

H.A.L.O.

Hall-effect Array for Live Observation

Magnetic Field Imaging

EEL4914 - Senior Design 1
Group 18



UCF

College of Engineering and Computer Science

UNIVERSITY OF CENTRAL FLORIDA

Authors

Ryan McCord

Gabriel Martin

Aidan Kennedy

Brogan Dougherty

Electrical Engineering

Computer Engineering

Electrical Engineering

Electrical Engineering

Review Committee

Dr. Enxia Zhang

Dr. Jaesung Lee

Dr. Piotr Kulik

Assistant Professor

Assistant Professor

Assistant Professor

Advisor

Dr. ChungYong Chan

Table of Contents

Table of Contents.....	1
List of Tables.....	5
List of Figures.....	5
1: Executive Summary.....	6
2: Project Description.....	7
2.1 Background and Motivation.....	7
2.2 Existing Products & Technologies.....	8
2.3 Overview of Project Goals.....	8
2.4 Project Features and Functionalities.....	9
2.5 Requirement Specifications.....	13
2.6 House of Quality.....	14
2.7 Block Diagram Representations.....	15
2.7.1 Hardware Block Diagram.....	15
2.7.2 Software Block Diagram.....	16
3: Research & Investigation.....	17
3.1 Analysis of Similar Existing Products.....	17
3.1.1 Commercial & Research Systems.....	17
3.1.2 Comparison and Justification.....	18
3.2 Hall-Effect Sensors.....	18
3.2.1 Texas Instruments DRV5053VA.....	19
3.2.2 DRV425RTJT.....	20
3.2.3 A1324LLHLX-T.....	22
3.2.4 SS495A1.....	23
3.2.5 Comparison & Selection.....	24
3.3 Data Storage and Handling.....	25
3.3.1 RAM.....	25
3.3.2 Flash.....	26
3.3.3 SD Card.....	27
3.3.4 External Hard Drive.....	28
3.3.5 Comparison & Selection.....	29
3.4 Microcontroller.....	31
3.4.1 MSP430FR6968.....	32
3.4.2 ESP32.....	32
3.4.3 STM32.....	33
3.4.4 Comparison & Selection.....	34
3.4.5 Development Board Selection.....	36
3.5 Sensor Addressing.....	36

3.5.1 Direct GPIO and ADC.....	37
3.5.2 MUX.....	37
3.5.3 Row/Column Multiplexing.....	38
3.5.4 ADC.....	40
3.5.5 I/O Expander.....	41
3.5.6 Binary Counter.....	42
3.5.7 Decoder.....	43
3.5.8 Comparisons & Selection.....	43
3.6 User Interface.....	44
3.6.1 RGB LED Display.....	45
3.6.2 Touchscreen GUI.....	46
3.6.2.1 Capacitive vs. Resistive Touchscreens.....	47
3.6.3 Comparison & Selection.....	48
3.7 Communication Protocols.....	49
3.7.1 I2C.....	50
3.7.2 SPI.....	51
3.7.3 UART.....	51
3.7.4 Comparison & Selection.....	52
3.8 Languages, Tools, and Development Flow.....	53
3.8.1 Embedded Development.....	53
3.8.1.1 C.....	53
3.8.1.2 C++.....	53
3.8.1.3 Comparison and Selection.....	54
3.8.2 Graphics Rendering.....	55
3.8.2.1 Python.....	55
3.8.2.2 Matlab.....	55
3.8.2.3 Open Source and Proprietary Graphics Libraries.....	55
3.8.2.4 Comparison and Selection.....	56
3.8.3 GitHub.....	57
3.9 Power Supply Unit.....	57
3.9.1 Batteries.....	57
3.9.1.1 Lead-acid Batteries.....	58
3.9.1.2 Lithium Batteries.....	58
3.9.1.3 Comparison & Selection.....	60
3.9.2 USB Power.....	61
3.9.2.1 USB-A.....	61
3.9.2.2 Micro-USB.....	62
3.9.2.3 USB-C.....	62
3.9.2.4 Comparison & Selections.....	62
3.9.3 Power Regulation.....	63

3.9.3.1 Linear Regulators.....	63
3.9.3.2 Switching Regulators.....	64
3.9.3.3 Comparisons & Selections.....	64
3.9.3.4 Analog Devices LTC3525 Switching Boost Converter.....	65
3.9.3.5 Texas Instruments LM2623MM Switching Boost Converter.....	66
3.9.3.6 Texas Instruments LM2596S Switching Buck Converter.....	66
3.9.3.7 Texas Instruments LP2985 Linear 3.3V Regulator.....	67
3.9.3.8 Final Component Selection.....	68
3.9.4 Battery Management Chips.....	68
3.9.4.1 Microchip Technology MCP73871.....	69
3.9.4.2 Analog Devices LTC4054.....	70
3.9.4.3 Texas Instruments BQ24610RGER.....	70
3.9.4.4 STMicroelectronics STC4054GR.....	71
3.9.4.5 Final Comparison and Selection.....	72
4: Standards & Design Constraints.....	74
4.1 Standards.....	74
4.1.1 IPC-A-610.....	74
4.1.2 IPC-2221: Standards in Circuit Board Design.....	75
4.1.3 IPC-7351: Surface Mount Design & Land Pattern Standard.....	77
4.1.4 IEEE 1149.1: JTAG Boundary Scan Testing.....	78
4.1.5 IEC 61000-4-2: Electromagnetic Compatibility (EMC).....	78
4.2 Constraints.....	78
4.2.1 Safety Concerns.....	79
4.2.1.1 Safety with the NED project.....	79
4.2.1.2 Safety with the HALO project.....	79
4.2.2 Feasibility Constraints.....	80
4.2.2.1 Feasibility with the NED project.....	80
4.2.2.2 Feasibility with the HALO project.....	80
4.2.3 Time Constraints.....	80
4.2.3.1 Time with the NED project.....	80
4.2.3.2 Time with the HALO project.....	82
4.2.4 Cost Constraints.....	82
4.2.4.1 Cost with the NED project.....	82
4.2.4.2 Cost with the HALO project.....	83
4.2.5 Layout Constraints.....	83
4.2.5.1 Layout with the HALO project.....	83
4.2.6 Available Resources.....	84
4.2.6.1 NED.....	84
4.2.6.2 H.A.L.O.....	85
5: ChatGPT vs. Similar Platforms.....	85

5.1 ChatGPT.....	86
5.2 DeepSeek.....	87
6: Hardware Design.....	88
6.1 Overall Schematics.....	89
6.1.1 Schematic Overview.....	90
7: Software Design.....	91
8: Fabrication & Prototyping.....	91
9: Testing and Evaluation.....	91
10: Administrative Content.....	91
10.1 Bill of Materials.....	91
10.2 Project Milestones.....	92
11: Conclusion.....	94
12: References.....	94

List of Tables

List of Figures

1: Executive Summary

Our senior design project is HALO, a Hall-Effect Array for Live Observation. The goal of our project is to develop a real-time magnetic field imaging system using an 8x8 array of Hall-Effect Sensors. This system will allow users to capture and visualize magnetic fields from everyday objects in real-time, providing a detailed and comprehensive map of magnetic field strength and polarity. The system is designed to be standalone and user-friendly, offering an interactive display where users can view various modes and display options. By detecting and mapping the strength of magnetic fields at various points, HALO will provide valuable insights into magnetic field distributions at an affordable cost, which can be applied in engineering and material science fields.

At the heart of our system is the STM32H735IGT6 microcontroller, which is responsible for processing data from all 64 Hall-Effect sensors in the array. Each sensor will measure the magnetic field strength at its respective location within the array. The data from the sensors will be multiplexed and transmitted through the microcontroller's internal ADC, which will enable data processing in real-time. The STM32H735IGT6 features a 32-bit ARM Cortex M7 core clocked at 550 MHz ensuring proper handling of sensor data, graphical rendering, and user input with a minimum refresh rate of 60 frames per second. The microcontroller's 564 KB of SRAM and 1 MB of embedded flash memory will allow us to seamlessly process data without the need for external memory. Our custom PCB houses 64 analog Hall-Effect Sensors, the STM32 microcontroller, and other necessary components, such as the capacitive touch screen and SD card module for data logging and visualization.

One of the key design considerations for HALO is power management, which ensures that the system can function independently from a wall power source and can be used for extended periods of time. The device is designed to be portable and robust, with a durable 3D-printed enclosure that protects the sensors and all other components from damage during use. The system has a graphical user interface (GUI), which allows users to view real-time magnetic field maps. The interface is also programmed with various modes for users who would like to adjust parameters.

HALO has been designed with scalability in mind, as an expanded array of sensors will allow for higher resolution when mapping magnetic fields. By combining multiple 8x8 arrays together, users will be able to view complex magnetic fields in more detail. Our primary goal for HALO is to create a system that is not only highly functional but also accessible and easy to use. With a focus on real-time data processing and visualization, the system can be used in

applications from observing magnetic materials to predicting electromagnetic interference within electronic components. By combining Hall-Effect technologies with real-time data visualization, HALO offers a significantly more affordable method to observe magnetic fields.

Our senior design project is part of Group 18, which consists of four team members: Ryan McCord, Aidan Kennedy, and Borgan Dougherty, all majoring in Electrical Engineering, and Gabriel Martin, who is majoring in Computer Engineering.

2: Project Description

2.1 Background and Motivation

Magnetic fields are present in nearly all electrical systems, but measuring and visualizing them often requires specialized equipment. In fields like electrical engineering, understanding magnetic fields is incredibly important. Traditional methods, such as single-point measurements or bulky sensors, can be quite expensive and don't provide a dynamic view of the magnetic environment.

One way to overcome these limitations is through the use of Hall effect sensors. These sensors generate a voltage that's proportional to the magnetic field they detect, offering a reliable and cost-effective way to measure magnetic fields. However, while Hall effect sensors are great for pinpoint measurements, there's still a significant need for a system that can map these fields in real time across a larger area.

This project aims to address this by creating a system that uses a 12x12 array of Hall effect sensors to generate a color-coded map of the magnetic field that refreshes at 60 Hz. By visualizing the magnetic field in real time, this system will provide users with a much deeper understanding of magnetic field behavior, especially in complex systems with oscillating magnetic fields. This can be especially useful in industry and for research, where it's vital to monitor the magnetic fields of electronics to prevent electromagnetic interference (EMI). Some potential applications for this system include monitoring power systems, detecting anomalies in machinery, or assessing sensor performance in different environments.

Our motivation behind this project comes from the need for a cost-effective, user-friendly way to study and visualize magnetic fields. With its practical

applications ranging from industry and research labs to educational use, users will be able to engage more effectively with the concept of magnetism.

2.2 Existing Products & Technologies

To better guide our project design, we analyzed two existing systems that align with our objectives and give insights into effective design strategies. These projects serve as useful references, helping us understand other design approaches and areas for improvement.

The first project, “A Magnetic Imager Tile” [1], uses an 8x8 array of Hall-effect sensors to visualize magnetic fields. This system detects field polarity but does not measure the field’s intensity. Data is transmitted using the I2C protocol to a laptop which displays the magnetic field. One of the key improvements we hope to make is transitioning from I2C to SPI for data transmission. SPI has higher bandwidth and lower latency, which is crucial since we want the refresh rate to be 60Hz. Also, this project relies on a development board and laptop to visualize the field. We hope to eliminate the need for external components by implementing a microcontroller chip and onboard LCD touchscreen to make the system function independently.

The second project, “A High-Resolution Magnetic Field Imaging System Based on the Unpackaged Hall Element Array” [2], has a different approach. It uses an array of 256 Hall-effect sensors attached to a motorized rail system to scan and render a high-resolution image of magnetic fields. Our project differs in that we aim to provide continuous real-time visualization rather than a static scan. However, we will incorporate data storage, allowing users to save static images or video recordings of 30 seconds or less of the magnetic fields onto an SD card. The motorized rail system is also a valuable concept that could be integrated into future iterations of our design.

By analyzing these existing projects, we aim to combine their strengths to create a system capable of high-speed magnetic field visualization.

2.3 Overview of Project Goals

The primary goal of this project is to develop a real-time magnetic field visualization system that can detect and display magnetic field strength using an array of Hall effect sensors. By providing a clear visual representation, we hope to make magnetic field analysis more accessible and user-friendly. The visual

representation will be comprised of a heat map of the magnetic field where different gradients correspond to different field strengths.

We want to create a compact, lightweight, and reliable device that can be used by engineers, researchers, and educators working with electronics. The device should be easy to use in a variety of environments, primarily in classrooms and laboratories, while maintaining high accuracy and responsiveness.

Goals:

- Display a visualization of the strength of a magnetic field relative to space and frequency
- The system should be lightweight and portable
- The device should be low-cost
- Magnetic field data must be captured by the sensors and transmitted to the MCU
- Implement a user-friendly interface for visualizing magnetic fields
- Create a standalone system with an integrated touchscreen for ease of use
- Implement effective shielding techniques to prevent electromagnetic interference
- Allow the system to be powered by a USB-C connection and lithium-ion battery
- The system should be able to store and export magnetic field data on an SD card

Stretch Goals:

- The system could be enhanced with automated scanning capabilities using motorized rails to analyze larger and more complex magnetic fields

Objectives:

- Hall-effect sensors must be properly calibrated for precise magnetic field detection
- The touchscreen interface should be user-friendly and provide real-time interaction
- EMI shielding techniques will minimize interference with the sensor array
- House the system in a durable 3D-printed enclosure
- The system operates independently, utilizing its own power source, built-in controls, and onboard display without relying on external hardware.

2.4 Project Features and Functionalities

Overall functionality

The overall functionality of the project will be to utilize Hall effect sensors to map magnetic fields and display them as a 2-dimensional heat map for intuitive readability. The device will also update the screen at a rate of 60 Hz to show any changes in the magnetic field in real time.

Sensor Matrix

The sensor matrix will be made up of 144 Hall effect sensors in a 12 by 12 configuration. We will be utilizing analog Hall effect sensors instead of digital ones because it will make visualizing the heat map more comprehensive and will allow for different field strengths to be represented instead of just having the sensors detect the presence or polarity of said field. When analog Hall effect sensors come into contact with a magnetic field the electrons and holes separate in them creating a voltage that is proportional to the strength of the magnetic field. The analog outputs of the sensors would feed into several multiplexers which would then feed into the analog-to-digital converter pins on the MCU.

MCU

The MCU will have built-in analog-to-digital converters which will allow us to take the analog data coming from the multiplexers and convert it to digital signals that can be processed by the MCU. The MCU will be reading the voltage value of the Hall effect sensors and assigning a field value to them. This information will then be transmitted to our display screen via SPI.

Display

The display will receive data from the MCU and show the 2D map of the field. It will be a colored display using red to represent a region of high intensity and green to represent a region of low intensity, dark red being the highest and light green being the lowest. The colors will work similarly to a weather radar showing the levels of rain using different colors to represent different intensity. The display will have 144 such regions that can change colors according to the Hall effect sensor from which it is receiving information. The display will have a refresh rate of 60 Hz which will allow for real-time updates from the sensors to show any changes in the magnetic field such as moving a magnet across its surface or rotating its polarity.

The display will be able to switch between mapping things such as the amplitude of the magnetic field, the frequency of the magnetic field, the polarity of the magnetic field or change to a display that shows the number associated with each color mT for field intensity, Hz for frequency, and 1 or -1 for field direction. The screen will display which view is currently being shown. These

views will also display a guide to show the user what value each color and intensity of said color represents.

The display will be a touchscreen to allow for user inputs. The user will be able to use the screen to change between displays, choosing which graph they want to see and how they want the data to be presented. The user will be able to change if they are viewing a graph showing amplitude, frequency, direction of field, or the raw data values. The screen will also allow the user to save the data from the past 30 seconds to an SD card.

Memory Card Connection

The device will come equipped with an SD card capable of storing 32 GB of information. This SD card will be ejectable so that a user can download information from the device and view it on their laptop or other devices capable of reading an SD card. The MCU will automatically save the last 30 seconds' worth of data to the SD. If the user does not save the data then it will be overwritten. The user can choose when to save 30 seconds worth of data and will be able to save multiple instances. The SD card will contain the information that could have been displayed on the screen such as the amplitude graphs, frequency graphs, field direction graphs, and the raw data.

Software

The software for the project will be responsible for acquiring sensor data, processing it, and updating the visual heat map on the onboard LCD display in real time. The software will be designed to meet a 60 Hz display refresh rate while sampling a 12 x 12 sensor array (144 sensors) with sufficient accuracy and speed. The microcontroller firmware will be developed in C, utilizing the STM32's existing libraries for LCD display interfacing (via SPI), expediting the development and testing process.

SPI is chosen for its higher data throughput, full-duplex capability, and lower latency compared to I2C or UART, advantages that are critical for achieving a smoothly refreshing, real-time display. An ADC (most likely 12-bit) will be used to convert the analog signals from the Hall effect sensors (routed through multiplexers) into digital values. Each sensor in the array is sampled individually. Given the need for real-time performance, it is critical that both the ADC conversion rate and the MCU's clock speed are sufficient to complete all 144 sensor readings within 1/60th of a second.

Initial calibration routines should allow for programmable offset and gain adjustments per sensor to account for base interference and component variances. Although the exact calibration method will be refined after sensor

integration on the PCB, a calibration routine will be included in the firmware to adjust for ambient noise and baseline offsets.

The raw ADC values should be directly mapped to a set of bins corresponding to intensity levels. Each bin is associated with a specific color, with 16-24 possible colors available to represent different ranges of magnetic field strength. The min and max sensor values will be determined either via calibration or dynamic scaling, ensuring that the bin sizes accurately represent the range of measured voltages. At this stage, minimal processing beyond binning and scaling is performed on the sensor data, as the main objective is to maintain a high refresh rate for the heat map. However, if future testing indicates that interpolation or smoothing is needed to enhance the visual representation, this functionality can be incorporated as an optional processing step before mapping values to colors. The values acquired from the ADCs at this stage are all that is required for the comprehensive development of the various graph types and display modes outlined in the display section above.

The software's main loop should be optimized to sequentially sample each sensor, process the ADC values, update the corresponding bins and values, save data to the SD card, and transmit the color-coded data to the display. The clock speed of the MCU should be configured to ensure that all sensor readings, storage operations, and display updates are completed within the time frame required for a 60Hz refresh rate. This approach guarantees that the heat map remains accurate and responsive to changes in the magnetic field. Data is sent from the MCU to the LCD via the SPI protocol. SPI is selected due to its high throughput, full-duplex operation, simplicity, and reliability. The MCU will use the binned ADC data to generate a color-coded heat map. Each sensor's ADC reading is mapped to one of the predetermined color bins. The display would be partitioned into 144 regions, each representing a sensor's reading, which together form the complete 2D visualizations of the magnetic field.

Early software testing will be performed using a reduced sensor array (potentially a 2 x 2 array on a breadboard) to verify the basic acquisition, calibration, and display update functionality before scaling up to the full array size. This phased approach should allow for early debugging and performance verification. Although we do not expect the need for a full user interface, a simple status and legend display on the LCD will be incorporated to indicate operational mode or diagnostic information.

2.5 Requirement Specifications

Table 1: List of Requirements & Stretch Goals

Parameter	Purpose	Value
Basic Features		
Number of Tiles in Heat Map	Magnetic field visualization	≥ 64 tiles
Field Size Detected	Detect a large portion of a magnetic field	$\geq 64 \text{ cm}^2$
Field Intensity Range Detected	Detect a range of magnetic field strengths	$[-10, 10] \text{ mT}$
Detection Distance	Accurately measure fields close to the sensor	$\leq 1 \text{ cm}$
Screen Refresh Rate	Display data through the screen quickly	$\leq 30 \text{ Hz}$
Amount of Data Saved	Store values for analysis	$\geq 30 \text{ s}$
Power Consumption	Consumer a small amount of power	$< 5 \text{ W}$
Power Efficiency	Efficient power use	$> 75\%$
Operating Temperature	Maintain a non-extreme operating temperature	$< 120^\circ \text{ F}$
Latency	Transmit data quickly	$< 200 \text{ ms}$
Weight	System is light weight for convenience	$< 5 \text{ lbs}$
Size	System is small for convenience	$< 10 \text{ cm}^2$
Cost	System is affordable	$< \$250$

Specification notes:

Items in Blue will be demonstrated on D-Day

For demonstration purposes, we will use a magnetic field source such as a permanent magnet, and hold it above the sensor matrix. The display will show a heat map showing areas of high intensity in red and areas of lower intensity in green. We will have the display show values at each sensor and compare this data with an external off-the-shelf device with the ability to measure magnetic fields to verify that our display is showing valid readings.

To demonstrate that the display is refreshing we will move the magnet across the sensor, further away from the sensor, and rotate the magnet causing the polarity to flip. The display will be refreshing in real time to accurately show how the field is changing. The display will also be able to show the refresh rate to confirm that it is refreshing at least 60 times per second.

To demonstrate the data logging feature, we will save a recording after 30 seconds to the onboard SD card and put it into another device to view the contents of the card. This will allow us to see the past 30 seconds of recorded data.

2.6 House of Quality

The House of Quality provides insights into the trade-offs between various engineering requirements in H.A.L.O. Many of these requirements are interconnected, meaning improvements in one area could come at the expense of another. For example, increasing the refresh rate of the magnetic field mapping system may demand higher processing power, which could lead to increased power consumption and have an impact on our battery life. Similarly, reducing the system's size would require components to be placed closer together, increasing the risk of electromagnetic interference, which could degrade overall performance. Understanding these relationships allows us to make better-informed decisions in our design to ensure we create a well-balanced system that meets technical and logistical goals.




