues, the plane was then physically complete, as shown in Figure 3.19:

Figure 3.19: Image of Completed Aircraft

3.2.2 Avionics Assembly

The assembly of the avionics system involved running wires in between and the securing of the individual components. This task was complicated by the fact that the avionics system had not been fleshed out in full detail at the time of final design, resulting in the exclusion of the motor ESCs, the GPS module, the BECs, the main power rail, the current sensors

and their hub, and the Power Cube. Luckily, the open nature of the body's monocoque construction provided ample space for mounting these additional components. Figure 3.20 showed the layout of the avionics system following real-world packaging, where Figure 2.20 provided the wiring diagram for the system:

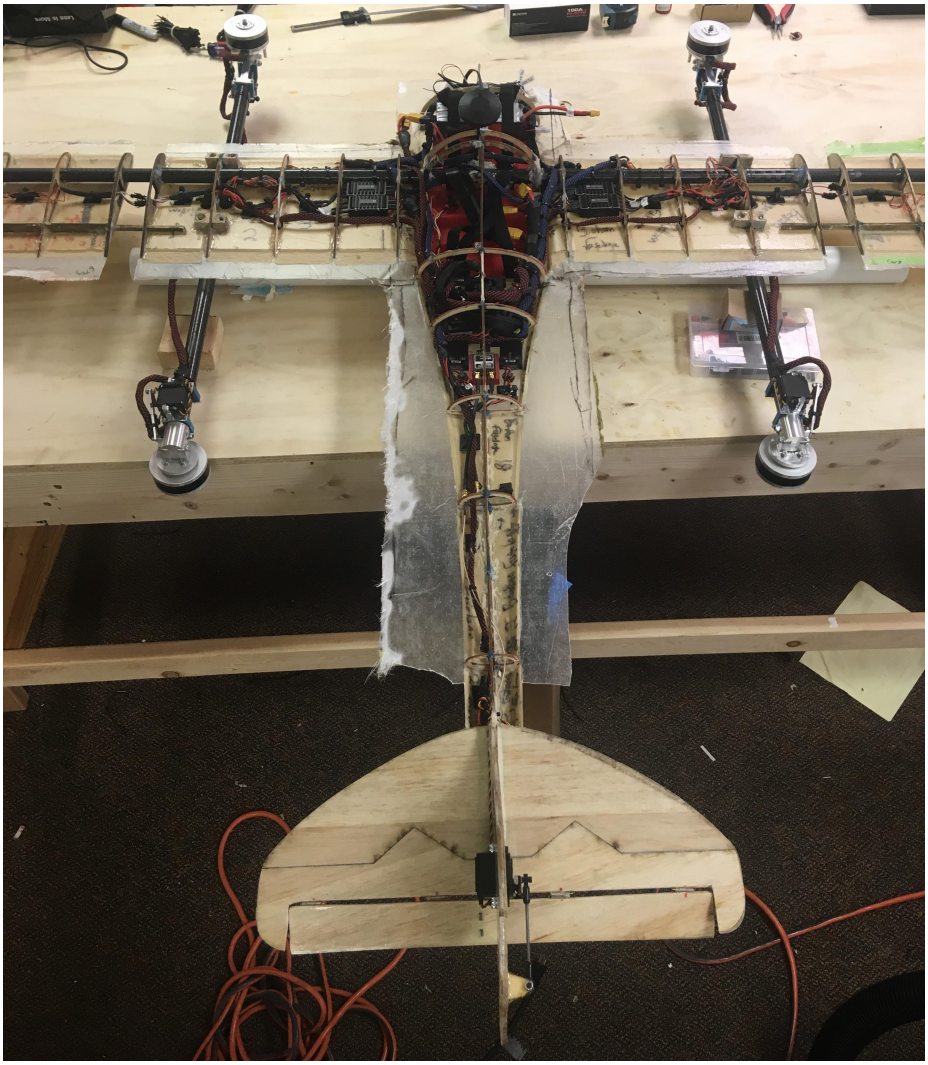

Figure 3.20: Image of Packaged Avionics System

The two banks of batteries were packaged as designed, with XT90 anti-spark connectors protecting against deterioration of the their individual wiring interfaces by preventing sparking at plug-in. The current sensors were small enough to be placed in-line to the thick wires coming off of the battery banks. However, the sensor hub and power supply board

were placed immediately aft of the rear battery bank, displacing the RFD900+ modem rearward to the next skeletal cubby. This modem's antennas were arranged with one being bonded to the vertical tail and the other inside the tail boom, just aft of the modem itself. The slew of signal and ground wires leading to the Pixhawk displaced the FrSky receiver to immediately above its original location at the flight controller's side, being bonded to one of the bulkheads. The airspeed sensor remained in place; its pitot tube required the drilling of a hole through the nose of the top skin, which was complemented by a flexible grommet in order to provide a more watertight seal. An external I^2C bus was bonded to the interior of the top skin at the nose; this easy-to-access bus would provide the ability to switch out multiple components for communication with the flight controller. A mini-USB extension cable likewise provided access to the Pixhawk for the ground station computer. The GPS module was mounted between the battery banks to a thin carbon fiber boom which raised it above possible electromagnetic interference from the batteries' DC. The ESCs were placed just inside the inner wings, in between the battery banks and the motors. The power, ground, and signal cables to each of the tilt motor and servo pairs were sheathed for protection and afixed along their respective struts to the aircraft's interior through grommets in the undersides of the wings. Those wires extending out to the ailerons and wing LEDs in the outer wings were provided with quick-disconnect connectors at the wing break to aid in dis/re-assembling of the plane for transportation.

Chapter 4

Testing

Testing of certain aspects of the aircraft was required to ensure the platform's successful operation. Potentially fatal flaws were addressed prior to their permanent implementation, and the reality of the platform's construction had to be reconciled with performance expectations. In any case, testing served as the final say regarding the operation of some of the aircraft's most essential systems.

4.1 Propulsion Testing

The ability to take off and land vertically as well as to transition to and from fixed wing flight was a core requirement of the UAS. Accordingly, special attention was ascribed in testing the real-world performance of the propulsion system prior to on-plane implementation to ensure successful operation. This testing was performed on a single tilt-motor assembly which was securely cantilevered off of a tabletop so as to allow complete freedom of motion as shown in Figure 4.1:

Figure 4.1: Image of Initial Propulsion Testing

A real-world understanding of the structural stability of the mounts and struts as they supported the rotation of the motor and propeller at incrementally higher throttle levels while static or transitioning was desired. The ability of the tilt servo to rotate the outboard assembly at the different throttle positions and the stiffness and responsiveness of the tilt linkage assembly were to be investigated as well.

The first test was run as follows: each throttle position was run while the propulsion system was held static in its fixed wing configuration. In the event of a successful test, the transitioning ability of the assembly would then be investigated at the same throttle percentage. Success of the transitioning then prompted a repeat at the next highest throttle percentage, with the increment being 25%. The tests ran well until the investigation of transitioning ability at 50% throttle. In this test, the assembly began violently oscillating about the strut axis upon reaching a transition angle of approximately 45 degrees to the horizontal. The test was repeated at 50% throttle while holding the motor-prop assembly static in its VTOL configuration to investigate whether the oscillations were specifically due to transitioning or due to orientation of the propeller's force vector beyond the horizontal; as the assembly exhibited the same violent behavior, the latter of these two possibilities was confirmed. Additionally, the motor and propeller were spun all the way up to full throttle while being held static in the fixed wing configuration, thereby confirming the torsional nature of the problem.

