HydroFleet

A MULTI-DRONE SYSTEM FOR AUTONOMOUS SOIL MOISTURE MAPPING

Olin College Robotics Lab

Contents

1 Project Intro

1.1 One Sentence Summary

We introduce **HydroFleet**, a multi-drone system that enables data-driven precision irrigation by generating high-resolution soil moisture mapping data, helping farms increase yields while reducing water and energy costs.

1.2 Ag Partner Overview

Our team was privileged to collaborate with two agricultural partners through the duration of this project. Their advice, feedback and insights significantly shaped the outcome of our work. We worked with **Tim Laird & Aubrey Dority** at Powisset Farm in MA, and **Clark McPheeters** at CBM Farms in Nebraska.

Tim & Aubrey, Powisset Farm (Ag expert advisors and Testing partners): Tim Laird, the Field Crop Manager, and Aubrey Dority, the Seasonal Assistant Field Crop Manager, of Powisset Farm in Dover, Massachusetts, served as our primary agricultural partners. Powisset Farm is a 109-acre community-supported agriculture (CSA) farm specialising in organic speciality crops, primarily vegetables. Tim and Audrey served as advisors and testing partners for implementing a multi-drone soil moisture measurement system on a real operational farm. Tim and Aubrey drew on their extensive experience as farmers to provide insights into the value of accurate soil moisture measurements, sharing stories with us about how inadequate knowledge of soil moisture levels can affect crop yields. As testing partners, they graciously opened their farm to us for flight and soil measurement tests throughout the project. While the commercial value of such a project may be limited for a farm of their scale, their hands-on, intentional involvement was instrumental in advancing our on-farm testing and validation.

Tim and Aubrey were just as excited as we were to be working together. Tim stated, *"An automatic robot system could be incredible. A dream of a lot of farmers is to take a day off and then come back and see all the work that's been done. That would be an amazing thing. It would make me feel like I could take a day off... most farmers during the season don't take days off; we work every day."* Tim believes that achieving this would be amazing.

Clark McPheeters, CBM Farms (Large-scale Ag expert advisor): Clark McPheeters, President of CBM Farms in Nebraska, provided key insights into the needs of large-scale farms for high-resolution soil moisture information, what ROI could look like for such a system and some of the strengths and limitations of our system. CBM Farms is a family-owned farm (spanning six generations) that operates over 3000 acres of row crop fields in Gothenburg, Nebraska, and produces food-grade corn, soybeans, and alfalfa. Clark provided expert practitioner insights on how pressurized irrigation works in large-scale farming, walking us through his process. He described his experience with different soil measurement technologies, their pros and cons, and the opportunities in this space. He provided critical feedback on the design of our proposed solution and guidance on moving forward. While distance limited our ability to visit and perform tests on his farm, he expressed a strong interest in continuing to support this work, and we are committed to cultivating this relationship.

1.3 Problem Statement

Soil moisture is a primary determinant of crop yield and irrigation efficiency. Variable-rate irrigation technologies, which modulate water application spatially across a field based on measured

Figure 1: Powisset Farm of Dover, MA, a 109-acre organic CSA farm, served as our off-campus testing site

soil water content, can reduce irrigation-related water usage by **25–40 percent** while increasing yields by up to **15 percent** for certain crops [Liakos et al., 2017, Edwards et al., 2025]. Realizing these gains requires high-resolution volumetric water content (VWC) maps that capture spatial variability at the sub-field scale (Figure 2). Currently, however, over **75 percent** of farmers rely on hand-sampling and rule-of-thumb practices to determine irrigation timing and intensity [Zhang et al., 2021, Evans et al., 2013]. With the average U.S. farm now spanning 463 acres and growing [U.S. Department of Agriculture, 2022], manual sampling at the required density—approximately one measurement per acre to depths of up to 18in (the full root profile of many cropping species) [Kashyap and Kumar, 2021]—is too costly and labor-intensive to support precision irrigation decisions.

Existing technologies for soil moisture measurement fall into two main categories. (1) **Stationary in-situ sensors** (e.g., capacitive sensors, time-domain reflectometry) are easy to deploy and provide reliable point measurements; however, achieving high-resolution sub-field mapping requires dense sensor networks that are prohibitively expensive and maintenance-intensive for most farms [Gasch et al., 2017]. (2) **Remote sensing approaches**, including satellite and drone-mounted ground-penetrating radar (GPR) systems, enable measurements without physical contact with the soil and scale well over large fields. However, they suffer from either insufficient spatial resolution for effective sub-field decision making or are unable to operate under vegetation cover, significantly degrading their effectiveness [Zhang et al., 2021]. As a result, current solutions fail to provide scalable, high-resolution, and robust soil moisture mapping at the sub-field level needed for precision variable-rate irrigation.

We propose **HydroFleet**, a multi-drone system for scalable, high-resolution soil moisture mapping to support precision irrigation. At the core of HydroFleet is a novel cable-suspended soil sampling module that enables drones to perform in-situ moisture measurements without landing, significantly improving probing speed and efficiency. By coordinating a fleet of small, modular drones, our system can adaptively sample soil locations in an information-efficient manner, leading to dense, sub-field moisture maps that are difficult to achieve with existing technologies. These high-resolution moisture maps provide actionable decision support for variable-rate irrigation, enabling more precise and efficient water application across farms.

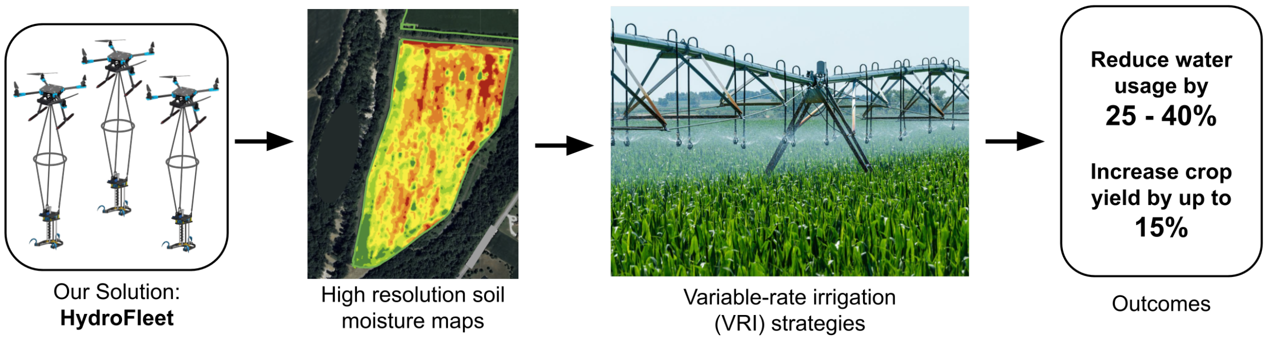

Figure 2: The solution flowchart for HydroFleet: enabling high-resolution moisture maps that inform variable-rate irrigation strategies which lead to higher crop yields and reduced water and energy costs [Liakos et al., 2017].

2 Design Explanation

2.1 Technical Description & Operations

The primary objective of HydroFleet is to autonomously coordinate a fleet of sampler-equipped drones to collect spatially distributed soil moisture measurements across a crop field and deliver actionable information to the farmer. To achieve this objective, the system comprises four subsystems: a drone platform, a soil probing module, multi-vehicle planner software, and a web-based farmer dashboard.

- **Drone Platform:** This subsystem serves as the mobile foundation, integrating an off-the-shelf drone airframe with a dedicated companion computer. It is responsible for managing custom flight behaviors by communicating with the drone's flight controller.
- **Soil Probing Module**: This subsystem serves as a soil moisture probing mechanism that measures a range of soil moisture level for a depth of about 5 inches. This mechanism consists of an anchoring, drilling, and sampling mechanism.
- **Multi-Drone Planner**: This subsystem serves as the offline path planner for multi-agent drone operations, ensuring that each drone's assigned route remains within its single-charge flight endurance while collectively achieving full coverage of the target area.
- **Farmer Dashboard**: This subsystem serves as an intuitive, easy to use interface between HydroFleet and the farmers. It provides a centralized platform for the farmer to plan sampling routes and visualize data collected by the drones.

Detailed descriptions and functionalities of each subsystem are provided in the respective sections below, and the full HydroFleet system can be seen in Figure 3 and 4.

2.1.1 Drone Platform

The drone platform is a multirotor UAV (HolyBro X500 V2) airframe equipped with two computing elements:

- A flight controller (PixHawk 6X) responsible for low-level attitude stabilization, position hold, and waypoint navigation. The flight controller accepts high-level commands (e.g., go to waypoint, hold position, ascend, descent) and executes them using onboard inertial measurement, barometric, and GNSS sensors (DroneCAN H-RTK F9P Rover GPS).

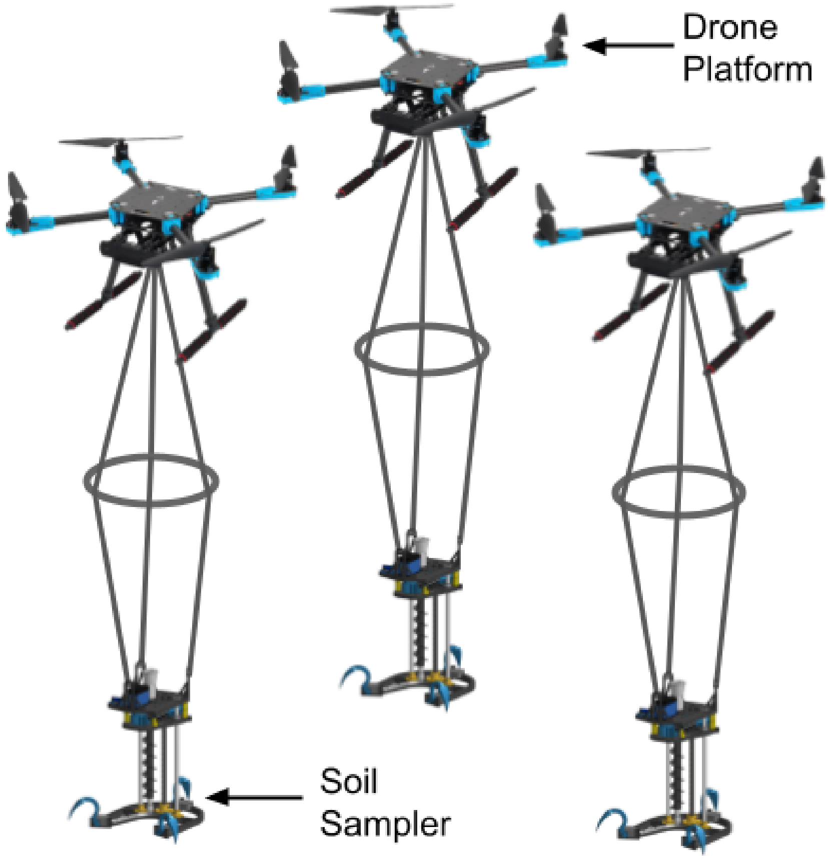

Figure 3: HydroFleet System Concept

Figure 4: HydroFleet operating at Powisset Farm, Dover, MA

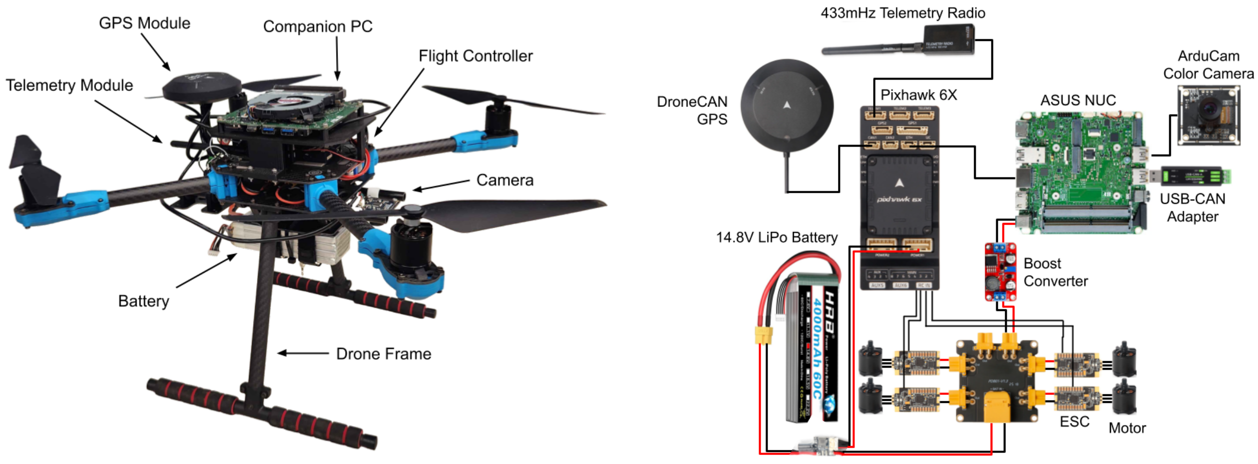

(a) (b)

Figure 5: (a) Physical drone with labeled components. (b) Drone platform wiring diagram

- An onboard companion PC responsible for autonomous mission-level decision-making, including the sequencing of measurement sites and coordination of the soil measurement cycle with the soil probing module.

Wiring details are shown in Figure 5b.

The drone platform also provides electrical power to the soil sampling module from its onboard battery, eliminating the need for a separate energy source on the module and minimizing suspended mass.

Standard ground control software and flight controllers are already capable of tasks like basic waypoint navigation, but they lack the native logic to coordinate multi-stage mission sequences. Tasks that require tight synchronization between flight maneuvers – such as precision descent, CAN-based sampler actuation and real-time data logging – demand a higher-level autonomy layer. For this project, we implemented this higher-level orchestration using ROS2 [Macenski et al., 2022].

ROS2 is a popular general purpose robotics library that has been deeply integrated with PX4, an open-source system for flight controllers [Meier et al., 2026], to the extent that custom flight modes can be developed in ROS2 that can be used like internal PX4 modes. Communication between ROS2 and PX4 is established via the XRCE-DDS protocol, which exposes PX4 uORB messages as ROS2 messages and types. This allows direct access to PX4 using ROS2 workflows and nodes.

Additionally, the PX4 ROS2 Interface Library is a C++ library that further simplifies interacting with PX4 from ROS2. The library provides the Control Interface, which allows for easy registration of custom modes written in ROS2. It implements the ModeExcutor() class, which acts as the orchestrator of multiple custom modes.

For HydroFleet, three custom modes are implemented and registered (Figure 6).

1. **GoToWaypoint**: Navigates the drone to specified sampling coordinate.
2. **SoilSample**: Descends the drone until the soil sampler is on the ground. Then, it actuates the sampler with a CAN signal and waits until measurement data are sent back.
3. **CustomRTL**: Executes a custom landing behavior where the drone descends to land the sampler first. Then, the drone moves forward 1 meter and lands fully.