Figure 1: House of Quality Diagram

2.7 Block Diagram Representations

2.7.1 Hardware Block Diagram

The H.A.L.O. system consists of two primary PCBs: the Sensor PCB and the Main PCB. The Sensor PCB houses the Hall-effect sensor array, which is used for measuring magnetic fields, as well as the multiplexer (MUX), analog-to-digital converter (ADC), binary counter, and decoder that work together to process and address the sensor readings. The Main PCB handles the system's power, processing, and output. It features a microcontroller (MCU), which manages tasks such as data acquisition, processing, and user interface

interactions. The power subsystem includes a USB-C port for charging, a 3.7V lithium-ion battery, a battery management chip, and 3.3V/5V voltage regulators to ensure stable voltage across the system. Additionally, the Main PCB is equipped with an SD card connector for data storage and is connected to a touchscreen LCD display for real-time visualization of the magnetic field data. To minimize magnetic interference and optimize overall system performance, the Sensor and Main PCBs are connected via a flexible printed circuit (FPC).




Figure 2: Hardware Block Diagram

2.7.2 Software Block Diagram

The software for the H.A.L.O. system is built to handle data collection, processing, and visualization. The firmware on the STM32H735IGT6 microcontroller controls the sensor array, manages data transfer using SPI/I2C, and processes the magnetic field measurements. Data flows through a real-time pipeline that ensures smooth communication between the sensors, signal

processing, and the graphical interface on the touchscreen. The software also supports data logging to an SD card, so the collected data can be saved for future analysis.

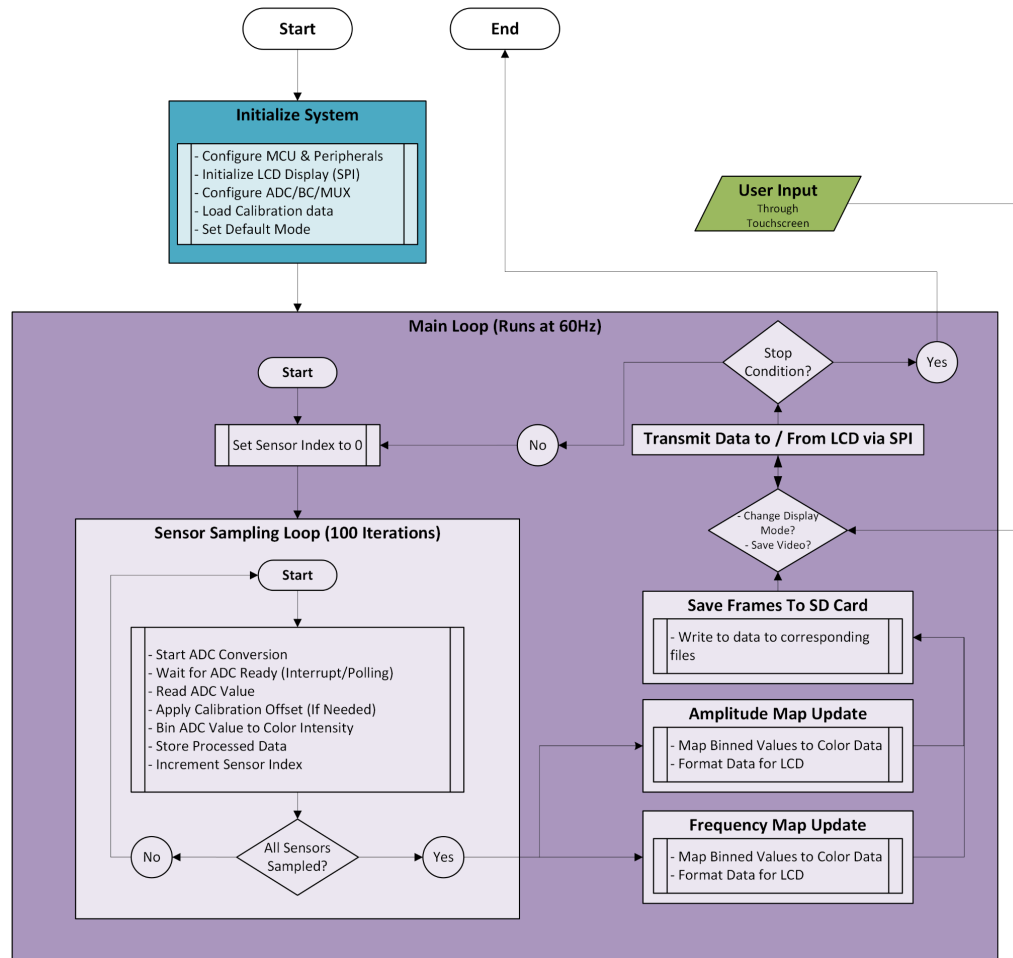


Figure 3: Software Block Diagram

3: Research & Investigation

3.1 Analysis of Similar Existing Products

3.1.1 Commercial & Research Systems

There are a variety of magnetic field imaging systems that already exist. Magnetic resonance imaging (MRI) is a mainstay in the medical field for its non-invasiveness and high-contrast images. [B8] In non-medical fields you'll find magnetic field imagers that operate on singular or multiple axes, that are handheld or stationary, that are on a rail system or free, and that are wired or wireless. The functions of these imaging systems encompass a large variety of

purposes. It can be used for research, educational purposes, or electronic design. It could even be used for something as simple as reading the Teslas around a permanent magnet.

3.1.2 Comparison and Justification

3.1.1 goes to show that there are numerous use cases for a magnetic field imager. The possible usages are so large that it's difficult to design a system for a specific purpose. In products like the HallinSight, Metrolab outlines the function and specifications of the device, but omits any text on how they recommend using it. [B9] The purpose of this long preamble is to show that the applications for a field sensor are broad and trying to design a sensor for a niche purpose would take away from the applicability of the system. Plus, the cases that are niche enough to warrant a specific design like an MRI machine are well beyond the scope of our project.

H.A.L.O doesn't seek to visualize fields in a new way, or find a niche application that has been capitalized on yet. The benefit and justification of H.A.L.O is convenience. Handheld sensors exist, and we have size and weight at the forefront of our minds when we make every design decision. Having a compact and lightweight imaging system will allow the user to contort the device and image magnetic fields from whatever angle or distance they prefer. The system being wireless only adds to this intention.

Having a light, compact, wireless magnetic imaging system are great foundation for our project, but the standout feature is having a screen incorporated into the device itself. I have been hard-pressed to find a single other magnetic imaging system that doesn't offload the display of the field to another device. By having a touchscreen on the device itself, measurements are very easy to obtain and the hassle of setting up another device and dealing with software is eliminated. The user won't have to shift their gaze from H.A.L.O to read the measurements. This reduces error. Having the screen directly on the device will make the device easier to use and more reliable.

3.2 Hall-Effect Sensors

Hall-Effect sensors are an essential part of HALO, as they are what allow for precise magnetic field measurements that are crucial for generating real-time maps of magnetic fields. These sensors operate on the Hall Effect principle, where the presence of a magnetic field will induce a voltage difference perpendicular to the current flow. Hall-Effect sensors are widely used in various fields and are available in two forms: digital and analog. Digital Hall-Effect sensors detect the presence of a magnetic field, while analog types provide a

continuous output proportional to field strength. We are considering various factors for sensor selection within HALO, including sensitivity, resolution, noise characteristics, power consumption, and ease of integration.

This section examines the Hall-Effect sensors we considered for HALO, each having potential strengths in the context of our project's requirements. These options include the DRV5053VA, a linear analog sensor with a ratiometric voltage output; the DRV425RTJT, a high-precision fluxgate sensor ideal for high-resolution applications; the SS495A1, a robust sensor designed for industrial settings, and the A1324LLHLX-T, a low-power sensor with moderate sensitivity. In this section, we will analyze each sensor's specifications and determine which is the most suitable for HALO's performance needs.

3.2.1 Texas Instruments DRV5053VA

The DRV5053 VA from Texas Instruments is a chopper-stabilized Hall-Effect sensor that provides a linear analog output voltage in response to magnetic flux density. This sensor operates with a supply voltage ranging from 2.5V to 38V, meaning it is a versatile option that can be used in various applications. The sensor's chopper stabilized design modulates its signal at a high frequency to minimize noise and DC offset, meaning it is capable of maintaining high accuracy over a wide range of operating conditions.

In the absence of a magnetic field, the DRV5053VA outputs a voltage of 1V. When we expose the sensor to a magnetic field, the output voltage changes linearly, with the direction of the magnetic field influencing the voltage polarity. This sensor is available in models with varying sensitivities, including -45 mV/mT, -23 mV/mT, -9 mV/mT, +23 mV/mT, and +45 mV/mT. For our application, the ± 9 mV/mT version will be the most suitable, as it provides a response to everyday magnetic fields such as desktop magnets or small electric motors.

One of the more notable features of the DRV5053VA is its capability to determine the polarity of an applied magnetic field. For this model, a south pole magnetic field decreases the output voltage below 1V, while a north pole magnetic field increases it above 1V. This polarity sensitivity is a great addition to this sensor, as we will also be able to observe the direction of magnetic fields.

This sensor also has excellent temperature stability, with its sensitivity only varying by $\pm 10\%$ over an operating temperature range of -40°C to $+125^{\circ}\text{C}$. While it is not likely that our system will be exposed to extreme temperatures, it

is helpful to note that this is a robust sensor that will not experience fluctuations when changing its surrounding environment.

The DRV5053VA is offered in a compact 3-pin surface-mount package, making it much easier to integrate into our custom PCB. Since our system involves integrating 64 of these sensors together, a compact packaging is crucial in minimizing the size of our PCB and simplifying routing.

It is worth considering, however, that this specific sensor is susceptible to errors caused by external magnetic fields unrelated to those we are trying to measure. If we do not properly manage this interference, we may receive inaccurate readings on our display. To mitigate this issue, we must be very careful when designing our PCB, ensuring that the sensor array is properly distanced from other components, such as the battery or voltage regulators. On the software side, incorporating algorithms to filter out noise would also help increase the overall reliability of our product.

In summary, the DRV5053VA Hall-Effect sensor offers a combination of high sensitivity, temperature stability, and ease of integration, making it an excellent choice for precisely measuring magnetic fields in our system. While it is important to mitigate interference when using this sensor, this limitation is minor compared to its significant advantages.




Figure 8-1. Typical Application Schematic - No Filter

3.2.2 DRV425RTJT

The DRV425RTJT, also developed by Texas Instruments, is a high-precision fluxgate Hall-Effect sensor that is used for applications requiring extremely accurate magnetic field measurements. Unlike traditional Hall-Effect Sensors,

the DRV425RTJT uses a fluxgate design, which enables it to detect small magnetic fields with high resolution.

One of the more notable features of this sensor is its ability to overcome noise, due to its very low noise floor. This characteristic allows it to detect very small changes in magnetic field strength, making it an ideal option for applications where high-resolution mapping of these fields is essential. This sensor outputs a differential analog voltage, which is beneficial for integration with an ADC to achieve precise digital processing of data.

In terms of operation, the DRV425RTJT operates within a supply voltage range of 4.5V to 5.5V, which is more narrow than the DRV5053VA but still offers flexibility. This sensor also operates from -40°C to +125°C, which makes it suitable for the operating environment of our device.

Despite its excellent resolution, accuracy, and temperature stability, the DRV425RTJT presents a few integration challenges. The sensor is housed in a surface-mount package, which is compact but more complex than the simpler designs of sensors like the DRV5053VA. The fluxgate design requires additional signal conditioning circuitry, which complicates the PCB layout. This added complexity may result in a larger PCB and more difficult integration, particularly for a system like HALO, which involves arranging 64 sensors in a compact array.

Another concern is the DRV425RTJT's increased sensitivity to external magnetic interference. While this higher sensitivity makes it ideal for detecting weak magnetic fields, it also makes the sensor more vulnerable to interference from nearby magnetic sources. In a tightly packed system like HALO, where sensors are placed close together, this could lead to inaccuracies if not managed properly. Ensuring proper shielding and a clean PCB layout would be essential to reduce these errors, adding extra complexity to the design process.

In summary, the DRV425RTJT offers high precision and stability but comes with challenges in integration, packaging, and interference management. While it could potentially offer superior performance with careful design, these factors make it a less ideal option for HALO, where a more straightforward and compact design is needed.




Figure 67. Simplified Schematic of a DRV425-Based Linear-Position Sensing Application

3.2.3 A1324LLHLX-T

The A1324LLHLX-T, developed by Allegro Microsystems, is a low-power, linear Hall-Effect sensor that provides an analog output in response to magnetic fields. Its main advantage is its low power consumption, which is essential for battery-operated systems like HALO. The A1324LLHLX-T operates with a supply voltage ranging from 4.5V to 5.5V, making it compatible with a wide range of power configurations.

One of the key features of this sensor is its ability to provide a ratiometric output, which varies linearly with the magnetic field strength. This ensures that the output voltage is proportionally related to the field strength, making it suitable for generating continuous magnetic field maps. With a typical sensitivity of 2.5 mV/G, it offers a moderate response to magnetic fields, which makes it ideal for general-purpose applications.

The A1324LLHLX-T's temperature coefficient is also quite favorable, with the sensitivity typically varying by $\pm 10\%$ over the temperature range of -40°C to $+125^{\circ}\text{C}$, similar to the DRV5053VA. This makes the sensor quite stable in varying environmental conditions.

In terms of integration, the A1324LLHLX-T comes in a compact 3-pin SIP package, which simplifies its mounting and reduces PCB space usage. However, like the DRV5053VA, external magnetic interference could affect its accuracy. Adequate shielding and PCB layout considerations will be necessary to ensure precise measurements.

While this sensor offers low power consumption, moderate sensitivity, and easy integration, its lower sensitivity may not provide the resolution needed for higher precision magnetic field mapping. However, it still offers a balanced solution for applications where lower power consumption is prioritized over fine resolution.




Figure X: Typical Application Circuit

3.2.4 SS495A1

The SS495A1 Hall-Effect sensor, developed by Honeywell, is a radiometric linear sensor designed for magnetic field measurement applications. This sensor provides an analog output proportional to an applied magnetic field. This sensor has a sensitivity of 3.125 mV/G at 25°C, with its resting voltage at 2.5V when no magnetic field is present. The output voltage increases or decreases depending on the polarity and strength of the applied field, allowing for precise measurements. The operational magnetic range extends from -600 Gauss to +600 Gauss, ensuring that the sensor is reliable in environments with strong magnetic fields.

One of the standout features of this sensor is its high temperature stability, allowing it to function in harsh environments ranging from -40°C to +150°C. This range is broader than the DRV5053VA, DRV425RTJT, and A1324LLHLX-T, making it more robust for extreme conditions. However, while the SS495A1 is designed to withstand such extreme temperatures, this feature will not be necessary for HALO, as it is a desktop device that will operate in a controller indoor environment.

In terms of performance, the SS495A1 has a linearity specification of $\pm 1.5\%$ of span from -600 Gauss to +600 Gauss, ensuring that the sensor delivers accurate readings across its operational range. Its null voltage at 25°C is 2.5V, with a drift of $\pm 0.04\%/^{\circ}\text{C}$ over the temperature range of 25°C to 125°C. This

makes the sensor stable and reliable, with minimal variation in its output voltage due to temperature fluctuations.

Another advantage of this sensor is its ease of integration into a custom PCB. The SS495A1 has a compact 3-pin packaging, which is advantageous for HALO where space is limited. Due to its simple design and ratiometric output, this sensor is a great option for HALO to create maps of magnetic fields.

In all, the SS495A1 has good stability, sensitivity, and ease of integration, which makes it a solid candidate for HALO. While its temperature tolerance is impressive, this is not a key factor for our application.

3.2.5 Comparison & Selection

After comparing the Hall-Effect sensors for HALO, the DRV5053VA stands out as the most suitable option due to its balance of sensitivity, ease of integration, and stability. While the DRV425RTJT provides higher resolution and low noise for detecting weak magnetic fields, its added complexity in terms of packaging and increased risk of interference makes it less ideal for the design of our system. The A1324LLHLX-T offers low power consumption, but it lacks the precision needed to map magnetic fields from common objects. The SS495A1 has a wide range temperature tolerance, but this feature is unnecessary for HALO, as it is designed to be a desktop device and will not be exposed to extreme temperatures.

The DRV5053VA's combination of linear output, high sensitivity, and ease of integration make it the best choice for HALO's sensor array. Its small footprint and temperature stability allow it to function effectively within our design, and its ability to detect a wide range of magnetic fields ensures the device can create precise, real-time maps of applied fields. It is also worth noting that this sensor is the most affordable option, which is a crucial factor to consider as we must integrate 64 of them into one printed circuit board. Overall, the DRV5053VA will simplify both the hardware design and integration process, making it the most efficient option for HALO.

Table X: Hall-effect Sensor Comparison

Component	Sensitivity	Temperature Range	Power Consumption	Cost	Risk of EMI
DRV5053VA	± 9 mV/mT	-40°C to +125°C	Low	Moderate	Medium

DRV425RTJT	± 2 mT	-40°C to +125°C	Moderate	High	High
A1324LLHLX-T	2.5 mV/G	-40°C to +125°C	Low	Low	Low
SS495A1	3.125 mV/G	-40°C to +150°C	Low	Low	Medium

3.3 Data Storage and Handling

Efficient data storage and handling is a highly critical component to both the performance and functionality of the HALO system. The project involves real-time data acquisition from the hall effect sensor array, writing to and reading inputs from the display, and saving data to physical storage; this requires careful consideration of storage solutions to balance speed, capacity, and efficiency. This section explores the various storage mediums and options, their implications on the project, and the final selections based on our requirements. Some key factors that influenced these decisions include the need for temporary data buffering, permanent storage of sensor readings / display output, display refresh rates, and power efficiency. The storage choices must also account for the MCU's internal and I/O memory limitations, the display's 480x272 resolution with a 60Hz refresh rate, and the frequency of sensor reading and logging required for accurate visualization and analysis.

3.3.1 RAM

Random Access Memory (RAM) is a crucial component of the HALO system, serving as the primary medium for temporary data storage, real time processing, and efficient data transfer between components on the PCB. RAM, unlike non-volatile storage options, enables high speed read and write operations that make it essential for handling the sensor data as well as managing the graphical display. The MCU includes 564KB of RAM, which will have to be carefully allocated to support various concurrent tasks, which as of this moment include sensor data buffering, display frame rendering, and system management.

One of the most significant demands on the RAM comes from the framebuffer for the 480 x 272 resolution display, which operates at a 60Hz refresh rate. Given that the display uses the RGB565 format (16 bits per pixel), a single full-frame requires approximately 261KB of RAM. Since we might need to employ a double-buffering scheme to prevent flickering and ensure smooth rendering, the

display alone could occupy over 520KB of RAM, consuming nearly all of the available internal memory. We hope to have the full functionality of our system at the desired specifications by using only the internal memory of the MCU (in order to ease PCB design), however we are purchasing additional RAM in the case that our system requires it. Fortunately, the STM32 we are using features Chrom-ART hardware acceleration, which reduces the computational load on the processor and optimizes memory usage by offloading graphical operations.

In addition to the display requirement, RAM is needed for buffering the readings from the Hall sensor array before processing and storage. HALO continuously collects readings from the array, which must be temporarily stored before being processed into meaningful visualizations or written onto the SD card. The exact RAM requirements for sensor data depend on the number of sensors, data resolution, and sampling rate. If each reading is stored as a 16-bit value and multiple samples are collected per cycle to improve accuracy, the total RAM demand could increase significantly, necessitating either additional storage or efficient memory management strategies.

Further RAM allocation is required for system tasks and operations, including real time data processing and potential filtering depending on testing results. The MCU supports Tightly Coupled Memory (TCM) which provides extremely low latency access to data, enabling further optimization strategies we could employ to optimize resources. The stack and heap memory allocations for the firmware must also be taken into account and carefully managed to prevent memory fragmentation, as well as to ensure system stability.

Given the aforementioned considerations, the internal 564KB of RAM must be allocated strategically so that we can most effectively balance display operations, sensor data processing, and execution. If additional memory is required, external SDRAM or PSRAM modules could be installed, although it would require reconsideration on the part of PCB design and power regulation. By optimizing RAM usage and utilizing any available hardware acceleration, our system can manage real-time data visualization and sensor data acquisition while maintaining our desired responsiveness and performance.

3.3.2 Flash

Flash memory will be the primary non-volatile storage medium for HALO, responsible for retaining critical data even when power is lost. The MCU includes 1MB of internal flash memory, which will handle the storage of firmware, calibration data, system configurations, and data logging. Unlike RAM, Flash memory retains information persistently, making it a great candidate for

storing essential system parameters, ensuring that the device can resume normal operation after power cycles or resets. However, Flash memory has inherent limitations like slower write and read speeds, as well as a finite amount of program / erase cycles, which must be factored into the design.

Our primary use of Flash memory is the storage of firmware, which includes all of the embedded software responsible for controlling the MCU, managing any peripherals, and executing the system's core functions. The size of the compiled software depends on the final functionality and complexity of our system. Given that the MCU features an ARM Cortex-M7 processor with floating-point and DSP capabilities, firmware can be relatively large when compared to that of a simpler MCU's implementation.

Beyond firmware, Flash will be used for calibration parameters and system settings, which define sensor thresholds, user preferences, and operational configurations. These parameters must persist between power cycles to ensure consistent performance without requiring recalibration on every boot. However, since flash has limited write endurance (100,000 - 1,000,000 cycles per sector), frequent updates to stored parameters must be managed carefully. One strategy to mitigate excessive wear is the use of wear leveling and data buffering using RAM, ensuring we only write to the flash when necessary.

While the Flash storage is technically capable of storing sensor data, it is not well-suited for high-frequency logging due to its slower write speeds and limited endurance. We will be requiring continuous data acquisition from the sensor array, which would quickly exhaust the available write cycles of internal flash if used for real time storage. Instead, a more suitable approach would be to buffer some of the sensor readings in RAM before dumping them onto the SD card.

Given these considerations, the 1MB of internal flash storage is allocated for firmware, calibration data, and critical system parameters / settings, while real-time sensor data and GUI frame buffering are reserved for RAM. By carefully managing write operations and leveraging RAM's temporary storage capabilities, we can preserve the Flash memory for long-term reliability while ensuring the system operates well over extended use.