It was postulated that the system's behavior was due to an alignment of the assembly's first natural frequency in the torsional direction with the vibrations induced by the motor and propeller's rotation and force, thereby leading to the exponential increase in amplitude of the assembly's angular displacement. In order to raise the assembly's natural frequency beyond the range of those vibrations to be expected from normal operation of the propulsion system, one of two options presented themselves: increasing the assembly's rotational inertia or increasing its torsional stiffness. The former option was quickly discarded as a possibility due to fears over adding additional cantilevered weight to a system

with demonstrated structural inadequacies; therefore, a solution involving an increase to system torsional stiffness was pursued.

The equation for angular displacement of a circular tube was given by Equation 4.1a (Gramoll 2019):

$$\theta = \frac{T(L)}{G(J)} \tag{4.1a}$$

$$J = \frac{\pi}{2}(r_o^4 - r_i^4) \tag{4.1b}$$

where J represented the cross section's polar moment of inertia, being defined by Equation 4.1b in which r_o^4 stood for the outer radius of the tube and r_i^4 the inner radius. From a geometrical standpoint, the amplitude of the propulsion assembly's angular oscillations could be reduced given equal torque inputs (T) and moment arms (L) by increasing the strut's shear modulus (G) through material selection or the strut's polar moment of inertia through an increase to tube diameter, wall thickness, or both. In the end, a 1" outer-diameter tube of $1/8$" wall thickness, representing an increase in both dimensions versus the original strut's 0.628" diameter and 0.064" wall thickness, was selected. Combined, these improvements would result in theoretical angular displacements 14% of those seen in the original strut, solely as a consequence of geometry. Additionally, superior material properties were selected for the new 1" tube including 45 degree plies within its layup schedule to better resist torsional loads. Going up in tube size would necessitate the re-manufacture of the motor pedestals, tube sleeves, strut clamps, and strut interfaces in order to accommodate the increase in diameter. However, doing so would allow the redress of additional issues encountered in the existing mounts such as pedestal-sleeve compliance, possible pedestal bending during transition, and failure of the fiberglass-reinforced lite ply servo platforms:

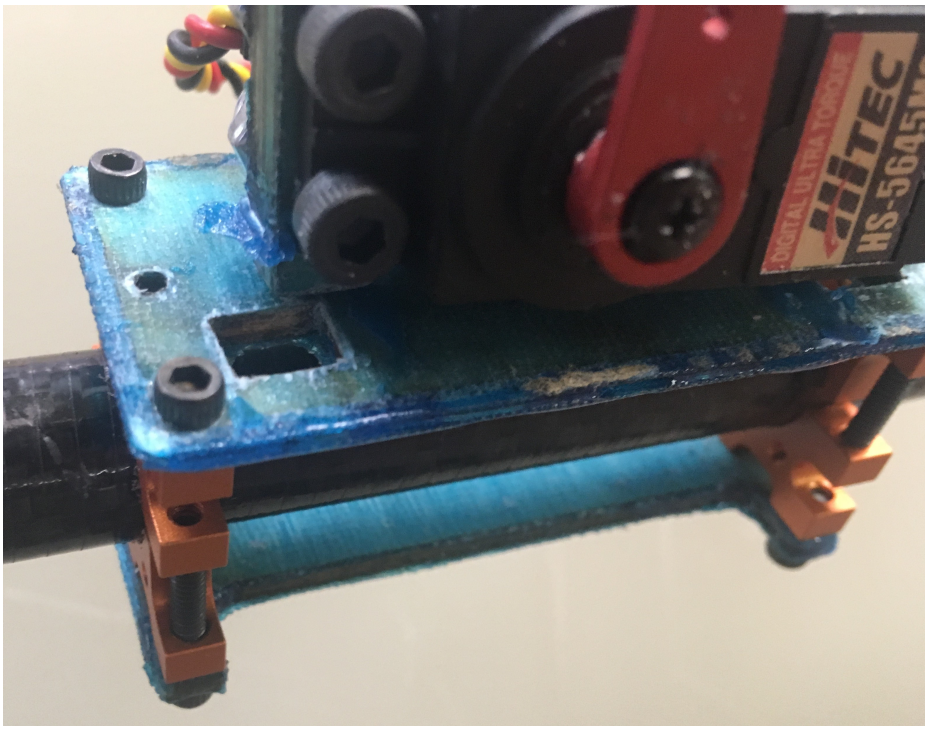

Figure 4.2: Image of Servo Platform Failure

to be remade in the same geometry but out of 7075 aluminum as with the other metal components of the propulsion system. The most significant physical changes to design came for the motor pedestals, illustrated for the original design in Figure 4.3a and the improved design in 4.3b:

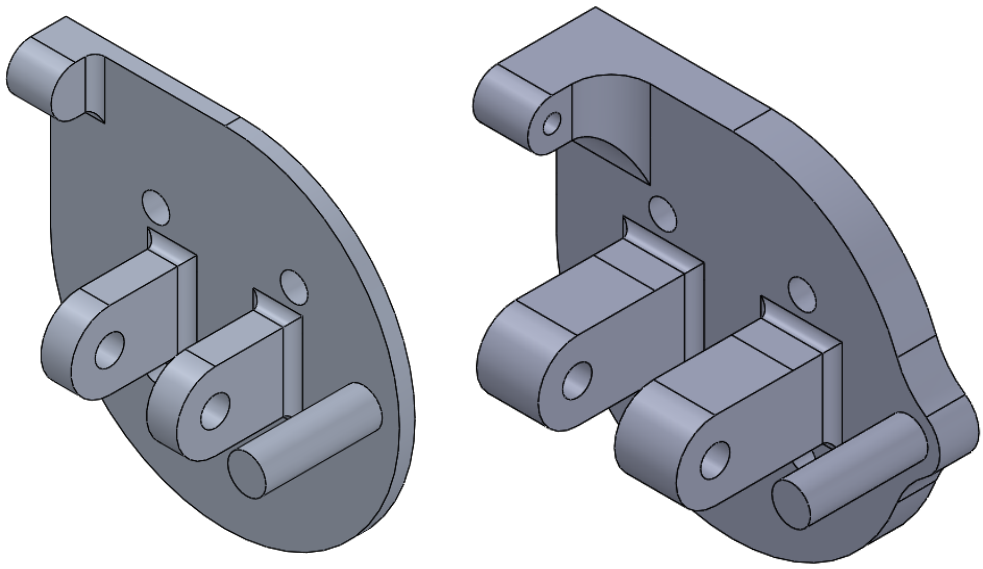

(a) (b)

Figure 4.3: CAD of Original and Improved Motor Pedestals

The platform itself was thickened to stiffen it against the possible bending observed in the initial propulsion testing. The tabs were thickened to combat compliance around the hinge bolt once mounted. These were also lengthened to accommodate the larger strut sleeve. The linkage's hardpoint fillet was increased to spread the force due to transition out over a greater area, and its distance from the pedestal's front face was increased in order to maintain its location relative to the hinge to preserve linkage geometry.

In order to simulate the effects of at least the strut's geometric improvements on system stiffness, the initial testing was redone using the same assembly mounted with as small a moment arm as possible. Though an L of 1.905" would have been required to achieve the analogous stiffness, 4.5" from the tabletop was the closest the motor-prop assembly could be mounted because of the servo platform, thereby achieving only 42% of the effect on θ predicted by the increase in tube size. Regardless, static testing at 45 degrees and 90 degrees up to and including full throttle exhibited none of the problematic behavior previously observed, lending confidence to the design decision. Transitioning was also tested at 50%

throttle. While successful, slop in the pedestal tab-sleeve connection dissuaded any further testing.

