Paradigm Engineering

B.O.A.T

May 15, 2023
30th Intelligent Ground Vehicle Competition
Memorial University of Newfoundland & Labrador

Paradigm

Aidan Clark
amhclark@mun.ca
Patrick Cleary
pcleary@mun.ca
Shiloh Burton
sjmburton@mun.ca
Ava Gogal
afgogal@mun.ca

Dylan Matthews
dmatthews19@mun.ca
Jumanah Babar
jbabar@mun.ca
Luke Cronin
lcronin@mun.ca
Noel Rowsell
ncrowsell@mun.ca

Andrew Nash
anash@mun.ca
Ian Carroll
icarroll@mun.ca
Thomas Porter
tgporter@mun.ca
Quinn Parsons
qbparsons@mun.ca

Dr. Oscar De Silva [Faculty Advisor]

____

I certify that the design and engineering of the vehicle by the current student team has been significant and equivalent to what might be awarded credit in a senior design course.

Table of Contents

1 Introduction

Paradigm Engineering is honored to represent Memorial University of Newfoundland & Labrador (MUNL) for the second consecutive year at the Intelligent Ground Vehicle Competition (IGVC). The team's latest creation, B.O.A.T., short for Built for Off-road Autonomous Transit, boasts advancements in mechanical, electrical, and software engineering designed for both parking lots and rugged terrains.

The mechanical team has innovated to create a design which is agile and efficient for all terrains. The electrical sub-team has developed three custom circuit boards, ensuring superior control, efficient power distribution, and safety through an effective emergency-stop (e-stop) mechanism. Lastly, the software team's strategy for the Auto-Nav Challenge leverages advanced computer vision algorithms and real-time data processing for precise obstacle detection and course navigation.

Building on the lessons learned from last year, Paradigm aims to make a significant impact at this year's IGVC with these innovations.

2 Team Organization

Paradigm employs a predominantly flat team structure, fostering effective collaboration and streamlined communication. The organization is divided into four distinct sub-teams, namely Business/Marketing, Electrical, Mechanical, and Software. The comprehensive roster of the design team and their respective departments can be found in Table 2-1.

Table 2-1 - Technical team members and departments

Aidan Clark	MUNL Mechanical Engineering
Patrick Cleary	MUNL Mechanical Engineering
Ava Gogal	MUNL Mechanical Engineering
Shiloh Burton	MUNL Mechanical Engineering
Andrew Nash	MUNL Computer Science (Graduate)
Ian Carroll	MUNL Computer Engineering
Thomas Porter	MUNL Computer Engineering
Quinn Parsons	MUNL Computer Engineering
Dylan Matthews	MUNL Electrical Engineering
Jumanah Babar	MUNL Electrical Engineering
Luke Cronin	MUNL Electrical Engineering
Noel Rowsell	MUNL Electrical Engineering

A combined total of approximately 4000+ hours have gone into the vehicle development at the time of submission on this report. Total vehicle component cost is approximately $20,000 USD.

3 Design Assumptions and Process

The design process at Paradigm Engineering is driven by the stipulations provided in the competition's rules document. These guidelines were carefully interpreted and converted into design criteria that guided the teams' efforts. The team's ideation process drew from a blend of tried-and-tested solutions and innovative approaches. These were systematically evaluated against the agreed-upon criteria, ensuring the design stayed on the right course.

Adding another layer of thoroughness, the team engaged Paradigm's extensive alumni network for design reviews. Their valuable insights and expertise brought an additional perspective, helping each sub-team refine their work. This balanced approach of peer and alumni reviews ensured the designs were robust, compliant with the criteria, and primed for the challenges of the competition.

4 Innovations in Vehicle Design

A standout feature of B.O.A.T., the team's entry in this year's competition, is its innovative skid-steer drivetrain. This setup, with two independent drivetrain modules, offers superior maneuverability, enabling the vehicle to easily navigate on both paved and off-road terrains.

On the software front, Paradigm's navigation system utilizes state-of-the-art machine learning techniques for semantic segmentation and bird's eye view (BEV) mapping. This approach allows B.O.A.T. to identify and map obstacles, such as lane lines and potholes, that are often challenging to discern using conventional lidar-based mapping techniques.

The electrical team recognized the importance of simplicity and reliability and streamlined the vehicle's architecture. The team consolidated key components onto three custom circuit boards, simplifying the wiring harness and enhancing overall system dependability. This strategic simplification improves manufacturing efficiency and ensures easier maintenance and troubleshooting.

5 Mechanical Systems

This section outlines the mechanical systems design and addresses the validation methods employed to ensure the vehicle can perform as anticipated.