3.3.3 SD Card

The SD card will serve as the primary long-term data storage solution for HALO, responsible for logging sensor data, system diagnostics, recordings of the display output, and any additional information that could be of use to the intended user of our system. Unlike RAM, which is volatile, and Flash, which has

limited write endurance, SD cards offer a high-capacity and removable storage medium capable of handling continuous data logging with minimal system constraints. These factors make an SD card the most practical choice, especially considering their inexpensive cost.

One of the key considerations in selecting an SD card is write speed, as data must be logged in real time without causing system bottlenecks. Given that the system operates at 60Hz, and sensor readings must be stored at a relatively high frequency, the SD card should be able to sustain continuous data writes without causing unnecessary latency, stuttering, or buffer overflow issues. A class 10 or UHS-I SD card is preferred, as these classifications ensure write speeds of at least 10MB/s, which is more than sufficient for the expected data throughput. We also plan to implement buffering strategies using RAM to temporarily hold sensor readings before writing them in batches to the SD card. These implementations should significantly reduce the risk of write latency affecting system performance.

Another advantage of using an SD card is the fact that it is removable. By storing sensor data onto an SD card, users can easily extract recorded information for offline analysis, debugging, or archival purposes without the need to directly interface with the MCU. This makes using an SD card an ideal solution for long-term data retention while maintaining flexibility, since in the case that the SD card malfunctions or becomes damaged, a new one can be swapped in. The system will interface with the SD card using SDIO (Secure Digital Input Output) mode, which offers higher performance when compared to SPI-based communication, lowering the CPU overhead and ensuring faster write operations.

Overall the SD card plays a crucial role in long-term data retention and management. It is the primary medium for recording data while offering removability, high capacity, and sufficient write speed for real time operation.

3.3.4 External Hard Drive

External hard drives (HDDs) and solid state drives (SSDs) offer substantial storage capabilities, in the ranges from gigabytes to terabytes, making them suitable for applications requiring extensive data storage. Integrating such drives for HALO presents several challenges, however, that outweigh their potential benefits.

The physical size and weight of external hard drives can be a disadvantage in embedded systems where space is limited. Since we plan to make this a

standalone and portable system, this is the most compromising factor when considering these larger drives. The next factor is interfacing and compatibility, since the MCU we are using lacks native support for the interfaces commonly used by external drives, such as USB Host or SATA. Implementing support for these interfaces would require additional hardware components and complex software development, which is incompatible with our target delivery date. These kinds of drives also consume considerable power due to mechanical components like spinning platters and read/write heads. Even a 2.5 inch external HDD would likely exceed the power budget of an embedded system like ours. This increased power demand can lead to reduced battery life and necessitate more robust power management strategies.

Considering these factors, external hard drives are not a practical choice for HALO's long term storage needs. The combination of interfacing difficulties, higher power consumption, and physical constraints makes them less suitable compared to alternatives like SD cards, which offer a more balanced solution for the needs of this project.

3.3.5 Comparison & Selection

Since HALO requires RAM for buffering, Flash for firmware storage, and an SD card for real time data logging, the selection process will focus on finding the most suitable components within each storage category. Below are comparisons of viable options.

Table X: Ram Comparison

Component	Type	Size	Speed	Features	Notes
Internal RAM (STM32)	Embedded SRAM	564KB	High (tightly coupled)	Directly accessible by MCU	Best for real time operations
IS66WV51216E BLL	External SRAM	8MB	108MHz	Parallel Interface	Compatible with STM32, but adds complexity
AP Memory APS6404L-3SQ R	PSRAM	4MB	133MHz	SPI / QSPI Interface	Lower power, useful for large buffers

The internal 564KB SRAM of the STM32 MCU we have selected is the primary choice due to its low latency, ease of access, and integration with the MCU. It will be used for sensor reading and display buffering, and computational tasks. If additional RAM is needed for larger buffers, PSRAM is a viable secondary choice due to its low power consumption and QSPI interface.

Table X: Flash Comparison

Component	Type	Size	Write Speed	Endurance (cycles)
Internal Flash (STM32)	Embedded NOR Flash	1MB	1MB/s	100,000
Winbond W25Q128JV	External NOR Flash	16MB	50MB/s	100,000
Micron MT29F2G08A BAEAWP	NAND Flash	2GB	20MB/s	10,000

Utilizing the internal Flash available on our STM32 is the most practical choice, similar to the RAM selection. Due to its low power consumption, high reliability, and direct accessibility, making efficient use of its available resources will ease development and reduce cost. After all, the MCU was selected in part due to the onboard Flash. External Winbond Flash could be considered if additional non-volatile storage is necessary.

Table X: SD Card Comparison

Component	Type	Capacity	Speed	Interface
SanDisk Ultra 16GB	SDHC	16GB	80MB/s	SDIO
Samsung EVO Select 32GB	SDHC	32GB	90MB/s	SDIO
SanDisk Extreme PRO 64GB	SDXC	64GB	170MB/s	UHS-I

A SanDisk Ultra 16GB is ideal due to sufficient storage capacity, good write speeds, and high availability. Higher capacity SDXC cards are unnecessary, as HALO's data logging will not generate excessive storage demands. Class 10 or UHS-I ensures smooth, real time writing.

These selections provide the best balance between performance, power, and simplicity, ensuring the storage needs for our system are met and smooth operation is achieved without unnecessary complexity.

3.4 Microcontroller

A microcontroller (MCU) is a compact integrated circuit designed to execute specific tasks within an embedded system. Unlike general-purpose CPUs, MCUs integrate a processor core, memory (RAM/Flash), and peripherals (I/O, timers, communication interfaces, etc.) into a single chip. Features like low power consumption, real-time processing capabilities, and adaptability make them ideal for applications requiring efficient deterministic control over hardware components.

Some key characteristics shared by most microcontrollers include:

- Power Requirements: MCUs generally operate at low voltages (3.3V or 5V) and have low power consumption, making them ideal for both battery-operated and hardwired applications.
- Programmability: Most MCUs support programming in C, C++, or assembly language, with various development environments available for ease of use.
- Cost-Effectiveness: Compared to CPUs, MCUs are generally much more affordable as they typically range from between \$5 to \$50, depending on model.
- Adaptability: The integrated peripherals available on most MCUs (GPIO, communication protocols, timers, for example) allow them to interface with external hardware like sensors, actuators, and displays.
- Analog-to-Digital Converters (ADCs): ADC functionality is critical for converting real-world analog signals, typically from sensor data, into a digital format for processing. The selection of an MCU often depends on the capability of its onboard ADC, factoring in resolution, sample rate, and available channels.

For HALO, a microcontroller serves as the backbone of the project, as it manages communication between system components and processes sensor data for visualization. The choice of our MCU will balance power efficiency, ADC capabilities, processing power, and development support.

3.4.1 MSP430FR6968

The MSP430FR6989, developed by Texas Instruments, is an ultra-low power 16-bit microcontroller designed for embedded applications requiring efficient processing, extended battery life, and integrated analog peripherals. It belongs to the MSP430 family, well known for low power consumption and flexible peripheral set.

Some key features of the MSP430FR6989 include a 16-bit RISC (Reduced Instruction Set Computer) CPU with a clock speed up to 16 MHz, 128KB of FRAM (non-volatile, low-power memory with fast writes) and 2KB SRAM, 83 general-purpose I/O pins, multiple 16-bit timers with PWM (Pulse-Width Modulation) support, and an operating voltage of between 1.8V to 3.6V. It offers support for I2C, UART, and SPI communication protocols to interface with peripheral devices, a 12-bit SAR ADC with up to 16 input channels and configurable sample and hold functionality, and various power modes to suit operational requirements.

In terms of form factor, the MSP430FR6989 is available in both QFP (Quad Flat Package) and VQFN (Very-Thin Quad Flat No-Lead Package), which are both surface-mounted packages, making it suitable for PCB design flexibility. For development and testing, the MSP-EXP430FR6989 LaunchPad is available, providing a USB interface for programming and debugging and breakout headers for easy prototyping. Texas Instruments developed Code Composer Studio (CCS) as the official development environment for the MCU, but there are some alternative compatible environments available such as IAR Embedded Workbench and Energia (an Arduino-style development environment).

The MSP430FR6989 is a strong contender for HALO due to its low power consumption, built-in ADC, flexible peripheral set, and FRAM memory (specifically advantageous for data logging). However, its 16-bit architecture and relatively low clock speed might limit performance for complex signal processing that is required for the project's real-time visualization tasks.

3.4.2 ESP32

The ESP32, developed by Espressif Systems, is a high-performance and low-power microcontroller with integrated Wi-Fi and Bluetooth connectivity. It features a dual-core Tensilica Xtensa LX6 32-bit processor, running up to 240MHz with 520KB of onboard SRAM. There are 34 general-purpose I/O pins with support for touch sensing and capacitive inputs, along with support for

various communication protocols, including I2C, SPI, UART, CAN, and PWM. It's also equipped with a 12-bit SAR ADC with up to 18 input channels. With an operating voltage of between 2.2V and 3.6V, the MCU is adaptable to a variety of applications.

The ESP32 is available in various form factors, most notably the surface-mount variant, the ESP-WROOM-32. There are various devkits available for the ESP32, with most offering USB support and advanced peripherals, as well as a RISC-V-based power-optimized variant. The ESP32 is extremely well-supported, with many of the most common development environments offering specific ESP32 IDEs including Arduino IDE, Espressif IDE, MicroPython, PlatformIO, as well as extensive Microsoft Visual Studio packages.

In regards to HALO, the ESP32 is a very compelling option due to its high-performance capability, versatile I/O, and built-in wireless capability. Although our project scope does not call for wireless connectivity at the moment, having the option available to our team could be valuable. The built-in ADC is another advantage of this chip, with the only downside being negative reviews of its performance at higher frequencies. However, paired with the performance of the processor and our application likely not requiring usage of the ADC in the MHz range, I believe this is not cause for concern.

3.4.3 STM32

The STM32 family, developed by STMicroelectronics, is a series of MCUs utilizing 32-bit ARM Cortex-M processors with a wide range of capabilities. The family consists of chips with various performance levels, from low-power to high-performance real-time. STM32 MCUs come in various flavors (depending on performance needs), with the STM32F3, STM32F4, and STM32L4 being the most relevant for our application.

They feature ARM Cortex-M (M0, M3, M4, M7, or M33) processors running at clock frequencies between 72 MHz and 180 MHz, with 64KB to 2MB of Flash storage and up to 512KB of SRAM. The higher end models have over a 100 GPIO pins, with most pins supporting advanced peripheral functions. Multiple 16 and 32-bit timers are available, with support for I2C, UART, SPI, USB, CAN, and Ethernet communication protocols. Depending on the series, the onboard ADCs (there are multiple) are either 12 or 16-bit, with most supporting up to 24 channels.

Despite the high-performance characteristics of the STM32, it runs at an operating voltage of between 1.8V and 3.6V, making it ideal for a portable

application. In terms of form factor, there are various package sizes available, from LQFP-32 (low pin count) to LQFP-144 (high-performance applications). A popular devkit is the STM32 discovery kit, a feature-rich board for rapid prototyping. Like the ESP32, the STM32 is widely supported in terms of development environments, compatible with Arduino IDE, PlatformIO, Visual Studio Code, and STM32CubeIDE (official IDE) to name a few.

The STM32 offers a strong balance between processing power, ADC performance, and efficiency. The high-resolution ADCs make it a strong candidate for use with precise sensor data, and the extensive peripheral set and power management capabilities provide flexibility. The STM32 will require more effort to set up compared to similar MCUs, and although being widely supported, lags behind the popular ESP32.

3.4.4 Comparison & Selection

Table X: Microcontroller Comparison

Feature	MSP430FR6989	ESP32	STM32
Architecture	16-bit MSP430	32-bit Xtensa LX6	32-bit ARM Cortex-M7
Clock Speed	16 MHz	< 240 MHz	< 550 MHz
Flash Memory	128 KB FRAM	4 MB (external)	64 KB - 2 MB
SRAM	2 KB	520 KB	16 KB - 1024 KB
GPIOs	40	34	100+
ADC Resolution	12-bit (16 channels)	12-bit (18 channels)	12 to 16-bit (24 channels)
Operating Voltage	1.8V - 3.6V	2.2V - 3.6V	1.8V - 3.6V
Wireless Capability	None	Wi-Fi, Bluetooth	None
Interfaces	I2C, UART, SPI, ADC, Timer, PWM	I2C, UART, SPI, ADC, DAC, PWM, CAN	I2C, UART, SPI, USB, CAN, Ethernet, LTDC

For HALO, we need a microcontroller that is powerful enough to handle real time data collection from the Hall sensor array, efficiently process and visualize this

data on a touchscreen display, and support additional functionality like data logging to an SD card, while maintaining a minimum refresh rate of 60 frames per second. We selected the STM32H735IGT6, an MCU from the STM32H7 family. This decision was driven by several factors, mostly related to bottleneck and performance.

First, the STM32H735IGT6 features a 32-bit ARM Cortex M7 core clocked at 550 MHz. This ensures that the MCU can handle ADC sampling from the sensor array while also managing display rendering and touch inputs without lag. Its 564 KB of SRAM is necessary for handling graphical data and sensor readings, helping to maintain smooth performance without requiring external memory. Also, it includes 1 MB of embedded Flash memory, allowing us to store firmware and graphical assets directly on the MCU, eliminating the need for external storage and simplifying PCB design.

The most important deciding factor was the LCD-TFT Display Controller (LTDC) and MIPI DSI interface, which provide a direct, hardware-accelerated path for driving a display. This allows the MCU to efficiently communicate with our selected Rocktech capacitive touchscreen while offloading graphical tasks from the main CPU. The LTDC ensures that rendering remains consistent and smooth, while the MIPI DSI interface optimizes data transfer between the MCU and the display, reducing the processing power required for screen updates. This combination guarantees that our project can achieve and sustain the 60 FPS requirement without bottlenecks.

The STM32H735IGT6 also includes multiple high-speed ADCs for collecting sensor data from the array. The I2C and SPI interfaces allow integration with external peripherals such as the capacitive touch controller and SD card module for data logging. The inclusion of an SDMMC interface enables us to save visualization data directly onto an SD card, a key specification of our project, without too much additional development.

Overall this MCU was selected because it meets all computational, memory, and display requirements while keeping the design fully self-contained (without the need for external RAM, Flash, or co-processors). This choice ensures that HALO remains a standalone, high-performance embedded solution, capable of delivering real-time, interactive visualization of magnetic field data at a smooth and responsive frame rate.

3.4.5 Development Board Selection

To assist in the development and validation of HALO's firmware and interface, we have chosen to go with our MCU's standard development board, the STM32H735G-DK Discovery Kit. This decision was motivated by the kit's close alignment with our final custom PCB design, while also providing an integrated and software-focussed development environment. It features the same MCU we have selected, as well as a very similar display that will allow us to hone in our UI design and functionality prior to transferring over to our final hardware, ensuring compatibility between the two. The board also includes native LTDC and MIPI DSI support, features we planned to use on the final product, allowing us to gain familiarity with STM's TouchGFX graphical library. An STLINK-V3 debugger is also onboard the devkit, which will help to streamline testing without requiring the use of breakout connectors or external hardware. While this board does include some extra components (external RAM and Flash that will not be available on the final hardware), it remains the best option for early prototyping, enabling us to validate ADC sampling, touchscreen response, and visualization.

3.5 Sensor Addressing

Sensor addressing will be an important part of how we will be collecting our data. Various addressing techniques exist, including direct GPIO connections, multiplexing, and digital protocols such as I2C and SPI. The choice of addressing method impacts system complexity, power consumption, and speed.

In this project, we are utilizing analog Hall effect sensors, which do not have built-in digital communication protocols such as I2C or SPI. This means that conventional digital addressing techniques cannot be applied, necessitating an alternative approach to efficiently read data from multiple sensors. Instead, we must implement an addressing scheme that allows the microcontroller to access and process data from each sensor individually without excessive wiring or hardware constraints.

To achieve this, we must convert the analog signals from the Hall effect sensors into a format that can be processed by the MCU. This requires an analog to digital converter to interpret the sensor outputs as meaningful data. Since the number of available ADC channels on the MCU is limited compared to the number of sensors in our system, we need a method to expand the MCU's ability to read multiple analog signals. This is where multiplexing becomes essential, as it enables the selection and reading of multiple sensors using a reduced number of ADC channels. Alternatively if enough GPIO pins are

available we could utilize ADCs for every sensor output to have a direct connection between the sensors and the GPIO pins.

3.5.1 Direct GPIO and ADC

GPIO stands for General Purpose Input and Output. If a microcontroller has enough available GPIO pins, each sensor can be connected directly to its own dedicated pin. By default, GPIO pins are typically configured for digital input, reading high or low voltage levels. This makes them unsuitable for reading analog signals directly. However, if the microcontroller lacks sufficient built in ADC channels, an external ADC can be used to convert analog signals into digital data that the MCU can process.

This approach offers the fastest data acquisition speed, assuming the selected MCU has enough GPIO pins but not enough ADC channels. Since each sensor has a dedicated connection, signals can be read in parallel without the delays associated with multiplexers. Additionally, this method simplifies circuit design by reducing the need for additional switching components.

However, the approach has some trade offs. It increases the number of required connections, which can impact PCB layout complexity, physical space, and cost. Power consumption may also rise due to the additional GPIO usage and external ADC requirements. In cases where space, cost, or power efficiency are critical factors, using multiplexers or shared ADC channels might be a more practical alternative.

3.5.2 MUX

A multiplexer (MUX) is a device that allows multiple input signals to be transmitted through a single output, which is particularly useful when working with a limited number of ADC channels on an MCU.

In this project, we will use multiplexing to read the signals from multiple analog Hall effect sensors to a single ADC input. This enables the MCU to read data from multiple sensors without requiring a separate ADC channel for each hall effect sensor.

How we will utilize MUX:

- Inputs: The input pins of the MUX are connected to the output signals of multiple Hall effect sensors. Depending on the size of the mux (8:1 or

16:1) a MUX could be handling the information from one or two rows of sensors.

- **Select Lines:** These control which input is connected to the output. If the MCU we choose has enough available GPIO pins, it will directly control these lines. Otherwise, a binary counter may be used to cycle through the inputs automatically.
- **Enable (or Encoder) Pin:** This pin is used to activate the MUX. The MCU controls this pin to enable data transfer when needed.
- **Output:** The selected sensor's signal is sent to the MUX's output, which is then connected to an ADC for digital conversion.

Handling the ADC Connection

If the MCU has an available ADC pin, the MUX output can be connected directly to it. If multiple MUXs are used or the MCU lacks sufficient ADC channels, an external ADC module may be used to handle the analog-to-digital conversion before sending the data to the MCU. The decision to use an internal or external ADC depends on the number of available ADC channels on the MCU and the number of MUXs required for the sensor array.

Table X: Multiplexer Comparison

MUX	Num Channels	R on (ohms)	Voltage	Switch Time (Ton, Toff) (Max)
CD74HC4067M96	16:1	160	2V ~ 6V	51ns, 49ns
ADG706BRUZ-REEL	16:1	4.5	1.8V ~ 5.5V	32ns, 16ns (Typ)
74HC4051D,653	8:1	120	2V ~ 10V	51ns, 42ns

3.5.3 Row/Column Multiplexing

Multiplexer row/column addressing is a method that uses two multiplexers to read data from a sensor matrix. One MUX is used for selecting rows, and the other is used for selecting columns. For example, in an 8x8 matrix, this method reduces the number of multiplexers to just 2, and each multiplexer would only need to be an 8:1 MUX, which requires only 3 select lines. However, the select lines for the row and column MUXs must be independent of each other, meaning they cannot be synchronized but must be coordinated to select the correct sensor at the row-column intersection.

With this method, the number of lines required to control the matrix would be 6 select lines for the two MUXs and 1 enable pin, totaling 7 lines. While this reduces the hardware required, the speed of reading sensors would also decrease, as only one sensor can be read at a time. This is in contrast to other multiplexing methods where multiple sensors could be read simultaneously, resulting in a slower overall reading process.

A key challenge in this setup is properly controlling the output of the sensors. Since each sensor will be outputting a voltage, sensors not currently selected must have their outputs disabled to avoid interference. This can be accomplished by using tri-state buffers to place the unselected sensors into a high-impedance state, effectively disconnecting them from the data line and isolating the outputs of the active sensor.

3.5.4 ADC

ADCs work by sampling an analog signal (a continuous range of voltages) at regular intervals and then converting each sample into a corresponding digital value. The resolution of the ADC is directly related to how precisely it can divide the analog signal's voltage range. Since this is done with bits, the number of discrete levels it can divide is 2^n where n is the number of bits. For instance, with an 8-bit ADC, the signal range is divided into 256 discrete levels, while a 10-bit ADC can divide the range into 1024 levels.

The higher the resolution, the more levels the ADC can use to represent the analog signal, which means that small variations in the input signal can be detected with greater precision.

The minimum detectable voltage change, known as the least significant bit, is determined by dividing the input range by the total number of levels.

$$\frac{\text{input range}}{2^n}$$

For example, with an 8-bit ADC and an input range of 1 volt, the voltage increment is 3.92 mV. This means that if the voltage of the input signal changes by less than 3.92 mV, the ADC will not be able to detect this change, and the value may stay the same. The table below represents the minimum voltage increments for different resolutions with an input range of 1 volt.

Table X: Digitizer step size vs. resolution [6]

Resolution	Ideal Dynamic range	Minimum Voltage Increment
8 bit	256:1	3.92 mV
10 bit	1024:1	0.98 mV
12 bit	4096:1	0.244 mV
14 bit	16384:1	61 μ V
16 bit	65536:1	15 μ V

The input range of the ADC is determined by its reference voltage. Typically, ADCs are designed to convert an input voltage that falls within a fixed range. For example, from 0 to 3.3V or 0 to 5V. If the analog signal exceeds this range, it will either be clipped or not properly detected.