Once the new metal components and tubes were received and assembled, the testing regime was repeated using the same test mounting setup as in the initial test, e.g., using the full moment arm. The assembly was first statically tested in the fixed wing configuration up to full throttle to confirm its baseline performance relative to the original assembly. This successfully passed, the test was repeated while holding the assembly static in the VTOL configuration; this too was accomplished uneventfully. The assembly was only then tested while transitioning, at 50% and 100% throttles. The mounts and strut remained visibly rigid, opening the door for their implementation onto the aircraft itself.

One aspect of this testing that did not pass satisfactorily was the performance of the tilt servo. Sized using the weight of the original motor mount, it struggled to rotate the motor-prop assembly while not spooled up. Therefore, larger 347.2 oz*in-rated servos were selected for use on the aircraft itself. The entire assembly, including the new servo, were displayed in Figure 4.4:

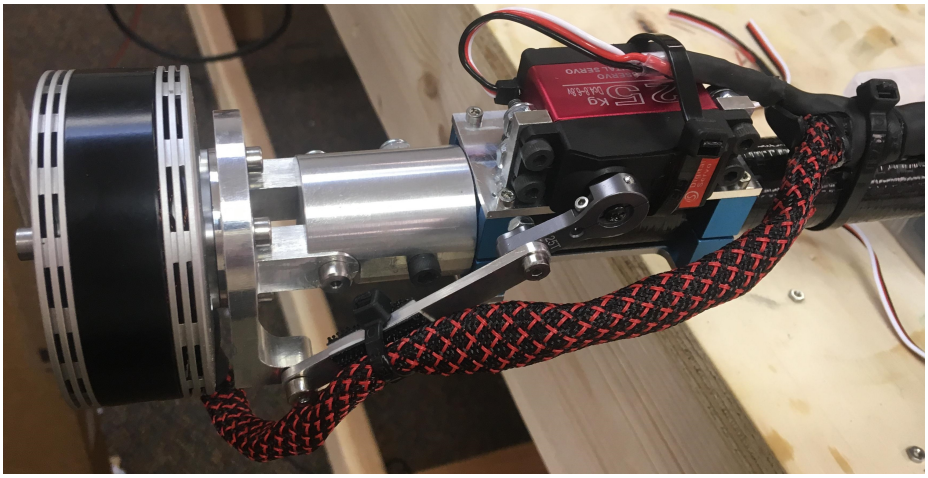

Figure 4.4: Image of Final Propulsion System Outboard Assembly

4.2 Weight Testing

The spectre of weight overages hung over the whole of the aircraft's manufacturing due to incomplete accounting during final design and assembly. The first effort at investigating this issue came just before sandwiching of the skins was performed but after the avionics components had been packaged in the form of a wing tip test. This test entailed supporting the whole of the aircraft's weight via supports place under the wing's tips. Doing so loaded the wing's structure in a sort of worst-case scenario with respect to actual flight conditions, the loading equivalent of a violent turn of just under 3g's. The as-manufactured aircraft appeared as in Figure 4.5 during this test:

Figure 4.5: Image of Initial Tip Test Sag

where the steady-state height was approximately 15 1/4," for a total sag of 3 3/4." A possible structural improvement to the wing spar involved placement of an additional spar sleeve through the body of the platform. Preliminary testing of this addition outside of the fuselage revealed that the outer wings could support themselves in bending even without their central

bolts with such a sleeve placed around the two spars' connection. Testing with a sleeve placed within the fuselage led to the following amount of sag:

Figure 4.6: Image of Spar-Reinforced Tip Test Sag

for a total sag of 1 1/2." With such clear improvement, the sleeve was implemented onto the aircraft.

After implementation of the new spar sleeve, the plane was weighed, including all components on or yet-to-be implemented. The result was 33 lb, representing a 65% overage in weight versus the all-up weight predicted by the mission simulation. A detailed audit of the sources of real-world aircraft weight revealed that while wiring and hardware had been adequately accounted for, the avionics system and aircraft structure had not been sufficiently modeled in enough detail to properly represent their real-world weight values. In the case of structure, the impact of the 1" strut redesign was a 150% increase in system weight, from 1.25 lb to 3.12 lb. The main wing spar had not been accounted for at all, nor

had the skeleton; altogether these parts alone contributed 2.43 lb in unaccounted weight. The avionics system weight model accounted for some components such as the motors and servos, but not others such as motor ESCs, the various power supply components, the BECs, or the propellers. Even in the case of the motors and servos, the calculators used to choose their weights based upon predicted power requirements still under-predicted their weights. The end result was 6.12 lb in unmodeled avionics weight.

Clearly this additional weight would greatly affect the aircraft's in-air behavior and performance. To investigate the effect on endurance specifically, XFLR5 was returned to. The SD7062 airfoil used in the main wing was retested at Reynolds numbers corresponding to the wing's known chord and flight velocities ranging from 49.21 to 131.2 ft/s. The peak C_L for each velocity was found and then discounted by 10% to account for the 2D nature of the analysis. The minimum velocity featuring a required C_L less than the corresponding maximum possible wing C_L was then determined for SLUF and for a 3g turn. The powers required for each scenario were calculated using Equation 2.3 and plotted in Figure 4.7 for a 33 lb aircraft:

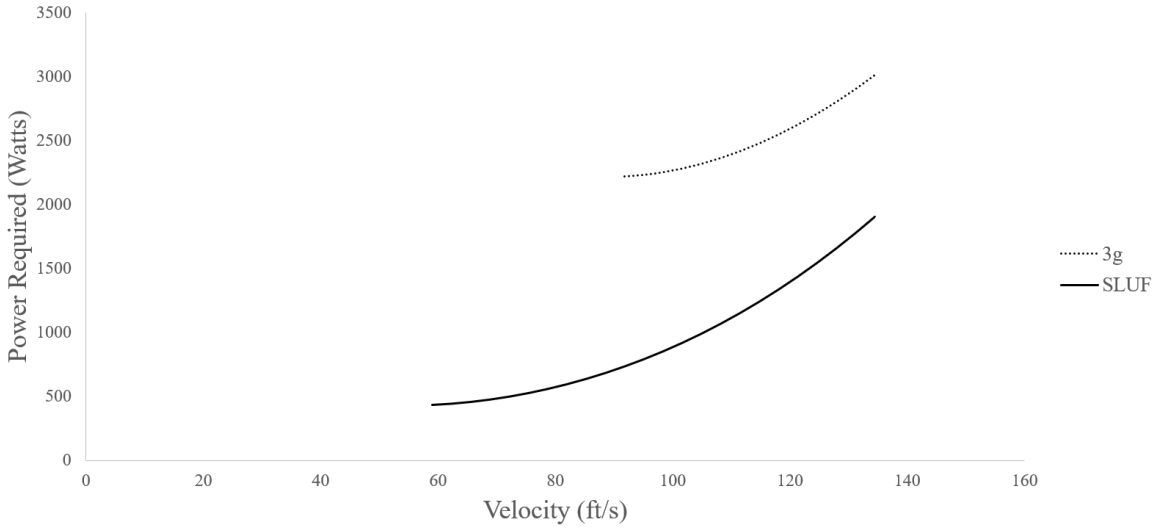

Figure 4.7: Updated Powers Required vs. Flight Velocity

Also visible in Figure 4.7 were the discrepancies in the minimum required velocities between the two flight states. This value was manually plugged into the mission simulation

code, along with updated values for weight, C_L, sampling velocity, and battery capacity, to produce updated sampling endurance numbers. Battery capacity was updated as this technique was applied not only to the 33 lb aircraft but also to cases of the aircraft with one and two batteries removed as well as a 38 lb plane for the 5 lb payload. An assumed flight path involving turning 60% of the time and SLUF 40% was utilized to produce a single endurance value. Table 4.1 summarized the updated endurance values for these cases:

Table 4.1: Updated Operational Endurance Values

Weight Case	**Sampling Endurance (hr)**
33 lb	0.79
33 lb + 5lb	0.64
33 lb - 1 battery	0.69
33 lb - 2 batteries	0.56

Though the endurances listed in Table 4.1 fell short of the 1 hr goal originally set, as well as the 2 hour value reported in the conceptual design report, CASS gave the go-ahead to proceed with the aircraft. Flight testing would serve as the ultimate test for determining real-world performance of the aircraft as a whole.