5.1 Design Overview

Table 5-1 - Vehicle overview

Total Vehicle Mass	120 kg
Vehicle Height	1195 mm
Vehicle Length	1141 mm
Vehicle Width	1110 mm
Minimum Ground Clearance	96 mm

5.2 Drivetrain

The skid-steer drivetrain consists of two modular packs that drive two wheels (left and right sides). The wheels are driven through a timing belt and pulley system powered by a stepper motor and gearbox arrangement. The wheel shafts are supported by two double row angular contact bearings. Figure 5-1 shows the drivetrain system.

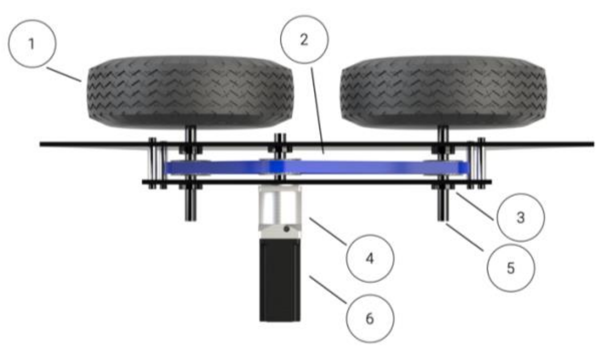

Figure 5-1 - Drivetrain module schematic

Table 5-2 - Drivetrain module part list

Number	Part
1	Wheel (x2)
2	Timing Belt
3	Double Row Angular Contact Bearing (x5)
4	10:1 Reduction Planetary Gearbox
5	Driveshaft (x3)
6	NEMA 34 Stepper Motor

5.3 Chassis

The vehicle chassis is designed to house and protect the electronics, battery pack, payload, and flight computer. This chassis design allows for easy access to all critical parts system, while providing sufficient mounting for all components and sensors. The chassis is comprised of 1”x1/8” 6061 Aluminum square tube, which is welded together. Figure 5-2 shows the vehicle chassis.

Figure 5-2 - Vehicle chassis

To ensure that the chassis could handle the loading inflicted during operation, an Ansys Mechanical simulation model was created to simulate the loading conditions to ensure safe operation is possible. The exterior of the vehicle is generally composed of 3mm white cast acrylic panels which provide protection from the outer elements. These panels also provide a location for the team’s sponsorship logos to be displayed.

5.4 Weatherproofing

The entire vehicle is shielded from the elements with the acrylic panelling. The intake and exhaust fans are water-resistant; this prevents key components from being exposed to moisture. Due to the competition timing coinciding with the natural shedding of the cottonwood seeds from local Michigan trees, fan dust covers were implemented to prevent intake of this or any other environmental debris. As an additional layer of protection against the elements, all exterior seams are sealed with a waterproof silicone to ensure the design is watertight.

6 Electrical and Power Systems

This section delineates the design of the electronic and power systems, including the power distribution system, electronic suite description, and safety devices.

6.1 Electrical System Overview

The vehicle’s electrical system consists of three main subsystems: power distribution, compute, and control. The power distribution subsystem ensures the safe and efficient delivery of power to all vehicle components. The compute subsystem is accountable for tasks such as perception, navigation, and control decisions. Meanwhile, the control subsystem is responsible for motor actuation and the implementation of safety circuitry.

6.2 Power Distribution Subsystem

The power distribution subsystem features an Off The Shelf (OTS) 48V 50Ah lithium-ion battery with a battery management system and an OTS 1500W power inverter for DC to AC conversion. A custom circuit board ensures power distribution and fault protection, while custom wire harnessing provides secure power routing. The setup meets Paradigm's vehicle-specific power and safety needs. See Figure 6-1 for the subsystem architecture.

Conservative calculations determined that the vehicle’s nominal power draw while driving is approximately 1200W, which corresponds to 25A of current at 48V. The lithium-ion battery is the most compact high-capacity battery available that was capable of providing the required current. The battery’s specifications are given in Table 6-1. With the battery’s capacity of 50Ah, the vehicle has a run time of 2 hours while driving. While the vehicle is stationary, the power consumption is reduced to around 600W, giving a standby run time of 4 hours.