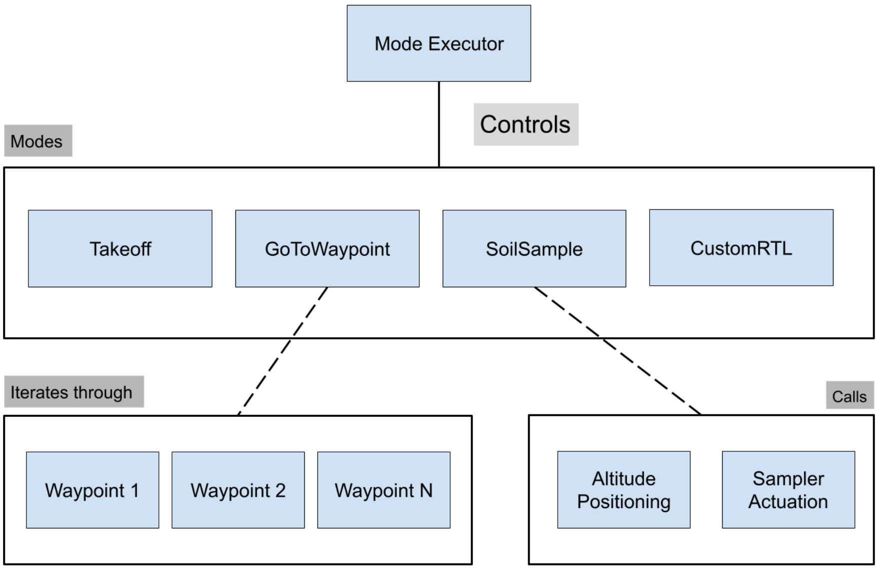

Figure 6: Illustration of custom mode structure

Simulation Pipeline (SITL)

To minimize the risk of damaging hardware, we made extensive use of a custom Gazebo simulation environment [Koenig and Howard, 2004] for testing autonomous behaviors. Building upon the standard PX4 single-vehicle simulation, the environment was expanded to support multiple vehicles within a customizable farm environment. This digital twin allowed for rapid iterative testing, allowing us to identify and resolve any bugs in a controlled setting before deploying to the physical hardware (Figure 7).

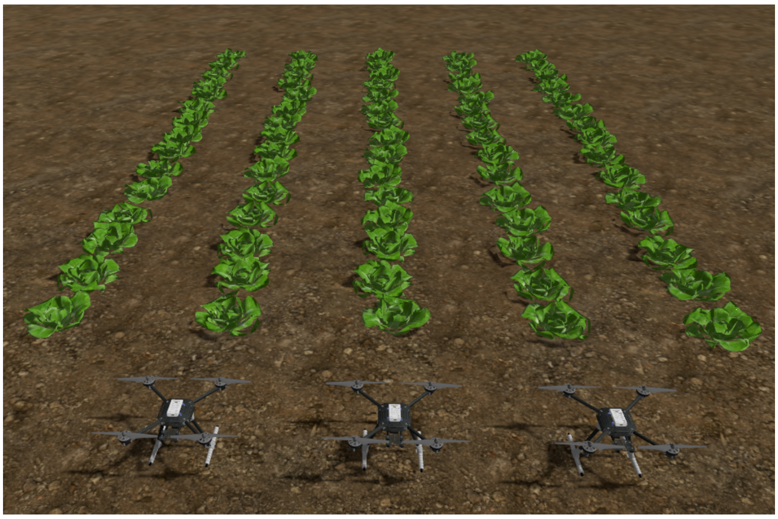

Figure 7: Multi-drone simulation in Gazebo Simulator

Cloud Workflow

To streamline the multi-vehicle development workflow, we adopted a cloud-based workflow using balenaCloud (Fig. 8). By grouping all drones into a single fleet, we can ensure that identical software environments are maintained across all units simultaneously. This centralized management eliminates

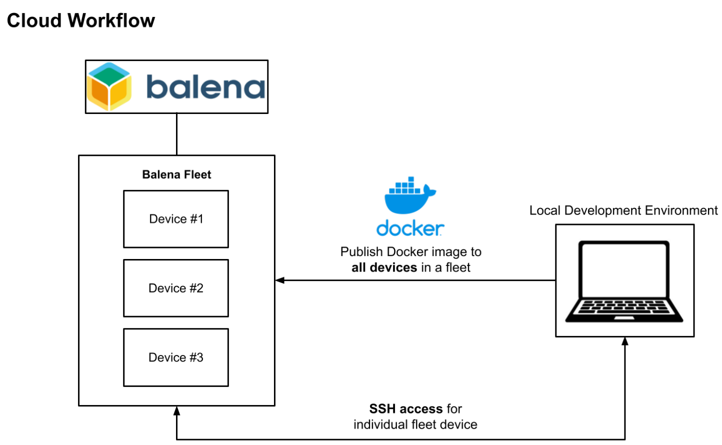

Figure 8: Balena Cloud fleet management workflow

the need for manual, device-by-device updates. Instead, any updates in the code or environment are pushed to the entire fleet in one action. This approach significantly reduces configuration mismatches and ensures that every drone operates on the same validated software version during field tests.

2.1.2 Soil Sampling Module

The soil probing module is the truly novel component of our proposed system and was designed from scratch throughout the project. It is a lightweight, self-contained electromechanical assembly that, once lowered to the ground surface, performs the complete measurement cycle autonomously under the direction of its onboard microcontroller. The module contains the following principal components (see Figure 9):

- **Sampling mechanism (Auger + Soil Moisture Sensor):**
  The sampling mechanism consists of the auger and the soil moisture sensor. We used a capacitive sensor (DSMM500 Precision Digital Soil Moisture Meter with Probe from General Tools) for soil measurements. This was selected because it can accurately measure the soil moisture of the dirt, while having the distinct cylindrical shape that other soil moisture sensors don't have. For the sensor to get deep into the ground, we went with the auger mechanism that could effectively penetrate the soil to the target measurement depth.

- **Anchoring Mechanism:**
  The anchoring mechanism utilizes a set of servos (MEUS Racing V2 Coreless RC Micro Servos–8.5KG) to deploy custom-designed anchors. These servos were selected for their high torque-to-weight ratio, providing the necessary force without adding significant mass to the assembly. To ensure efficient soil penetration, the anchor geometry features a tapered profile that begins with a narrow tip and widens progressively. The anchor's circular cross-section reduces drag during penetration and maximizes the surface area acting perpendicular to the auger's rotation. This also minimizes the amount the sampler lifts during anchoring (Figure 10). This design minimizes initial resistance while maximizing holding force once fully engaged. Each anchor is keyed directly to the servo horn for precise control and structural integrity. Additionally, the base of the system incorporates several spikes that improve stability and enhance initial ground contact upon landing, making it easier for the anchors to penetrate the soil.

- **Actuation Mechanism (lead screw + gear system):**
  The actuation mechanism utilizes a continuous servo to drive a high-torque gear train. This

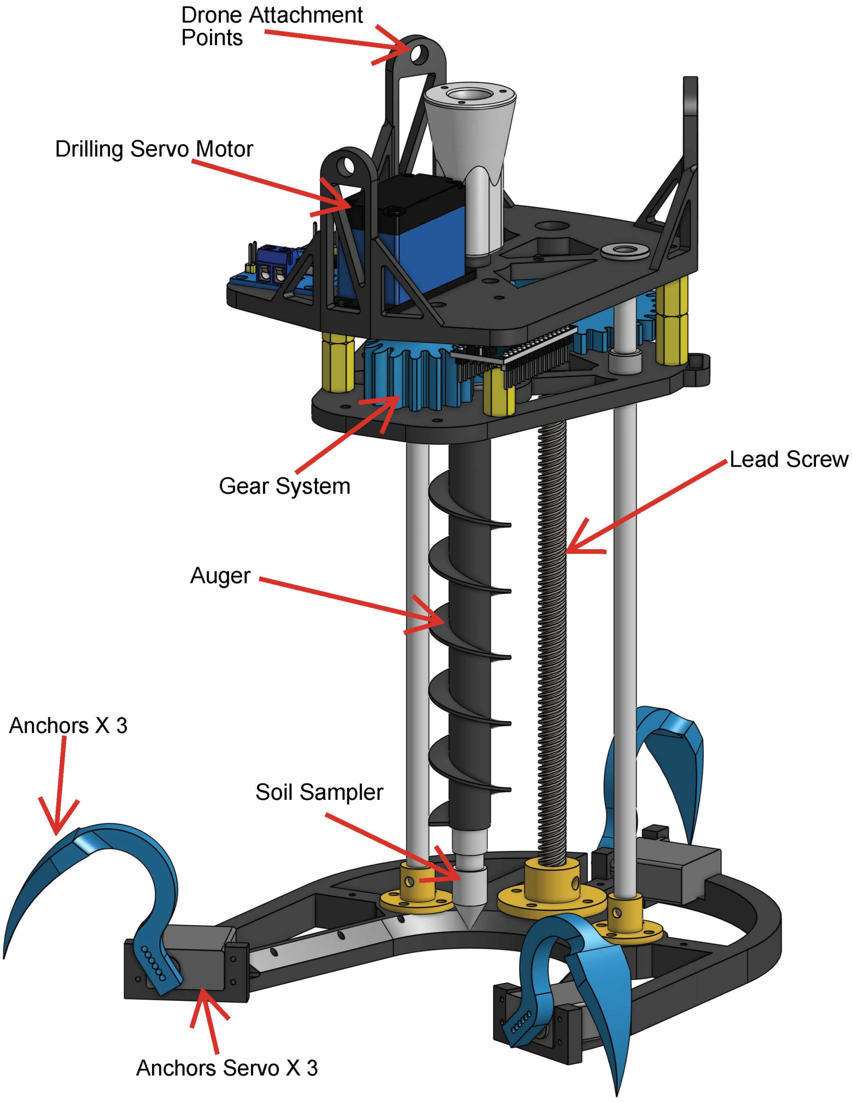

Figure 9: Soil sampling module with labeled components

primary gear interfaces with a secondary gear that rotates both the auger and the soil moisture sensor, enabling them to penetrate the ground effectively. To prevent wire tangling during rotation, the sensor is integrated with a slip ring on its top. The secondary gear subsequently drives a final gear containing an internal lead screw nut and as this gear rotates around a fixed lead screw, which translates the entire assembly linearly. Through iterative testing, a 1:1:2 gear ratio was established to synchronize the auger's rotational speed with its descent rate. The platform remains stable during travel by sliding along two lateral support rods with nylon bushings for smooth movement. To maintain a low weight while ensuring durability, the assembly uses an aluminum shaft and 3D-printed gears, powered by a 25kg/cm servo capable of penetrating varying soil densities.

Figure 10: (Left) The anchoring and stabilization system during operation. Three servo-driven anchor arms, each terminating in a curved digger hook, are arranged at near 120° intervals around the perimeter of the base frame. (Right) Fixed nails on the underside of the base frame to resist rotational motion during both anchor embedding and auger operation.

- **Onboard computing and communication:**
  In order to drive the sampler servos and read values from the sensors, the sampler is controlled by an Arduino Nano mounted to the top of the sampler. The entire sampling routine — driving the anchor servos, lowering the probe into the soil, reading moisture values, retracting the probe, and retracting the anchors — is executed from a script on the Arduino, which simplifies the communication between the drone and the Arduino to a simple set of commands and the logging of data. By default, the sampler script is set to an 'OFF' state where the servos are retracted. Upon receiving a specific command, the sampler either enters a 'PAUSED' state (waiting for specific commands to drive individual servos, used for debugging and tuning) or the start of the autonomous sampling routine, which begins with 'CLAW_EXTENDING'. The hardware diagram is in Figure 11. The sampler routine is as follows (see Fig. 12):
  - **Claw extension**: For a set time, extend the anchor servos to dig into the soil and hold the sampler in place.
  - **Main extension**: Drive the continuous servo forwards, extending the moisture probe into the soil. Check the ultrasonic sensor periodically to get the distance between the top and bottom sampler plates and calculate the travel of the moisture probe. If this travel is greater than the minimum specified, start sampling; otherwise, abort after a set time.
  - **Moisture sampling**: For a set time, read the moisture sensor pin to get multiple readings for the soil moisture.
  - **Main retraction**: Drive the continuous servo backwards, retracting the moisture probe from the soil. After the extension as reported by the ultrasonic sensor is past a threshold or a set time interval elapses, proceed to the next state.
  - **Claw retraction**: For a set time, retract the anchor servos from the soil. At the end,

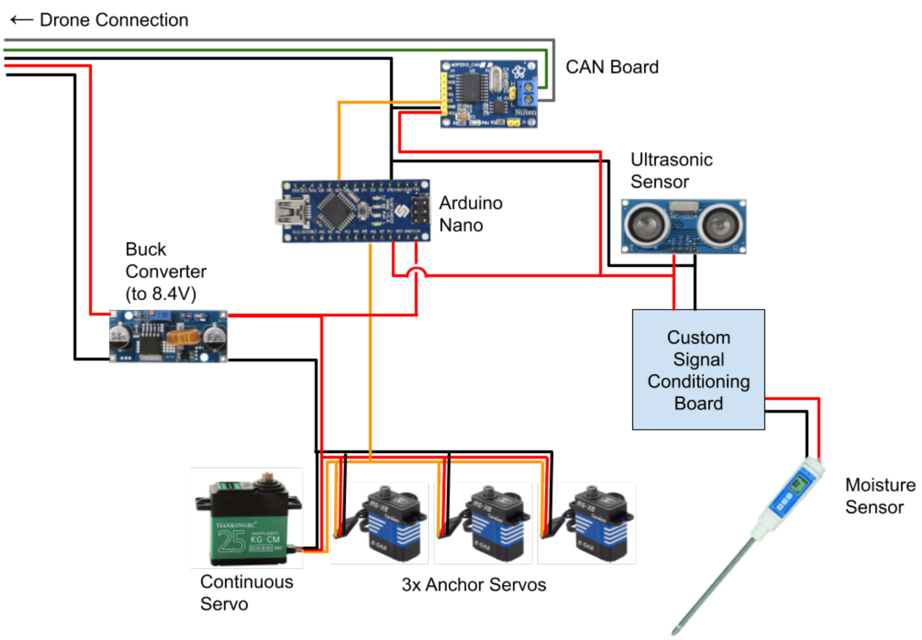

Figure 11: Hardware diagram of the soil sampler

if the sampling succeeded, transmit the recorded data to the drone computer. Otherwise, transmit a signal indicating sampling was unsuccessful.

Due to concerns of electrical noise from the motors affecting the communication between the drone computer and Arduino, we use a CAN bus for the communication. CAN consists of a twisted pair of two wires for data, CAN-high, and CAN-low, where signals are encoded as the difference between the two wire voltages. This negates the effect of most electrical noise, which would affect both lines equally, at the cost of requiring an adapter on each end to interpret the CAN signal. For this, we use the Waveshare USB-CAN-A adapter on the drone computer, which converts a serial connection on a USB device to a CAN signal, as well as the MCP2515 board, allowing us to interpret the CAN signal on the Arduino side.

The drone battery provides 14.8V, which is routed to the sampler and stepped down to 8.4V to allow the continuous servo and anchor servos to achieve maximum torque. This voltage can power the Arduino Nano (via the built-in linear regulator on the VIN pin), which internally steps it down to 5V. The 5V from the Arduino powers the MCP2515 board, the ultrasonic sensor, and the board that communicates with the moisture sensor. Figure 11 shows the full hardware diagram of the soil sampling module.