4.3 Flight Testing

4.3.1 Hover Testing

A simple hover test using three batteries was planned to serve as a baseline test for the aircraft's ability to fly. This would confirm synergy of the platform's various individual systems and reveal any kinks lingering in its implementation. Two batteries were excluded from the full complement to ensure takeoff given the plane's over-weight, bringing the all-up weight to 27.8 lb. The test was performed at The University of Oklahoma's Kessler Atmospheric and Ecological Field Station. Ardupilot's Q_STABILIZE mode was chosen to control the aircraft with for the first few test flights.

There existed some concern over the lack stiffness of in the UAS' landing gear. The two front struts attached to the rear face of the motor pedestals were designed to be long enough to accommodate the conceptual payload pod shown in Figure 2.39. As no pod was to be used during the first test flights, an additional set of legs were bolted to the rear payload mounting points to serve this purpose for landing. It was also noted that the vertical tail rested on the rudder on the ground; therefore, a landing peg was added to the vertical stabilizer to protect the rudder and its hinges.

The plane was successfully hovered over two flights. Takeoff was achieved at a throttle position of 49.6%. It was postulated that this low required-takeoff-throttle was due to the absence of the two batteries as well as to the peak voltage being outputted by the freshly charged batteries, being closer to 4.2 Volts per cell rather than the average value of 3.7 used throughout the performance calculations.

The flights were short but already they revealed a number of issues. The epoxy bonding the front landing struts to their respective aluminum pegs failed. Additionally, one of the struts snapped as a consequence of normal landing operation. Struts with thicker walls would be procured and more rigorously fastened to address these issues. After having been in the air for a number of seconds during the first flight, the craft was seen to weathervane into the wind despite the pilot and autopilot's attempts to counteract this, though careful attention to this during the second flight allowed the pilot to account for it. Mild oscillations about the yaw axis were observed as he did this in a behavior known as "tail wag." The gain of the yaw control in Ardupilot, being the default quadPlane value, did not accurately represent the comparatively large yaw inertia of the given aircraft and would require some tuning during future testing. Overall, the aircraft succeeded in its hovering capacity; Figure 4.8 captured the aircraft midway through the second flight:

Figure 4.8: Image of Hummingbird Mid-Hover

A second round of hover testing was performed to facilitate tuning of the autopilot's quad-specific control gain values as well as to collect current usage values in order to calculate projected hovering endurance values. This round of testing featured a notional payload pod to more accurately match the weight distribution of the craft during operation and to serve as an improved method for supporting the aircraft on the ground. The aircraft's yaw, roll, and pitch characteristics were manually investigated one at a time in hover by the pilot whose feedback was used to alter each axis' control gains in Ardupilot. Following confirmation of each new gains' successful operation, the battery amount in the aircraft was increased and the gains were affected as necessary to maintain favorable handling. In this way, the control gains were manually determined so as to provide stable but responsive control outputs to the autopilot software. Electric current usage was determined using the flight logs from each battery layout. Approximate average current and peak current values were taken from the strictly hovering segments of each flight; this information is summarized for the three flight layouts in Table 4.2:

Table 4.2: Hovering Current Draw and Projected Endurance

Layout	All-Up Weight (lb)	Avg. Current (A)	Peak Current (A)	Avg. Current Endur. (hr)	Peak Current Endur. (hr)
2-Battery	25.8	52	59	0.23	0.20
3-Battery	27.8	53	64	0.34	0.28
4-Battery	29.8	58	67	0.41	0.36

It was noted that, as the 4-battery layout was difficult to control, a 5-battery setup would be impossible to safely land, necessitating its exclusion from this testing. At this point the focus of testing shifted towards the transitioning ability of the aircraft and forward flight.

4.3.2 Transition Testing

Before the actual tests of transition could begin, the lingering issues affecting tilt servo burnout had to first be addressed. An investigation into the state of the burned out servos' internals revealed that one of each servo's MOSFET transistors were charred, leading to the hypothesis that the servos were pulling too high a current at some point during their operation. The peak current draw of one of the remaining operational tilt servos was measured while holding its prop-motor assembly statically in the VTOL and fixed wing configurations as well as during transition in both directions. The highest current draw was consistently seen while simply holding VTOL. This produced the conclusion that, though care had been taken to set the VTOL endpoints of the tilt servos to prevent bottoming out the motor pedestal tabs against the mechanical stops of the strut sleeves, the VTOL endpoints remained too close to the mechanical stops and were causing the tilt servos excess strain. With new tilt servos implemented, the VTOL endpoints were set with an appreciable and equivalent gap and the issue ceased.

Transition testing was initially performed using two batteries, positioned in the two cubbies closest to the wing spar. Additional hover testing of this layout was performed to determine the amount of flight time that would be available for controllable VTOL flight in order to gauge how long fixed wing flight would need to be to allow for a safe landing. Loss of upward mobility was observed after 3 minutes of hovering, so two fixed wing laps

of the test area at Kessler were chosen to be attempted. With two fresh batteries, the aircraft was vertically brought to 262.5 feet above the ground. The transition was initiated via the transmitter, proceeding autonomously until the craft achieved forward flight. The transition was smooth and undramatic, taking approximately 10 seconds to complete, with only a few feet of altitude being lost during the process. However, the aircraft almost immediately entered an oscillatory pitch behavior featuring audible autonomous motor throttling in sync with the cycles of the motion. The craft completed the two laps in this manner before being transitioned back to its vertical configuration; this too was achieved in a surprisingly simple fashion. Sufficient battery capacity had been depleted by this point to cause a hard landing, with the notional payload pod absorbing all of the force of the impact. It was noted that this would not be ideal during real-world implementation with sensitive sensors in such a pod and would require additional safety provisions in the construction of the pod and/or in the battery capacity-remaining margins for landing. Regarding the oscillatory forward flight behavior, the aforementioned field observations regarding abnormal throttling of the motors led to an investigation into the Ardupilot values dealing with automatic stall prevention. The "Q ASSIST SPEED" parameter set the minimum forward flight velocity below which the software would utilize motor tilting to aid in keeping the aircraft aloft and was set to 92 ft/s; the "ARSPD FBW MIN" parameter referred to the minimum acceptable forward flight speed and was set to 59 ft/s. It was postulated that the autopilot was confused by this conflict in velocity values, resulting in the observed behavior. The tilt-motor stall protection was disabled to remove this conflict, and the aircraft was relaunched with two fresh batteries. A good vertical climb to altitude was followed by successful transition to forward flight. However, within a matter of seconds following transition, the aircraft was seen to re-engage in its oscillatory behavior. In the absence of the automatic stall protection from the autopilot these perturbation rapidly grew until the platform pitch up vertically, flipped inverted, and fell to the ground. The aftermath of the crash was shown in Figure 4.9:

Figure 4.9: Image Depicting Aftermath of Crash

Video of the crash was available to use in examining the behavior of the aircraft leading up to the crash as were the Pixhawk's log files. The oscillatory behavior, rather than being caused by the flight controller's stall protection feature, was due to simple pitch instability which had been actively counter-acted by the autopilot during the first test flight.