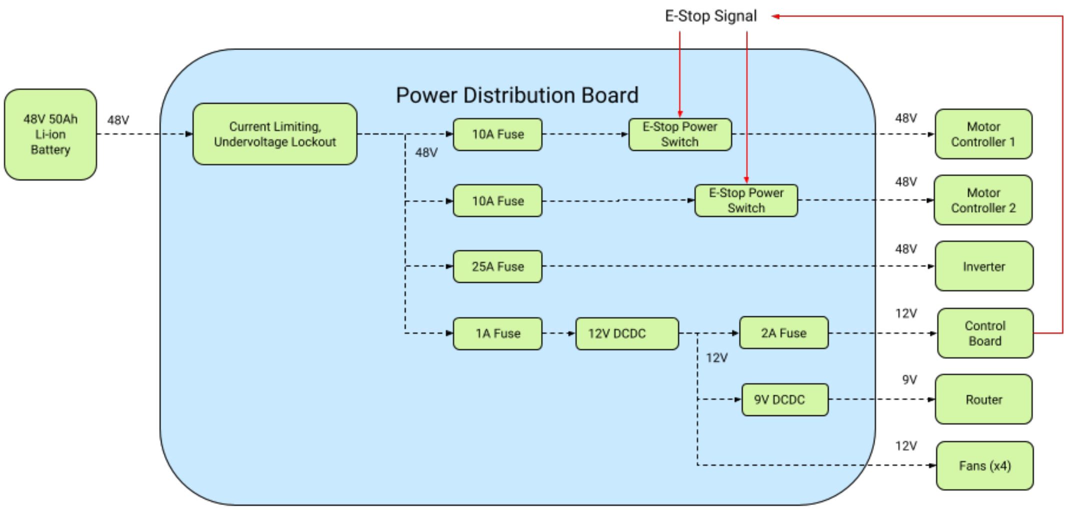

Figure 6-1 – Block diagram of power distribution subsystem

Table 6-1 - Battery specifications

Specification	Value	Unit
Battery Chemistry	NMC Lithium-Ion	N/A
Nominal Voltage	48	Volts
Charge Time	~5	Hours
Run Time (Full Load)	~2	Hours
Battery Capacity	50	Amp-Hours
Maximum Continuous Discharge Current	50	Amps

The power distribution board was designed to perform the following functions:

- Regulate 48V battery power to other voltages.
- Distribute 48V, 12V, and 9V power to vehicle's electrical components.
- Implement protection circuitry (fuses, inrush current limiting, undervoltage lockout).
- Interface with e-stop signal from control board to cut motor power.
- Power and control fans to cool power components.

A hot swap controller is used at the battery's input to the board to implement inrush current limiting and undervoltage lockout functions. Switching DC-DC converters are used to efficiently convert 48V battery power to 12V and 9V rails. A render of the power distribution board is shown in Figure 6-2.

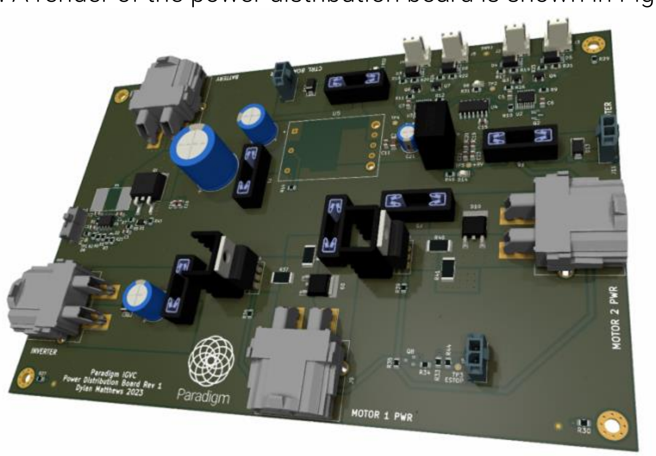

Figure 6-2 - Custom power distribution board

6.3 Control Subsystem

The control subsystem uses two OTS NEMA 34 stepper motors for high-torque output, interfaced with OTS drivers for smooth operation. A custom control circuit board ensures optimal performance, while custom signal wire harnessing facilitates data transmission. This configuration aligns with Paradigm's vehicle operational requirements.

The main component of the control subsystem is Paradigm's custom control circuit board. The control board was designed to perform the following functions:

- Communicate with the compute subsystem over CAN (Controller Area Network) bus.
- Control the vehicle's motors.
- Implement circuitry for required IGVC safety functions (wireless and wired e-stop, safety lights).
- Regulate 12V power to other voltage rails.

Communication with the compute subsystem, and control of the motors and safety lights, are all achieved with the use of an ESP32-S3 microcontroller running custom firmware on the control board. The logic for the vehicle's e-stop is implemented in hardware using logic integrated circuits. A render of the control board is given in Figure 6-3, while Figure 6-4 illustrates the architecture of the control and compute subsystems.