- **Soil moisture sensor integration and calibration:**
  To measure the moisture level of the soil we utilized the DSMM500 Precision Digital Soil Moisture Meter with Probe, manufactured by General Tools (Figure 14aa). This is a capacitive sensor which measures dielectric constant of the surrounding medium. Since water has a significantly higher dielectric constant (approximately 80) than dry soil (typically 3 to 5), introducing water increases the probe's overall capacitance.

  Hydrofleet requires a continuous data stream that could be processed and stored onboard our drone's flight computer. To minimize payload weight, we discarded the original housing and retained only the metal probe. To interface the raw probe with our microcontroller, we designed a custom signal conditioning circuit using an LMC555 CMOS Timer configured in astable mode.

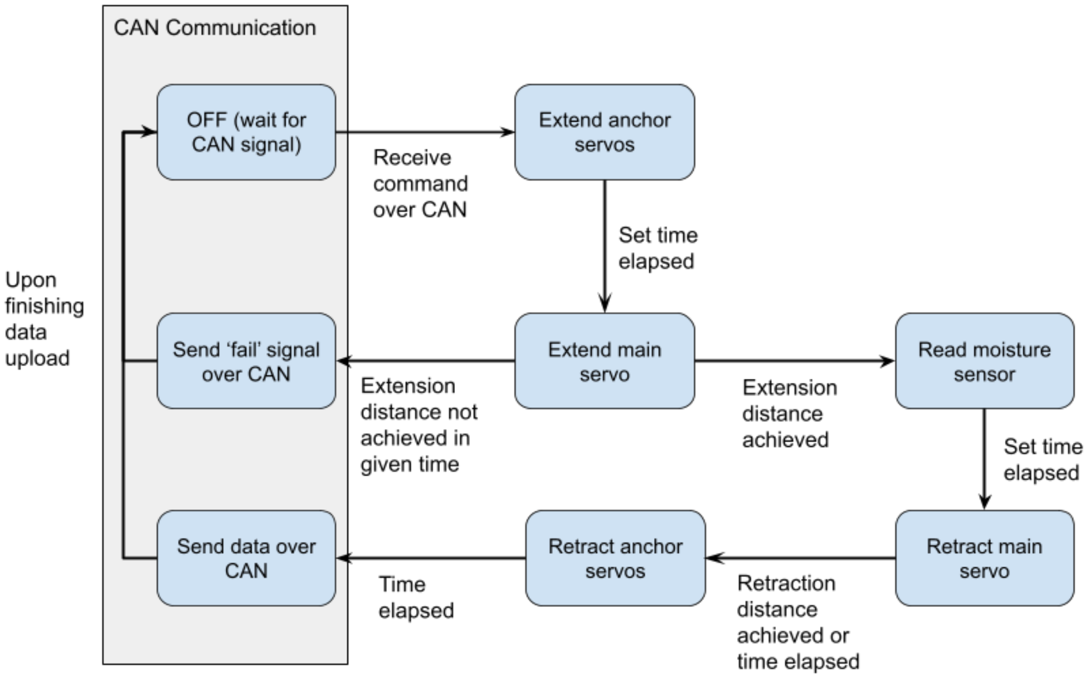

Figure 12: State diagram of the sampling routine

This circuit works by utilizing the soil probe as the primary capacitor within the timer's RC (resistor-capacitor) oscillator network. As the soil moisture increases, the probe's capacitance increases, which acts as an electrical dampener. This dampening effect alters the charging and discharging timing of the circuit, which inversely shifts the voltage of the output pulses. This setup allows us to translate a physical change in soil capacitance into a measurable analog voltage reading that the drone can easily log.

From this dry baseline, we added specific masses of water to achieve our target Gravimetric Water Content (GWC) ratios of 0%, 15%, 20%, 25%, 30%, and 50% (Figure 13). For each increment, the soil and water were mixed and packed to a uniform density, creating a homogeneous mixture. Both our custom astable circuit probe and an unmodified factory DSMM500 probe were inserted into the soil at each moisture level to record readings. By plotting the sensor's output voltage against the known GWC percentages, we generated a calibration curve (Figure 14b). Our data show a strong linear relationship between voltage and moisture percentage, with $R^2 = 0.9969$. Notably, the unmodified factory probe exhibited the same linear trend across the targeted percentages, which successfully validated that our custom signal conditioning circuit was functioning correctly and accurately reflecting the soil's true dielectric shift.

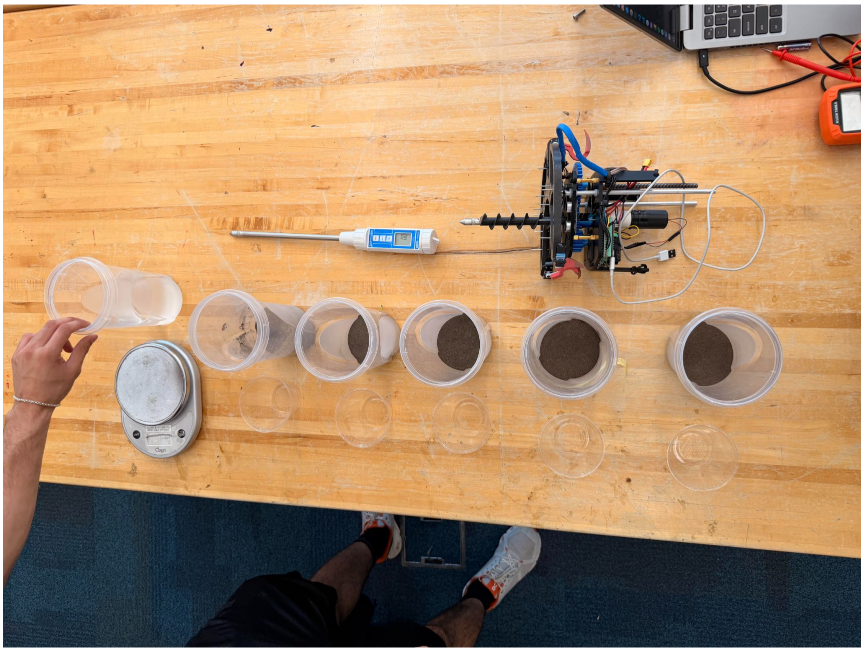

Figure 13: Soil moisture sensor calibration process using 6 calibrated points FROM 0% to 50% of gravimetric water content (GWC)

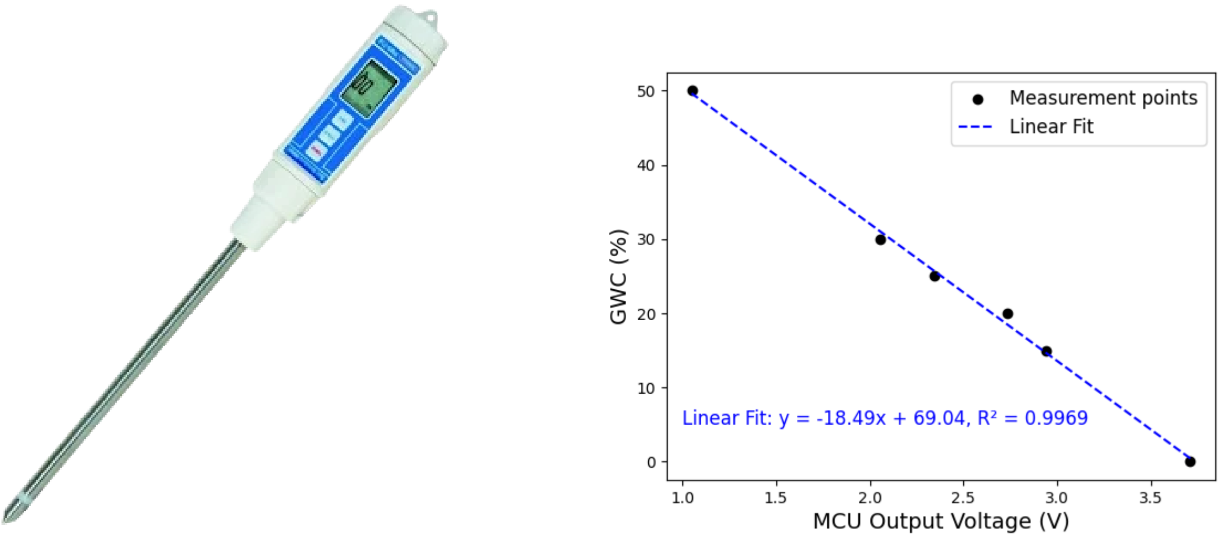

(a)

(b)

Figure 14: (a) DSMM500 Precision Digital Soil Moisture Meter (b) Sensor calibration curve

2.1.3 Multi-Drone Planner

In order to safely and efficiently fly a swarm of drones with payloads over any farm field, robust path planning is necessary to ensure efficient labor and time distribution. HydroFleet integrates a path planning framework that takes into account the bounds of the drone's flight endurance and the temporal and spatial constraints of any given farm field.

We define endurance by a fixed energy capacity that undergoes constant depletion during operation, precluding mid-mission recharging. Each mission entails a takeoff from a home depot, completion of a set number of probing points, and a return to the starting point before the battery depletes. The duty cycle is therefore limited by the energetic costs of both spatial transit and the hover duration required for soil probing.

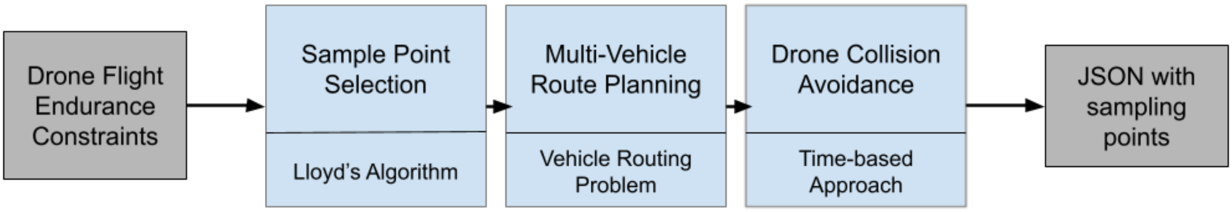

Figure 15: Software Architecture for determining optimal sampling points for drone system

We break this problem into four parts.

1. Defining the number of cells to sample from based on arguments like polygon (of the block of land) area, probing time, flight speed, total mission time for each drone based on energy constraints, and number of agents to cover the entire block
2. Generating uniformly-spaced probing points (Voronoi iterations)
3. Generating optimal paths between the probing points (Vehicle Routing Problem) based on the number of drones available
4. Adding time delays to routes that overlap to ensure there are no collisions

Finally these routes are exported as a json file that interfaces with the aforementioned ROS2 missions. Figure 15 shows the software architecture for the multi-drone planner.

Defining Number Of Cells:

In order to ensure the flight endurance for each drone, we calculated an optimal number of sampling points for a given polygon area and number of drones. For visualization, consider one of Powisset Farm's crop block and the 3 drones in our swarm (Figure 16a).

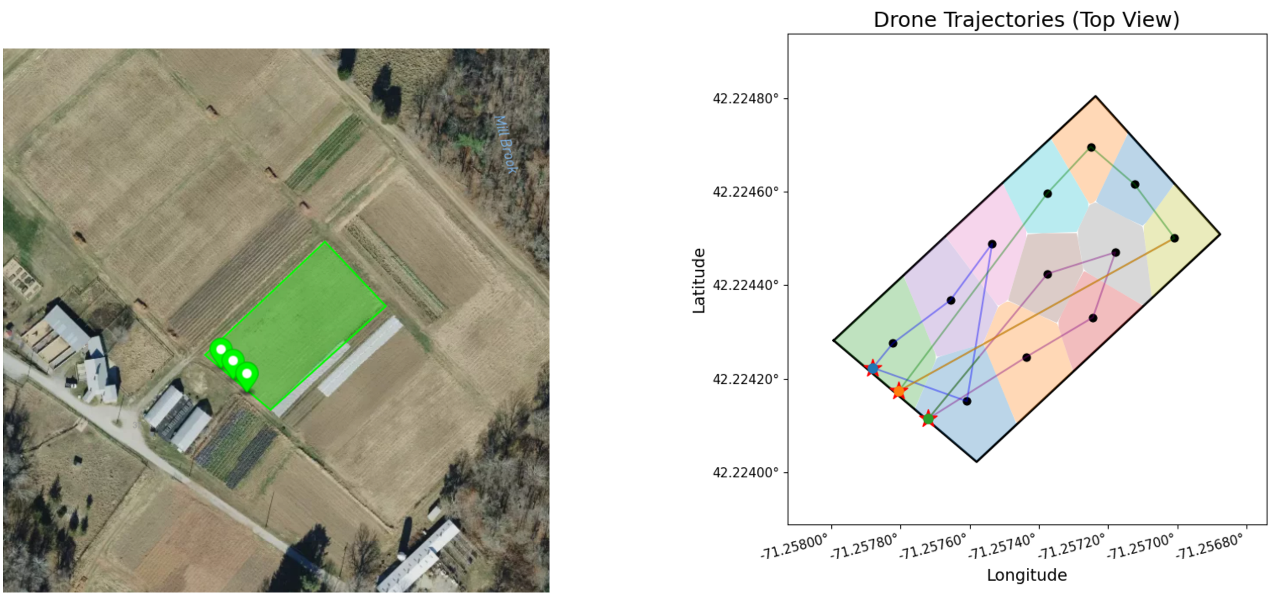

(a) (b)

Figure 16: (a) Image of Powisset Farm with selected crop block and depot points (b) VRP generated optimal paths for three drones on Powisset Farm

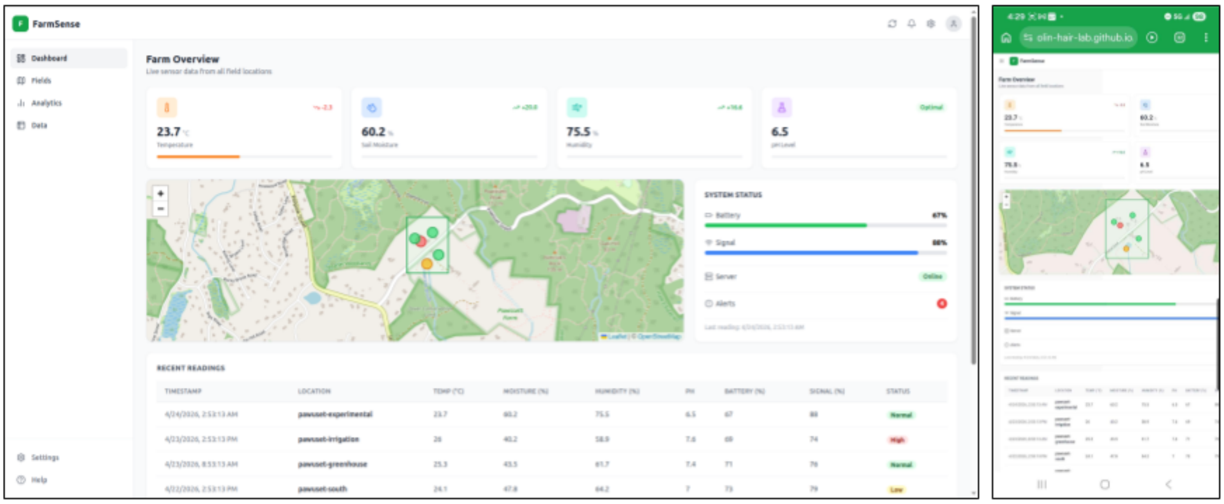

Figure 17: Farmer dashboard web-view (left) and mobile view (right)

The maximum possible time the drones can collectively spend probing can be defined as the sum of the cumulative flight time for each drone divided by the sampling time. The actual flight time for each drone is calculated based on the area of the crop block, thus accounting for the time spent traveling between sampling points. The number of sampling points is incremented and returned when the actual flight time is less than the maximum possible time and divisible by the number of drones. Using this method, we get 12 total sampling points across the Powisset crop block (Figure 16b.