Chapter 5

Conclusion

5.1 Lessons Learned

The failure of the aircraft's first iteration was not solely limited to its crash. Many mistakes and oversights were made throughout its design and manufacturing, each of which took the platform's implementation further and further from the configuration laid out in the mission simulation code. The result was a platform which, even if it had not crashed, would never achieve the performance desired by CASS. With the most up-to-date predictions in Table 4.1 putting the aircraft's endurance on the order of half an hour at most, the Hummingbird would have been firmly within the performance range of several of the commercially available rotorcraft included in Figure 1.1. This showing would not have justified the platform's comparatively large size, the cost involved with its development, or its regular usage in light of the relative difficulty of repair in the event of a crash.

Perhaps the single most impactful cause for failure in the Hummingbird's development was insufficient book-keeping of the platform's mass properties. The extreme final weight overage taxed the propellers thereby requiring much higher powers during vertical flight and thereby greatly diminishing the battery capacity remaining for forward flight and even limiting the maximum possible battery load-out, in turn affecting the placement of batteries with respect to the overall aircraft center of gravity. Going hand-in-hand with the issue of incomplete weight accounting was incomplete center of gravity determination. Examining Table 2.5, the single most massive line item was the battery weight; any change to this weight source's value and/or center of gravity location would have an outsize effect on the location of the overall aircraft's resultant center of gravity. With the use of incomplete battery load-outs, both of the values were subject to variations which were not adequately

accounted for during testing and which almost certainly contributed to the unstable pitch behavior which ultimately brought the platform down.

Though adequate center of gravity determination only emerged following the stability analysis of the detailed design phase, the ultimate source of the platform's weight overage extended back to incomplete weight modeling in the mission simulation code. This issue manifested itself in two ways: complete exclusion of weight sources, as in the case of the wing spar, and under-prediction of certain components' weight, as in the case of the motors and servos. As discussed in Section 4.2, these oversights mainly materialized in the avionics and structure of the platform as evident in comparing Tables 2.6 and 2.5. Part of the failure of the avionics system was due to an attempt to implement all desired components onto a platform which was ultimately a prototype to be used as a developmental stepping stone towards an eventual production-grade version, allowing for analysis of the feasibility of including the various features. Additionally, the motor and servo weights, though determined based upon real-world values, did not draw from the T-motor and Hitec product lines which would eventually be used to choose the implemented components. Other sources, such as those due to the strut's redesign, had been accounted for but undersized. The remaining unaccounted weight appeared as individually minuscule additions in the course of the aircraft's manufacturing such as joint epoxy, small sources which eventually added up.

Several options were considered regarding the final direction of the project. Among these was the implementation of the hummingbird's surviving autonomous capability and VTOL hardware onto a commercially available fixed wing platform. While investigated at length, this course of action would necessarily serve as a stop-gap measure, leading to an under-performing aircraft satisfying only the barest interpretation of the project's intended goal. More prudent would be the application of the lessons learned throughout implementation of Hummingbird version one towards the design of a second iteration, and it was this course which was taken as the project's final direction.

5.2 Mission Simulation Improvements

The application of the lessons learned from the first version of the aircraft came in the alteration of the mission simulation code. More accurate modeling of the aircraft and its mission requirements would produce a more feasible configuration from the outset. This would leave the code in its best operating condition in case of future application and would also allow the production of a new recommended configuration.

The two overarching failures of the mission simulation code were incomplete system modeling and over-simplified mission design. In addressing the first of these oversights, the lessons learned through the course of the Hummingbird's weight audit were applied in the redefinition of existing models and in the creation of new ones. The linear density of the VTOL strut was updated to reflect that used on the Hummingbird, and its length was redefined to be a function of a given player's chord length and propeller diameter. The choices of electric motors were recompiled to include T-motor's entire catalogue, with the exception of those eight motors of 18S and 22S voltage; the motor selection calculations were updated to reflect these new options. The energy density and maximum continuous discharge of those battery options available from Gens Ace and Tattu were tabulated according to voltage. The maximum possible power output from the selected motor and continuous discharge from the resulting battery selection were used as checks against a given player's feasibility. The choices of servos were drawn from Hitec's catalogue, with control surface and tilt assembly requirements being calculated separately to allow for two different servo selections. The standard 90 Amp Pixhawk power module was substituted to provide power and battery voltage data to the flight controller. All composite skin densities were updated to reflect measured values, and a span-dependent spar was implemented matching the linear density of the Hummingbird's inner spar. The aluminum motor pedestals, strut sleeves, servo platforms, strut clamps, front and rear wing interfaces, and spar inserts were added as constant weight values. Finally, internal skeleton weight was calculated as a proportion of overall weight equivalent to that seen in the implemented platform.

Mission design was altered to better approximate the actual usage of the platform as informed by the Hummingbird's flight testing. In practice, the platform vertically climbed to an altitude of 197.0 feet before transitioning rather than climbing to altitude in fixed wing mode, so this behavior was implemented instead. As found during weight testing, the aerodynamic requirements of turning lead to much higher-required forward velocity values than that required for steady, level, unaccelerated flight, with the resulting increase in power draw producing lower but more realistic operational endurance values. A relatively severe bank angle of 30 degrees was selected to represent the turning condition, corresponding to a g-loading of 1.154. The same theoretical turning/SLUF time distribution, being 60/40, was initially implemented. To find the turning coefficient of lift and velocity, the table of Reynolds number-dependent C_Ls for the SD7062 airfoil was expanded to accommodate a wide range of design possibilities. Each configuration progressively tested each coefficient of lift-velocity pair starting with the lowest value and continuing until a flight state was reached which supported the aircraft. The calculated on-board battery capacity available for use in determining a configuration's operational endurance was decremented to be 0.85 of its actual value as platforms were typically landed with some such safety margin. Finally, the atmospheric values of density and viscosity were altered from their standard sea level values to those matching Kessler's altitude (Toolbox 2003).

Preliminary results of the revamped mission simulation code fell well beneath the desired endurance range. In an effort to better understand the relative effects of a number of the code's new features on operational endurance, tests were run under equivalent conditions with the exception of the parameter being examined. The average operational endurance of each test's top five configurations were averaged and compared against that of the unchanged code. The results of this comparison were plotted in Figure 5.1:

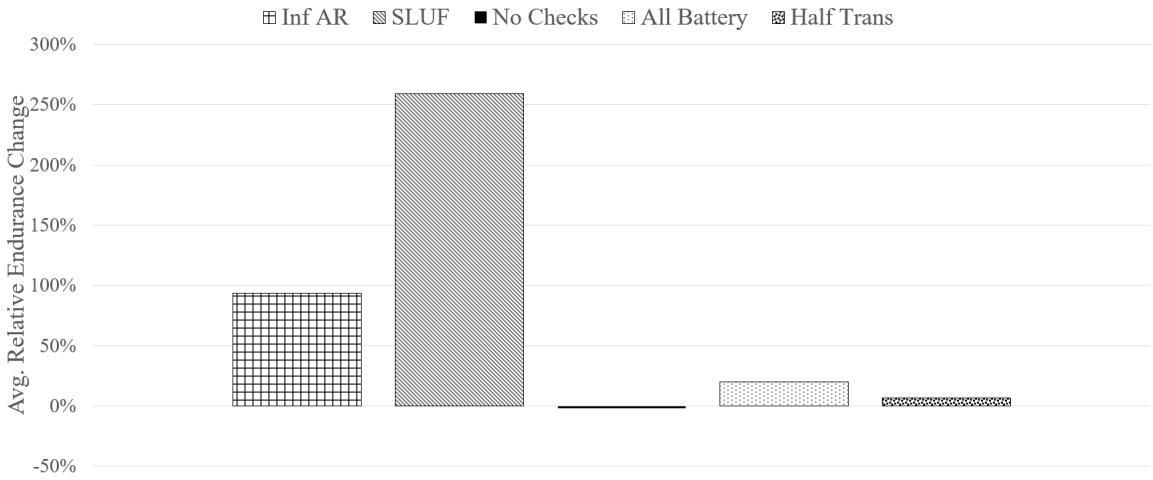

Figure 5.1: Plot of Relative Effects of Mission Simulation Features

The parameter "Inf AR," standing for "Infinite Aspect Ratio," eliminated the provision for a maximum allowed wing aspect ratio of 10. The "SLUF" parameter returned the forward flight power required to solely steady, level, unaccelerated flight. The temporarily eliminated checks referred to in "No Checks" were those related to maximum allowed motor power required and battery discharge. "All Battery" allowed for the use of 100% of the on-board battery capacity, and "Half Trans" reduced the transition altitude from 197.0 feet to 98.43 feet. It was clear from this comparative analysis that the implementation of turning power required had an oversize effect on operational endurance.

An examination of the intended profiles to be flown by the platform during sampling flights revealed that the vast majority of forward flight time would be spent in SLUF. This observation informed the decision to alter the ratio of time spent in turning/SLUF from 60/40 to 10/90. Tests confirmed that this change would still have a significant effect on operational endurance. With this ratio implemented, a new series of tests were run for a range of maximum allowed overall weights. The operational endurance values of the top five configurations from four such tests were plotted in Figure 5.2 against those originally predicted during Preliminary Design:

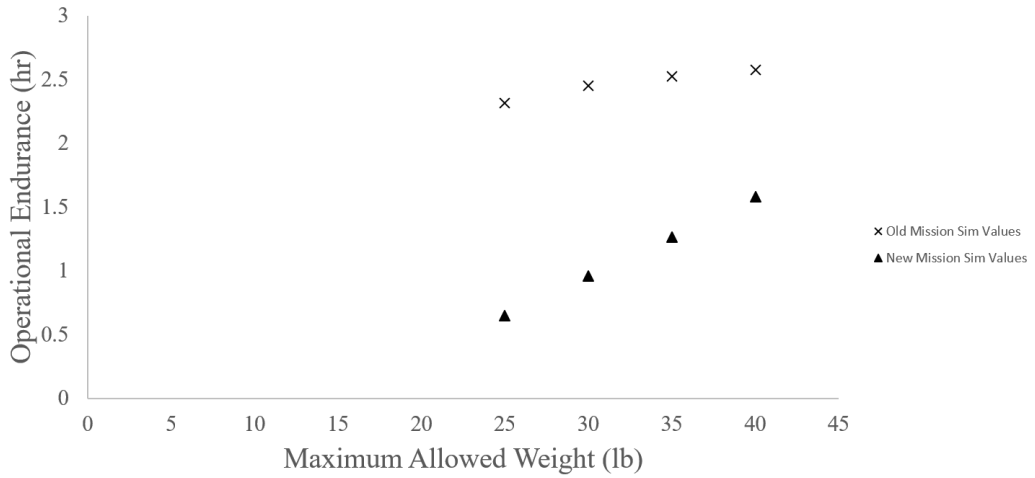

Figure 5.2: Plot of Operational Endurance vs. Maximum Allowed Aircraft Weight

The configuration corresponding to 40 pounds, having a theoretical operational endurance of 1.58 hours, was selected as the recommended design to maintain some margin beyond the desired 1 hour mark.

5.3 Newly Optimized Design Configuration

The final recommended design configuration corresponded to the maximum allowed overall weight of 40 pound plotted in Figure 5.2. This configuration achieved an operational endurance of 1.58 hours while supporting a payload of 5 pounds, thereby theoretically fulfilling two of the three performance requirements of Table 1.1; wind resistance was not modeled was part of the mission simulation and would need to be confirmed in flight testing. As the 40 pounds naturally included the 5 pound payload, it was interesting to note that the resulting 35 pound airframe was similar to the 33 pound as-implemented weight of the Hummingbird. Examining the differences in the two design states, it was clear that a much higher percentage of this weight would be in the batteries given the increase in battery capacity from 642.6 for the previous top design to 1371 Watt-hours, with much of the decreases in the other parameters working towards saving enough weight to allow for this additional capacity. Within the top five designs, wing area ranged from 5.644 to 6.312 ft^2

and wingspan falling in line according to the maximum allowed aspect ratio of 10. That of the Hummingbird's configuration was 6.848 ft^2. The result were higher sampling velocities, of between 61.58 ft/s and 64.90, and 3g-turn velocities of 110.8 to 117.1 ft/s versus 55.77 ft/s and 91.86 ft/s, respectively, predicted in Figure 4.7.

The decision variable values of the top-performing 40 lb configuration were compiled in Table 5.1:

Table 5.1: Final Recommended Design Configuration

Wing Area (ft^2)	**Wingspan** (ft)	**Battery Capacity** (WH)
5.785	7.605	1371
VTOL Takeoff Vel. (ft/s)	**VTOL Landing Vel.** (ft/s)	**Prop Diameter** (ft)
13.47	6.181	0.8924

A simplified CAD model based upon that of the Hummingbird was generated according to the geometrical decision variable values of Table 5.1 as shown in Figure 5.3:

Figure 5.3: Render of Simplified Final Configuration CAD

5.4 Recommendations for Future Work

Having applied the lessons learned throughout the development of the platform's first iteration in improving the mission simulation code, the program was set up to produce more feasible design configurations and performance predictions. This improved capability positioned the project well for its next stage of work. However, there remained a number of points for which additional attention could be paid.

The current battery model of the mission simulation code allowed for continuous battery weights based upon the desired value of capacity and energy densities derived from real-world battery options as dictated by the system voltage. Though more accurate than that utilized previously, this model could be improved if it were to feature discrete battery options. Additionally, a better model would optimize selection of the battery to be used based upon current draw and maximum possible discharge as well as power draw.

Much of the platform's mechanical design could be better optimized as well. Much of the aircraft's structure was redundant. The monocoque skin of the wing was load bearing, but so was the carbon tube-spar. Depending upon the ultimate size and weight of future iterations, the spar could be thinned or even eliminated. The layup schedule of the tail surfaces, being two plies of two ounce fiberglass on either side of a quarter-inch thick balsa core, was overbuilt for the expected loads; thinner core or a different core material better optimized for the aerodynamics loads expected on the tail surfaces could lead to appreciable weight savings. The aluminum hardware associated with the tilt motor assembly served as constant, significant weight sources during the design's re-optimization as did the cross sectional weight of the strut tube; smaller scale UASs could certainly utilize lighter versions of these components. If a similar motor hinge design were to be reutilized in following platform, the implementation of sleeve or ball bearings at the connection between the hinge and the aluminum components would be useful to prevent possible wallowing of the bare aluminum holes as well as to provide the least rotational resistance to the tilt servos. The use of a turnbuckle linkage arm could provide the fine adjustment necessary to account for manufacturing tolerances in the tilt assembly which resulted in some slop in the linkage connection. A more durable landing solution would be required for full use of the platform. Though the payload pod could be constructed with a durability capable of protecting its sensitive sensor payload and supporting the aircraft during takeoffs and landings, the additional weight to do so would be counterproductive. Were the payload to

be packaged elsewhere on the platform, the monocoque skin could be designed to support belly landings without much additional weight.