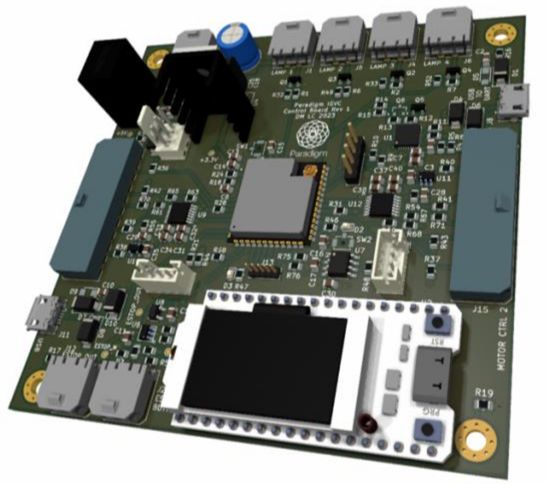

Figure 6-3 - Custom control board

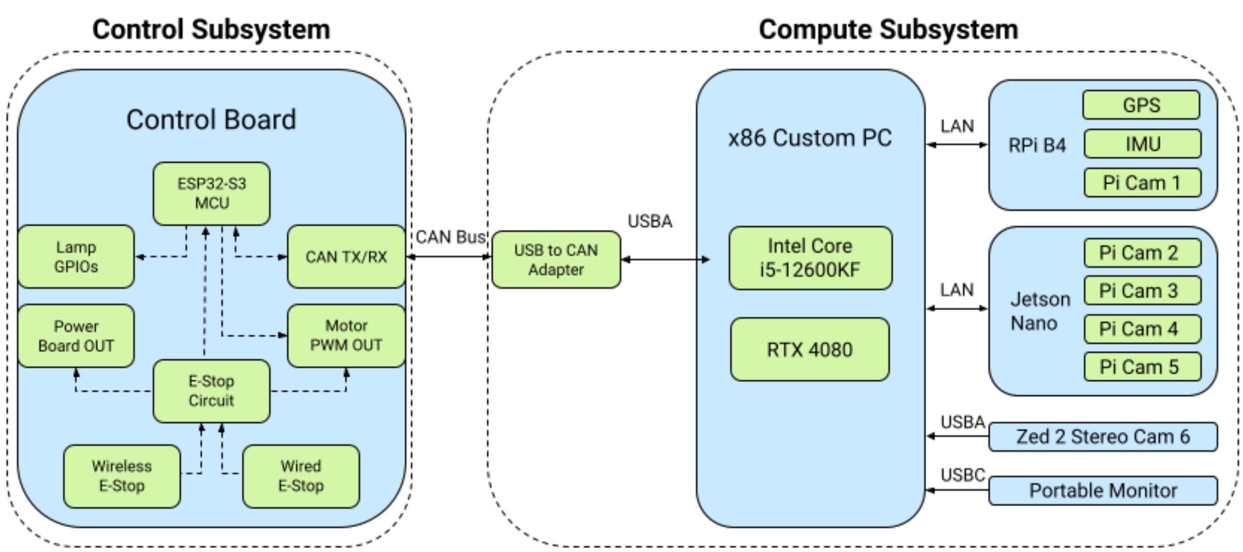

Figure 6-4 - Block diagram of control and compute subsystem architecture

6.3.1 CAN Message Summary

The CAN communication protocol, commonly used in vehicles, is being used to control the B.O.A.T. effectively. The CAN bus is configured to transmit/receive at a rate of 500 kbps, resulting in less than 200 μs reaction speeds to control messages. The Compute and Control subsystems must agree regarding the exchange of CAN messages. Due to this, a message map has been implemented aboard both subsystems, which details behavior according to what type of message has been transmitted/received. For instance, to control the B.O.A.T. to go around a right turn, a CAN message would be sent detailing motor actuation. This message would have an ID that controls the motors, with a payload detailing a fast velocity on the left motor and a payload detailing a slower velocity on the right motor to make the right turn.

6.4 Compute Subsystem

The team's compute subsystem is a robust assembly of high-performance computing modules and camera arrays. The central component is a custom-built computer, the "flight computer," with an RTX 4080 Graphical Processing Unit (GPU) for real-time machine-learning tasks.

Supplementing the flight computer, a Jetson Nano with an Ardu Camera Hardware Attached on Top (HAT) live-streams four cameras over the Robot Operating System (ROS) network. A Raspberry Pi B4 module, equipped with a BerryIMU-GPS HAT for localization, supports Inertial Measurement Unit (IMU) and Global Positioning System (GPS) modules that communicate over the ROS network. The Raspberry PI's native Camera Serial Interface (CSI) is also utilized for a fifth Ardu camera.