Sample Point Selection using Lloyd's Algorithm:
Voronoi Iteration, also known as Lloyd's algorithm, divides a plane into regions where each region contains exactly one seed point, and every location within that region is closer to its seed than to any other. With each iteration, the seeds spread apart and converge toward a uniform, equidistant distribution across the space. This approach is well-suited for soil moisture probing for two reasons. First, the uniform distribution ensures that samples are spread evenly across the entire plot, giving the farmer a complete picture of soil conditions rather than data concentrated in one area. Second, because the points are approximately equidistant, the distance each drone needs to fly between stops is minimized, which makes the most of the available battery life.

Multi-Drone Route Planning using Vehicle Routing Problem:
The Vehicle Routing Problem (VRP) is an optimization method for finding the most efficient routes for a fleet of vehicles to visit a set of locations. In our case, it takes the Voronoi-generated sampling points and distributes them across the available drones, giving each drone an ordered sequence of waypoints that covers the plot within its battery limit.

2.1.4 Farmer Dashboard

The Farmer Dashboard serves as the primary interaction point between HydroFleet and the farmer. It was designed as a **mobile-friendly, intuitive user interface** that provides readings of measured values and plots trends across all current fields. This is presented in a user-friendly format that does not require the farmer to have a deep technical understanding of the process. These design decisions stemmed directly from conversation with farmers where they emphasized clear visuals and wanting to add custom sampling points. Our current main dashboard is shown in Figure 17 below. Within the main dashboard, the core features and information farmers need are displayed in a simplified format. We also developed the software to accommodate future additions, such as humidity, pH,

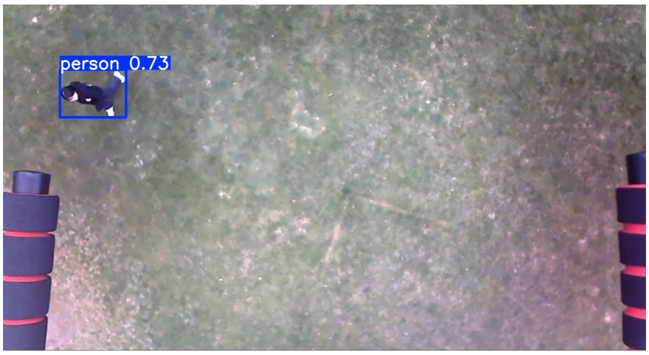

Figure 18: Person detection from drone's downward camera

and temperature measurements. The subpages then provide greater depth when needed, allowing the farmers to focus on a particular field, trends over time, or the raw data itself.

The farmer interface runs as a React app through Github Pages using a Python FastAPI backend hosted on Railway. Incoming drone and sensor data is posted to the API and stored in a PostgreSQL database via Supabase, and then received through the dashboard. The dashboard can be viewed through a mobile device, allowing for easy access outdoors.

2.2 Safety Features

When operating on the farm, HydroFleet must be aware of its surroundings, since farmers could be working around the crop beds where it operates. Especially with an additional suspended load at the bottom, it is crucial that those in the drone's vicinity are aware of it to prevent injuries. To address this, we implemented three safety measures: one on the drone, another on the ground station, and the third as a handheld system control override.

2.2.1 Person Awareness

The drone is equipped with a downward-facing camera that continuously monitors the area below it while running an existing deep learning model for person detection to check for humans. When a person is detected, the drone holds its position in place until the area below is clear (Figure 18).

2.2.2 Audio-Visual Safety Alerts

In addition to the person-awareness safety feature, the ground station has a safety light tower that provides both visual and audio alerts to those in the vicinity (Figure 19). When all drones are on the ground, a solid green light is activated to indicate that it is safe to approach the area. Upon activation of the drones, a flashing red light accompanied with the buzzer is used to indicate that it is dangerous to be nearby. This is triggered whenever the drone is taking off or landing.

2.2.3 Built-in Safety

The drone's flight controller also provides various built-in safety features to handle unexpected issues that may arise during flight. When the drone reaches critical battery levels mid-flight, an automatic return-to-launch (RTL) is triggered. Similarly, RTL is triggered when connection with the base station is lost. This is particularly helpful in a farm environment where connection can be unstable. Lastly, the drone can be manually overridden by a human pilot at any point mid-flight in case of an emergency (Figure 20.

Figure 19: Audio-visual safety alert at the system Ground Station using a safety lighthouse

Figure 20: Manual control override with controller

2.3 Performance Factors

The performance of HydroFleet is influenced by various factors that require careful consideration to optimize its effectiveness for farming operations of various sizes and contexts.

2.3.1 Weather Condition

For autonomous drone operations, weather conditions are a primary performance factor on the sampling mission's success and safety. In high-wind conditions, the oscillation of the slung load can amplify, inducing major instability in the flight controller and creating a severe safety hazard with a swinging mass. During heavy precipitation, the integrity of the onboard hardware is at risk, particularly the companion computer mounted at the top. While a custom protective enclosure can be mounted, it adds significant weight on a system already operating near its maximum takeoff weight, directly reducing flight endurance and mission range. Users must take these risks into account when determining the best conditions to activate HydroFleet. We recommend using HydroFleet during clear weather with minimal wind for most reliability.

2.3.2 Soil Condition

Another factor impacting HydroFleet's operations is soil condition, particularly density and moisture. Soil that is too dry, wet, loose, or compact can hinder machinery movement and complicate soil interaction. For the soil sampler, semi-compact soil is optimal, as it allows the anchors to fully penetrate while keeping the sampler stable during auger operation. Overly compact soil prevents full anchor penetration, causing the sampler to lift and become unstable, which reduces the moisture sensor's depth of travel and could cause it to fall over. Conversely, dry and brittle soil provides insufficient support, allowing anchors to be easily pulled out during digging.

The sampler performs well under the common farm conditions, though it is not designed for rarer extremes where soil moisture sensing would be unnecessary in any case. Users should weigh these trade-offs when deciding the best conditions for deploying the HydroFleet, keeping in mind that its ideal operation environment is semi-compact soil where anchors can fully penetrate and digging is unobstructed.

2.3.3 Crop Type

The success of HydroFleet's operation relies heavily on the robustness of its soil sampler landing sequence. In that regard, the types of crops, especially their height and spacing, can significantly impact the landing procedure. HydroFleet operates approximately over 6 meters above ground level, with the sampler hanging 2 meters below its frame. Upon arriving at a sampling point, the suspended sampler experiences a mild oscillation due to the momentum from flight before stabilizing. While this is acceptable for short crops, for crops like corn, which can grow up to 2 to 3.5 meters tall and are often planted in close spacing, the oscillation can cause the sequence to fail if the sampler gets tangled in the crop.

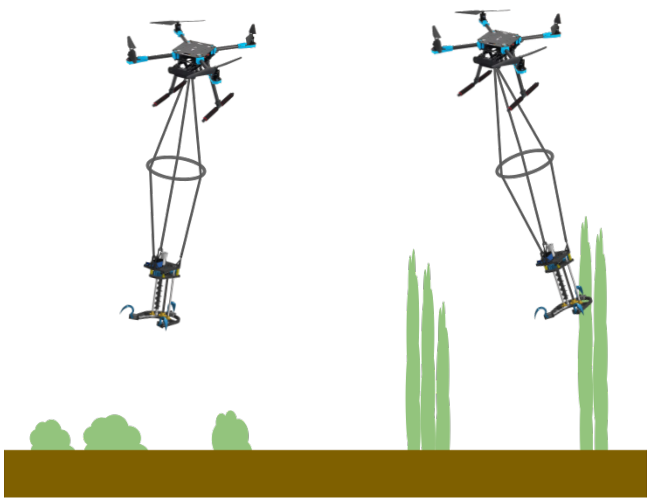

Figure 21: Demonstration of landing outcomes under different crop heights

2.3.4 Drone Flight Time and Farm Size

The energy budget, based on battery capacity, remains the largest operational constraint for drone-based systems and affects how our solution scales with farm size. For our current prototype drone platform build, we estimated a flight time of approx. 10 minutes per battery charge, with each sample sequence taking about 30 seconds to complete. While this flight time can be further optimized with more custom drone frame designs and larger battery capacities, energy constraints still remains an important limiting factor in large farm operations. Another important factor related to farm size is spatial resolution. This represents the number of samples to take per unit area. Higher spatial resolutions (e.g., 10-100 meters per sample, depending on field variability) yield higher-resolution moisture maps, but also require more sample points per unit area and more time for drones to complete.

We have considered several interesting solutions to the aforementioned constraints. First, using our multi-drone planner, we integrate the ability for the drones to autonomously return to their ground station to recharge or swap batteries while in operation. In practice, this would require fixed ground stations (with charging) at locations spread out around the operating field. The second solution we seek to implement is the integration of adaptive sampling algorithms into our system. Adaptive sampling improves efficiency in large fields by iteratively selecting new sampling locations that maximize information gain—targeting areas of high uncertainty or variability—based on the current state of the evolving soil moisture map. We plan to continue working in this direction during the summer.

2.4 Cost of System Design

Table 1 below shows the cost breakdown of a single drone plus sampling module in the fleet. Most parts were purchased off the shelf, but some were custom parts manufactured in-house. The soil sampling module and drone platform are calculated separately for readability. The Bill of Materials (BOM) for a single unit of HydroFleet comes to **$1,846.83**.

Item Name	Cost	Purpose
TD-8825MG Digital Servo 25kg	$19.99	Drives Actuation System
MEUS Racing Servo 8.5KG x3	$101.97	Anchoring Mechanism
200mm Tr8x8 Lead Screw and Nut	$8.19	Actuation System
DSMM500 Soil Moisture Meter	$221.64	Soil Moisture Sensing
Arduino Nano ATmega328P	$3.79	Sampler Microcontroller
HiLetgo Nano Terminal Board	$2.93	Nano Wiring
MCP2515 CAN Bus Module	$2.66	Drone - Sampler Comms
HiLetgo XL4015 Buck Converter	$3.16	Voltage Stepdown
HC-SR04 Ultrasonic Distance Sensor	$1.99	Depth Sensor
Arducam 2MP Color 50fps	$59.99	Sampler Camera
Misc. Hardware	$28.79	Connectors and Stock
Soil Sampling Module TOTAL	**$455.11**	
HolyBro X500 V2 Drone Frame	$260.99	Drone Frame
Companion Computer (ASUS NUC)	$289.78	Offboard Computer
PixHawk 6X Flight Controller	$309.99	Flight Control
DroneCAN F9P Rover GPS	$296.99	Noise-Free GPS
HolyBro Telemetry Radio	$58.99	Telemetry
FlySky RC Controller & Receiver	$64.99	Safety Manual Fallback
Arducam Global Shutter Camera	$59.99	Safety Person Detection
4S 4000mAh Lipo Battery	$50.00	Battery
Drone Platform TOTAL	**$1,391.72**	
System TOTAL	**$1,846.83**	

Table 1: Bill of Materials for Soil Sampling Module and Drone Platform

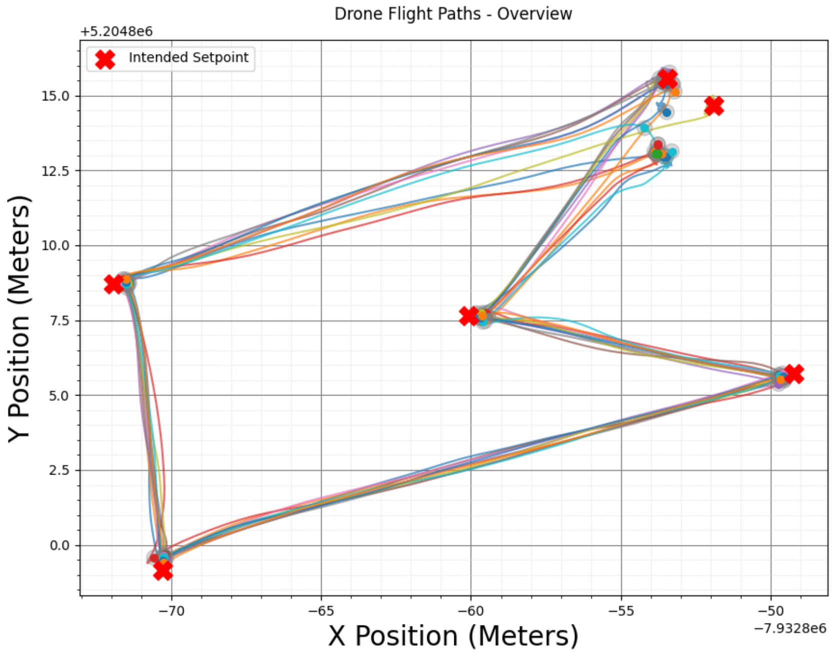

Figure 22: Results of 12 flights each visiting 4 waypoints. Top right point is the launch point and should be ignored.

3 Design Evaluation

In order to evaluate the performance of our system and design, four primary tests were conducted across navigation accuracy, soil measurement accuracy, sampler landing accuracy, and multi-drone performance analysis.

3.1 Navigation Accuracy

In order to sample reliably at target waypoints, the drone needs to have precise waypoint navigation. To test navigation accuracy, the drone was instructed to fly to a repeatable set of 4 waypoints for 12 flights. The flight logs were then saved and plotted against the ground truth waypoint locations.

For a quantitative measurement of the navigation accuracy, the **Circular Error Probable (CEP50)** was utilized. For each waypoint, we observed a consistent **CEP50** value of **0.43 meters**, meaning that half of the experimental landings concluded within this radius from the target center (Figure 22-23).

To further improve the system's navigation capabilities, future iterations will integrate a Real-Time Kinematic (RTK) base station to reduce GNSS-related drift.

3.2 Soil Moisture Measurement Accuracy

Since inserting the sensor using the drone differs from the hand-held method it was originally intended for, we sought to validate that our mechanically instrumented sampling module would not materially alter the sensor measurements. To test this, we took 10 measurements using our drone. For each drone sample, we then manually inserted the soil sampler by hand very close to, but not in contact with, the same piece of soil. We compared the results from both the sampling module and hand-held tests.