To ensure that the analog signal falls within the ADC's input range, you can amplify the signal using an operational amplifier (op-amp) or other types of signal conditioning. This can be particularly useful if the signal's voltage range is too small compared to the ADC's input range.

The sampling rate determines how frequently the ADC samples the analog signal. The faster the sampling rate, the more frequently the ADC will update the digital representation of the signal, allowing for a smoother and more accurate tracking of fast changing signals which is an outcome that we are hoping for with our project.

According to the Nyquist theorem, to capture a signal without aliasing, the sampling rate must be at least twice the highest frequency component of the signal. This ensures that each waveform is sampled sufficiently to capture its full shape. If you sample too slowly, the ADC may miss important variations in the signal, leading to distortion.

For example, if you are sampling a sine wave with a maximum frequency of 1 kHz, the sampling rate should be at least 2 kHz to avoid aliasing. For the purposes of measuring changing magnetic fields we will try to ensure that our ADCs can over compensate to avoid aliasing.

Many MCUs come with built in ADC channels, however if we do not have enough ADC channels or if the ADC channels are not a high enough resolution to provide meaningful information then we may need to add external ADCs.

Table X: ADC Comparison

ADC	Number of Bits	Sampling Rate (Per Second)	Data Interface	Voltage - Supply, Analog
ADC081C021CIMKX /NOPB	8	188.9k	I2C	2.7V ~ 5.5V
AD7940BRJZ-REEL 7	14	100k	SPI, DSP	2.5V ~ 5.5V
MCP3021A5T-E/OT	10	22.3k	I2C	2.7V ~ 5.5V

3.5.5 I/O Expander

I/O expanders can be used to increase the number of available control pins, reducing the burden on the MCU. This is particularly useful when the MCU does not have enough GPIO pins. In our case, using four 16:1 multiplexers requires a total of 4 select pins, 1 encoder pin, and 4 ADC pins. Since all MUXs will share the same select and encoder signals, there is no need for independent enable/disable control for each MUX.

If we were to utilize multiplexor row/column addressing then the number of GPIO pins would increase, but the number of ADC pins would be decreased to 1. This is because the select pins on the MUXs would need to be independent of each other bringing the number of GPIO pins for selection from 4 to 6. With the enable pin that would require 7 pins intotal dedicated to multiplexing.

I/O expanders communicate with the MCU using various protocols, each affecting how many GPIO pins are needed. The most common protocols are I2C and SPI:

I2C uses two lines:

- SDA (Serial Data Line) for bidirectional data transfer and
- SCL (Serial Clock Line) for synchronizing communication.

While I2C is efficient in minimizing pin usage, it is relatively slow.

SPI is faster but requires four lines:

- CS (Chip Select): Selects which sub-device the main MCU communicates with.
- SCLK (Serial Clock): Carries the clock signal from the main.
- MOSI (Main Out, Sub In): Data sent from the main to the sub.
- MISO (Main In, Sub Out): Data sent from the sub to the main.

SPI has the advantage of supporting daisy chaining, allowing multiple I/O expanders to be connected while still using the same four lines. However, not all I/O expanders support daisy chaining, and some may require separate CS lines for each device, increasing pin usage.

Table X: I/O Expander Comparison

I/O Expander	Number of I/O	Interface	Clock Frequency	Voltage - Supply
TCA9555PWR	16	I2C, SMBus	400 kHz	1.65V ~ 5.5V
MCP23S17T-E/SS	16	SPI	10 MHz	1.8V ~ 5.5V
TCA6408ARSVR	8	I2C, SMBus	400 kHz	1.65V ~ 5.5V
MCP23S08T-E/SO	8	SPI	10 MHz	1.8V ~ 5.5V

3.5.6 Binary Counter

Binary counters or decoders are others that could be utilized to select the select pins on our MUXs if we were to run out of suitable pins on our MCU.

Binary counters can output a count in binary corresponding to the number of available output pins. For example, 3 output pins count from 0 to 7, and 4 output pins count from 0 to 15. A binary counter can be used to activate the select pins on a MUX if the MCU does not have enough available GPIO pins to do so. The counter is controlled by a CLK pin, which connects to a clock to control the counting, and a CLK reset pin to reset the count.

Table X: Binary Counter Comparison

Counter	Bits	Count Rate	Voltage - Supply
MC14516BDR2G	4	8 MHz	3 V ~ 18 V
SN74HC590ADR	8	24 MHz	2 V ~ 6 V
SN74LV163ADR	4	90 MHz	2 V ~ 5.5 V
74HC161D,653	4	48 MHz	2 V ~ 6 V

3.5.7 Decoder

A decoder is a combinational logic device that converts a binary number into one of several output lines. A decoder takes a binary input and activates one of its output lines based on the input value. For example, a 2-to-4 decoder would have 2 input bits and 4 output lines. This means that for any combination of the 2-bit input (00, 01, 10, or 11), one of the 4 outputs will be activated, while the others remain inactive.

This means that we could utilize a decoder to ensure that only one of our MUXs is active at a single time. This could be accomplished by connecting the enable pins of our MUXs to the outputs of the decoder.

Table X: Decoder Comparison

Decoder	Circuit	Voltage - Supply
SN74HC138DR	1 x 3:8	2V ~ 6V
74LVC138APW,118	1 x 3:8	1.2V ~ 3.6V
SN74HC139DR	1 x 2:4	2V ~ 6V
74HC139D,653	1 x 2:4	2V ~ 6V

3.5.8 Comparisons & Selection

Given the STM32 MCU's 100+ GPIO pins and 24 ADC channels, we have more than enough resources to directly connect the sensor array without needing additional components like I/O expanders or binary counters. It is unlikely that expansion will be needed for other parts of the project since the addressing will not utilize too many pins. This simplifies the design and avoids unnecessary hardware.

With 100+ GPIO pins it is possible to utilize the direct GPIO and external ADC method however we have decided against this for a number of reasons. Having a dedicated ADC for each sensor would take up far too much space on the board and more than double the size of the sensor array. Having to solder every sensor onto the board already presents a challenge where we could introduce error, but nearly doubling the total components that need to be soldered on creates potential for error that is undesirable. Each of the ADCs would also greatly increase the power consumption which would put us further away from our goal of a low power device. While it may improve the speed that we could collect data, the cost for that many ADCs would inflate our budget by more than what the team would like.

When comparing the two multiplexing methods, row/column addressing offers some advantages, but the drawbacks make it less ideal for our application. Row/column addressing would reduce the number of MUXs needed to read the sensor data, but it would also require more GPIO pins and would take longer to read all the sensors. While this method would allow us to use only a single ADC pin, we have 24 available ADC channels, so this is not necessary. Additionally, isolating the outputs of the sensors becomes more complicated, requiring tri-state buffers or other methods to prevent interference, ultimately adding complexity to the design.

Instead, the optimal choice is to use 4 16:1 MUXs, with each MUX reading two rows of the sensor matrix. This solution uses 4 ADC channels, simplifies sensor selection (requiring only 4 GPIO pins for selection and 1 for enable), and offers faster sensor reads. Synchronizing the select pins is also easier, which streamlines the process of selecting sensors to read. Because we will be using an MCU with enough ADC channels, all of the MUXs will be able to transmit data at the same time meaning they can all be enabled simultaneously. This means that a decoder would not be required.

If we need even faster sensor reads in the future, we could consider using 8 8:1 MUXs and utilizing 8 ADC channels. However, this would require more space and increase the complexity, making it a less appealing option unless speed becomes a significant concern.

3.6 User Interface

When determining a display and output solution for HALO, we've narrowed down our selection to two primary options: an onboard RGB LED display and an onboard embedded touchscreen display and GUI combination. Each of these

presents us with unique advantages and disadvantages when considering functionality, cost, complexity, ease of implementation, ease of use, and overall suitability for the project's needs, especially as throughout development our needs may change. Through analysis of these factors we will determine the selection of the most appropriate solution that suits our projects current and future needs.

3.6.1 RGB LED Display

Utilizing an RGB LED display is a simple and cost-effective method that will allow us to visualize heat map data, acquired from the hall-effect sensor array, by using an array of light-emitting diodes that uses a combination of red, green, and blue light sources in a single package. These kinds of displays range from small, segmented indicators to large matrix panels capable of displaying simple text, symbols, and even low resolution graphical interfaces. For HALO, an LED display could be used to represent real-time magnetic field intensity using various color schemes or graphical patterns, providing a highly visible, low power, and simple means of data representation.

One of the key advantages that makes an LED display an attractive option is its cost-effectiveness. Basic models can be acquired for as little as \$20 depending on size and resolution, with custom solutions being even cheaper and more customizable (although likely less dense). Power consumption would be very manageable, especially if considering an OLED or low-brightness LED matrix. Unlike LCDs that require backlighting, an LED display does not require additional illumination, increasing its efficiency.

From an implementation perspective, RGB LED displays are generally easy to integrate with microcontrollers, as many models use serial communication protocols including SPI, I2C, and UART, which simplifies the wiring and minimizes the number of pins required. Some displays can also be used with direct GPIO control allowing for rapid updates, but may require some additional external hardware to ensure the LEDs are driven effectively. Popular models like the Adafruit 64x64 RGB LED Matrix provide well-supported libraries for microcontrollers like the STM32 and ESP32, reducing development complexity.

However, RGB LED displays have significant limitations in terms of resolution and interactivity. Even with a dense LED matrix, the level of detail that can be conveyed is very limited. This makes it more difficult to display complex numerical values, intricate visualizations, or real-time graphs. If a high level of detail were to be implemented on an LED matrix display, it would have to be physically large, increasing both power consumption and cost. Another

drawback is that LED displays lack native interactivity support, meaning that any user input would need to be handled through external physical buttons, rotary encoders, voice commands, or some other kind of interface. This would add additional complexity to the overall system design, as an independent input method must be integrated alongside the display.

Another key factor that must be considered is scalability (or adaptability). If future modifications to HALO were to require more detailed visualization of data, or different display modes, transitioning away from an LED matrix would be necessary. While their simplicity is useful for initial implementation, LED displays do not provide a long-term scalable solution especially if interactive data visualization or high resolution output is required.

3.6.2 Touchscreen GUI

In contrast to an RGB LED display, an embedded touchscreen display offers a far more sophisticated and interactive means of visualizing data. These displays range from small TFT touchscreens to larger full-color capacitive screens, enabling real-time graphical representations, user interactions, and enhanced control capability. For HALO, an embedded touchscreen could be used to visualize the heatmaps of magnetic field intensities, provide real-time graphs, allow for the addition of various display modes, and even potentially allow for users to interact with the data in a dynamic fashion.

Likely the biggest advantage of using a touchscreen over an LED matrix display is its native support for full interactivity. Unlike an LED display, which is strictly an output medium, a touchscreen allows for direct user input, eliminating the need for development of external control interfaces and streamlining the integration of hardware and software. This significantly simplifies system design as the user controls are built directly into the interface, and if needed can be completely overhauled to meet the needs of our design as it matures. Additionally, interaction like zooming, selecting data points, or switching between different display modes is made significantly more intuitive through the use of a touch-based interface, and allows for the overall system to fit in a smaller footprint.

However, this added functionality comes at the cost of increased cost and complexity, on top of greater processing requirements. Embedded touchscreens typically cost between \$30 - \$200 depending on size, resolution, display technology, and touch technology. The power consumption of an LCD with backlighting is significantly greater than that of an LED matrix, which could be a concern if HALO's final power budget is highly constrained.

The implementation complexity of a touchscreen display is another important consideration, which presents a potentially critical bottleneck in our development cycle. Unlike an LED display, which can be updated with simple framebuffer write operations, a touchscreen requires a graphical user interface (GUI) framework to manage rendering, touch input, and real-time updates. Libraries such as LVGL (Light and Versatile Graphics Library), TouchGFX, or Nextion's HMI tools can significantly aid in UI development, but again require additional processing power as well as sufficient RAM to handle framebuffer storage. Microcontrollers working with limited processing resources may struggle or fail entirely in rendering high-resolution visuals while simultaneously managing sensor data acquisition and system control.

To address this challenge, some touchscreen displays come with integrated processing units that handle the UI rendering independently. Displays such as the Nextion HMI Series or WaveShare TFT Touchscreens offload rendering tasks, reducing the computational burden on the microcontroller. However, these solutions can increase cost and require additional development to configure/customize the interface.

When selecting an embedded touchscreen, some of the key specifications we are considering include:

- Screen size and resolution - Higher resolution allows for more detailed visualization but at the cost of increasing data processing requirements significantly.
- Touchscreen type - Resistive screens work with any input (stylus, gloves) but are less responsive. Capacitive screens provide a smoother experience but require direct skin contact.
- Communication interface - SPI-based displays are common but have limited refresh rates. Parallel RGB, LVDS, or HDMI connections are better suited for high-speed rendering.
- Development support - Some screens have pre-built firmware and GUI tools, making implementation much easier.

From a long-term perspective, using an embedded touchscreen would provide far more flexibility over using an LED display. It would expand our options to potentially allow for the development of interactive features, real-time data visualization, and future expansions without requiring significant hardware changes. However, the cost of this added flexibility includes higher computational demands, increased power consumption, and a more complex software development process.

3.6.2.1 Capacitive vs. Resistive Touchscreens

When considering touchscreen technology, two varieties are available in the form of capacitive and resistive displays. Capacitive touchscreens are more responsive and support multi-touch gestures, making them the better choice when it comes to smooth and intuitive interactions. They rely on the conductive properties of the human body, meaning that direct skin contact or a capacitive stylus is required for them to function. Their high transparency and lack of mechanical layers allow for better color accuracy and brightness. Capacitive touchscreens also fair better over time as they are more durable and do not suffer from mechanical wear due to repeated presses. However, they are more expensive than their counterparts, and do not work well with non-conductive inputs like gloves.

On the other hand, resistive touchscreens rely only on pressure applied to the screen, making them compatible with any input method including gloves, styluses, or any other firm object. They are generally more affordable and function in a wider range of environmental conditions. But, they are less responsive and do not support gestures requiring multiple inputs. The mechanical layers situated above the display also reduce the clarity and brightness of the screen compared to capacitive touchscreens. For our needs, since the touchscreen needs to function mostly as a data output with limited interactivity (mostly for navigation through menus and simple 'button presses'), a resistive touchscreen is the better option as it will help to reduce our overall costs and reduce the environmental constraints in which our device can operate.

3.6.3 Comparison & Selection

The choice between an RGB LED display and an embedded touchscreen ultimately depends on the current and potential future requirements of HALO. If the primary need was to display simple status indicators of the basic magnetic field intensity changes, the LED display would be the most cost effective and efficient choice. Since it is relatively easy to implement, requires minimal computational resources, and has an abundance of support for various microcontroller configurations and communication protocols, it would have the advantage if its lack of interactivity and limited resolution were not a hindrance to the goals of the project.

Since we require real time, high resolution data visualization in tandem with the capability for user interaction, the embedded touchscreen display is the more suitable choice. It enables intuitive controls, supports detailed graphical representations, and offers a scalable interface that can dynamically grow and

contract to fit the needs of HALO. While it is the more capable solution, it comes at the cost of additional hardware and software complexity, demands greater processing power, and increases the overall system cost.

Below is a comparison of some of the displays we considered, along with their features and our final selection.

Table X: Screen Comparison

Feature	Best LED Option	Best Touchscreen Option
Cost Efficiency	Adafruit 32x32 RGB Matrix (\$40)	Rocktech 4.3" RK043FN66HS-CTG (\$45)
Best Resolution	Waveshare 64x64 RGB (\$70)	NHD-5.0-800480AF-ASX P-CTP 5.0" Capacitive Touchscreen (\$80)
Fastest Refresh Rate	Pimoroni Interstate 75 (>120 FPS)	Nextion 5.0" Enhanced HMI (>80 FPS)
Easiest Development	Pimoroni (Prebuilt drivers)	Nextion (GUI pre-built tools)
Interactive UI	No interactivity	Nextion or Rocktech
Lowest MCU Load	Pimoroni (integrated driver)	Nextion or Rocktech (offloads rendering)

Since our MCU has sufficient RAM, Flash, and graphics processing capabilities, we chose to go with the Rocktech 4.3" as it offers a visually pleasing, bright and responsive display that is capable of achieving our desired frame rate, and has a highly customizable UI to accommodate the various GUIs we need to incorporate for data visualization. It also has an LTDC (LCD-TFT Display Controller) that can integrate with the same functionality on the MCU, offloading some of the graphics processing and allowing for the most efficient usage of CPU resources.

3.7 Communication Protocols

Communication protocols are the rules for sending information between devices. In the case of our project, we will need to select a communication

protocol that will be able to facilitate communication between our touchscreen and our MCU.

Protocols can be described as being simplex, half duplex or full duplex. A simplex protocol is one that only allows for a single direction of communication, meaning device A could send information to device B, but device B could not send it back to device A. Half duplex allows for communication in both directions but not at the same time. This means that device A could send information to device B and device B could send information to device A, but while one of them sends information, the other must be on standby and hold the transmission of their information. Full duplex allows for simultaneous transmission of information. This means that device A and device B can both be sending and receiving at the same time without the need to hold.

There is also synchronous and asynchronous transmission. Synchronous transmission uses a clock signal to ensure that the devices are in sync. This requires an extra line for the clock. Asynchronous communication does not use a clock line but instead relies on start and stop bits in the transmission to signal the beginning or end of a transmission. It also requires that both the devices have the same baud rate. Synchronous communication is used for high speed continuous data transfer while asynchronous is used for low speed intermediate transfer.

3.7.1 I2C

I2C (Inter-Integrated Circuit) is a widely used communication protocol for low-speed communication with peripheral devices. It operates in half-duplex mode and utilizes synchronous communication with only two lines.

- SDA (Serial Data Line), which is used for bidirectional data transmission
- SCL (Serial Clock Line), which is used to synchronize the transmission of data using the clock

I2C typically uses a 7-bit address, though in rare cases, a 10-bit address may be used for applications that require more devices on the bus. The standard speed of I2C is 100 kbit/s (standard mode), but it can also operate at 400 kbit/s (fast mode) for faster communication. In some cases, I2C can reach speeds of 3.4 Mbit/s in high-speed mode.

In I2C, data is transmitted one byte at a time. After each byte, the receiver sends an acknowledgment (ACK) bit. This acknowledgment confirms that the byte has been successfully received. If the sub device cannot receive the data, it sends a

negative acknowledgement (NACK) to indicate an issue. If the main does not receive an ACK from the sub, it knows there was an issue with the transmission, and it may retry or take corrective action.

If a device does not respond with a NACK, the main device can decide whether to resend the data or handle the error in another way. Conversely, if the sub device receives a NACK from the controller, it knows that the main has stopped requesting more data, and it should halt further transmission.

I2C also allows for multiple masters and multiple slaves. Only one main will be able to use the bus at a time, and they will be ranked on priority so that the highest priority main will use the bus first.

3.7.2 SPI

SPI (Serial Peripheral Interface) is a widely used communication protocol for high-speed communication with peripheral devices. It operates in full-duplex mode and utilizes synchronous communication with four lines.

- CS (chip select) which allows the main to choose which sub to send the signal to
- SCLK (serial clock) carries the clock signal from the main
- MOSI (main out, sub in) data sent from the main to the sub
- MISO (main in, sub out) data sent from the sub to the main

The four lines enable full-duplex transmission, allowing data to be sent and received simultaneously on both MISO and MOSI. When transmitting data, the MOSI on the main connects to the MOSI on the sub, and the MISO on the main connects to the MISO on the sub. Data is transmitted with the most significant bit first. All subs share the same MISO and MOSI lines but each has its own CS line. If additional subs are needed, the IO must be expanded to accommodate the extra CS lines.

SPI does not have a maximum clock speed and it doesn't use addresses for sending information. Instead, it uses the CS line to choose the sub to send information to. It also does not have built-in error correction.

Because its signal degrades over long distances it is typically only utilized for short distance applications like connecting devices on the same board.

3.7.3 UART

UART (Universal Asynchronous Receiver-Transmitter) is a hardware communication protocol used for asynchronous serial communication. It can operate in three modes: simplex, half-duplex, and full-duplex. UART uses up to two lines for communication:

- TX (Transmit) is used to send data to a receiver
- RX (Receive) is used to send data back to the transmitter

In its simplex configuration, only the TX line is needed because communication occurs in one direction. In half-duplex and full-duplex configurations, both TX and RX lines are required. However, in a half-duplex, only one line can transmit information at a time. These lines only support a connection between a transmitter and a receiver meaning the protocol does not allow for the connection of multiple devices.

When UART data is transmitted, it is done using framing so that the receiving device knows when a transmission starts and stops. This is done by changing from a logical high (1) to a low (0) for the start bit. Following the start bit are the data bits. In some cases, a parity bit follows the data bits to check for errors in transmission. Finally, the stop bit switches back to logical high.

Because UART is asynchronous, it does not require a clock signal. However, this means other parameters must be synchronized between the transmitting and receiving devices, such as voltage levels, baud rate, parity bit, data bit size, and stop bit size.

Additionally, flow control mechanisms (either hardware-based with RTS/CTS or software-based with XON/XOFF) may be used to manage data flow, especially in more complex systems.

3.7.4 Comparison & Selection

Table X: Communication Protocol Comparison

Spec	I2C	SPI	UART
Speed	Slow	Fast	Slow
Distance	Long	Short	Short
Clock	Synchronous	Synchronous	Asynchronous
Configuration	Half duplex	Full duplex	Simplex, Half

			Duplex, or Full Duplex
Connectable devices	Multi	Multi	Single
Lines	2	4	2

For our project, we need a touch screen that can communicate with an MCU, and vice versa, to process user inputs and update the screen in real-time. This requires full-duplex communication, where both the MCU and touch screen can send and receive data simultaneously. We also want the screen to be updated as quickly as possible, which means we'll need a high-speed communication protocol. Given that the device will be small and doesn't require long-range communication, we've chosen SPI as our communication protocol. SPI offers the fast data transfer speeds we need while maintaining simplicity and reliability for short-range communication.