Together, these five cameras, and a standalone ZED 2 camera, continuously feed real-time visual data into the team's computer vision models. This data integration is achieved through a ROS network via a router.

6.5 Safety Devices

6.5.1 Emergency Stops

The mechanical e-stop is a latching, normally open, push button placed at the top rear of the vehicle. This button is connected to the control board to generate an active high signal when pressed.

The wireless e-stop uses two ESP32 LoRa 915MHz development boards; one is used as a transmitter and the other as a receiver. A custom PCB, shown in Figure 6-5, is used to interface a button and power with one of the development boards for the transmitter. The second development board is located on the vehicle's control board, acting as the receiver. The transmitter uses the same button as the mechanical e-stop, and when the button is pressed, a signal is sent to the receiver. The receiver then receives this signal and generates an active high signal on the control board.

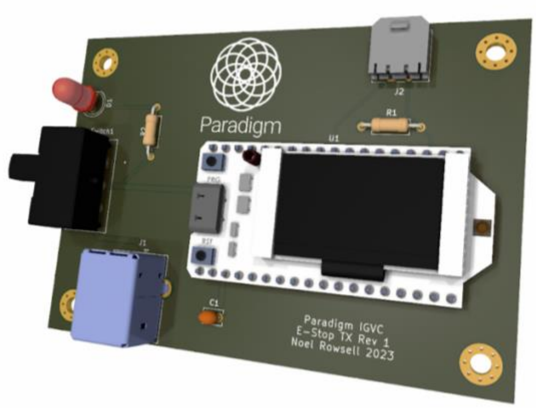

Figure 6-5 - Render of e-stop transmitter board

For this competition, the e-stop circuit was designed to be completely digital. The team observed interference issues with the analog e-stop circuit during last year's competition, so a digital circuit was used to reduce noise susceptibility. The e-stop system was designed to perform three actions:

1. Cut power to both motors (hardware based).
2. Disable the motor controllers (hardware based).
3. Stop the ESP32-S3 on the control board from driving the motors (firmware based).

Only one of these three actions needs to be performed to stop the vehicle, so even if two actions fail, the vehicle will still stop. This greatly increases the reliability of the circuit. The circuit works by taking the logical OR of the active high mechanical and wireless e-stop signals. The output signal is inverted to make both active high and active low e-stop signals available. The active low signal is sent to the motor controllers as the controllers will be disabled by a low signal. The active high signal is sent to the ESP32-S3, as well as to the power board through a cable. On the power board, a signal-level MOSFET is used to convert the e-stop signal to a gate signal for power MOSFETs which switch motor power on the low side. When the e-stop signal is low, the MOTOR_FET_GATE signal is high, and the motor power is on. When the e-stop signal is high, the signal-level MOSFET pulls MOTOR_FET_GATE low, and the motor power is switched off. An overview of the e-stop circuit is shown in Figure 6-6.

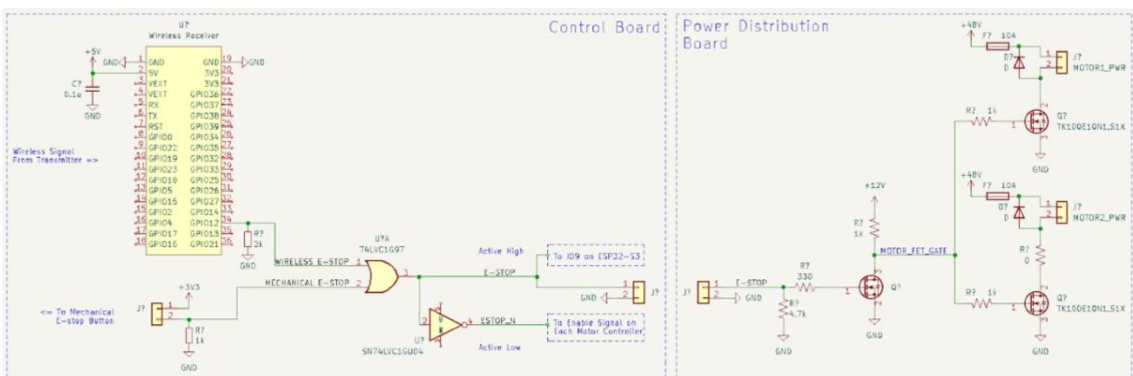

Figure 6-6 - Overview of e-stop circuit

6.5.2 Safety Lights

To go above and beyond the imposed safety regulations on the B.O.A.T, the control subsystem contains an internal state machine residing in the firmware. This state machine accommodates actions for the safety light requirements and acts as a safety net for the hardware e-stops. This firmware runs on the control board's ESP32-S3 chip and actuates the safety lights, which are connected to the control board, using low-side NMOS transistor switches.