The ergonomics of internal access could also benefit from a rework. As they stood, the platform's use of plastic nuts and bolts, while lightweight, were slow to use; small scale quarter-turn fasteners could be a feasible alternative. Additionally, the size of the hatches relative to the size of the batteries was small enough that battery changes became an involved process, and some components such as the flight controller lacked hatches completely. Though the addition of more and larger hatches could solve this issue, it would greatly decrease the stiffness of the load-bearing monocoque skin. Re-packaging of the platform's avionic components into a removable module would have the benefits of total access as well as a completely continuous monocoque except for its single access point.

Finally, it was recognized that many of the additions to the mission simulation code, while increasing the accuracy of the modeling of the given platform, also narrowed the scope of possible platform concepts. Its implementation was a compromise between over-burdening its operation with additional features, such as wind resistance, and maintaining a sufficient degree of reality that the testing metrics would lead to a well-performing platform even following real-world implementation. The feasibility of altogether different aircraft concepts should not be ruled out in the case of future iterations; for example, significant weight savings could be realized were a tilting strategy leveraging the shape of the airframe and flight controller programming to be utilized rather than additional hardware, and thus weight, solely for this purpose as in the case of the Hummingbird's tilt struts and motor mount assemblies. So long as the mission simulation code were to be accordingly altered to capture the major requirements of the different concept, comparative analysis of its performance versus that of the original layout could validate such a design decision.

Bibliography

Afework, B., J. Hanania, K. Stenhouse, and J. Donev, 2018: Sterling engine. University of Calgary, URL https://energyeducation.ca/encyclopedia/Stirling_engine, [Accessed 10 August 2018], online.

Allred, J., and Coauthors, 2007: Sensorflock: An airborne wireless sensor network of micro-air vehicles. ACM, Sydney, Conference on Embedded Networked Sensor Systems.

Alonso, J., 2009: Simple performance estimation. URL http://adl.stanford.edu/sandbox/groups/aa241x/wiki/e054d/attachments/31ca0/performanceanddrag.pdf, [Accessed September 2018], online.

APC Propellers, 2017: Apc propeller performance data files. URL https://www.apcprop.com/technical-information/performance-data/, [Accessed October 2018], online.

ARPRO, 2018: Hardness "type c" comparison for arpro epp vs. arpak epe vs. arpex xepe molded from jsp expanded bead foam. URL http://www.arpro.com/tech-docs/tech-docs.php, [Accessed 13 August 2018], online.

Atomics, G., 2018: Predator xp rpa. URL http://www.ga-asi.com/predator-xp, [Accessed 9 August 2018], online.

Avery, A., 2013: Design and development of a storm research uas. M.S. thesis, Oklahoma State University, Stillwater.

Banis, D., J. Marceau, and M. Mohaghegh, 2015: Design for corrosion control. *Aero Magazine*, **(7)**, doi:http://www.boeing.com/commercial/aeromagazine/aero_07/corrosn.html.

Brady, J., M. Stokes, J. Bonnardel, and T. Bertram, 2016: Characterization of a quadrotor unmanned aircraft system for aerosol-particle-concentration measurements. *Environmental Science Technology*, **50**, 1376–1383.

Brandt, J., R. Deters, G. Ananda, and M. Selig, 2015: Uiuc propeller data site. University of Illinois at Urbana-Champaign, URL https://m-selig.ae.illinois.edu/props/propDB.html, [Accessed 3 December 2018], online.

BRP Rotax GmbH and Co., 2016: Engine type 914 115 hp (ul/f). URL https://www.flyrotax.com/files/Bilder/Produkte%20Rotax/Datasheets/Produktdatenblatt_914_115hp_rev.BRP-Rotax_20160823.pdf, [Accessed 24 August 2018], online.

Budapest University of Technology and Economics, 2018: Fuel properties: Aviation engineer/specialist. URL ftp://ftp.energia.bme.hu/pub/AviationEngineerSpecialist/PowerPlant/AvEng_Fuelproperties.pdf, [Accessed 24 August 2018], online.

Buonassisi, T., 2013: Fundamentals of photovoltaics. Massachusetts Institute of Technology, URL https://ocw.mit.edu/courses/mechanical-engineering/2-627-fundamentals-of-photovoltaics-fall-2013/, [Accessed 12 August 2018], online.

Carpenter, P., 2018: Model airplane kits-what's available? URL https://www.rc-airplane-world.com/model-airplane-kits.html, [Accessed 14 August 2018], online.

Chilson, P., and Coauthors, 2009: Smartsonde: A small uas platform to support radar research. American Meteorological Society, Boston, 34th Conference on Radar Meteorology.

Delacourt, C., P. Allemand, M. Jaud, P. Grandjean, A. Deschamps, J. Ammann, V. Cuq, and S. Suanez, 2009: Drelio: An unmanned helicopter for imaging coastal areas. *Journal of Coastal Research*, **(SI 56)**, 1489–1493.

Drela, M., S. Hall, P. Lagace, and I. Lundqvist, 2006: Lab 8 notes – basic aircraft design rules. URL https://ocw.mit.edu/courses/aeronautics-and-astronautics/16-01-unified-engineering-i-ii-iii-iv-fall-2005-spring-2006/systems-labs-06/spl8.pdf, [Accessed October 2018], online.

Ellwood, J., 1990: Design and construction of a composite airframe for uav research. M.S. thesis, Naval Postgraduate School, Monterey.

Elston, J., B. Argrow, M. Stachura, D. Weibel, D. Lawrence, and D. Pope, 2015: Overview of small fixed-wing unmanned aircraft for meteorological sampling. *Journal of Atmospheric and Oceanic Technology*, **32 (1)**, 97–115.

Fleming, J., W. Ng, and S. Ghamaty, 2004: Thermoelectric-based power system for unmanned-air-vehicle/microair-vehicle applications. *Journal of Aircraft*, **41 (3)**, 674–676.

Forbus, K., 2018: How do batteries work? URL http://www.qrg.northwestern.edu/projects/vss/docs/power/2-how-do-batteries-work.html, [Accessed 8 August 2018], online.

Glover, J., 2013: Foam!...types of foam that are used in rc fpv mini uav airframes. URL http://kiloohm.com/foam-types-of-foam-that-are-used-in-rc-fpv-mini-uav-airframes/, [Accessed 13 August 2018], online.

Gramoll, K., 2019: Torsion: Mechanics - theory. URL https://www.ecourses.ou.edu/cgi-bin/ebook.cgi?doc=&topic=me&chap_sec=02.1&page=theory, [Accessed April 2019], online.

Griffis, C., T. Wilson, J. Schneider, and P. Pierpont, 2009: Unmanned aircraft system propulsion systems technology survey. Tech. rep., U.S. Department of Transportation, Washington D.C.

Hays, T., 2009: High energy density propulsion system dynamometer for small engines. M.S. thesis, Oklahoma State University, Stillwater.

Holland, G., and Coauthors, 2001: The aerosonde robotic aircraft: A new paradigm for environmental observations. *Bulletin of the American Meteorological Society*, **82 (5)**, 889–901.

IAI, 2002: Pioneer. URL http://www.iai.co.il/2013/36945-41973-en/Older_UAV.aspx, [Accessed 9 August 2018], online.

Insitu, 2017: Scaneagle. URL https://www.insitu.com/images/uploads/pdfs/, [Accessed 31 July 2018], online.