3.8 Languages, Tools, and Development Flow

3.8.1 Embedded Development

3.8.1.1 C

C is a widely used language for embedded systems and microcontroller programming because it is efficient, portable, and interacts with hardware at a low level. For our system, C is a strong candidate for the MCU because it provides direct memory access and deterministic execution, which are crucial for real-time applications like reading from an ADC and updating the touch screen display. It is well-supported by STM32CubeIDE (our chosen development environment), ensuring that we have robust toolchain compatibility. Also, most STM32 HAL (Hardware Abstraction Layer) and LL (Low-Layer) libraries are written in C, making it the most compatible option. C also benefits from an extensive community of developers, which can ease the development process for us as we will have a wealth of documentation and troubleshooting resources. While C is procedural and lacks object-oriented features, it still allows for us to develop highly optimized and well structured code. Given our real time constraints and reliance on embedded peripherals, C remains the most practical choice.

3.8.1.2 C++

C++ builds on C with object-oriented programming features, which could provide modularity and scalability benefits. One of its main advantages is the ability to implement abstraction for our peripherals, which can result in a cleaner and more reusable codebase. The language also supports the Standard Template Library (STL), which may not be the most ideal for embedded systems due to memory constraints, but can still improve development efficiency by providing some useful data structures.

However, C++ introduces additional complexity and runtime overhead, such as virtual function calls, which can negatively impact performance on an embedded system. While STM32CubeIDE does support C++, the available STM32 HAL and LL libraries are primarily written in C, necessitating a mix of C and C++ should we choose to utilize embedded programming features. Given our needs, sticking with C for core firmware development is the most practical choice, though we may use C++ selectively for higher-level abstractions like managing the display interface.

3.8.1.3 Comparison and Selection

Table X: Programming Language Comparison

Feature	C	C++
Performance	High, deterministic execution	Slightly lower due to OOP overhead
Memory Usage	Low, efficient memory management	Higher due to vtables, OOP features
Hardware Interaction	Direct access, low-level control	Can use abstraction, but requires mixing with C for HAL/LL compatibility
Toolchain Support	Fully supported in STM32CubeIDE	Supported, but HAL/LL libraries require C interop
Code Modularity	Procedural, structured programming	Supports OOP, better for abstraction and reuse
Community Support	Large embedded systems community	Smaller but still viable for embedded systems
Real-time Suitability	Excellent, predictable execution	Potential overhead from dynamic features

Ease of Development	Simple but requires careful structuring	More structured but adds complexity
---------------------	---	-------------------------------------

C is the preferred language for our embedded development due to its efficiency, direct hardware access, and full compatibility with the STM32CubeIDE toolchain we will be using. It ensures deterministic execution, which is crucial for real-time tasks such as ADC readings and touchscreen updates. While C++ offers modularity through object-oriented programming, the additional runtime overhead and compatibility challenges make it less suitable for core firmware development. However, C++ may be used selectively for higher-level abstractions, such as display management.

3.8.2 Graphics Rendering

3.8.2.1 Python

Python is widely used for data visualization and analysis, making a potential candidate for generating the heat map. However, several challenges present themselves when considering Python for real time embedded applications. It is an interpreted language that requires significant processing power and memory, making it unsuitable for the resource-constrained environment provided by our MCU. While MicroPython exists it isn't optimized for real-time tasks like ADC data acquisition and display rendering. If Python is used at all, it would be for external post-processing of sensor data for debugging or analysis purposes. Overall, Python will not be a feasible option for real-time heat map rendering.

3.8.2.2 Matlab

Matlab is a powerful tool for mathematical modeling, data visualization, and signal processing. It offers Simulink and embedded code generation tools, but several barriers exist when considering using Matlab for real time embedded execution. While Matlab can generate C code using Embedded Coder, this process is often complex and not optimized for real time performance. Also, Matlab is not designed for embedded systems where low-level hardware access and deterministic timing are required. However, Matlab could still be useful for offline simulation and algorithm development, particularly in testing heat map processing techniques before implementing them in C on the MCU. Despite the advantages Matlab offers, it is not an entirely viable option.

3.8.2.3 Open Source and Proprietary Graphics Libraries

Since Python and Matlab are unsuitable for real time heat map rendering on the STM32, alternative solutions must be explored. One option is to use LVGL (Light and Versatile Graphics Library), which is an open source graphics library optimized for embedded systems that support touchscreens and real time visualization. Another option is to use the STM32 graphics framework TouchGFX, which is specifically designed for use with STM32 MCUs. TouchGFX provides efficient rendering capabilities and integrates directly with STM32CubeIDE, which will likely result in us using it for the majority of GUI development.

3.8.2.4 Comparison and Selection

Table X: Programming Language Comparison

Feature	Python	Matlab	LVGL	TouchGFX
Performance	Low (interpreted, high overhead)	Moderate (code generation possible but inefficient)	High (optimized for embedded systems)	High (designed for STM32)
Memory Usage	High (requires significant RAM)	High (not optimized for MCUs)	Low (optimized for embedded)	Low (efficient STM32 integration)
Real-time Suitability	Poor (interpreted, non-deterministic)	Poor (not designed for real-time execution)	Good (designed for embedded graphics)	Excellent (real-time STM32 graphics support)
Ease of Development	High (simple syntax, many tools)	Moderate (powerful but complex)	Moderate (requires embedded knowledge)	Moderate (STM32-specific learning curve)
Hardware Interaction	Limited (not MCU-friendly)	Limited (mainly for simulation)	Direct framebuffer access	Optimized for STM32 hardware
Integration with STM32	Poor	Poor	Moderate	Excellent (native STM32 support)
Best Use Case	Post-processing, debugging	Algorithm simulation, testing	Lightweight GUI on embedded devices	Primary GUI development for STM32

Python and Matlab are unsuitable for real-time heat map rendering due to their high resource requirements and lack of embedded support. While they may still come in handy for offline analysis and simulation, they cannot be directly used

on the STM32. For embedded graphics, LVGL and TouchGFX are the most viable options. TouchGFX is the preferred choice due to its seamless integration with STM32CubeIDE, optimized rendering performance, and direct hardware compatibility. LVGL remains an alternative if a lightweight option is needed, but TouchGFX is expected to be the primary graphics framework for our project.

3.8.3 GitHub

GitHub will be used as the version control system for our software development. The repository will be structured to include firmware, visualization, and documentation. The firmware directory will contain the C and/or C++ source code responsible for MCU operation. The visualization section will house code related to heat map rendering, and documentation will include reports, design documents, and notes. The repository may also contain binary files for precompiled firmware and relevant data. Using GitHub ensures that all team members can collaborate effectively, track changes, and resolve conflicts systematically. Version tracking will allow us to maintain a history of modifications and revert to previous versions if necessary.

3.9 Power Supply Unit

The power supply unit (PSU) of the system is comprised of the battery, battery connection, power regulators, and battery management chips. The PSU is dependent on every component in the system, so a table is included below that outlines each component's supply voltage range and maximum current drawn.

Table X: Circuit Components

Component	Supply Voltage Range (V)	Maximum Current Drawn (mA)
Hall-effect Sensors	2.5 - 38	3.6
SD Card Reader	< 17	500
MCU	1.62 - 3.6	360
Multiplexers	2.5 - 6	25
Decoders	2 - 6	.001
Touchscreen	3.0 - 3.6	40

3.9.1 Batteries

Every system needs a power source. It's just a matter of choosing the right type for the project. For HALO, we're making a handheld magnetic sensor, so maneuverability, weight, and size are the most important concerns. A system that weighs too much or is too unwieldy, loses its convenience, and thus our project which is focused on making a convenient way of imaging magnetic fields loses much of its benefit. Secondly, battery life and cost need to be considered as well. It is worth mentioning that a direct outlet connection is a possibility, but being tethered to a wall infringes on the tenants of our project as stated above.

3.9.1.1 Lead-acid Batteries

Lead-acid batteries were developed in the 19th century and have not undergone much change in design since. They are very dependable but aren't widely used in electronics because of their weight, size, and low energy densities. Lead-acid batteries have one of the lowest gravimetric energy densities as well as one of the lowest volumetric energy densities. Lead-acid batteries will often find utilization in systems that value dependability above all else or systems that have a large power storage requirement. Examples of such systems include emergency systems such as sirens and backup power. They're also utilized by renewable energy farms and automobiles. It's also worth noting that lead-acid batteries are quite affordable which lends to their continued use. [B1]

While the affordability of lead-acid batteries are alluring, their low energy densities make them non-ideal for our project as our magnetic field imager has a handheld component, and convenience of use is of high importance. For that, a lightweight and compact system is desired which lead-acid batteries don't lend themselves well to.

3.9.1.2 Lithium Batteries

There are many types of lithium batteries. The ones that will be considered below are: lithium-polymer (Li-po), lithium-ion (Li-ion), and lithium-iron-phosphate (LFP). These batteries share lithium in their namesake because they all use lithium in some part in the composition of their cathode. Each lithium battery type sports a high energy density resulting in less weight and size.

Lithium-polymer

Compositionally, the lithium-polymer battery mainly differs from its lithium battery counterparts in terms of its electrolyte. Where the typical lithium-ion and lithium-iron-phosphate battery utilizes an organic solvent, the lithium-polymer battery uses a polymer electrolyte from which it derives its name.

Lithium-polymers garner a wide variety of benefits such as the many varieties they are built in. A specific battery can be found for almost any need, and they are able to maintain lower temperatures under high rates of discharge; however, the biggest concern with lithium-polymers is safety. Fire and explosion are not uncommon due to mishandling. Overcharging and improper storage can more easily compromise the integrity of the lithium-polymer over other lithium battery types. In addition to these factors, the casing of the lithium-polymer is non-rigid. This leads to better affordability, but the batteries themselves become easier to damage. Their fragility combined with their volatility is a dangerous combination. The safety factor of lithium-polymer batteries is too prevalent for them to be used in this project.

Li-ion

The typical lithium-ion battery shares a lot of similarities with the lithium-polymer battery, so it's important to note the differences as an uneducated purchase can see a lithium-polymer battery acquired instead of the intended lithium-ion. This is furthered by the fact that a lithium-polymer is a type of lithium-ion battery, so many online listings will have lithium-ion in the title even if it is truly a lithium-polymer. Lithium-polymer batteries are usually flat while lithium-ion batteries are more often than not cylindrical. They both often have heat shrink over multiple cells, but the cells that make up a lithium-ion are rigid which is one of the factors that contributes to them having a better degree of safety. Lithium-ion batteries have a comparable gravimetric energy density to the lithium-polymer, but it has a better volumetric energy density.

LFP

The last battery type that will be examined in this section is the LFP battery. The LFP battery type has marked advantages over other lithium batteries such as: thermal regulation, safety, and cycle life. Like the Li-po, it is able to maintain lower temperatures under higher output. It's even safer than the lithium-ion as its casing is the most rigid of the lithium types discussed here and thus it is the most impervious to damage. These features however lead to a higher price point. [B1]

If it weren't for the increased price, the LFP battery would be the prime choice for this project; however, the improvements that the LFP provides over the lithium-ion aren't worth the price increase. The lithium-ion is already plenty safe

and provides adequate temperature regulation for the scope of this project. The cycle life might sway this decision if we were truly making a product to be delivered to market, but this project is for demonstration purposes and will not be in use for the amount of time that the better cycle life of the LFP would make itself apparent.

3.9.1.3 Comparison & Selection

Since battery types are being compared here and not specific battery models. The categories and their rankings will have to deal in approximations and ranges.

Table X: Battery Type Comparison

	Lead-acid	Lithium-polymer	Lithium-ion	Lithium-iron-phosphate
Safety	great	terrible	good	great
Volumetric Energy Density (Wh/l)	25-100	300-375	350-450	140-330 [B3]
Gravimetric Energy Density (Wh/kg)	10-40	125-190	125-190	90-120 [B2]
Affordability	good	great	good	bad
Nominal Voltage (V)	12 [B4]	3.7	3.7 [B4]	3.2 [B5]

Now that the battery chemistry has been decided. We can choose a specific battery.

Table X: Battery Comparison

	Option 1	Option 2	Option 3
Minimum Order Quantity	1	1	4
Arrival Date via	03/26/25	04/09/25 -	04/02/25

Cheapest Shipping Option		04/25/25	
Listed Price	\$7.99	\$8.52	\$17.34
Capacity (mAh)	500	2500	2600
Nominal Voltage (V)	3.7	3.7	3.7
Brand	Bancroft	Qisuo	YDL
Connector	Tamiya	XH2.54	XH2.54

Although Option 1 has the lowest capacity and the least functional connector, it has the lowest price and the best arrival date. Option 2 would be ideal if its range for arrival wasn't so large, but it's shipping from China. Option 3 is a non-starter as its minimum order quantity is 4 which results in around double the price.

3.9.2 USB Power

Now that a battery type has been selected, the matter of how it connects to the system must be discussed. Finding a decently priced lithium-ion battery that also has a connector that easily interfaces with typically available PCB ports has proven to be quite the challenge. We put more stock in the cost of the battery than the effort of adapting a niche connector to typical ports, so that's led us to select a battery with a Tamiya connector.

We'll have to cut off the Tamiya connector and solder a new connector as a replacement. One of the most universal options is a USB connector, but the decision of which USB type to go with requires analysis. This section will look at USB-A, micro-USB, and USB-C.

3.9.2.1 USB-A

The USB-A standard was created in 1996 [B11], and you can continue to see it in common use today. It has undergone revisions and improvements, but those changes are centered around data speed. We are not concerned with the data transfer of USB for this analysis as this cable will only be used for power delivery. The USB-A has four wires, ground, V_{cc} , data+, and data-. [B12] The

data wires won't even be soldered when we change the connection on the battery.

USB-A cables deliver 5 volts at a range of 0.5 - 2.4 A for a maximum wattage of 12 W [B13]. A downside of the USB-A is its size. Amongst the other USB derivations being examined here, it is the largest. A port could still be incorporated onto the PCB without much issue, but in a project like H.A.L.O where size is a constant consideration, it does play a role in the decision.

3.9.2.2 Micro-USB

The Micro-USB was introduced in 2007 [B11], 11 years after USB-A, and like its predecessor, it still sees widespread use to this day. Micro-USBs are rated for 5 V as well, but they're only rated for 2 A, so their maximum wattage is 10 W [B14]. That's 2 W less than USB-A. It has the same four wires as USB-A, but there could be a potential fifth wire used for identification depending on if it's a Micro-A or Micro-B [B12]. It does see the smallest size of the lot with dimensions of approximately 7.40 x 2.35 mm [B15].

3.9.2.3 USB-C

Lastly, we will be taking a look at USB-C. It's the newest addition of the bunch, seeing introduction in 2014. With its rotational symmetry, bidirectionality, and high data transfer speeds, it's poised to phase out the other USB types, a factor that shouldn't be disregarded as it affects the longevity of the product.

USB-C sees the same four-wire scheme as USB-A. Most pertinently, USB-C stands out in power delivery. It can support a voltage range of 5 - 20 V at an amperage range of 0.5 - 5 A [B13]. This means the maximum wattage of the USB-C is 100 W. This dwarfs the maximum wattage of the USB-A and micro-USB standards. The USB-C standard also allows for bi-directional power flow, but this isn't a pro or con in our case as our circuit will not be supplying voltage back to the power source. In terms of size, USB-C finds itself in between USB-A and micro-USB at approximately 8.34 x 2.65 mm [B16].

3.9.2.4 Comparison & Selections

It is worth noting that the sizes compiled above and in the table below are only approximations, as the precise dimensions vary by manufacturer and the figures were compiled from specific examples around the internet. Also all three variations of USB see similar price ranges, so price will not be included in the table or analyses.

Table X: USB Type Comparison

	USB-A	Micro-USB	USB-C
Voltage (V)	5	5	5 - 20
Current (A)	0.5 - 2.4	2 A	0.5 - 5
Maximum Wattage (W)	12	10	100
Dimensions mm	12 x 4.5 [B17]	7.40 x 2.35	8.34 x 2.65

After compiling the data, USB-C is the option that we will be moving forward with. The benefits of longevity, rotational symmetry, increased power delivery, and smaller size all at the same price point of USB-A or micro-USB makes it the clear choice.

3.9.3 Power Regulation

No power supply is complete without power regulation. Even the best voltage source doesn't supply a perfectly constant voltage level, and as the voltage source discharges, it strays further and further from its nominal voltage until it can no longer supply current to the system. Power regulation fixes this issue by implementing a circuit that corrects the discrepancies in the voltage source's output and keeps the voltage to the load constant even as the voltage supplied by the battery decreases during the cycle.

3.9.3.1 Linear Regulators

Linear regulators are the most common types of voltage regulators. Their simplicity makes them easy to implement and cheap to produce. Thus, their cost is very affordable. The linear regulator is purely a step-down mechanism. Since HALO will split the 3.7 V voltage source into a 3.3 V channel and a 5 V channel, this project has to use a switching regulator for the boosting of the 5 V channel, but a linear or switching regulator can be used for the step-down to 3.3 V. This section will go over three basic linear regulator types.

Standard

When it comes to regulators, two of the most prevalent qualities to be concerned with are the dropout voltage and the ground pin current. The dropout voltage can be thought of as the excess voltage that is consumed by the regulator. If a regulator outputs at 3.3 V and has a dropout voltage of 1 V, the

voltage source will need to be able to supply 4.3 V. Since HALO is using a 3.7 V power source and using a step-down regulator to achieve an output of 3.3 V, this feature is very important when choosing the proper regulator.

The standard linear regulator has the most dropout voltage at around 1.7 V to 2.5 V. However, its ground pin current is the lowest of three types examined here at less than or equal to 10 mA. The numbers included here and for the two other types of regulators are for a typical regulator, hence why a general range is being used. Before a specific model can be chosen, these ranges need to be examined to determine the proper regulator type.

Low Dropout (LDO)

The LDO regulator varies from the standard regulator in that it only contains one PNP transistor in its feedback control loop. This results in less voltage being consumed by the regulator and results in the low dropout voltage the regulator receives its names from.

The LDO regulator has the lowest dropout voltage of the three types examined in this section, but as a tradeoff it has the most ground pin current. A higher ground pin current is less desirable as the ground pin current is analogous to wasted current.

Quasi Low Dropout (Quasi LDO)

The quasi LDO regulator is a modified version of the standard regulator. Its dropout voltage is higher than that of the LDO, but lower than the standard regulator that it is derived from. The ground pin current is comparable to the standard regulator, but the quasi LDO's ground pin current is larger. [B6]

3.9.3.2 Switching Regulators

Linear regulators dissipate voltage through resistors. The excess electrical energy is converted to heat. This is why the linear regulator can only be used to step-down voltages. The switching regulator operates differently. Depending on the type of switching regulator, there's an amount of switches that cycle on and off to regulate the voltage. The excess power isn't dissipated through heat like in a linear regulator. It isn't delivered to the source in the first place. This results in better efficiency and a cooler system, but the switching mechanisms introduce more noise than linear regulators and they're more expensive as well. [B7]

Although switching regulators are more efficient, we've decided to use linear regulators for the step-down regulator in our project. The simpler design and lower cost of the linear regulator are the main driving factors. The lower noise is

a bonus, but it wasn't a determining factor as we're already taking measures to limit the noise from the voltage source, and it wouldn't be a hassle to apply the same techniques to the voltage regulator (distance and shielding).

3.9.3.3 Comparisons & Selections

After choosing a linear regulator for the step-down regulator the question becomes which type. The dropout voltages and pin currents are compiled in the table below. As was the case in the text describing the options, the ranges look at a typical regulator from each type and are thus approximations.

Table X: Voltage Regulator Type Comparison

	Standard	Low Dropout	Quasi-Low Dropout
Dropout Voltage (V)	1.7 - 2.5	0.1 - 0.7	0.9 - 1.5
Ground Pin Current (mA)	≤ 10	$\leq 20 - 40$	≤ 10

For our project, the dropout voltage is the most prevalent feature as there is only 0.4 V of difference between our V_{in} and our V_{out} . The costs between the different types of regulators is negligible which allows us to choose the LDO solely for its low dropout voltage.

3.9.3.4 Analog Devices LTC3525 Switching Boost Converter

The LTC3525 from Analog Devices is a boost switching converter that provides a fixed 5V output, which makes it a strong candidate for power regulation within HALO. With an efficiency of up to 95%, this converter is optimized for applications where power loss must be minimized. This regulator is capable of delivering 5V from a 1-4.5V input, making it suitable for power loads such as lithium-ion batteries.

A major advantage of the LTC3525 is its ultra-low quiescent current of 7 μ A, which significantly reduces power consumption and helps to maximize battery life. Also, it only requires three external components, which helps to simplify PCB design. Additionally, this converter features output disconnect and inrush current limiting. This feature ensures that there is no current flow during startup or shutdown, helping to prevent damage to other sensitive components.

As stated before, this converter can operate with input voltages ranging from 1V to 5V, which will work effectively with out 3.7V Lithium-Ion battery. Once power is supplied to the converter, it can maintain regulation even if the input voltage drops as low as 0.5V.

The LTC3525 is clearly an excellent choice for power regulation due to its high efficiency and low quiescent current. The only specification that could be considered a drawback is its size being significantly smaller than older LM26xx models. While this might make it more difficult to integrate into our final PCB, it is a great option since we are dealing with limited space.