In Paradigm's vehicle state diagram, the vehicle transitions from boot to standby state with solid safety lights. Manual or autonomous control prompts flashing lights. If not in these control states, possibly due to an emergency stop, the vehicle rejects motor commands, ensuring safety. The return to standby solidifies the lights again. The transition between these states and their corresponding safety light behaviors are illustrated in Figure 6-7.

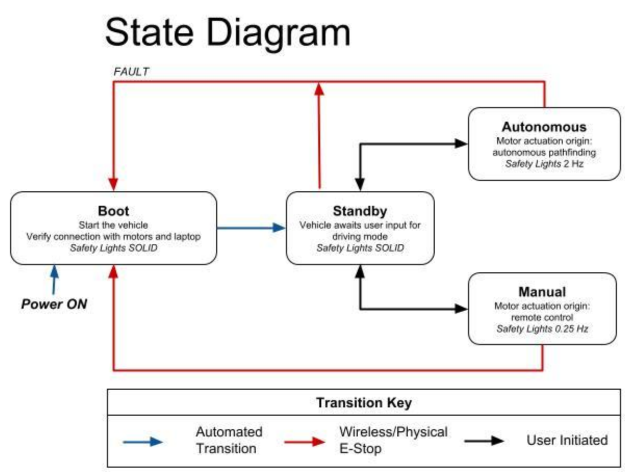

Figure 6-7 - System safety and control state diagram

6.5.3 Battery Management System

The vehicle's OTS battery has an integrated battery management system (BMS) to protect itself. This BMS offers protection against undervoltage, overvoltage, short circuits, and reverse polarity, and it also performs cell balancing for the battery. In addition to this, a hot swap controller is placed at the battery input on the power distribution board to limit inrush current and offer additional protection against undervoltage and short circuits.

6.5.4 Fuses

Automotive blade fuses are placed in the path of every high-power output connector on the power distribution board. These fuses provide protection against faulty components or wire damage and will help protect components in fault conditions and prevent wire fires. These fuses also simplify debugging by making it easier to find faulty components if issues arise. All fuses are connected to the board via fuse holders and can be easily replaced.

7 Software Strategy and Mapping Techniques

7.1 Overview

Paradigm's autonomous vehicle software strategy is especially designed to tackle the Auto-Nav Challenge. Paradigm's software system uses advanced computer vision algorithms to map the Auto-Nav course with varying obstacles. By processing real-time visual data, Paradigm's navigation system swiftly identifies obstacles, assesses their positions, and makes real-time decisions to adapt the vehicle's path. A full system diagram can be seen in Figure 7-1.