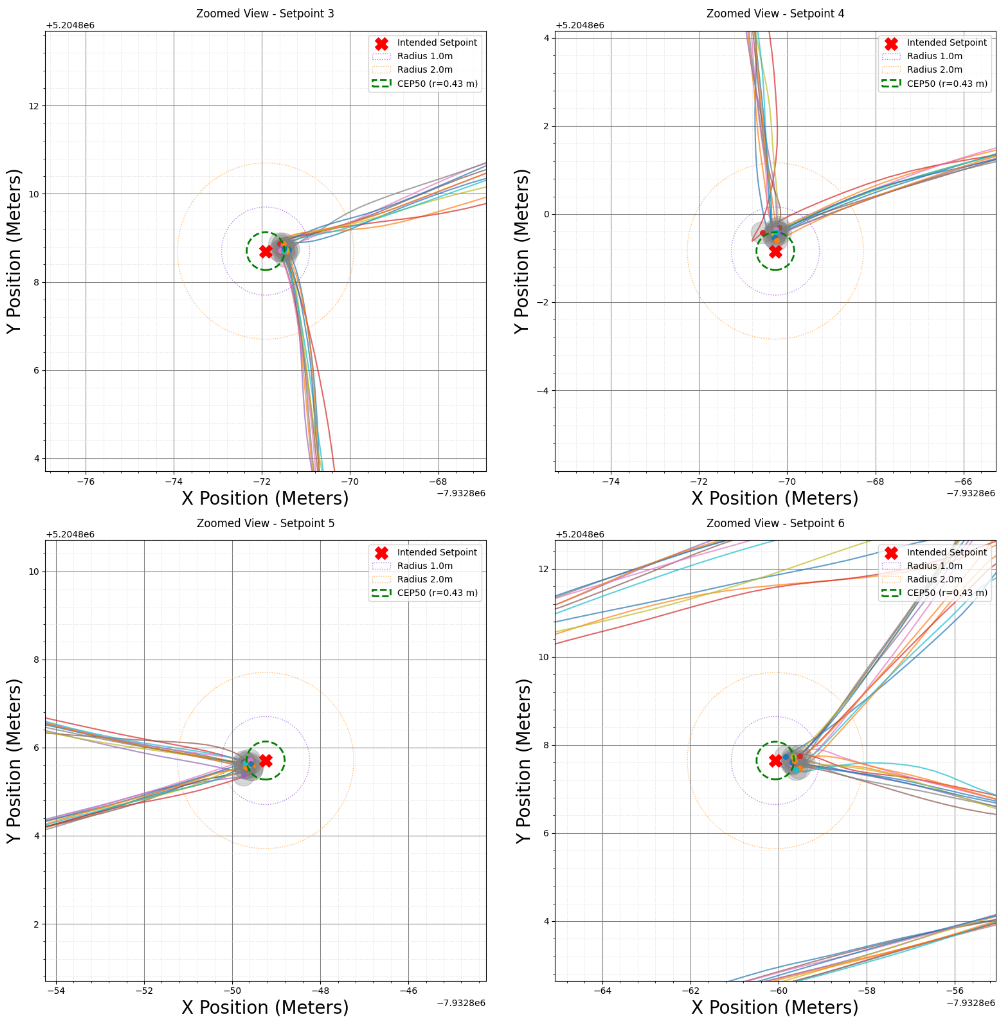

Figure 23: Zoomed in results at each waypoint for the navigation accuracy experiment

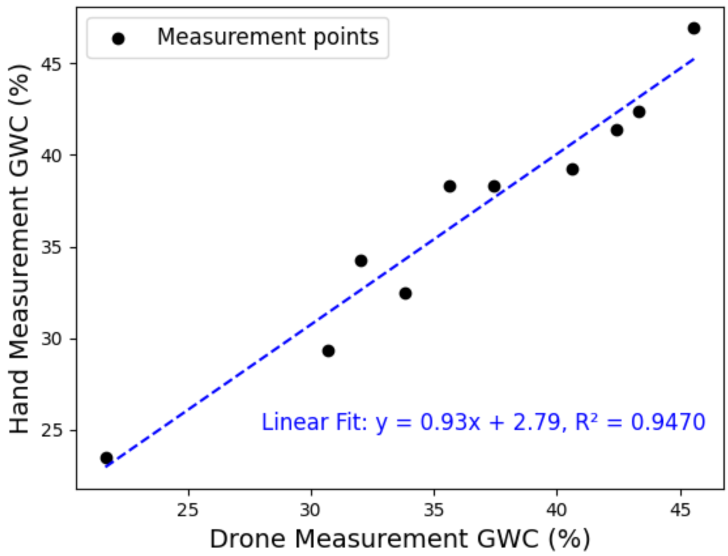

Figure 24: Soil moisture measurement accuracy comparing measurements obtained manually and using the HydroFleet sampling module

As shown in the Figure 24, the measurements from the drone-based sampling module and the hand-held sensor were essentially identical, with an R^2 **value of 0.947.** This is especially convincing when you consider the slight noise the sensor has on its own. This confirms that our drone measurements are valid when compared to the traditional hand measurements.

3.3 Sampler Landing Accuracy

Even if the drone's navigation is accurate, there is still a large, underactuated payload beneath the drone, subject to its own dynamics. To assess the accuracy of the sampler landing, the drone was flown over a known ground position marked with a grid and the precision of the sampler positions was logged.

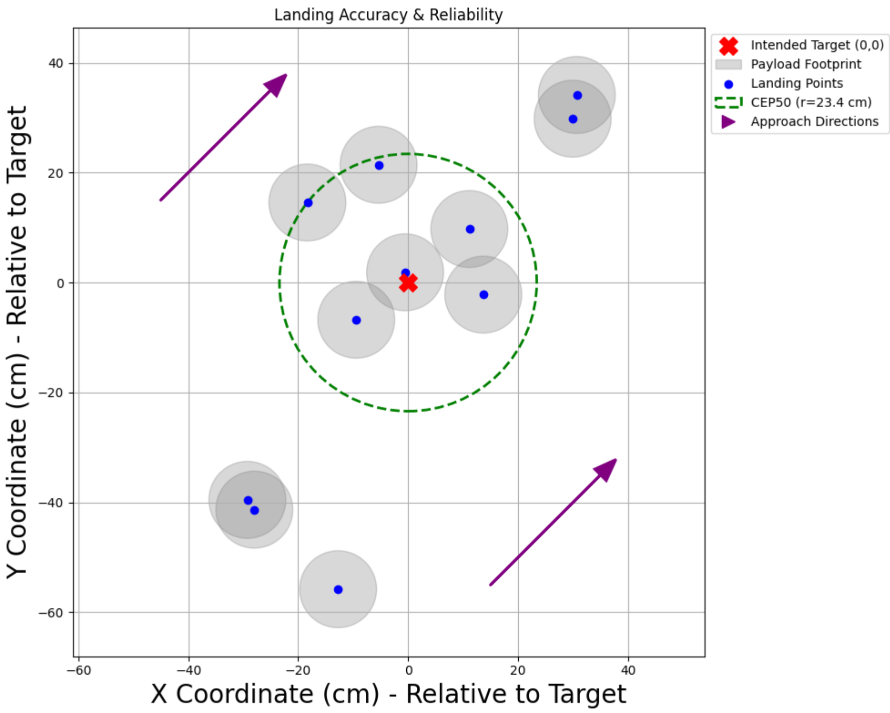

Figure 25: Results of sampler landing over 11 flights

The results showed a **CEP50 of 0.234 meters** (Fig. 25). While this establishes a strong baseline for autonomous operations, further refinements are necessary for high-precision operations to minimize crop damage. Future iterations will integrate a RTK base station and a vision-based landing algorithm to mitigate the impact of flight oscillations during the final descent phase.

3.4 Multi-Drone Performance Analysis

Ensuring efficient, collision-free route planning for each drone in the HydroFleet across the field is crucial to the successful deployment of this system. To evaluate our proposed multi-drone planning and coordinating subsystem, we ran a simulation using our testing farm (Powisset Farm) as a case study. We evaluated the mission completion time relative to farm size and the number of drones operating in the fleet. For our analysis, we assumed a fixed sampling spatial resolution of one sample per 10 meters (which is significantly higher than most farmers currently use). Other input parameters are: drone speed of 5 m/s between sampling points, a 30-second dwell period per sample, 10 minutes of flight endurance per charge, and an 8-minute battery swap time. It is important to note that this model assumes a battery swap station is always within negligible travel distance of the drone's last sampled point. Hence this analysis focuses on sampling and recharge time rather than transit home.

For a farm the size of Powisset which serves as the upper bound of the graph (Figure 26) at 109 acres (440,000 m²), the time savings between one and two drones is significant. However, the difference between four and five drones seems negligible, and a farmer may opt for four drones to be part of their total fleet. This demonstrates that this swarm system may be tuned to sampling density needs and practical time constraints.

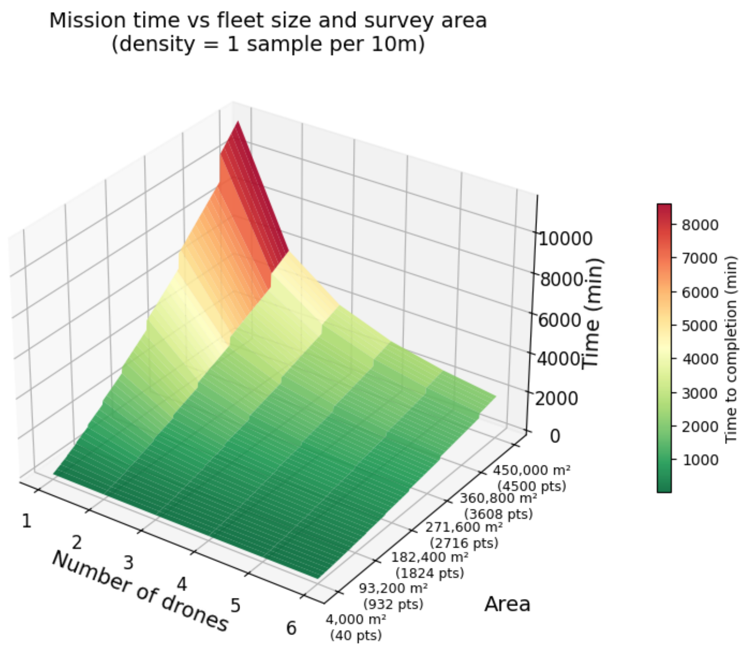

Figure 26: Mission time analysis based on farm area and fleet size

3.5 Summary Table of Evaluation Metrics

Metric	Description	Values/Results
Navigation Accuracy	Evaluation of the drone's ability to navigate to a global waypoint	**CEP50**: 0.43 meters
Measurement Accuracy	Evaluation of the sampler's measurement abilities compared to a manual measurement	R^2: 0.947 between sampler and manual measurement
Sampler Landing Accuracy	Evaluation of the drone's ability to land the sampler at a specified location	**CEP50**: 0.234 meters
Multi-Drone Performance	Evaluation of optimal fleet size for a small farm like Powisset	4 drones

Table 2: Summary Table of Evaluation Metrics

4 Design Story

4.1 Drone Platform

The success of the drone platform depended on its ability to reliably transport the soil sampling module across the farm and land it accurately. During our earlier integration tests, the sampler was connected to the drone at a single point, as shown in Figure 27a.

(a) (b)

Figure 27: (a) Single-point mounting (b) Multi-point mounting with bridle ring

This mounting position, however, transferred the rotation of the sampler directly to the drone, causing it to continuously rotate throughout its flight. To isolate the rotation of the sampler from the drone, we added a fishing swivel at the connection point between the drone and the sampler. In addition, we increased the number of mounting points from one to three for increased stability, and added a bridle ring in the middle to prevent twisting of the ropes (Fig. 27b).

For controlling the drone autonomously, we started with a ROS2 offboard control approach that was supported by PX4. While it was easy to integrate into the pre-existing drone system, we quickly realized the limitations of using offboard control for our task.

Traditional offboard control treats the flight controller as a "black box" that accepts external setpoints, but this often leads to a disconnect between high-level ROS2 logic and the drone's internal safety state machine. To address these gaps in reliability and integration, we pivoted to the PX4 ROS2 Interface Library. This modern approach allowed us to register distinct autonomous behaviors as "flight modes", allowing the flight controller to natively understand our custom logic. This also improved visibility on the ground control software, since it would display the name of the custom mode rather than just "offboard" (Fig. 28).

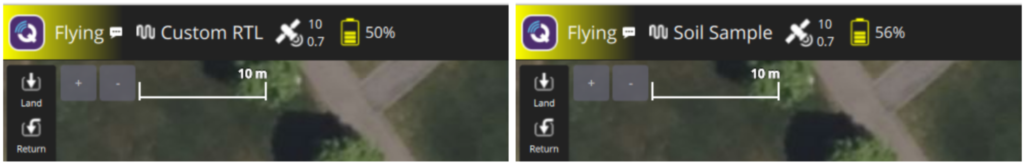

Figure 28: Custom mode names displayed on ground control software

4.2 Soil Sampler Module

For the initial phase of the soil sampler, we started by defining the problem. Our problem description was the following:

"Farmers need a reliable way to make informed decisions about irrigation. However, the current process is slow and inefficient, taking significant time and money to collect data. We want to create an attachment to the drone that takes soil moisture samples and makes consistent, reliable data collection at a farm field scale. Current solutions are either too expensive for the average farmer or don't get enough data to be useful."

After defining the problem description, we then listed all of the requirements and design constraints. The key requirements are listed below.

- Probe soil moisture across a depth of up to 6 inches
- A sensor able to insert and retract from the ground without breaking
- Within the drone payload capacity
- Probe must be able to sample without damaging plants
- Get precise moisture monitoring readings (margin of error of 5% on proper insertion)

With our requirements and design constraints, we then came up with 3 different ideas (see Fig. 29).

1. **Embedded Drill Probe:** In this design, the probe is embedded with the drill, where 1 motor turns the lead screw to lower the top base, while the second motor turns the drill itself. This makes the lead screw push the drill and the probe. With this, it will stop and measure every increment (1") to get a full soil moisture profile.
2. **Harpoon:** In this first harpoon idea, we use the spring's force to drive the harpoon into the ground. As the motor attached to the helical cam rotates, the harpoon structure will rise, compressing the spring. When it reaches the peak, the compressed spring will be uncompressed, sticking the harpoon inside the ground. In this section, the harpoon idea, we use a compressed spring meshed with a helical cam. As the motor attached to the helical cam rotates, the structure for the harpoon will go up, compressing the spring. When it reaches the peak, the compressed spring will be uncompressed, sticking the harpoon inside the ground.
3. **Sliding Hammer:** In this sliding hammer design, the hammer is released and hits the probe, which is connected to the rod and a wire. This will push it to go on the dirt. This will repeat itself until it goes fully down into the ground.