3.9.3.5 Texas Instruments LM2623MM Switching Boost Converter

The LM2623 from Texas Instruments is a general purpose boost DC-DC converter designed for efficient voltage conversion in battery powered systems. With an operating input voltage ranging from 0.8V to 14V, it is able to be used with many different power supplies. This converter has an efficiency of up to 90%, which is ideal for HALO since we are looking to minimize power loss. It also uses Pulse Frequency Modulation (PFM) to effectively regulate output voltage while reducing output ripple. The LM2623 provides output voltages from 1.24V to 12V, covering a wide range suitable for many different applications

One of the major advantages of the LM2623 is its low start up voltage being 1.1V, meaning it can operate with minimal input. Its high switching frequency also allows for the use of smaller passive components. This converter has a built in N-channel MOSFET with a low on-resistance, improving thermal performance and overall efficiency. While the LM2623 is not the most modern boost converter, its larger packaging makes it much easier to integrate onto a PCB, as its wider pin spacing simplifies routing and soldering. As long as we have the board space, this model may be a better option when compared to smaller modern regulators.

3.9.3.6 Texas Instruments LM2596S Switching Buck Converter

The LM2596S from Texas Instruments is a step-down (buck) switching voltage regulator designed to efficiently convert higher input voltages to lower output voltages. With an input voltage range of up to 40V and the ability to provide output voltages of 1.2V to 37V, the LM2596S is suitable for a wide variety of applications, including the HALO project. The device can deliver up to 3A of

continuous output current, making it ideal for powering components that require a stable and reliable power source.

The LM2596S uses a fixed internal oscillator that operates at 150 kHz, allowing it to use smaller filter components compared to lower-frequency regulators. With an efficiency of up to 90%, the LM2596S helps minimize power loss and maximize battery life, a critical factor for portable applications like HALO. The regulator is available in fixed output versions of 3.3V, 5V, and 12V, with an adjustable version that can range from 1.2V to 37V.

A key benefit of the LM2596S is its simple design, requiring only four external components for operation. This simplicity reduces design complexity and simplifies PCB layout. The device also features line and load regulation, thermal shutdown, current-limit protection, and low quiescent current of around 80 μ A in standby mode, which ensures efficient operation even under low power conditions.

However, its larger size compared to newer, more compact converters may present challenges in highly constrained spaces. Either way, its robust design and ability to handle various input loads make it a solid option when space allows, especially in applications like HALO that demand reliable voltage regulation for various subsystems.

3.9.3.7 Texas Instruments LP2985 Linear 3.3V Regulator

The LP2985 from Texas Instruments is a low-noise, low-dropout (LDO) linear voltage regulator that provides a stable 3.3V output, making it a strong candidate for powering low-power components in HALO, such as the microcontroller. The LP2985 can easily be powered with a 3.7V lithium-ion battery, as it can handle input voltages ranging from 2.5V to 16V. It offers a $\pm 1\%$ output voltage accuracy over load and temperature, which is essential for maintaining stable power delivery to sensitive subsystems.

LP2985 features a very low dropout voltage, typically around 40mV at a 150mA load. This ensures that the regulator maintains a consistent 3.3V output even when the input voltage is near the output, which is important for maximizing battery life in a system like HALO. The LP2985 also features low quiescent current and ultra-low shutdown current of just 1.12 μ A, which helps to extend operation time when in standby mode.

The regulator's low output noise (30 μ V_{RMS} with a 10nF bypass capacitor) and high power supply rejection ratio (PSRR) of 70dB at 1kHz make it ideal for

applications where noise-sensitive analog components, such as the Hall-Effect sensor array, are in use. This performance reduces the need for additional filtering, simplifying the design and enhancing reliability. Also, the LP2985 includes built-in protection features like current limiting and thermal shutdown, ensuring that the device remains safe under various operating conditions.

However, it is important to note that, like all LDO regulators, the LP2985 is less efficient than switching regulators when the input voltage is significantly higher than the output. For HALO, this may not be a concern as long as the device is used to power lower-power components where noise and size are primary considerations. With its compact SOT-23 package, the LP2985 can easily be fit into the HALO design, making it a viable solution for specific low-power subsystems.

3.9.3.8 Final Component Selection

Table X: Voltage Regulator Comparison

Component	Input Voltage	Output Voltage	Size	Key Feature
LTC3525	1V to 4.5V	Fixed 5V	Small	Low Quiescent current
LM2623MM	0.8V to 14V	1.24V to 12V	Large	Low start-up voltage
LM2596S	4.5V to 40V	Fixed 3.3V	Large	High current capability
LP2985	2.5V to 16V	Fixed 3.3V	Small	Low noise, high PSRR

After reviewing the available options, we have selected the LTC3525 and LP2985 for the HALO project. The LTC3525 is ideal for providing efficient 5V regulation from a 3.7V lithium-ion battery, offering high efficiency up to 95% and low quiescent current, which is crucial for maximizing battery life. Its compact size is a key advantage, though it may require careful PCB integration due to its small form factor. On the other hand, the LP2985 is well-suited for powering low-power components like the microcontroller with its low dropout voltage, low noise, and minimal quiescent current. Its size and low noise characteristics also make it a perfect fit for noise-sensitive subsystems in HALO, providing stable and reliable performance for smaller, lower-power circuits. Together, these two

components provide the necessary power regulation across different parts of the system while maintaining efficiency and reliability.

3.9.4 Battery Management Chips

When working with rechargeable batteries such as Lithium-Ion, it is important to have an effective battery management system to ensure safe and efficient operation. Battery management chips, also known as power management integrated circuits, regulate the charge and discharge cycles of batteries, controlling the voltage and current to prevent overcharging or discharging. Many of these common chips offer features such as thermal management and overvoltage protection to increase safety and extend the lifespan of a battery.

For our application in HALO where power reliability is a key factor, selecting the right battery management chip is crucial for creating a stable system. A well chosen power management integrated circuit will not only protect the battery but also optimize charging for smooth operation.

3.9.4.1 Microchip Technology MCP73871

The MCP73871 from Microchip Technology is an integrated solution for both system load sharing and Li-Ion/Li-Polymer battery charge management. It is designed to power a system and simultaneously charge a single-cell Li-Ion battery using either an AC-DC wall adapter or a USB port. One of the key features of the MCP73871 is its Voltage Proportional Current Control (VPCC), which ensures that the system load takes priority over the charge current to the battery, helping prevent overcharging.

This charger operates with a constant current/constant voltage (CC/CV) algorithm and offers selectable charge termination points, with voltage options of 4.10V, 4.20V, 4.35V, or 4.40V, all with a regulation tolerance of $\pm 0.5\%$. The device also supports thermal regulation, adjusting charge current based on temperature conditions to protect the system and optimize the charge cycle time.

The MCP73871 can handle up to 1.8A of total input current, making it suitable for portable applications where power requirements can vary. Charge current control is resistor-programmable, ranging from 50 mA to 1A, providing flexibility for different battery configurations. Additionally, it includes essential safety features such as reverse discharge protection, undervoltage lockout (UVLO), and a safety timer with enable/disable control.

To keep users informed of the charging status, the MCP73871 provides several indicators, including a low battery status indicator (LBO), a power good status indicator (PG), and charge status/fault condition indicators. This makes it easy to monitor the charging process, and the device is designed to function within a wide temperature range of -40°C to +85°C, ensuring reliable operation in various environments.

Due to its small physical size and low number of external components, the MCP73871 is ideal for portable devices like GPS systems, PDAs, digital cameras, and Bluetooth headsets, where both system power and efficient battery charging are required.

3.9.4.2 Analog Devices LTC4054

The LTC4054 from Analog Devices is a complete constant-current constant-voltage linear charger specifically designed for single-cell Lithium-Ion batteries. Its compact ThinSOT package and low external component count make it a great choice for portable applications where space and simplicity are key. This charger is particularly well-suited for devices powered via USB ports, such as cellular phones, PDAs, MP3 players, and Bluetooth devices.

The LTC4054 is capable of charging Li-Ion batteries with a preset charge voltage of 4.2V, providing $\pm 1\%$ accuracy. It offers programmable charge current up to 800mA, which is set with a single external resistor, giving flexibility for various battery configurations. The charger operates with a constant-current/constant-voltage (CC/CV) charging algorithm, and it includes thermal regulation to prevent overheating during high-power operations or high ambient temperatures. This thermal regulation ensures that the charge rate is maximized without compromising the safety and longevity of the battery.

The charger also includes several useful features, such as automatic recharge and C/10 charge termination, which automatically ends the charge cycle when the current drops to 1/10th of the programmed charge value after the final float voltage is reached. Additionally, the LTC4054 includes a charge current monitor output for gas gauging, which provides real-time data on the battery's state of charge. It also has a status output pin that indicates the charge termination and the presence of an input voltage.

For low power consumption, the LTC4054 enters a low current state (less than 2 μ A) when the input supply is removed. It can also be placed in a shutdown

mode that reduces the supply current to just 25 μ A, helping to conserve battery power when the charger is not in use.

This charger is ideal for compact, low-power devices that require efficient and reliable charging, such as charging docks, cradles, and Bluetooth applications.

3.9.4.3 Texas Instruments BQ24610RGER

The BQ24610RGER from Texas Instruments is a highly integrated, stand-alone battery charger controller designed for Li-Ion and Li-Polymer batteries. This device uses a synchronous buck converter to efficiently manage charging, supporting up to 6 cells in series with an operating input voltage range from 5V to 28V. With a charge current capability of up to 10A, it is suitable for a variety of applications, including mobile devices, industrial equipment, and ultramobile PCs.

The BQ24610RGER is known for its high-accuracy voltage and current regulation, with a charge voltage accuracy of $\pm 0.5\%$ and charge current accuracy of $\pm 3\%$. It offers dynamic power management (DPM), which helps reduce battery charge current when the input power reaches its limit, preventing overloading of the AC adapter while simultaneously powering the system and charging the battery. This makes it an excellent choice for devices that require efficient power management and high-performance battery charging.

This device also includes several safety features to ensure the protection of both the battery and the system, including overvoltage protection, battery thermistor sensing for hot and cold charge suspends, battery detection, reverse input protection, and thermal shutdown. Additionally, it supports a programmable safety timer and includes charge overcurrent protection and battery short protection, which makes it a reliable choice for applications where safety is a priority.

The BQ24610RGER offers excellent status monitoring capabilities with three status output pins (STAT1, STAT2, and PG) to indicate the charging status and adapter presence. This allows for easy integration with a host processor or system for real-time monitoring of the charge cycle. It also features low quiescent current consumption, ensuring minimal drain on the battery during operation.

In addition, the BQ24610RGER is optimized for mobile and portable applications, such as netbooks, PDAs, handheld terminals, and medical

equipment, providing a complete solution for Li-Ion/Li-Polymer charging with flexibility and safety.

3.9.4.4 STMicroelectronics STC4054GR

The STC4054GR from STMicroelectronics is a standalone linear charger designed for single-cell Li-Ion batteries, offering constant-current/constant-voltage operation with thermal regulation to ensure optimal charging without risking overheating. This makes it an excellent choice for portable applications where efficient charging is required without the need for complex external components. The device charges Li-Ion batteries directly from a USB port and operates within USB power specifications, making it suitable for a wide range of consumer electronics.

The STC4054GR features programmable charge current of up to 800mA, which is set using a single external resistor. It also offers a fixed 4.2V charge voltage with 1% accuracy, ensuring precise charging for optimal battery life. The device automatically terminates the charge cycle once the current drops to 1/10th of the programmed value, and it supports automatic recharge once the battery voltage drops below a certain threshold.

Key features include an internal current regulation block that adjusts the charge current to prevent overheating during high-power or high-ambient temperature operation. The device also provides a charge current monitor output, which can be used for gas gauging, allowing the system to monitor the battery's charge level accurately. Additionally, the low battery voltage detection triggers precharge mode, and the soft-start function limits inrush current during the charging process, ensuring a smooth start-up without damaging the device.

The STC4054GR is packaged in a compact TSOT23-5L package, making it ideal for space-constrained applications like cellular phones, PDAs, Bluetooth devices, and other battery-powered electronics. The device also includes a shutdown mode that reduces the supply current to just 25 μ A, further enhancing energy efficiency when not in use.

With its simple design, low external component count, and advanced features like thermal regulation and automatic recharge, the STC4054GR is an excellent choice for portable, energy-efficient applications.

3.9.4.5 Final Comparison and Selection

In choosing the appropriate battery management chip for the HALO project, several factors must be considered, including the power source, charging current, thermal management, and ease of integration with the system. The following table summarizes the key specifications of the four battery management chips evaluated:

Table X: Battery Management Chip Comparison

Component	Charge Current	Charge Voltage	Thermal Regulation	Packaging	Additional Features
MCP73871	Up to 1.8A	4.1V, 4.2V, 4.35V, 4.4V	Yes	20-Lead QFN	System load sharing, automatic recharge
LTC4054	Up to 800mA	4.2V	Yes	5-Lead SOT-23	Charge status output, soft-start
BQ24610R GER	Up to 10A	Customizable based on the number of cells	Yes	24-Lead VQFN	Dynamic power management, battery detection
STC4054GR	Up to 800mA	4.2V	Yes	TSOT23-5L	Charge current monitor, undervoltage lockout

Each of the battery management chips offers robust features for different applications. The MCP73871 provides high flexibility in terms of selectable charge voltage options (4.1V, 4.2V, 4.35V, and 4.4V), making it suitable for a wide variety of Li-Ion batteries. It also supports simultaneous system load sharing with charging, which could be advantageous for ensuring system performance during battery charging. However, its larger 20-lead package may not be as space-efficient as some of the other options.

The LTC4054, with its compact SOT-23 package and low external component requirements, is another excellent choice for portable applications. Its charge current limit of 800mA makes it suitable for devices with lower power needs, and its thermal regulation and soft-start features provide protection in high-power or high-temperature environments.

For systems requiring higher charge currents, the BQ24610RGER stands out with a maximum charge current of 10A, making it ideal for devices with larger battery capacities or power demands. Its dynamic power management ensures that it can handle both the battery charge and system load efficiently without overloading the power source. However, its more complex features and larger package size may be overkill for simpler applications.

Lastly, the STC4054GR is a great option for applications that require a straightforward, efficient charging solution in a compact package. With a preset 4.2V charge voltage and programmable charge current up to 800mA, it is well-suited for battery-powered devices with moderate charging requirements. Its low quiescent current and shutdown mode also help to save power when the system is not in use.

Considering the requirements of the HALO project, which involves real-time data collection and processing, the MCP73871 is the most suitable choice due to its flexible charge voltage options, system load sharing, and automatic recharge capabilities. These features ensure a reliable power supply and efficient battery charging during the operation of the system, making it the best fit for the project's needs.

4: Standards & Design Constraints

4.1 Standards

4.1.1 IPC-A-610

IPC-A-610 is one of the most widely referenced standards for electronic assemblies. This of course includes PCBs. It's obvious now why this standard deserves to be discussed and looked into further in this report.

IPC-A-610 discusses component layout and orientation. The orientation of the components needs to be consistent with the documentation, and the layout of

the components must ensure proper spacing. Proper spacing avoids interference between components and allows for proper testing to occur.

IPC-A-610 also includes requirements for the through-hole and surface-mount component types. For through-hole, the standard demands insurance of the condition of the component as well as proper protocol when attaching it to the assembly. The leads of the component can't be bent or damaged in any capacity. Looking for cracks in the leads of the component is essential in ensuring that the component complies with standard IPC-A-610. The leads of the component also need to be of the proper length to ensure that a sufficient and proper soldering job can be accomplished. To comply with IPC-A-610, the manufacturer needs to carefully examine the components before they are used, and they need to keep a diligent mind when securing the component.

IPC-A-610 also deals with surface-mount components. Just like with through-hole components, the condition of the surface-mount component needs to be evaluated before it can be used in the construction of the assembly. The standard once again requires cracks to be searched for. For the attachment of a surface-mount component, the manufacturer must ensure that the component is flat and securely pressed against the PCB. No tilting or misalignment is permitted under this standard. The standard also requires pressure be applied to the component throughout the entire soldering process. Only once the process is finalized, can the manufacturer release the pressure on the component. Whether it be through-hole or surface-mount, the component should be properly labelled so that identification is easy. This responsibility lies on the manufacturer of the component, but we as buyers and component selectors need to be aware of this when purchasing components so that we comply with IPC-A-610.

IPC-A-610 governs many facets of the soldering process as well. The shape of the solder is important. Excess solder resulting in a bubble is unacceptable. There should be a clear concave shape of the solder that goes from the borders of the solder to its peak, and the leads of the component need to be trimmed to proper length once the soldering has been completed. The edges of the solder cannot touch the edges of other solder, pins, or leads as that would result in a short. Not only could that affect the functionality of the product, but it could affect the safety as well. Lastly, the condition of the solder itself must be inspected as well. The solder must be clean. No residue can be present. IPC-A-610 also provides rules for coating processes once the PCB is completed to ensure longevity of the product and protection of the board.

4.1.2 IPC-2221: Standards in Circuit Board Design

IPC-2221 as indicated by its title is about printed circuit board design. The actual construction is covered more in standards like the above IPC-A-610. IPC-2221 has a hierarchical structure with other relevant standards. It's important to discuss how they relate to each other to fully understand IPC-2221.

IPC-2221 covers the generic design of PCBs. There are four standards below it that each cover a specific type: IPC-2222 covers rigid circuit boards. This standard definitely applies to our project as we are most definitely going to be utilizing rigid circuit boards. IPC-2222 goes over spacing, material selection, board thickness, routing, etc. IPC-2223 covers flexible circuit boards. This one has possible bearing on our project as a flexible circuit board is a potential solution to deal with interference the battery may induce in the rest of the system. Like IPC-2222, IPC-2223 addresses material selection. It also covers minimum bend radius, filling of vias, and more. IPC-2224 deals with laminated multi-chip modules or MCM-L. A laminated multi-chip module uses laminated PCB as the substrate and has multiple chips on the PCB [B19]. The last in the IPC 2220 series is IPC-2226. IPC-2226 covers high density interconnect PCBs, also known as HDI. These PCBs are characterized by their high wiring density. There are significantly more traces and connections than on your typical PCB. It is possible that our project will delve into this territory as our hall effect array will have 64 Hall Effect sensors alone, not counting any regulators, multiplexers, or other components. Although, our project will definitely dive into territory covered by IPC-2222 and possibly subject matter covered by IPC-2223 and IPC-2226, this paragraph was only included in this report to discuss the hierarchy of the 2220 Series so that IPC-2221's role as a standard could be better understood and the below discussion of IPC-2221 can be analyzed with its proper context in mind.

Two words that are seen often in IPC-2221 are clearance and creepage. Clearance is the distance between conductors measured through the air, and creepage is the distance between conductors measured along the surface of the PCB. IPC-2221 has tables that provide minimum clearances based on a variety of voltage levels from 0 to 500 V. The maximum voltage that our system will see is 5 V, so only the 0 - 15 V range of these tables is relevant to H.A.L.O. The applicable data from these tables is compiled below.

Table X: Board Clearances

	Minimum Clearance for Bare Board	Minimum Clearance for Assembled Board
--	-------------------------------------	--

Internal Conductors (mm)	.05	
External Conductors, uncoated (mm)	0.1	0.13
External Conductors, coated (mm)	0.1	0.13

After seeing this table, it is worth defining the terms “bare board” and “assembled board.” A bare board is a PCB that does not have any electrical components or vias. It consists of the PCB material itself and traces [B20]. The assembled board is the PCB, after all of the electrical components have been soldered on [B21].

IPC-2221 provides the minimum creepage and clearance values, but also acknowledges that increasing distances can be challenging for designing PCBs where size is often a chief constraint of the project. So, it provides a few alternate solutions to achieve these values. One of the offered solutions is to physically mar the circuit board by carving specific patterns in between the components such as v-grooves or slots. Although this doesn’t increase the clearance between the conductors. It does increase the creepage.

Another route to achieve these values is by installing physical insulation on the board itself, and IPC-2221 provides guidance on selecting insulator material. Their recommended basis on making this decision is by evaluating different materials’ Comparative Tracking Index or CTI. The CTI is essentially a measurement of how much voltage it takes for an insulator to begin tracking which is the process of conducting pathways forming in the insulator [B23].

Other solutions for achieving acceptable creepage and clearance values are implemented in the design portion of the PCB. If possible, components with large voltage differences between them can be placed further away from one another or can even be placed on opposite sides of the board if the board type permits [B22].

4.1.3 IPC-7351: Surface Mount Design & Land Pattern Standard

IPC-7351 is a critical standard that governs the design and placement of any surface mount components. This standard provides land pattern dimensions to

make sure that components are placed with proper spacing for soldering. Since our project integrates an 8x8 array of Hall-Effect sensors and other surface mount components, following IPC-7351 guidelines would help us optimize our PCB design for when we manufacture and assemble it.

IPC-7351 highlights the need for correct pad sizes and tolerances so that one can avoid soldering issues such as misalignment. This is important for the design of HALO, where all sensors and peripheral components must be correctly integrated. Also, this standard provides guidelines on thermal relief patterns, which help to maintain heat dissipation to preserve signal integrity and reduce thermal stress.

4.1.4 IEEE 1149.1: JTAG Boundary Scan Testing

IEEE 1149.1 is the Joint Test Action Group (JTAG) standard, which is essential when testing and debugging printed circuit boards. We are integrating a STM32H735IGT6 microcontroller into HALO's design, which has capabilities to support JTAG for in-system debugging and testing. By integrating JTAG testing into the design process of our project, we may be able to ensure better fault detection while reducing reliance on physical probes.

Boundary scan testing is very helpful for identifying soldering defects, shorts, and open connections without requiring powered functional testing. This will be especially useful in our design, where the PCB has many components integrated close together making it more difficult to probe individual connections. Overall, utilizing IEEE 1149.1 as a design consideration could improve quality of life during manufacturing and testing.

4.1.5 IEC 61000-4-2: Electromagnetic Compatibility (EMC)

Given that we are looking to measure magnetic fields in real time, it is critical that we ensure HALO has electromagnetic compatibility. IEC 61000-4-2 outlines testing procedures and design strategies that protect systems from electrostatic discharge (ESD) events, which can damage sensitive components like Hall-Effect sensors, microcontrollers, and more.

To comply with this standard, our design for HALO's PCB should incorporate ESD protection mechanisms such as ground planes and shielding techniques. By having proper PCB layout with controlled impedance traces and isolation of

high-speed signals, we will be able to mitigate the risk of signal degradation due to ESD events.