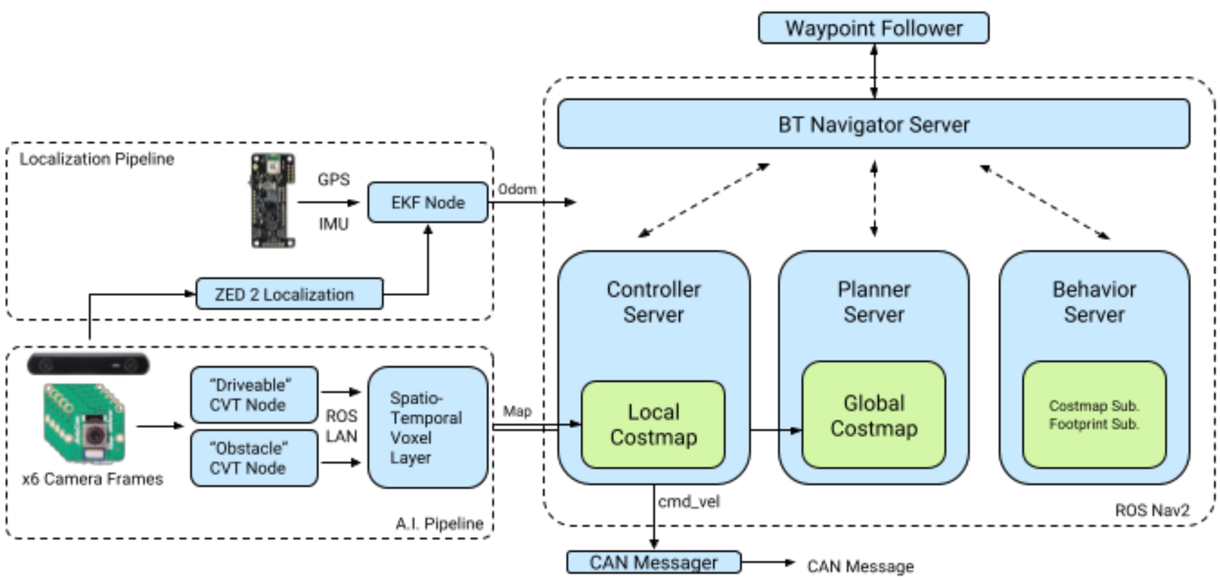

Figure 7-1 - Paradigm navigation software overview

The system diagram in Figure 7-1 highlights the seamless integration of Paradigm's custom AI and localization pipelines with the ROS2 Nav2 system. This fusion enables precise real-time decision-making and path adjustments, underlining the robustness of Paradigm's autonomous vehicle software in tackling the Auto-Nav Challenge.

7.2 Obstacle Detection and Avoidance

The Paradigm navigation system applies state-of-the-art research from the autonomous vehicle domain, to construct real-time BEV environmental maps surrounding the robot. Paradigm's perception system utilizes the Cross-View Transformer (CVT), an attention-centric machine learning model proposed by Zhou et al. [1]. These BEV predictions feed into the ROS2 navigation system for path planning.

7.2.1 Vision-Based Obstacle Detection

Six high-resolution cameras, capturing at a rate of 60 frames per second, are strategically stationed at the top of the vehicle, offering a 360-degree view of the robot's vicinity. Two distinct CVT instances process the real-time visual data from these cameras, each generating a 256x256 matrix of confidence levels. These levels indicate the model's confidence in the presence of an obstacle in an approximately 1 in^2 area, for a total perception area of approximately 21 ft^2. These outputs are subsequently incorporated into a ROS costmap. An example prediction from the model can be seen in Figure 7-2 below.

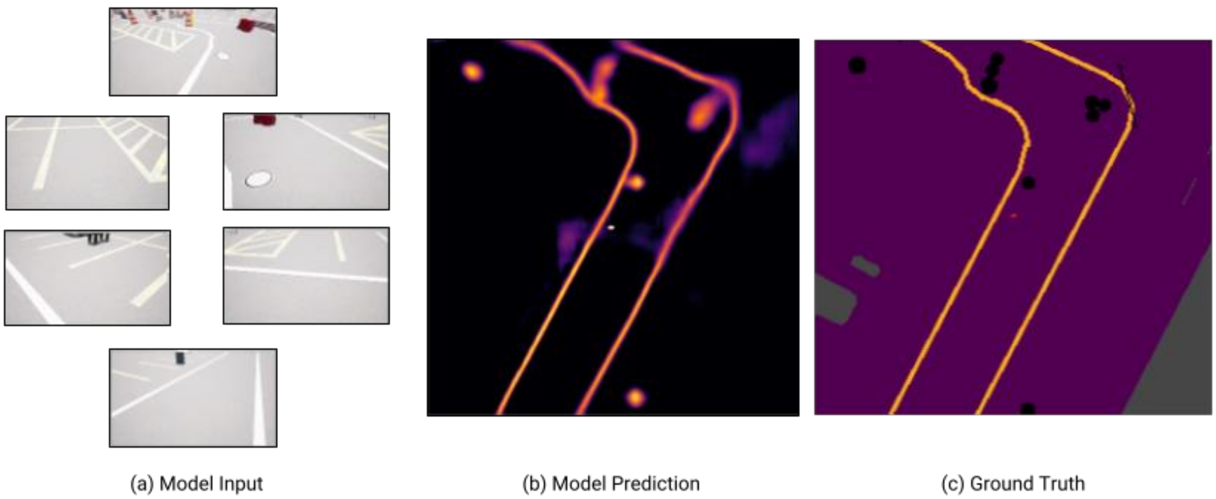

Figure 7-2 - Example perception situation; (a) Six camera images input to model (b) CVT obstacle prediction confidence map - lighter is higher confidence (c) Ground Truth Semantic Segmentation

The CVT employs a cross-view, cross-attention mechanism to aggregate the six camera image features into a shared map-view representation. The cross-view attention module learns the geometric structure of each camera to link camera-view pixels to bird's eye view locations. All cameras share the same image-encoder but use a positional embedding dependent on their individual camera calibration.

7.2.2 Obstacle Avoidance

The two CVTs collaborate to form detailed BEV of the environment. The "obstacle" CVT instance detects various obstacles, including artificial barrels, traffic cones, humans, cars, lane lines, and simulated potholes. The second CVT instance, the "driveable" model, predicts the feasible navigation area around the robot, marking the space between the lane lines and the course ramp as driveable.