After our 3 ideas, our next steps were to downselect to one. We ended up choosing the idea of embedded probes. The main reason for this was the level of consistency that the other two designs didn't offer. In the harpoon idea, there were questions about the spring and whether it could effectively penetrate 6 inches of dirt. In the sliding-hammer idea, much of the force came from the drone, which

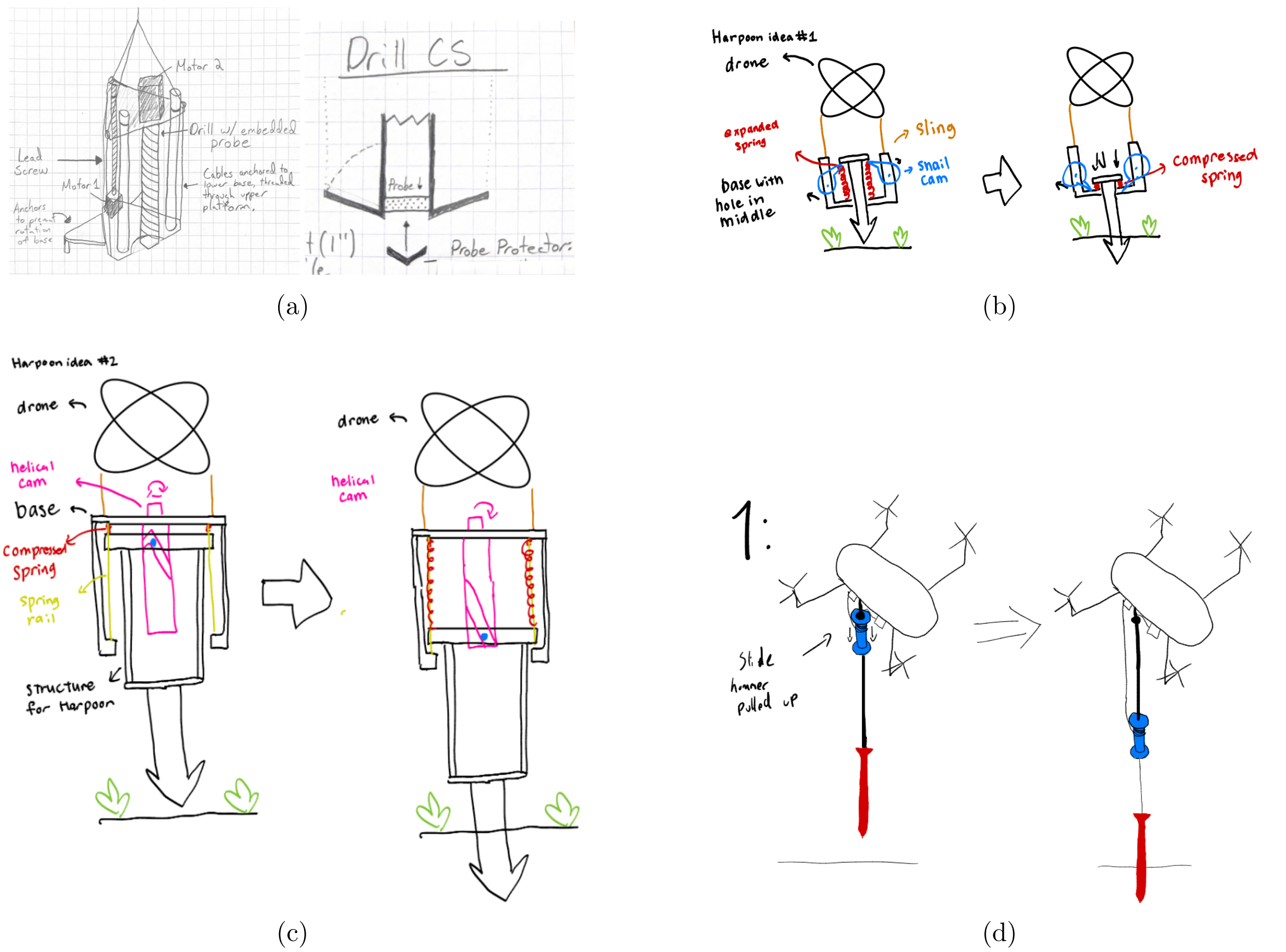

Figure 29: Design concept sketches. (a) Embedded drill probe (b) Harpoon with spring probe (c) Harpoon with helical cam probe (d) Sliding hammer probe

was unreliable for testing. This left us with the embedded probe idea, which provided a consistent method for testing and iterating.

Initial Physical Mockup

Before building a detailed CAD model, we created a simple physical mockup to test the basic architecture of the embedded probe concept. This low-fidelity prototype allowed us to quickly explore the two-plate frame layout, vertical support structure, and how the system would hang when suspended. It also helped us check whether the central opening and overall geometry made sense before investing more time into a refined mechanical design. After confirming that this overall concept was promising, we moved on to the first CAD-based iteration.

This was our first CAD iteration of the embedded probe idea. In this design, we modeled the probe with the lead screw on a stepper motor. The lead screw would be fixed with a driven lead nut that moves the top plate down. In the middle is the auger and our soil sampler. We decided to get the DSMM500 Precision Digital Soil Moisture Probe from General Tools because of its cylindrical shape and measurement accuracy/resolution of $\pm5\%/0.1\%$. During this iteration, we wanted the auger to be fabricated with steel for more reliability, and most other components to be 3D printed to minimize the weight and to allow for rapid iterations.

Our initial problem was how we wanted to drive the drill and the soil sampler. We had mechanical concepts of chains, belts, and pulleys, and gear systems. We ended up choosing the gear system

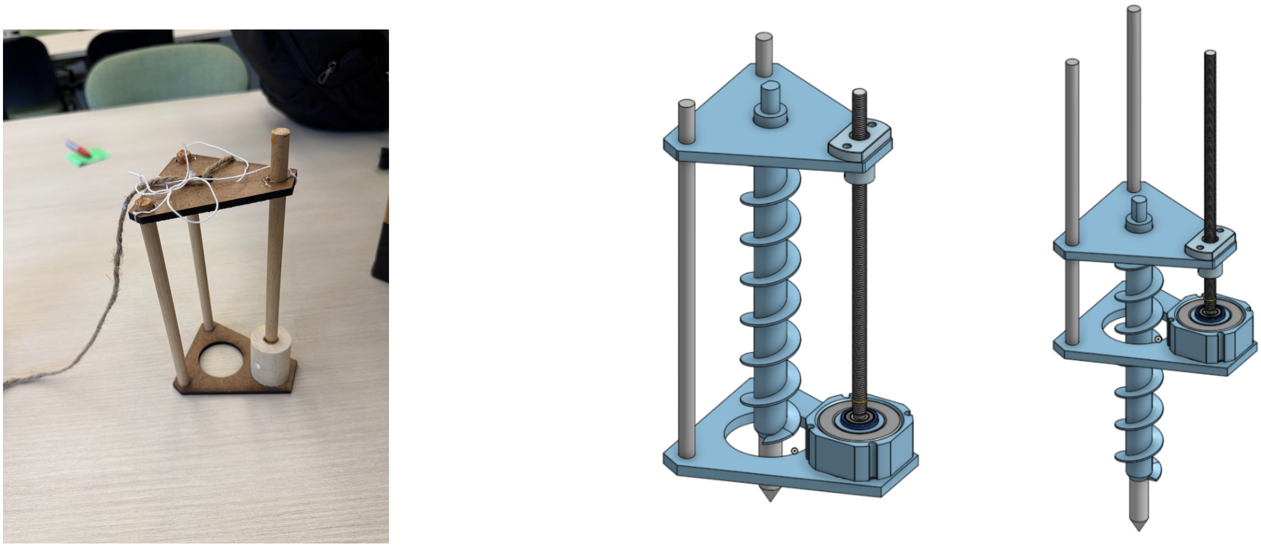

(a) (b)

Figure 30: (a) Initial physical prototype of soil sampler (b) CAD Iteration of Initial Design with Embedded Probe

because of the ease of changeability. The chain would be significantly too heavy, and the belt system would require exact ratios and a tensioning system. The gears, however, would be easy to change, lightweight because we could 3D-print them, and would not require a tensioning system.

Two motors vs. one motor: With the CAD architecture figured out, our next question was how we would drive the lead screw and the auger/soil sampler. We decided by prototyping two separate designs. Our first design was to drive the lead screw and the auger with separate motors.

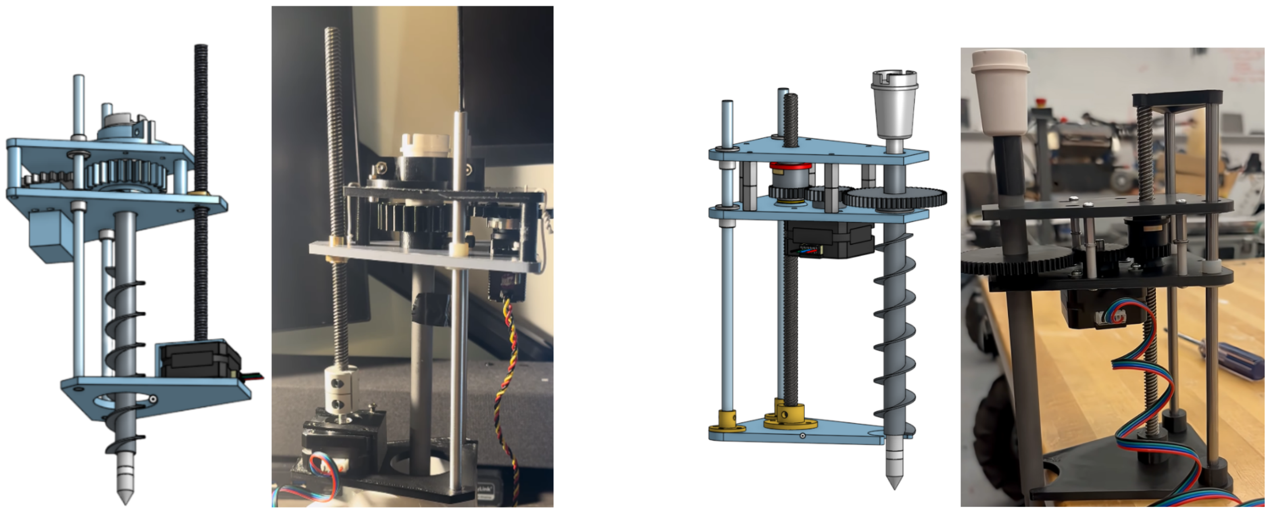

(a) (b)

Figure 31: (a) Two-motor soil sampler (b) One-motor soil sampler

In this design, the stepper motor drives the lead screw, while the servo motor drives the auger. The prototype weighed around 610 grams. The pros of this design were that it was simpler to test individual motors and allowed for more power to be driven for the auger. It was also less prone to single-point mechanical failure. However, this design was heavier because of the extra motor.

The second design was driving the lead screw and the auger with the same stepper motor. We actuated the lead screw by rotating the lead nut rather than the screw itself, which kept the screw fixed in place. The prototype weighed 626 grams (509 grams without the probe), with the increased weight primarily from the use of steel rods instead of the aluminium rods used in the initial design. One advantage of this design is that it theoretically reduces weight by requiring only a single motor. However, a key drawback is that it results in a higher center of mass, which can negatively affect stability.

After prototyping both designs, we decided to proceed with the second design, in which the stepper motor drives both the lead screw and the auger. We recognized that weight would be a major constraint, so saving it by removing the extra motor was the best step for us. This configuration became our Version 1 design.

Coupling the Sensor to the Auger: Early on, we planned to keep the soil moisture sensor stationary while the auger rotated around it. The thinking was that the sensor would just sit in the dirt, taking readings while the auger did the digging work around it. In practice, this didn't hold up. The friction between the rotating auger and the stationary sensor was higher than we expected, and, worse, placing the sensor in the middle of the auger's path made the auger less effective at clearing dirt. The auger was basically fighting the sensor instead of digging.

We pivoted to rotating the sensor along with the auger. To do this, we used a single set screw to pinch the gear, the auger, and the sensor together so that all three would spin as one unit. This fixed the digging performance immediately, but it created a new problem: the sensor's wires would now twist and tangle every time the assembly rotated. To solve that, we added a slip ring above the sensor, which allows the sensor body to rotate freely while keeping a continuous electrical connection to the wires running up to the Arduino.

Position Feedback for the Probe: Switching from a stepper motor to a continuous servo earlier in the design saved us weight and simplified our electronics, but it cost us position feedback. We considered two options to address this: a time-of-flight sensor and an ultrasonic sensor, both mounted on the upper plate to measure the distance to the lower plate as the assembly traveled down the lead screw. We tested both and went with the ultrasonic. The time-of-flight sensor was more precise on paper, but in field conditions, it was easily thrown off by leaves, grass blades, and other debris that could drift between the two plates during a sample. The ultrasonic was less affected by that kind of interference and gave us reliable readings in the actual environments where the sampler would be working.

Anchor Designs: We placed a strong emphasis on finding the right anchor design because we knew we had to get a firm grip on the ground to reliably penetrate the soil. We started out with the passive spikes design (Figure 32a). While it did help by preventing the soil sampler from twisting, it didn't prevent the sampler from lifting as it went down into the dirt. We then explored active anchors. Based on our testing, we decided to move forward with two separate designs: a 3-prong anchor (Figure 32c) and a thick anchor (Figure 32d). The idea behind these designs was to have a wider contact area when going into the dirt. The biggest problem was the anchor servo. The 9-gram servo, although light, was too weak for the anchors to reliably pierce the ground. We ended up switching to the MEUS Racing Servo 8.5KG. Although it was around 22 grams per servo, the performance was exponentially better than that of the 9-gram servo. It was able to pierce the dirt effectively. However, we still had the challenge of the anchors being too flimsy, unable to hold the ground effectively, and breaking as soon as they hit the dirt.