4.2 Constraints

Constraints have played a significant role in shaping our project, leading us to redefine its scope multiple times. At the start of the semester, our initial concept was a photodiode-based radiation detector. However, after consulting with the Florida Space Institute (FSI) to assess its feasibility and applications, we secured sponsorship for radiation hardened components and a nuclear event detector on the condition that we shift our focus to something that the sponsors would be interested in. This was an EMP protection system. Following our initial divide and conquer report and discussions with Dr. Chan and Dr. Weeks, the team ultimately decided to abandon the original idea and instead develop a magnetic field imager.

4.2.1 Safety Concerns

4.2.1.1 Safety with the NED project

In our initial proposal, we planned to incorporate a nuclear event detector manufactured by Micross as the key component for detecting high doses of radiation. However, a critical limitation of this component was its minimum dose rate threshold of 50 krad per second, meaning that testing would need to be conducted in a high radiation environment that could pose significant health risks to individuals involved. Our team lacked direct experience in radiation testing, but we were working alongside Dr. Zhang and industry experts at Honeywell who had expertise in the field. With their guidance and proper safety procedures, we were confident in our safety.

When we initially presented our concept to Dr. Chan and Dr. Weeks, we had not yet developed a formal safety plan. As a result, they expressed strong concerns about the feasibility and risks associated with our proposed testing procedures. One of their primary objections was that all radiation testing would need to take place off UCF's campus, beyond their direct supervision. Given these concerns, Dr. Weeks suggested an alternative approach, using electrostatic discharge (ESD) as a means to trigger the device instead of radiation. This suggestion significantly alleviated their safety concerns and allowed us to explore a safer and more controlled method of testing.

4.2.1.2 Safety with the HALO project

Although safety considerations were a major factor in prompting our team to reconsider the project's direction, they were not the sole reason for our decision to pivot. Ultimately, after evaluating the technical complexities, logistical challenges, and long-term feasibility of our design, the team decided to move away from radiation detection entirely. Instead, we opted to focus on measuring low level magnetic fields, a direction that eliminated safety risks while still providing a suitable project. Our current project is entirely centered around magnetic field imaging, ensuring that all testing can be conducted safely and efficiently on the UCF campus without exposing team members or collaborators to hazardous conditions.

4.2.2 Feasibility Constraints

4.2.2.1 Feasibility with the NED project

Feasibility was a challenge that the team initially underestimated. When we were introduced to the idea of using the NED to trigger clamping in order to protect critical components from an EMP, we understood that it would be a difficult task. However, with the confidence expressed by our sponsors and the FSI, we believed it was an achievable goal. It wasn't until we spoke with Dr. Weeks that we fully understood that the project was beyond our abilities.

The NED operates with a response time in the nanosecond range, and for clamping to effectively mitigate the effects of an EMP, it would need to activate within a similarly short timeframe. The level of high speed circuitry required for this type of system was well beyond our team's current expertise. Designing and implementing such a system would have demanded knowledge of high speed electronics, specialized components, and testing environments that were simply outside our skill set.

4.2.2.2 Feasibility with the HALO project

This lack of feasibility was the primary reason we ultimately decided to pivot to a magnetic field imager, a project that has proven to be a much more attainable goal. With the imager, our initial target was to achieve a refresh rate of 60 frames per second. However, after conducting further research into the microcontroller we selected, we discovered that it might not have enough processing power to sustain that rate. We explored potential solutions, including upgrading to a more powerful MCU or offloading some of the processing to a computer. After weighing the trade-offs, we concluded that the best course of action was to

adjust our expectations and reduce the target frame rate. This compromise allows us to maintain a functional and reliable design while staying within our technical limitations.

4.2.3 Time Constraints

4.2.3.1 Time with the NED project

Senior design places a strong emphasis on completing the project within the designated time frame. If we fail to do so, we are required to retake the course, we do not have the option to extend the project until it is finished. Because of this, it is crucial that our project remains feasible within the time allotted. For our team, this means completing our work within the spring and summer semesters. While some teams choose to spread their work across the spring and fall semesters, allowing them to utilize the summer as extra development time, our team's goal has always been to graduate by the end of summer.

At the start of our project, we devoted a significant amount of time to refining our idea. We scheduled meetings with professors, organizations like the FSI, and potential sponsors to explore different possibilities. When we were introduced to a project concept from Micross and Vorago, our team was eager to work with companies that could provide both technical expertise and critical hardware components. Our interactions with Vorago were smooth and efficient. They quickly answered our questions, confirmed their support, and even offered to provide a development kit along with a standalone MCU. However, our experience with Micross was much slower, and unfortunately, their hardware was a key part of the project. By the time we reached our D&C meeting, Micross had only just confirmed their interest in moving forward.

During our D&C meeting, Dr. Weeks and Dr. Chan raised concerns about working with these sponsors, many of which we had not previously considered. The components we were relying on were highly specialized, with limited publicly available information. This meant that if we encountered any technical issues, we would be entirely dependent on our sponsors for guidance and troubleshooting. Initially, both professors were completely opposed to us working with these sponsors, not only due to safety concerns regarding the NED but also because of the feasibility challenges associated with EMP clamping. However, once we determined that ESD could serve as a safer alternative to radiation, Dr. Weeks and Dr. Chan agreed to let us move forward, provided that we found a different application for the components beyond EMP protection.

After exploring alternative ways to integrate these parts into our project, we ultimately decided to abandon the idea of working with sponsors altogether. By this point, we had already spent a significant portion of the semester pitching ideas and searching for projects that sponsors might support. Although we had found a sponsor interested in a project similar to our original concept, moving forward with them would have required us to pitch a new idea and refine it until it aligned with their interests. Given that our initial D&C meeting had not been approved, we were already behind schedule, and we could not afford the additional delays that would come with negotiating and troubleshooting a new sponsored project.

In the end, we chose to proceed without a sponsor to avoid being slowed down by dependence on an external company with its own priorities. While working with industry partners can provide valuable resources, the uncertainty surrounding sponsors' responsiveness and the potential time investment required to secure a viable project ultimately led us to take full control of our design. This decision has allowed us to move forward at a more efficient pace, ensuring that we remain on track to complete our project within the required timeframe.

4.2.3.2 Time with the HALO project

Scrapping our initial idea had significant consequences. Most notably, we had no completed pages for our final report and had to make rapid progress on our new project to stay on track. To ensure we met our deadlines, the team established a goal for each member to write a minimum of five pages per week while maintaining clear communication about which sections had been completed. This structured approach allowed us to make steady progress and avoid falling further behind.

When selecting components for our new project, we made a strategic decision to ignore lead times entirely. Most components had lead times of several weeks. On the time frame that we were working on having to wait 7 weeks for a component vs having to wait 20 weeks didn't make that much of a difference. Both would likely result in project failure. To mitigate this risk, we chose to exclusively use widely available components with ample stock that could ship immediately. Additionally, for any relatively inexpensive components, we ordered spares to account for potential failures or mishaps during development.

Beyond that, our research also focused on identifying potential backup components. In the event that a critical part became unavailable, we ensured that we had alternative options that could be quickly sourced without

significantly altering our design. By prioritizing availability and flexibility in our component selection process, we positioned ourselves to avoid unnecessary delays and maintain steady progress toward completing our project on time.

4.2.4 Cost Constraints

4.2.4.1 Cost with the NED project

Cost was a significant concern when working with sponsors, particularly because the components they provided were extremely expensive, costing thousands of dollars each. If we accidentally damaged any of these parts, it was highly unlikely that we would have been able to replace them, adding another layer of risk to our project. This financial constraint played a role in our decision to move away from working with sponsored components and instead focus on a more cost effective, self managed design.

4.2.4.2 Cost with the HALO project

When selecting components for our revised project, we prioritized building a strong foundation with our microcontroller. The MCU was the most expensive single component, and its capabilities dictated many aspects of our overall design. Early on, we encountered a potential issue with the processing power of our selected MCU. We explored alternative options, including upgrading to a more powerful microcontroller that could comfortably handle our initial project requirements. However, after evaluating the trade offs, we determined that the increased cost of a high performance MCU wasn't justified for our specific needs. Instead, we made a strategic decision to scale back certain aspects of our project. Such as lowering our target refresh rate to ensure that our existing MCU could still perform effectively within our revised design constraints.

Cost also played a major role in the selection of our sensors, as we needed to purchase a large quantity of them. Since our design relied on multiple sensors, we had to strike a balance between performance and affordability. Choosing an overly expensive sensor would have dramatically increased our budget, while opting for a cheaper, lower-quality sensor could have compromised the accuracy of our system. Additionally, we anticipated that some sensors might fail or get damaged during development, so we needed to account for the cost of ordering spares. Ultimately, our approach was to find a sensor that met our performance requirements while remaining cost-effective, ensuring that we could afford the necessary backups without exceeding our budget.

4.2.5 Layout Constraints

4.2.5.1 Layout with the HALO project

Since one of our primary design goals was to make the device as compact as possible, we focused on optimizing the layout by minimizing the number of components and ensuring that each component was utilized as efficiently as possible. Initially, we planned to use a 10x10 sensor array, mainly because the idea of having 100 sensors sounded nice. However, once we began researching the implementation, we realized that using a power of two arrangement would be more practical, especially since we planned to use MUXs to manage the ADC channels on our MCU.

By structuring the array as an 8x8 grid with 64 sensors, we could optimize our use of MUXs and simplify the routing of signals. This arrangement allowed us to dedicate one MUX to each row or pair of rows, improving efficiency. When selecting our MUX configuration, we weighed the trade offs between speed and space. Using eight 8:1 MUXs would have provided faster switching, but it would have also required more board space. Instead, we opted for four 16:1 MUXs, which, while slightly slower, reduced the physical footprint of our circuit.

Another critical consideration in our layout design was the potential for interference between components and the sensors. Every component in our system generates its own magnetic field, which could interfere with the readings from our magnetic sensors. To mitigate this issue, we carefully planned the placement of components, ensuring that any elements generating significant electromagnetic interference were positioned away from the sensor array. By balancing efficiency, compactness, and signal integrity, we aimed to create a design that was both practical and high-performing within our constraints.

4.2.6 Available Resources

It is important to keep in mind what is available to us for the duration of this project, both in terms of available equipment and available guidance. By default, we have the senior design lab and the instruction of Dr. Chan and Dr. Weeks. Additional equipment and guidance can be sourced through additional effort on our part, but that is the starting point.

The senior design lab comes equipped with Tektronix oscilloscopes, Tektronix dual arbitrary function generators, Tektronix DMM 4050 digital multimeters, Keithley 2230-30-1 triple-channel power supplies, Dell Precision 3420 computers, an SMD rework station, soldering and desoldering stations, a digital

microscope inspection station, and many additional instruments not listed on the website [B24]. If processes that require equipment not contained on this list, we must be prepared to acquire it ourselves.

4.2.6.1 NED

The NED project highlighted what was available to us as our initial scope was way beyond what we had access to; however, we were able to acquire access to additional facilities and equipment through the aid of sponsorships and connections: including access to the FSI laboratory and equipment from Micross and Vorago. Gaining access to these resources wasn't the end of the dilemma as it shined a spotlight on the non-tangible resources we would need: guidance.

Dr. Weeks is dedicated to helping students succeed in senior design through hands-on help as well as providing the knowledge that he possesses. He made aware of these benefits by telling us we would not have access to them if we were to conduct most of our work in outside laboratories or with niche, specialized equipment.

We were taking a gamble on how much help the companies would be able to provide us and how much time they would be able to devote. This point was especially concerning since their proprietary equipment would be an integral part of our project and we would therefore be dependent on them for documentation and professional guidance. We decided that the dedicated guidance from an experienced professional like Dr. Weeks was too valuable a resource to neglect, and thus, we decided to proceed with an entirely different project.

4.2.6.2 H.A.L.O

There are two aspects to the resources at our disposal, the equipment and the guidance. By choosing H.A.L.O, we chose to forego equipment for guidance, and thus the constraints of the senior design lab are again prevalent. We have already corresponded with Dr. Weeks about some of these limitations that will shape our project. He has informed us that UCF does not have the resources to solder ball grid array (BGA) parts. This constraint and constraints like it have shaped our selection of components.

5: ChatGPT vs. Similar Platforms

Artificial Intelligence (AI) has become increasingly popular in recent years, with a growing presence in educational and industrial fields. From natural language processing to machine learning, artificial intelligence tools have proven to be valuable assets for automation of tasks and solving complex problems. Among the tools available to the public, ChatGPT has grown in popularity due to its advanced language generation and user-friendly interface. Within this chapter, we will observe the benefits and drawbacks of ChatGPT in comparison to other AI platforms. Through this comparison, we aim to assess the effectiveness of artificial intelligence in applications like our senior design project, HALO.

The use of natural language processing (NLP) is what creates the overall functionality of ChatGPT, combining computer science and AI to help machines understand the human language. NLP techniques enable tools like ChatGPT to comprehend user prompts and generate responses. Through the utilization of datasets, ChatGPT is able to analyze text, recognize context, and provide coherent answers. This process is crucial for complex tasks. With NLP, ChatGPT can process and analyze entire documents or websites, offering insights based on the context the user provides. For our project, this feature is key to helping with research, documentation, and technical problem-solving. It allows the team to make the most of AI for smoother communication and more efficient task management as we develop HALO.

5.1 ChatGPT

ChatGPT is the juggernaut of the AI landscape. To measure a product's adoption, there exists a statistic that tracks how much time the product took to acquire 100 million users. ChatGPT snatched the record and claimed the title of the fastest growing consumer product after acquiring 100 million users in only 2 months. TikTok took 9. Instagram took 2.5 years [B27]. Since then, it has been the name that people think of when AI is said. Being many people's first foray into the modern AI world, it has garnered a solid reputation amongst its contemporaries. This isn't without basis as its dedicated teams and years of development give it a leg up against its competition.

This reliability is one of ChatGPT's biggest strengths. It is solid across the various purposes you might need AI for, but that's also one of its downsides. ChatGPT is a jack of all trades and a master of none. It isn't weak in any areas, but it isn't the strongest either. It has strong reasoning capabilities, but it's outperformed by AIs like Grok and DeepSeek.

Being a jack of all trades comes with its own set of benefits as you can use ChatGPT for nearly any reason you would seek out an AI in the first place. It

doesn't neglect the creative types as it offers image generation and has decent proficiency in long-text responses. On top of this, it has low AI detectability. In other words, its text output is less likely to be immediately identified as the product of an AI tool [B26]. All of these are significant advantages of ChatGPT, but for our project, the creative side of ChatGPT is superfluous.

Cost is one of if not the most important factor in choosing products. ChatGPT does offer a free version, but it also offers a Plus and Pro plan at \$20 a month and \$200 a month respectively. The free version includes but is not limited to GPT-4o mini, (ChatGPT's cost optimized model), limited access to GTP-4o and o3-mini (flagship and reasoning models respectively), and limited access to file uploads, data analysis, image generations, and ChatGPT's voice mode.

The Plus plan lessens the limits imposed by the Free version, provides access to Sora as well as more reasoning models, and allows users to use experimental features and models. The Pro plan effectively allows access to all substantial models ChatGPT offers and pro mode, which uses more computational power for better answers [B25].

5.2 DeepSeek

DeepSeek is the new kid on the block and has been making monumental waves throughout the AI landscape. With all of ChatGPT's prominence and reliability, DeepSeek recently overtook ChatGPT in the Apple app store, showing that ChatGPT's place on the top isn't as certain as some would believe; however, it has an air of controversy due to its country of origin: China.

DeepSeek's ties to China has led to concerns not only about security but censorship as well. Since DeepSeek is located within China, the company has to comply with the CCP's censorship policies. This leads to certain prompts around forbidden topics in China yielding lackluster results. DeepSeek also doesn't have the language capabilities of models like ChatGPT or Google's Gemini. Less globally popular languages are prone to grammatical and translational errors.

Then there's the host of problems that come with the data storage being located in China. TikTok's current plight is indicative of how these concerns could prove problematic for DeepSeek. At best, international audiences may be reserved about fully placing their trust in DeepSeek, leading to lower use. At worst, governments may outright ban the platform in their respective nation. TikTok's place in the United States is still uncertain, a fact that does not bode positively for DeepSeek's longevity in foreign markets. Anecdotally, I wasn't able to access

DeepSeek on UCF's internet due to "connection privacy" concerns on Google Chrome. Only once I was on my home network was I able to successfully access the website.

Even with these setbacks, there are multiple benefits that DeepSeek provides leading to its surpassing of ChatGPT on the Apple app store. The platform is the most open-source out of all its competitors. It has made its DeepSeek-R1 model open-source, an unprecedented move in the AI landscape that could force other companies to do the same. Grok has already hinted at allowing open-source access to older models.

Once again, the importance of cost cannot be overstated in consumers' decision making process. DeepSeek uses cheaper GPUs than its competitors, allowing for the service to be offered at a cheaper price point than Grok or ChatGPT's higher tiers. Yet, it doesn't follow the typical subscription based model. It instead uses "tokens" relegated on the basis of quantity inputted and outputted from the model. For reference, one English character equates to 0.3 token while one Chinese character equates to 0.6 tokens.

Like ChatGPT, DeepSeek offers two pricing plans, but its only free option is in the form of a demo. DeepSeek Chat is its more affordable model and DeepSeek Reasoner is the premier tier. Chat is priced at 7 cents per 1 million token inputs with cache hit, 27 cents for 1 million token input with cache miss, and \$1.10 for 1 million token output. Cache hit occurs when the answer has already been compiled by DeepSeek and therefore less computing power has to be expended on the company's part to provide the user with their answer. Reasoner is priced at 14 cents for 1 million token input with cache hit, 55 cents for 1 million token input with cache miss, and \$2.19 for 1 million token output [B28]. The two tiers are mainly set apart by Reasoner's titular reasoning abilities.

DeepSeek's increased affordability has not only allowed it to surpass ChatGPT on the Apple store, but could see its widespread use in foreign markets, especially regions like Europe where the cost barrier to AI is more significant than in the United States. It's unclear whether its setbacks will outweigh its benefits as its future is uncertain. It's important to remember that it hasn't been around for long enough to truly cement itself in the industry. Unlike ChatGPT, which has undergone years of development in the public eye, DeepSeek has yet to fully attain a dependable reputation. The questionable longevity of DeepSeek in foreign markets may affect our adoption of it for our senior design project, and it may not. It's simply too early to tell if security concerns will impact our access to it during our senior design process.

6: Hardware Design

The hardware design of the Hall-effect Array for Live Observation (HALO0 system revolves around a single printed circuit board that integrates all the necessary components to ensure system functionality. The PCB houses the sensor array of 64 Hall-Effect sensors, which is the primary mechanism responsibly for detecting magnetic fields. Alongside the sensors are multiplexers, which are responsible for combining the sensor data into readable data for the microcontroller. The STM32 microcontroller unit on the PCB plays a critical role in the system for handling data acquisition, processing, and user interface management. The power subsystem is also integrated into the PCB, featuring a USC-C port for charging, a 3.7V lithium-ion battery, a battery management chip, and 3.3V/5V voltage regulators that provide stable voltage supply across the system. The SD connector provides data storage, while a touchscreen LCD displays magnetic field data to the user. To minimize magnetic interference, we will incorporate specific PCB layout techniques that maximize sensor performance and avoid interference issues.

The hardware design for HALO involves a collaborative approach, where each team member must take responsibility for specific components and subsystems based on their individual strengths. Ryan and Aidan, both majoring in Electrical Engineering, are responsible for creating the sensor array, which consist of the Hell-Effect sensors and addressing circuitry for magnetic field measurements. Gabriel, majoring in Computer Engineering, is responsible for programming the microcontroller, focusing on data acquisition, processing, and handling interactions with the user interface. Brogan and Ryan are then responsible for the power subsystem, which will ensure proper voltage levels are distributed throughout the systems. Having a division of responsibilities will make sure to maximize the strength of each team member and enhance our system's performance and reliability.









Figure X: Hardware Block Diagram



6.1 Overall Schematics

6.1.1 Schematic Overview



Our schematic design for HALO serves as a blueprint for the systems electrical architecture, outlining all connections between components on our PCB. the overall schematic is broken down into key subsystems: the sensor array, microcontroller section, power subsystem, and user interface. Each subsystem is designed so that the system will efficiently collect and process magnetic fields data to be interpreted by the user.










One Section of the Hall Effect Array - 8 sensors and their multiplexer





MCU and MUX connections

6.2 Schematic Justification

6.2.1 Power Supply Subsystem

The power supply subsystem of HALO is responsible for providing voltage

6.2.2 MCU

6.2.3 User Interface

6.2.4 Sensor Array

7: Software Design

7.1 Overview of Software Architecture

HALO implements a bare-metal firmware architecture on the STM32H735IGT6 microcontroller. The software is structured in a modular fashion to accommodate sensor acquisition, user interface rendering, and data logging in a responsive manner. The overall architecture follows a super loop paradigm along with interrupt service routines (ISRs) for time critical events such as ADC completions, touch input, and SD card write operations. This design minimizes overhead and ensures deterministic execution, which is critical for consistently updating the 8×8 sensor array in real time. The software architecture is divided into the following logical modules:

- **Sensor Acquisition:** Coordinates the GPIO-controlled multiplexer select lines and reads sensor outputs via the MCU's ADC peripherals.
- **Signal Processing:** Converts raw ADC readings into normalized intensity values suitable for graphical display. This may include optional smoothing or mapping to a color scale.
- **User Interface:** Manages display updates via the TouchGFX rendering framework. This includes drawing the sensor grid, updating visual elements, and handling user interactions.
- **Storage Manager:** Facilitates writing either raw sensor data or heatmap snapshots to the onboard SD card, using the STM32 FATFS middleware.
- **System Control:** Orchestrates global behavior such as state transitions, power-up initialization, and potential expansion features (e.g., battery monitoring or system status reporting).