Li, Y., J. Deng, C. Mu, Z. Xing, and K. Du, 2014: Vertical distribution of co_2 in the atmospheric boundary layer: Characteristics and impact of meteorological variables. *Atmospheric Environment*, **91**, 110–117.

LiquidPiston, 2018: Power reimagined compact, quiet, low-vibration, high-efficiency rotary engines. URL http://liquidpiston.com/wp-content/uploads/2018/07/LiquidPiston-X-Specifications-2018-05-14-3.pdf, [Accessed 24 August 2018], online.

Ma, S., H. Chen, G. Wang, Y. Pan, and Q. Li, 2004: A miniature robotic plane meteorological sounding system. *Advances in Atmospheric Sciences*, **21 (6)**, 890–896.

Northrop Grumman, 2018: Global hawk. URL http://www.northropgrumman.com/MediaResources/Pages/Photo.aspx?pid%3DGL-10001_015%26rel%3D%2F%26name%3DPhotos, [Accessed August 2018], online.

Northwestern, 2018: Thermoelectrics: The science of thermoelectric materials. Northwestern Materials Science and Engineering, URL http://thermoelectrics.matsci.northwestern.edu/thermoelectrics/, [Accessed 12 August 2018], online.

PBS Aerospace, 2014: Tj20 turbojet engine. URL http://www.pbsaerospace.com/our-products/tj-20-turbojet-engine, [Accessed 24 August 2018], online.

Raymer, D., 2012: *Aircraft Design: A Conceptual Approach*. 5th ed., American Institute of Aeronautics and Astronautics, Reston.

Raytheon, 2018: Manta uas long-range unmanned aircraft system. URL https://www.raytheon.com/capabilities/products/manta, [Accessed 30 July 2018], online.

RCV Engines Ltd., 2016: Heavy fuel uav engines. URL http://www.rcvengines.com/documents/rcv%20flyer.pdf, [Accessed 24 August 2018], online.

Reyes, C., 2018: Introduction to rc airplane foams. URL http://rcadvisor.com/introduction-rc-airplane-foams, [Accessed 13 August 2018], online.

Rhinehart, R., 2019: Leapfrogging: A multiplayer optimization direct search algorithm. URL http://www.r3eda.com/leapfrogging-optimization-algorithm/, [Accessed September 2018], online.

Scholz, D., 2017: Empennage general design. URL http://www.fzt.haw-hamburg.de/pers/Scholz/HOOU/AircraftDesign_9_EmpennageGeneralDesign.pdf, [Accessed July 2018], online.

Schrod, J., and Coauthors, 2017: Ice nucleating particles over the eastern mediterranean measured by unmanned aircraft systems. *Atmospheric Chemistry and Physics*, **17 (7)**, 4817–4835.

Scitech, 2009: How things work: The wankel engine. *The Tartan*.

Selig, M., J. Donovan, and D. Fraser, 1989: *Airfoils at Low Speeds*. H.A. Stokely, Virginia Beach.

Stierwalt, S., 2015: Einstein's legacy: The photoelectric effect. *Scientific American*.

T-Motor, 2017: T-motor products. URL http://store-en.tmotor.com/category.php?id=11, [Accessed October 2018], online.

Tian, L., 2013: Miniature homogeneous charge compression ignition free-piston engine compressor. Ph.D. thesis, University of Minnesota, Minneapolis.

Toolbox, E., 2003: U.s. standard atmosphere. URL https://www.engineeringtoolbox.com/standard-atmosphere-d_604.html, [Accessed October 2019], online.

TuffWing, 2018: Tuffwing products. URL http://www.tuffwing.com/products/products.html, [Accessed 13 August 2018], online.

UASUSA, 2018: The tempest unmanned fixed wing drone. URL https://www.uasusa.com/products-services/aircraft/the-tempest.html, [Accessed July 2018], online.

Wang, R., X. Xu, S. Jia, L. Ran, Z. Deng, W. Lin, Y. Wang, and Z. Ma, 2017: Lower tropospheric distributions of o_3 and aerosol over raoyang, a rural site in the north china plain. *Atmospheric Chemistry and Physics*, **17**, 3891–3903.

XFLR5, 2014: Xflr5 v6.02 guidelines. URL http://www.xflr5.com/xflr5.html, [Accessed October 2018], online.

Appendix A

Tables of Compiled Datasets

Table A.1: Table of Commercial Platform Performance Data

Platform	GTOW (lb)	Max Payload (lb)	Max Endurance (hr)	Max Wind Resistance (knots)
Fixed Wing				
SUMO	1.279	0.3086	0.5000	Unreported
DataHawk	1.543	0.4409	1.000	Unreported
Tuffwing	4.400	1.320	0.6667	Unreported
MASC	8.818	2.205	1.500	Unreported
Aerolemma-3	7.937	2.646	0.2500	Unreported
M^2AV	12.35	3.307	0.7500	Unreported
NOVA	17.00	3.700	1.500	26.10
CU NexSTAR	11.02	4.409	0.6667	Unreported
Tempest	17.00	7.000	1.500	52.10
RPMSS	28.66	8.820	8.000	38.90
Aerosonde	33.07	11.02	40.00	Unreported
ScanEagle	48.50	12.35	11.00	Unreported
Manta	61.07	22.05	6.000	Unreported
UMARS 2	66.14	22.05	4.000	Unreported
Cruiser	77.16	22.05	6.000	Unreported
MARIA	35.00	10.00	8.000	69.50
Rotorcraft				
Phantom 4 PRO v2	3.031	0.000	0.5000	19.44
Typhoon H	4.300	0.5620	0.4167	Unreported
Inspire 2	9.370	1.790	0.4500	19.44
MATRIX-i	10.00	6.200	0.7000	34.76
H920	11.00	2.650	0.4000	Unreported
MATRICE 200 Series	13.54	5.160	0.6333	23.33
MATRICE 600 Pro	34.17	13.23	0.6333	15.55
DRELIO	37.48	13.23	Unreported	27.00
Alta 8	40.00	20.00	0.5833	Unreported
Agras MG-1	54.01	22.05	0.4000	Unreported

Table A.2: Motor Repository

Motor	**KV**	**Internal Resistance (Ω)**	**Idle Current (*A*)**	**Weight (lb)**	Max Continuous Power (W)	Max Continous Current (*A*)
'U11KV90'	90	0.09900	0.7000	1.797	3000	60
'U11KV120'	120	0.05700	0.7000	1.746	4000	80
'U11II KV120'	120	0.08000	1.500	1.702	2783	57.30
'U13KV85'	85	0.02300	1.250	2.822	3120	65
'U13KV100'	100	0.01800	1.640	2.866	3848	80
'U13II KV130'	130	0.01800	1.500	2.183	5659	118
'U15IIKV80'	80	0.01700	3.800	3.836	8580	143
'U15IIKV100'	100	0.01200	4.500	3.836	9942	143
'P80KV100'	100	0.05100	1.600	1.433	2800	60
'P80KV120'	120	0.04100	1.950	1.433	3000	70
'U10 PLUS KV170'	170	0.02100	1.900	1.127	2000	60
'U12 KV90'	90	0.03600	1.200	1.750	2500	50
'U12 KV100'	100	0.02100	1.300	1.746	2500	50
'U12II KV120'	120	0.02200	1.500	1.737	4560	95
'MN705-S KV260'	260	0.01600	2.800	0.9921	2000	80
'MN801-S KV150'	150	0.05000	1.400	1.058	2800	60
'MN805-S KV120'	120	0.04800	1.600	1.367	3200	65
'MN805-S KV150'	150	0.03200	2.200	1.378	3600	75
'MN805-S KV170'	170	0.02800	2.600	1.378	4000	83