We solved the challenge by making the anchor mounting screws thicker, which helped with the breaking

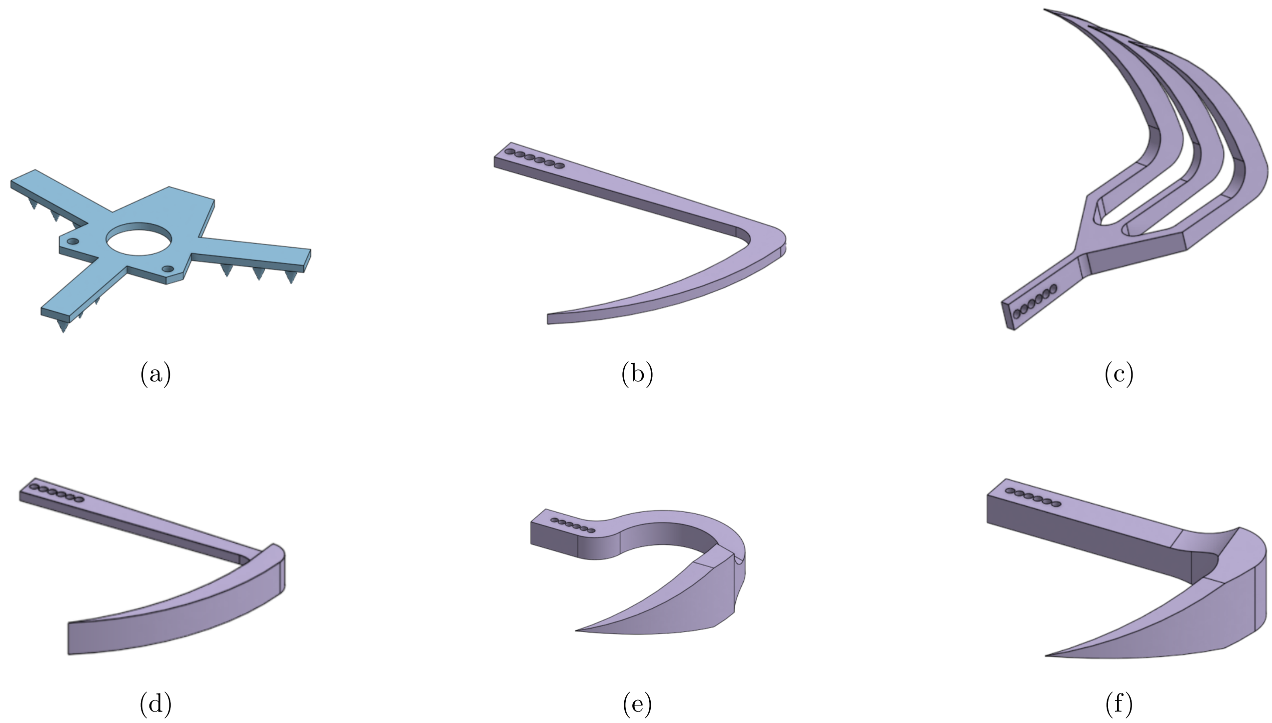

Figure 32: Anchor design iterations, in chronological order (a) Passive anchor (b) Initial design (c) Three-pronged anchor (d) Thicker anchor (e) Tapered anchor (f) Tapered circular anchor

and flimsiness issues. We also had the anchors' geometry features as a tapered profile that begins with a narrow tip and widens progressively. We decided to go with that to make the initial contact with the dirt at a singular point, making it easier to initially dig in and then widen for more contact with the dirt. Although this worked, it was stopped by the ground when the tapered profile ended. This made it more inconsistent and more prone to getting pulled out. To prevent that, we made the anchor circular (Figure 32f). With that, the taper profile continued down, where a circular arc helped the tapered part go deeper into the ground, so that it was perpendicular to the normal force when the auger is digging.

Final design:

A key change and driving factor for all of these changes was to reduce the weight while creating opportunities to be flat on the ground. A major change we made for our final design (Version 4) for the base (Figure 33 & 34). In the base, we pocketed everything to make it lighter. We also added small holes throughout the base to add spikes. Through our testing, this proved effective, as it provided good initial contact between the soil sampler and the soil before the anchors went in. Another way that we made the soil sampler weigh less was by pivoting away from the steel auger idea. As mentioned earlier, we initially wanted the steel auger idea because we thought it would be more effective for displacing the dirt. However, we realized that the 3D printed version was able to dig dirt effectively while not weighing as much as the steel auger.

The biggest single weight save in Version 4 came from the soil sensor itself. The sensor came with a long probe and a heavy housing that was carrying a lot of weight unnecessary for our use case. We cut the probe shorter, TIG-welded it back together, and rewired it to interface directly with our custom signal conditioning circuit (Fig. 35). This one change saved 47 grams. Combined with the lightened base, the 3D-printed auger, and the holes we cut throughout the structure, Version 4 ended up substantially lighter than Version 3 while still meeting our 5.25-inch sampling depth.

(a)

(b)

Figure 33: (a) Final design of soil sampler (b) Fabricated soil sampler

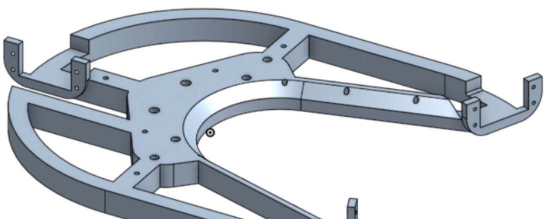

(a)

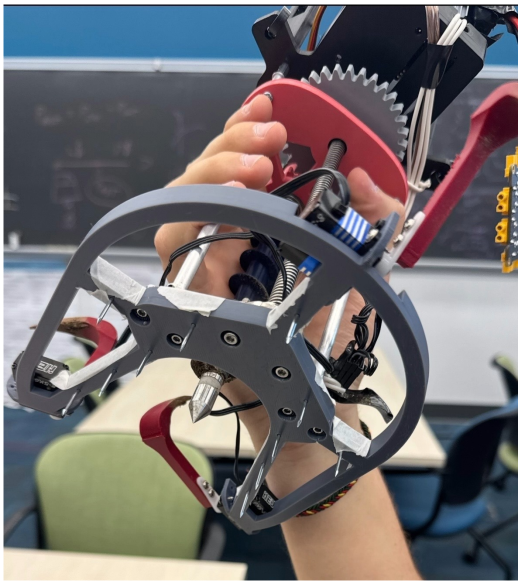

(b)

Figure 34: (a) Final design of base plate (b) Spikes on sampler base

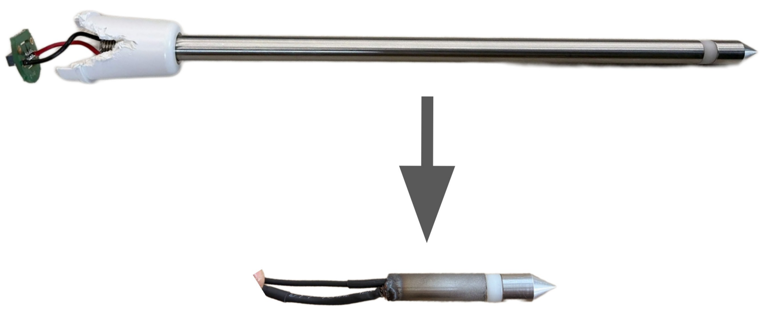

Figure 35: Custom lightened sensor probe

By Version 4, the sampler had converged to a significantly lighter and more stable design. The anchor geometry had been refined through several iterations to reliably grip the soil, the sensor had been mechanically coupled to the auger and rewired for weight savings, the ultrasonic sensor gave us reliable position feedback for the continuous servo, and the suspension and base improvements meant the sampler made better contact with the ground on landing. These iterations together let us balance penetration, stability, and payload constraints while preserving reliable soil moisture sampling.

Sampler System Integration

For the initial phase of this project, we didn't specifically designate 'integration' as a key subsystem of the project, and we therefore found that while both the sampler and the drone system were being developed nicely, there was little effort in how these subsystems before the end of the project. Thus, near the start of winter, we pivoted by specifically delegating 'soil sampler and sensor integration' as a discrete component of the project, and some of us pivoted to directly working on this integration.

The initial vision for the code that would run on the sampler was a straightforward wrapper over the functionality that was manually controlled by the sampler team — we would have functionality for driving the anchor servos and driving the main servo, and would combine these behaviors in a script to execute the full sampling routine. However, after drafting this code and testing the functionality of the sampler and communications, we realized that a blind extend–sample–retract loop would only be valid in ideal cases. In particular, we encountered cases where the quality of the soil would affect how deep the probe could extend into the soil, and envisioned cases where, perhaps due to a rock or a poor landing, the probe might not be able to penetrate the soil at all. In these cases, a naive sampling routine would end up sampling junk data, rather than sampling the depth of soil that would be valuable, and so we needed a way of detecting if the sampler had been able to properly extend.

To address this, we purchased an ultrasonic distance sensor, tested it to ensure that it could properly report on the distance between the top and bottom plates of the sampler, and integrated it into the sampler design. Mechanical integration of the ultrasonic sensor into the upper plate of the sampler was thankfully not challenging, and the wiring of the sensor was not difficult, making this a fairly simple in-place change to make. This did complicate the flow of the sampling routine, as requiring that the sampler extend far enough to sample added a second possible execution path, but drawing up a state diagram for the sampler made implementing this straightforward.

Our initial idea for the communication between the soil sampler and drone was a CAN bus, as we had concerns over noise, though we generally had little experience working on CAN communication. When first testing CAN as a proof of concept, we found that the MCP2515 board we were using was simple to set up electrically, thanks to numerous guides online, but we initially struggled to get communication working due to a misconfigured clock rate in the code.

4.3 Multi-Vehicle Planner

In typical robot sensor coverage path problems (CPP), lawnmower paths are generated as a systematic, parallel approach. In our initial algorithmic strategy, we used cell decomposition and grid CPP and dropping points along the generated path to create our sampling points (Fig. 36- 37). We also initially posited that spatially distributed drones would be necessary given the drones have a payload that could cause unexpected collisions.

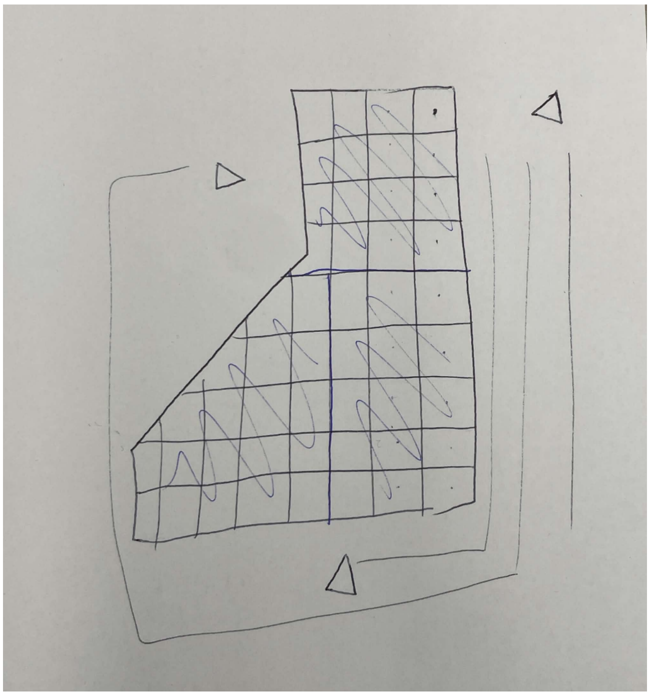

Figure 36: Initial sketch of multi-agent path planning

We realized that our approach of generating points based on a lawnmower path left gaps in the crop block unsampled, namely around the borders.

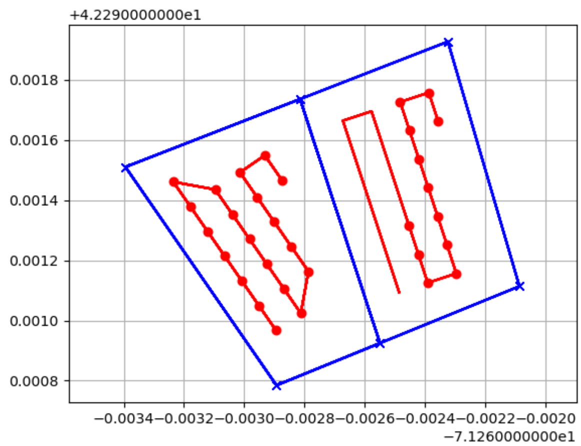

Figure 37: Initial simulation output with lawnmower paths and sampling gaps near borders

Additionally we weren't accounting for flight endurance in our initial simulations, so we pivoted to deciding the number and location of sampling points prior to generating paths. This led us to using the VRP algorithm to generate our routes in which we also learned that we could temporally distribute paths to ensure no collisions during takeoff or landing.

4.4 Summary Table of Design Story

Design Challenges	Solutions and Adaptations
Heavy tangling of the slung load	Increased mounting points with bridle ring.
Rotational transfer from slung load to drone	Fishing swivel for isolating rotation
Black box approach of traditional offboard control	PX4 ROS2 Interface Library for direct mode registering
Firm grip of the ground from the sampler	Iterated various anchor designs and integrated active & passive anchors
Tight weight constraints of the drone platform	Fabricated custom lightened sensor probe to cut 47 grams
Soil sampler may overextend and get damaged	Integrate an ultrasonic distance sensor as a failsafe
Noise and interference can affect signal communication over 2 meters between the drone and sampler	Integrate CAN communication between the drone and sampler module
Ensure maximal field coverage [under limited battery life] while ensuring no collisions	Optimize VRP with battery life and predetermined sampling points

Table 3: Summary of Implementation Challenges and Solutions

5 Impacts

5.1 Project Summary

HydroFleet is a novel multi-drone autonomous system designed to provide scalable, high-resolution soil moisture mapping to support precision irrigation through the integration of a specialized drone platform, an automated soil sampler, a multi-vehicle planner, and a user-friendly farmer dashboard. Our team has achieved **Technology Readiness Level (TRL) 5-6** (prototype development and demonstration in relevant environment) and, through testing, have validated core aspects of our system design. The drone platform achieves high-precision navigation with a **Circular Error Probable (CEP50) of 0.43 meters** during transit and **0.234 meters** during landing, while the soil sampler maintains exceptional fidelity with an R^2 **value of 0.947** compared to manual ground-truth data. The multi-vehicle planner computes optimized waypoints for each drone that is collision-free. Lastly, the farmer dashboard provides a central platform for the farmers to easily monitor moisture measurements. Beyond technical performance, HydroFleet offers a scalable and sustainable solution that can help farms increase crop yields (up to %15) while reducing water and energy costs (up to %40) and improving overall soil health. Due to the novelty of certain aspects of this system, we have, in collaboration with our MIT partners, completed filing a **provisional patent application** with the US Patent and Trademark Office. Moving forward, we are committed to continuing to develop this system, progressing along the TRL through additional field tests and design iterations over the summer and beyond.

5.2 Commercial Potential

5.2.1 Target Customer Base

HydroFleet is designed specifically to serve **mid to large-scale (500+ acre) pressurized-irrigation farms** in water-scarce regions, where irrigation is both a major cost driver and key determining factor for crop yields. These kinds of farms are increasingly adopting variable-rate irrigation (VRI) as a solution, but remain constrained in how much value they can extract from these systems due to insufficient soil moisture data. *(As an anecdote, Clark, our large-scale Ag advisor, noted that while he owns and operates capable centre-pivot irrigation systems in his farm, he is unable to maximise them due to sparse and limited soil moisture data).*

We believe early adopters of this solution will be farmers operating irrigation-dependent farms who have some prior exposure to precision agriculture technology (e.g., utilizing yield and soil maps, auto guidance and precision weeding systems). Our system is particularly well-suited to large-scale fields where existing soil-monitoring approaches are insufficient to provide the necessary sub-field insights for VRI decision-making.

5.2.2 Market Size

According to the 2022 USDA census, The **total irrigated agricultural farm area** in the United States is **53 million acres** [U.S. Department of Agriculture, 2022]. Considering a service price range of **$5-$20 per acre** (this range was advised by our Ag advisor, benchmarking with existing solutions), we use a mid-value of $15 per acre for our analysis. This yields an estimated **Total Addressable Market (TAM) of approximately $800 million annually** for soil moisture mapping services in irrigated farms in the US alone.