The architectural goal is to keep display rendering and data sampling decoupled through buffering and DMA transfers, so that GUI responsiveness does not interfere with sensor readout. In this configuration, the MCU continuously cycles through sensor scan routines while conditionally updating the display and checking for touch input or save commands. Interrupts are used sparingly and only for latency-sensitive events.

7.2 Development Environment

The HALO firmware is developed within a Windows environment using STMicroelectronics' integrated STM32Cube toolchain. This suite offers a streamlined workflow for embedded development, combining code generation, peripheral configuration, graphical UI design, and debugging.

7.2.1 Integrated Development Environment (IDE)

STM32CubeIDE serves as the main environment for firmware development. Based on Eclipse and GCC (arm-none-eabi toolchain), it provides a unified interface for writing code, managing projects, and debugging via the onboard ST-Link debugger. It integrates seamlessly with STM32CubeMX for peripheral initialization and middleware configuration.

7.2.2 Peripheral Configuration

STM32CubeMX is used to graphically configure hardware peripherals, middleware (e.g., FATFS, TouchGFX), clock trees, and pin assignments. Once configured, CubeMX generates the corresponding initialization code, which serves as the base of the firmware project.

7.2.3 Graphics Framework

TouchGFX Designer is used to develop the graphical user interface. This tool provides a drag-and-drop canvas for designing screens, adding widgets, and mapping interactions. The output is C++ code that integrates directly with STM32CubeIDE projects. TouchGFX supports hardware acceleration (DMA2D) and partial framebuffer rendering, making it ideal for real-time embedded displays.

7.2.4 Compiler and Debugger

The ARM GCC toolchain compiles C/C++ firmware into binary output compatible with the STM32H735IGT6. Debugging and flashing is done through ST-Link v3, integrated into the Discovery Kit hardware. It enables breakpoints, live variable inspection, and peripheral register views.

7.3 MCU and Firmware Framework

At the heart of the HALO system is the STM32H735IGT6 microcontroller, this MCU is built around an Arm Cortex-M7 core running at 550–600 MHz, offering a strong balance of computational power, memory, and peripheral support for graphics-driven applications. It features 1 MB of Flash memory and 564 KB of SRAM, which is more than sufficient to handle HALO's display buffering, sensor data acquisition, and control logic without requiring external memory.

7.3.1 Firmware Architecture

HALO uses a bare-metal firmware approach — meaning no operating system or RTOS (real-time operating system) is employed. Instead, the application runs in a deterministic, single-threaded super loop structure, with key functions triggered by interrupts. This approach offers:

- Minimal system overhead
- Predictable control over timing
- Simplified debugging and system flow

The project follows a layered firmware model:

- Hardware Abstraction Layer (HAL): Auto-generated by STM32CubeMX, these drivers abstract low-level peripheral access and provide functions for initializing and interacting with GPIOs, ADCs, DMA, SDMMC, LTDC, and more.
- Application Layer: Implements core logic for sensor scanning, buffer management, user interface updates, and SD card interaction.
- Middleware: ST's libraries such as TouchGFX (graphics), FATFS (file system), and potentially USB Device Stack (if needed in future extensions).

7.3.2 Memory Management

The internal SRAM is divided into several regions, including:

- DTCM (128 KB): Fastest access; ideal for computation-critical variables.
- AXI SRAM (320 KB): Used for general-purpose buffering, display framebuffer, and TouchGFX-related memory.
- SRAM4 (64 KB): Can be used for I/O buffering or SD card data staging.

Special attention is given to DMA compatibility and cache coherency, as both ADC and display subsystems rely on DMA transfers. The framebuffer for the touchscreen display must be placed in a non-cacheable or cache-coherent region to avoid visual artifacts.

7.3.3 Startup and Initialization

- Clock tree configuration ensures high-speed operation of the core, ADCs, and display peripherals.
- GPIOs are initialized to control MUX select lines.
- ADCs are configured for high-speed sampling of analog input from the Hall sensors.
- TouchGFX is initialized with framebuffer allocation and UI setup.
- SD card is mounted via FATFS and initialized for logging.

Overall, the STM32's high clock speed, ample SRAM, and display-specific peripherals (LTDC, DMA2D, Chrom-ART Accelerator) make it an excellent match for HALO's performance and memory requirements, especially while keeping the software stack lightweight and real-time capable.

7.4 Peripheral and Driver Integration

HALO leverages several key peripherals and external components to fulfill its core functionality, and integrating these into the software system requires careful configuration and driver management. The STM32 offers a rich set of built-in peripherals, and these are leveraged to connect to the various sensors, display, and storage systems used in the project.

The Hall-effect sensors, DRV5053VA, are the primary input to the system. These sensors output analog voltage signals that represent the intensity of magnetic fields. To read these values, HALO uses the onboard ADC (analog-to-digital converter). The CD74HC4067M96 multiplexer is used to handle the selection of individual sensor channels, with GPIO pins acting as select lines to choose which sensor's output is routed to the ADC at any given time. This multiplexing technique allows for an 8x8 grid of sensors, providing the system with 64 distinct sensing points. The ADC is configured for fast sampling, ensuring that data from all sensors can be acquired at a sufficient rate for the display to be updated in real time.

Once the ADC readings are acquired, they are processed and mapped to appropriate values for graphical representation. This processing is minimal, simply normalizing the raw ADC values and converting them into an intensity scale that the display can use to visualize the data. This enables the system to maintain a real-time update rate of the display without significant delays, despite the continuous sampling of the sensor array.

The RK043FN66HS-CTG touchscreen display is used for visual feedback, rendering the heatmap of the sensor array. This display is controlled via the LTDC (LCD-TFT display controller), which interfaces directly with the screen. In combination with the DMA2D (Direct Memory Access 2D engine) and Chrom-ART Accelerator, it ensures smooth and efficient rendering, even with limited MCU resources. The TouchGFX library handles the GUI framework, which allows for quick and easy management of graphical elements like circles representing each sensor, color-coded based on their intensity. TouchGFX also handles touch input, enabling users to interact with the interface to save data to the SD card or switch between different display modes.

For data storage, HALO interfaces with an SD card via the SDMMC (Secure Digital MultiMedia Card) interface. The FATFS middleware library facilitates file system management, enabling the system to write raw sensor data or heatmap snapshots to the SD card for later analysis. This operation is integrated smoothly with the rest of the application, triggered by user interaction with the touchscreen interface. Data logging is efficient and non-blocking, with buffer management ensuring that no system tasks are delayed during the write operation.

Finally, the integration of the GPIO pins for managing the multiplexer select lines and triggering other control signals is straightforward. These pins are configured through the STM32 HAL and managed within the application logic to ensure the correct sensor channel is selected at any point in time. Interrupts are used for time-sensitive operations, ensuring that sensor reads and display updates happen synchronously without interfering with each other.

Overall, we integrate peripherals and drivers in a manner that ensures both real-time sensor data acquisition and responsive display performance. By

using direct hardware access through STM32's HAL libraries and leveraging middleware like TouchGFX and FATFS, the system is able to meet its performance goals with minimal overhead and complexity.

7.5 Graphics and UI Design

The user interface (UI) of HALO is designed to be both functional and efficient, providing real-time feedback on the sensor data in a simple and intuitive manner. The UI is driven by TouchGFX, a graphical user interface framework designed specifically for STM32 microcontrollers. This framework was chosen for its ability to render complex visual elements efficiently, even on resource-constrained systems.

The display itself is a 4.3-inch TFT touchscreen (model RK043FN66HS-CTG), which offers sufficient resolution (480x272 pixels) for the simple, grid-based heatmap visualization of the Hall-effect sensors. The graphical representation of the sensors is a grid of circles, each of which corresponds to a sensor in the 8x8 array. The color of each circle is dynamically adjusted based on the sensor's reading, with a color gradient representing the intensity of the magnetic field detected.

The TouchGFX Designer tool was used to design the UI layout, providing a drag-and-drop environment for placing visual elements like the sensor grid, buttons, and text fields. This graphical interface makes it easy to design and modify the screen layout without delving too deeply into the underlying code. TouchGFX outputs C++ code that is then integrated into the firmware project, where it interacts with the main application logic.

To manage the display's performance, DMA2D is used to accelerate pixel operations such as drawing the sensor grid and updating the color of each circle. This allows for efficient, non-blocking updates to the screen while the system continues to sample sensors and handle user input. The Chrom-ART Accelerator further enhances graphics performance by offloading certain image processing tasks from the CPU, helping to keep the MCU's resources available for other real-time operations like ADC sampling.

In addition to the sensor grid, the UI includes several buttons and interactive elements, allowing the user to initiate actions like saving the heatmap

data to the SD card. The touchscreen interface uses interrupts to handle user input, ensuring that the UI remains responsive and that actions like button presses are detected promptly. Since the display is relatively simple, the application avoids complex animations or transitions, focusing instead on smooth, real-time updates of the sensor grid and minimalistic navigation controls.

The decision to use TouchGFX ensures that the display remains responsive while consuming minimal processing power. Since the sensor acquisition and display update processes are both handled within a super loop, the UI design ensures that these tasks do not block each other. By using DMA for display updates and keeping the UI simple, HALO strikes a balance between graphical fidelity and real-time performance.

7.6 Task Scheduling and Concurrency

Given the simplicity and real-time requirements of HALO, a complex task scheduler or operating system is not necessary. Instead, the software operates on a bare-metal framework, where control is passed from one function to another through a super loop, with interrupts used to handle specific tasks that require immediate attention. This approach minimizes overhead and simplifies the software architecture, ensuring that the system is as efficient as possible while maintaining precise timing for critical operations like sensor sampling and display updates.

The primary loop in the firmware constantly handles the acquisition of data from the Hall-effect sensors. Through the use of a multiplexer (MUX), the 8x8 sensor array is scanned by sequentially selecting each sensor's output through GPIO-controlled MUX select lines. This process involves rapidly sampling the analog data from the sensors using the STM32H735IGT6's ADC channels, followed by digital conversion and storage in memory. Once all sensor data is collected, the system updates the graphical representation on the touchscreen display. This process is repeated continuously to maintain real-time feedback of the magnetic field readings.

Interrupts play a crucial role in handling concurrency, particularly when interacting with external hardware like the touchscreen and ADC. ADC interrupts are used to signal when a new sample is available, allowing the software to

process the data and update the system state without delay. Similarly, touchscreen interrupts are used to detect user input in real-time, allowing for actions like saving data to the SD card or navigating menus. By using interrupts, the firmware ensures that sensor data acquisition and user input processing can happen simultaneously without blocking each other.

However, careful attention is given to ensure that time-sensitive tasks are prioritized. For example, reading from the Hall-effect sensors and updating the display must be done frequently and within strict timing constraints to provide a smooth user experience. To prevent blocking or delays, all non-urgent tasks, such as SD card saving or system status checks, are performed in the background, allowing the primary functions to proceed uninterrupted.

Overall, this scheduling approach strikes a balance between simplicity and efficiency, ensuring that HALO's core functions—sensor data acquisition, display updates, and user interaction—are performed with minimal latency and maximum reliability.

7.7 Software Challenges and Considerations

Despite the straightforward nature of the system, several software challenges had to be addressed to ensure that the application runs reliably and efficiently. These challenges primarily revolve around optimizing the system for real-time performance, managing memory usage, and ensuring that interrupts and non-blocking operations do not interfere with critical tasks like sensor data acquisition and display updates.

HALO requires continuous sampling of the sensor array, real-time updates to the touchscreen display, and quick response to user input via touch gestures. One of the main challenges was ensuring that these operations occur with minimal delay while using the STM32H735IGT6's limited resources effectively. The use of a bare-metal approach and a super loop structure was critical in achieving the required performance, but careful attention had to be paid to task prioritization. Interrupts were essential for managing time-sensitive operations like ADC sampling and touch detection, but ensuring that interrupt service routines (ISRs) were kept as short as possible was key to maintaining system responsiveness. Longer ISRs could block other critical operations, so all

non-essential tasks were moved to the main loop or handled in the background to avoid timing issues.

Memory management was another significant consideration, especially since the STM32H735IGT6 is equipped with 564 KB of SRAM, which must be shared between various tasks, including sensor data storage, display buffering, and interrupt handling. To optimize memory usage, the system must allocate memory carefully between the sensor buffer, the framebuffer for the display, and any other temporary storage, such as the SD card write buffer. The framebuffer, in particular, requires a considerable amount of memory, and placing it in the appropriate non-cacheable region of SRAM was necessary to ensure correct rendering without data corruption. Additionally, DMA was used to transfer sensor data and update the display without blocking other operations, further optimizing memory usage by reducing CPU involvement in these tasks.

The use of interrupts is crucial for handling concurrent operations, especially for real-time tasks like reading from the ADCs and detecting touch inputs from the touchscreen. However, improper management of interrupts can cause timing issues, such as missed ADC samples or unresponsiveness to user input. Therefore, careful attention was paid to designing the interrupt service routines to be as fast and efficient as possible. Time-consuming tasks such as data processing and UI updates were moved out of the ISR and into the main loop to prevent blocking other important interrupts. The system also needed to ensure that tasks such as SD card logging did not interfere with the primary sensor sampling and display update loops. By delegating non-urgent tasks to lower-priority functions, the system could maintain smooth operation without unnecessary delays.

Although power management was not a primary concern for the initial iteration of HALO, it is an area that may need attention in future versions. The STM32H735IGT6 supports low-power modes that could be used to extend battery life, particularly when the system is idle or not actively sampling sensors. Optimizing the firmware to enter low-power states during inactivity, such as when the display is not being updated or when the sensor array has completed its cycle, could help improve overall energy efficiency. Additionally, the system could be designed to scale its performance based on user demand, lowering

clock speeds or reducing peripheral activity when maximum processing power is not needed.

Debugging embedded systems, particularly those with real-time constraints like HALO, can be challenging. One issue encountered was ensuring that timing constraints were met in the presence of debug breakpoints or logging statements. The debugging process required extensive use of hardware breakpoints and serial communication to monitor system state without introducing additional delays. Testing was also complicated by the absence of a real-time operating system (RTOS), which would have provided built-in tools for managing concurrency and scheduling. As a result, much of the testing focused on ensuring that the timing of sensor reads and display updates remained within acceptable bounds, even as new features were added or hardware changes were made.

While HALO is designed to work with an 8x8 grid of sensors, the system is flexible enough to support future expansion, either through increased sensor counts or more complex display features. If needed, the system could be adapted to handle additional sensor arrays or more complex user interfaces, although this would require careful attention to memory and processing power. Future enhancements might also include wireless communication for remote monitoring or control, although this would introduce additional software complexity, particularly around power management and data synchronization.

8: Fabrication & Prototyping

9: Testing and Evaluation

10: Administrative Content

10.1 Bill of Materials

Table X: Bill of Materials

Component	Name	Price	Quantity
-----------	------	-------	----------

<u>Electrical</u>			
Hall-Effect Sensors	DRV5053VA	\$0.55	256
STM32 Microcontroller	STM32H735IGT6	\$14.70	1
Development Board	STM32H735G-DK	\$87.91	1
Linear 3.3V Regulator	LP2985-33DBVR	\$0.48	1
Switching 5V Regulator	LTC3525ESC6-5#TRPBF	\$6.75	1
3.7V Lithium-Ion Battery	PRT-13853	\$5.50	1
Battery Management Chip	MCP73871	\$2.51	1
USB Type C Port	USB4155-03-C	\$1.61	1
16:1 Multiplexer	CD74HC4067M96	\$0.58	16
Analog-to-Digital Converter	AD7940BRJZ	\$13.52	1
8-Bit Binary Counter	SN74HC590ADR	\$0.48	4
2-to-4 Decoder	SN74HC138DR	\$0.26	4
Memory Card Connector	10067847-001RLF	\$2.00	1
Touchscreen LCD Display	RK043FN66HS-CTG	\$44.80	1
Flat-Flex Cable	0150150233	\$2.02	2
Misc Components		~\$20.00	
<u>Mechanical</u>			
3D-Printed Housing	ABS Filament	\$10.00	
Shielding Metal (Optional)	Mumetal Foil		
Total		\$182.86	

The total estimated budget for this project is \$179.78. This estimate accounts for all necessary components, with some flexibility for potential adjustments or

additional expenses. While the budget does not currently include spare parts, we plan to order extra Hall-effect sensors and other critical components to mitigate the risk of damaged or defective parts. Given the high number of sensors in our design, we will be sure to test each unit before integrating them into our PCB to ensure they function reliably. If unexpected changes must be made, we will be sure to look at cost-effective alternatives or adjust our component selection to stay within a reasonable budget.

10.2 Project Milestones

Table X: Senior Design I Milestones

Item	Start Date	End date	Anticipated Duration
Brainstorming	1/10/25	1/16/25	1 week
Divide and Conquer (D&C) Report	1/16/25	1/23/25	1 week
Divide and Conquer (D&C) Group Meeting	1/27/25	1/27/25	1 day
Update and upload D&C report onto group website	2/1/25	2/6/25	< 1 week
Parts selection	2/7/25	2/15/25	1 week
Acquire parts	2/7/25	3/10/25	Around 1 month*
Initial PCB design	2/16/25	3/2/25	2 weeks
Breadboard prototyping	3/11/25	4/14/25	1 month
Midterm Report	3/16/25	3/23/25	1 week
Midterm Report Group Meeting	3/26/25	3/28/25	1 day
Update and upload midterm report onto group website	3/28/25	4/3/25	1 week
SD1 Final Report & Mini Demo Video	4/15/25	4/23/25	1 week

Table X: Senior Design II Milestones

Item	Start Date	End date	Anticipated Duration
Finalize Design and Parts Inventory	05/12/2025	05/14/2025	2 days
PCB Assembly	05/14/2025	05/28/2025	2 weeks
Finalize Firmware Development	05/28/2025	06/25/2025	4 weeks
Integration Testing	06/25/2025	07/02/2025	1 week
System Optimization	07/02/2025	07/09/2025	1 week
Final Testing and Validation	07/09/2025	07/16/2025	1 week
Final Documentation and Presentation	07/23/2025	07/30/2025	1 week

11: Conclusion

12: References

[1] Jansen, Peter. "A Magnetic Imager Tile." Hackaday.io, 29 July 2017, <https://hackaday.io/project/18518-iteration-8/log/64376-a-magnetic-imager-tile>.

[2] Cai, Jiangwei, et al. "A High-Resolution Magnetic Field Imaging System Based on the Unpackaged Hall Element Array." *Applied Sciences*, vol. 14, no. 13, 2024, p. 5788. <https://doi.org/10.3390/app14135788>.

[3] S. Akka and M. Missous, "Design of a portable, two dimensional magnetometer, using 2 dimensional quantum Hall effect sensor array, optimised for low magnetic field applications," The 10th IEEE International Symposium on Electron Devices for Microwave and Optoelectronic Applications, Manchester, UK, 2002, pp. 177-182, doi: 10.1109/EDMO.2002.1174951. keywords: {Sensor arrays;Magnetometers;Hall effect devices;Design optimization;Magnetic sensors;Voltage;Data acquisition;Spinning;Magnetic fields;Magnetic switching},

[4] N. Fischer, J. Kriechbaum, D. Berwanger and F. Mathis-Ullrich, "Compliant Hall-Effect Sensor Array for Passive Magnetic Instrument Tracking," in *IEEE Sensors Letters*, vol. 7, no. 3, pp. 1-4, March 2023, Art no. 2500404, doi: 10.1109/LSSENS.2023.3250971.

keywords: {Sensor arrays;Sensors;Magnetic sensors;Bending;Target tracking;Magnetic flux;Shape;Magnetic sensors;passive magnetic sensing;flexible sensors;Hall-effect sensors;magnetic tracking;self-sensing;wearable sensors},

[5] McGill University, "Analog to Digital Conversion - Virtual Lab on Biomedical Signals," McGill University. [Online]. Available: https://www.medicine.mcgill.ca/physio/vlab/biomed_signals/atodvlab.htm. [Accessed: 22-March-2025].

[6] Spectrum Instrumentation, "ADC and Resolution," [Online]. Available: https://spectrum-instrumentation.com/support/knowledgebase/hardware_features/ADC_and_Resolution.php. [Accessed: 22-March-2025].

[B1] "Types of Batteries." *Monolithic Power Systems*, <https://www.monolithicpower.com/en/learning/mpscholar/battery-management-systems/introduction-to-battery-technology/types-of-batteries>. Accessed 16 March 2025.

[B2] "What is the Energy Density of a Lithium-Ion Battery?." *Flux Power*, 21 Aug. 2020, https://www.fluxpower.com/blog/what-is-the-energy-density-of-a-lithium-ion-battery?utm_campaign=Newsletter&utm_medium=email&_hsmi=93911746&utm_content=93911744&utm_source=hs_email.

[B3] Murden, Dave. "LITHIUM IRON PHOSPHATE (LIFEPO4) BATTERY ENERGY DENSITY." *Eco Tree Lithium*, 14 Apr. 2023, <https://ecotreelithium.co.uk/news/lifepo4-battery-energy-density/>

[B4] "An Introduction to Batteries: Components, Parameters, Types, and Chargers." *Monolithic Power Systems*, <https://www.monolithicpower.com/learning/resources/an-introduction-to-batteries-components-parameters-types-and-chargers>. Accessed 16 March 2025.

[B5] "Understanding the Voltage of LiFePO4 Cells: A Comprehensive Guide." *EV lithium*, <https://www.evlithium.com/Blog/lifepo4-cell-voltage-importance-and-impact.html>. Accessed 16 March 2025.

[B6] Simpson, Chester. "Linear and Switching Voltage Regulator Fundamentals." *Texas Instruments*, <https://www.ti.com/lit/an/snva558/snva558.pdf?ts=1741541072059>. Accessed 16 March 2025.

[B7] Castaldo, Adam. "Switching regulator fundamentals." *Texas Instruments*, Feb. 2019, <https://www.ti.com/lit/an/snva559c/snva559c.pdf?ts=1741524827804>