We believe that this estimate is conservative. The benchmark pricing reflects typical row crop systems (e.g., corn and soybean), where margins and willingness to pay are relatively lower. In higher-value specialty crops (e.g., fruits, vegetables, and permanent crops), precision irrigation delivers greater economic returns, and thus supports a significantly higher willingness to pay. Additionally, HydroFleet's platform has strong expansion potential: the same infrastructure can be extended to include **soil nutrient measurement and mapping**, broadening the TAM beyond irrigation into the larger precision agronomy and overall input optimization markets.

To estimate the **Serviceable Available Market (SAM)**, we focus on farms most likely to adopt advanced precision technologies. A USDA report indicates that **68% of large-scale farms** and **55% of mid-size farms** already utilize precision agriculture tools, including soil mapping, yield monitoring, and variable-rate technologies [Lim et al., 2024]. Restricting our analysis to these segments and cross-referencing with the total number of irrigated farms, we estimate that approximately **27.6 million acres** (21 million large-scale and 6.6 million mid-size) are strong near-term candidates for HydroFleet deployment. Applying the same pricing assumption yields a **near-term SAM of approximately $415 million annually**. We expect this serviceable market to grow as the technology continues to mature and adoption becomes more widespread.

We determine our **Serviceable Obtainable Market (SOM) of 5% of SAM** by setting a realistic capture percentage of **5% of SAM**, translating to **$20 million annually** in revenue. Further market research and assessment will inform a more detailed bottom-up estimate of our reasonably obtainable market. Figure 38 shows the market sizing and assessment for our project.

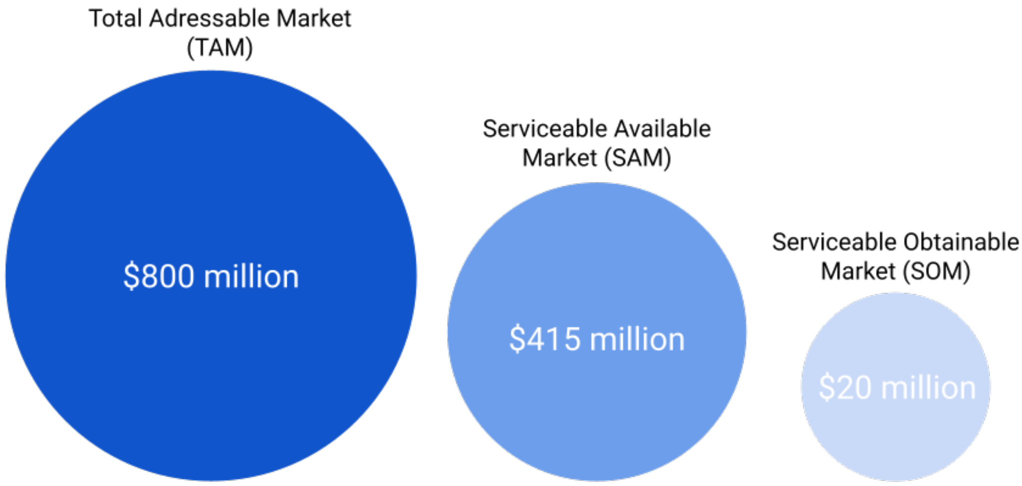

Figure 38: Preliminary market sizing and assessment (using TAM-SAM-SOM model) to determine commercial potential of HydroFleet

5.2.3 Value Proposition and Market Differentiation

HydroFleet offers value to farmers by closing the gap between precision irrigation equipment and the information needed to effectively utilize them by providing sub-field, high-resolution soil moisture data, enabling farms to increase yields while reducing operating costs. Our system extends the capabilities of existing systems by enabling scaling more efficiently to large farm areas, including non-contiguous fields.

The primary economic value drivers are: (1) **Crop yield optimization** (up to 15% in some crops) due to effective irrigation throughout the growth stages of the crop, (2) **Water and energy cost savings** (up to 40%) due to reduction in over-irrigation and more efficient use of water and energy resources [Edwards et al., 2025, Liakos et al., 2017].

As highlighted earlier, existing technologies for soil moisture measurement fall into two main categories. (1) **Stationary in-situ sensors** (e.g., capacitive sensors, time-domain reflectometry) are easy to deploy and provide reliable point measurements; however, achieving high-resolution sub-field mapping requires dense sensor networks that are prohibitively expensive and maintenance-intensive for most farms [Gasch et al., 2017]. (2) **Remote sensing approaches**, including satellite and drone-mounted ground-penetrating radar (GPR) systems, enable measurements without physical contact with the soil and scale well over large fields. However, they suffer from either insufficient spatial resolution for effective sub-field decision making or are unable to operate under vegetation cover, significantly degrading their effectiveness [Zhang et al., 2021]. This is where **HydroFleet** comes in, enabling a fully modular and scalable system to generate high-resolution, robust soil moisture maps at the sub-field level needed for precision variable-rate irrigation.

5.2.4 Production Costs and Business Model

Following the Bill of Materials (BOM) outlined in Section 2.4, the production cost for a single drone + sampler module unit is around **$1,900**. This does not take into account the engineering labor costs as we are not able to estimate that at this current time. Therefore, a HydroFleet system with 3 units, as we have showcased in this project, would cost about **$5,700**. Based on feedback from Clark, our Ag expert, the reasonable pricing range for these soil moisture measurements on his farm is $5-$20 per acre. We would love to conduct further market discovery to better understand the willingness to pay for such a service across a variety of farm operations. While still in the preliminary phase of exploring the business and commercial frameworks for this system, we anticipate that a **system-as-a-service model** may better align with the unit economics and support operational efficiency.

5.3 Solution Impacts

5.3.1 Impact on Workforce

Soil moisture sampling is a **physically demanding** and **repetitive task** in modern farm operations. As noted earlier, over 70% of farms still rely on some form of hand-sampling [Zhang et al., 2021, Evans et al., 2013]. Farm workers traverse large fields, manually probing the ground at predetermined intervals. Automating this process redirects labor toward higher-value activities — crop scouting, equipment maintenance, and agronomic decision-making — rather than the rote mechanical work of pushing a probe into the ground hundreds of times a day which some farmers still do primarily. The time savings compound across the growing season, particularly when sampling frequency increases to support precision agriculture practices that would otherwise be impractical with manual labor alone. Beyond efficiency, autonomous soil moisture sampling addresses a meaningful safety concern. Field workers regularly encounter hazards ranging from venomous snakes and aggressive livestock to ticks, fire ants, and wild boars, depending on the region. Deploying an autonomous system, such as the HydroFleet, in place of a person reduces the frequency of human exposure to these threats. This is especially relevant in pasture and rangeland contexts, where soil monitoring may be required in areas actively occupied by cattle or other large animals capable of causing serious injury.

5.3.2 Environmental & Agronomic Impacts

The two most direct environmental benefits of our proposed system are **reduced water waste** and **improved soil health**. Agricultural irrigation accounts for approximately **70% of global freshwater withdrawals** [Food and Agriculture Organization of the United Nations, 2021] (Fig. 39). Continuous, spatially resolved moisture data enables variable-rate irrigation, in which water is applied only where and when soil conditions actually require it rather than uniformly across an entire field. Under conventional irrigation schedules, significant portions of a field receive water that has already been adequately supplied, contributing to runoff, leaching of fertilizers into groundwater, and unnecessary draw on increasingly stressed aquifers. Tying irrigation decisions to real measurements rather than calendar-based schedules or worst-case assumptions substantially reduces these losses. The same data supports better soil management. Frequent, granular measurements help identify compaction, drainage problems, and moisture-stressed zones early enough to act on them, rather than discovering issues only after yield has already been affected. Over time, this improves the soil's resilience: better-managed fields retain more organic matter, hold water more effectively, and require fewer corrective inputs. The suspended sensor system is also extensible, broadening its potential environmental impact beyond moisture sensing alone. The same deployment mechanism can carry probes for pH, nitrate, and other soil chemistry parameters, allowing growers to fine-tune fertilizer application in the same way irrigation can be tuned to moisture. With modest modifications, the platform could be applied to use cases beyond agriculture, including geotechnical surveying, environmental monitoring of contaminated or remote sites, and exploration of terrain that is unsafe or impractical for humans to access directly. Each extension preserves the core advantage of the slung design: decoupling a low-cost, replaceable sensor from a more expensive carrier vehicle, while applying it to problems where ground-level data is currently expensive or sparse.

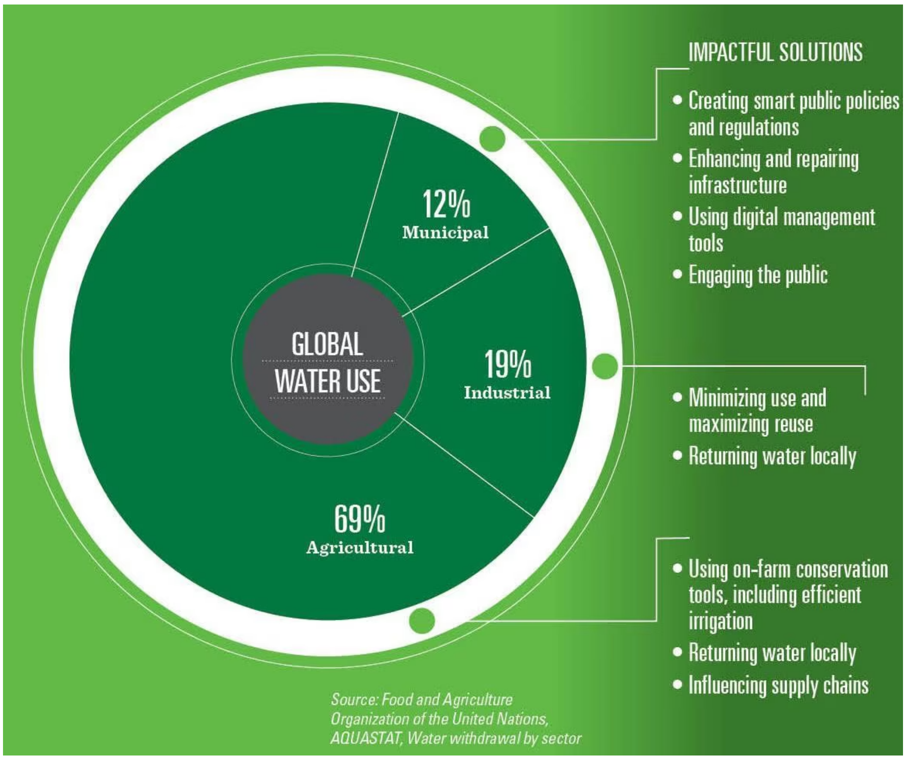

Figure 39: Drivers of global freshwater water use and how to meaningfully reduce them. Image source: [World Economic Forum, 2018]

6 Acknowledgments

The success of HydroFleet was made possible through the support of our Ag partners and collaborators. We extend our gratitude to Tim Laird and Aubrey Dority of Powisett Farm, as well as Clark McPheeters of CBM Farms, for providing invaluable field-specific insights and farm access that supported our development under real-world requirements. We also extend our gratitude to Harvey Merton from the BioIntrumentation Lab at Massachusetts Institute of Technology for mentoring all our undergraduate students. We also recognize the significant contributions of our student colleagues at Olin College: Pauline Petersen, whose expertise in videography was vital to creating our competition video, and Chang Jun Park, for his prior work on the farmer dashboard. Finally, we thank Lauren Hayes, the administrative assistant at Olin College, for her support in the purchasing process, ensuring the timely acquisition of all parts.

7 Team Information

- Dexter Friis-Hecht (**Team Co-lead**, Electrical Engineering Senior)
- Dokyun Kim (**Team Co-lead**, Engineering: Robotics Senior)
- Swasti Jain (Engineering: Robotics Senior)
- Benjamin Ricket (Engineering: Robotics Sophmore)
- Cian Linehan (Mechanical Engineering Junior)
- Heesung Han (Mechanical Engineering Sophomore)
- Zaraius Bilimoria (Engineering: Robotics Sophmore)
- Kilan Rougeot (Mechanical Engineering Sophomore)
- Grant Rechtin (Engineering: Robotics Sophmore)
- Pia Swarup (Mechanical Engineering Freshman)
- Mehmet Firat (Mechanical Engineering Sophomore)
- Harvey Merton (MIT Graduate student collaborator/mentor)
- Kenechukwu C. Mbanisi (**Project advisor**)

Figure 40: Team Photo with a few team members not pictured: Pia Swarup, Cian Linehan, Zaraius Bilimoria and Harvey Merton

References

[Edwards et al., 2025] Edwards, P., Vellidis, G., Porter, W., Mallard, J., Hall, D., and Bedwell, E. K. (2025). More crop per drop: An introduction to the cropfit app.

[Evans et al., 2013] Evans, R. G., LaRue, J., Stone, K. C., and King, B. A. (2013). Adoption of site-specific variable rate sprinkler irrigation systems. *Irrigation science*, 31(4):871–887.

[Food and Agriculture Organization of the United Nations, 2021] Food and Agriculture Organization of the United Nations (2021). The state of the world's land and water resources for food and agriculture – systems at breaking point (solaw 2021).

[Gasch et al., 2017] Gasch, C., Brown, D., Campbell, C., Cobos, D., Brooks, E., Chahal, M., and Poggio, M. (2017). A field-scale sensor network data set for monitoring and modeling the spatial and temporal variation of soil water content in a dryland agricultural field. *Water Resources Research*, 53(12):10878–10887.

[Kashyap and Kumar, 2021] Kashyap, B. and Kumar, R. (2021). Sensing methodologies in agriculture for soil moisture and nutrient monitoring. *IEEe Access*, 9:14095–14121.

[Koenig and Howard, 2004] Koenig, N. and Howard, A. (2004). Design and use paradigms for gazebo, an open-source multi-robot simulator. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2149–2154.

[Liakos et al., 2017] Liakos, V., Porter, W., Liang, X., Tucker, M., McLendon, A., and Vellidis, G. (2017). Dynamic variable rate irrigation–a tool for greatly improving water use efficiency. *Advances in Animal Biosciences*, 8(2):557–563.

[Lim et al., 2024] Lim, K., McFadden, J., Miller, N., and Lacy, K. (2024). America's farms and ranches at a glance: 2024 edition.

[Macenski et al., 2022] Macenski, S., Foote, T., Gerkey, B., Lalancette, C., and Woodall, W. (2022). Robot operating system 2: Design, architecture, and uses in the wild. *Science Robotics*, 7(66):eabm6074.

[Meier et al., 2026] Meier, Lorenz, and The PX4 Contributors (2026). PX4 Autopilot v1.17.0.

[U.S. Department of Agriculture, 2022] U.S. Department of Agriculture (2022). 2022 census of agriculture. Accessed: 2026-04-29.

[World Economic Forum, 2018] World Economic Forum (2018). Companies can help solve water scarcity. here's how.

[Zhang et al., 2021] Zhang, J., Guan, K., Peng, B., Jiang, C., Zhou, W., Yang, Y., Pan, M., Franz, T. E., Heeren, D. M., Rudnick, D. R., et al. (2021). Challenges and opportunities in precision irrigation decision-support systems for center pivots. *Environmental Research Letters*, 16(5):053003.