7.3 Software Strategy and Path Planning

Paradigm's software strategy integrates a semantically segmented costmap into the ROS2 navigation system to ensure efficient path planning. The Smac Hybrid-A* Planner, being cost-aware, uses this costmap to help the robot navigate effectively around obstacles and within drivable areas.

Driveable area prediction is central to Paradigm's strategy. Specifically, for scenarios like leaving the "No Man's Land", the robot needs to find and follow lane lines to re-enter the track. This is achieved by mapping drivable areas such as between lane lines and assigning them a negative cost value, making them attractive paths for the robot.

7.4 Map Generation

Paradigm's perception system feeds its BEV predictions into the ROS Spatio-Temporal Voxel Layer, a powerful tool for generating cost maps. This package provides an efficient mechanism for integrating BEV predictions with ROS costmaps while also accounting for the inherent noise in AI-based perception systems.

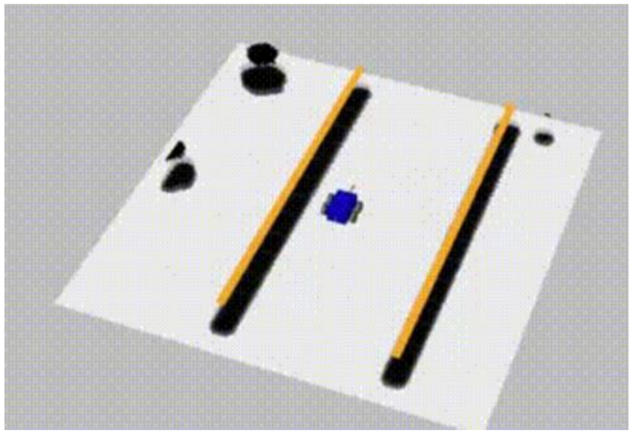

Figure 7-3 - Local costmap generation, showcasing obstacles detected

The costmap is a two-dimensional grid overlaid on the environment, where each cell represents a certain level of "cost" or difficulty for the vehicle to traverse. This cost value is based on the outputs of the CVT instances. Obstacles detected by the "obstacle" CVT instance increase the cost, while areas marked as drivable by the "driveable" CVT instance decrease the cost. An example obstacle cost map can be seen in Figure 7-3.

Paradigm's software team has carefully fine-tuned the costmap generation process. Low-confidence predictions, which are likely to be erroneous, are discarded. The system also applies a temporal filter to discard non-persistent predictions, reducing the likelihood of false positives.

7.4.1 Localization

In Paradigm's autonomous vehicle system, the ZED2 stereo camera and the BerryGPS-IMU work within the ROS2 network to enable precise localization. The ZED2 camera provides accurate visual odometry, while the BerryGPS-IMU contributes acceleration, angular velocity, and baseline geolocation data. These diverse data inputs are seamlessly integrated and communicated in real-time over the ROS2 network, facilitated by a router. Then, an Extended Kalman Filter (EKF) consolidates these inputs to estimate the vehicle's position and velocity. This robust localization approach, leveraging the ROS2 network, ensures reliable vehicle navigation under various conditions.

7.5 Goal Selection and Path Generation

While competition-provided waypoints guide Paradigm’s system, the robustness of its pathfinding capabilities allows for liberal tolerance on these waypoints, which account for any GPS errors. This flexibility ensures safe navigation, allowing minor deviations from the waypoints while keeping the vehicle within drivable areas.

Path planning on a rich costmap, including drivable areas and obstacles, enhances the team's navigation system resilience. When receiving a waypoint, the Smac Hybrid-A* Planner generates an optimal path based on the costmap. The DWB controller then refines this path using local data, with the Constrained Smoother ironing out any irregularities for more straightforward navigation. Finally, the SimpleGoalChecker verifies goal attainment, and the next waypoint is traveled to.

7.6 Additional Creative Concepts

Paradigm's B.O.A.T design goes beyond safety norms through rigorous testing of software, firmware, and electrical components, especially for e-stop and safety light functions. The team adopts the real-world industry standard CAN communication protocol, enabling fast and reliable vehicle communication. Six cameras stream data into advanced computer vision models for real-time environment mapping, enhancing obstacle detection and avoidance. This innovative approach bolsters both safety and performance, marking a milestone in the team's autonomous vehicle design.

8 Failure Identification and Mitigation

8.1 Vehicle Failure Modes and Resolutions

Potential failure modes were analyzed as a team and addressed throughout the design process. Table 8-1 demonstrates the critical failure paths addressed in the system design.

Table 8-1 - Failure modes and resolutions

Failure Mode	Repercussions	Resolution
Tip-over	Loss of maneuverability	• The vehicle's design has been strategically chosen to minimize tipping hazards by positioning the center of gravity towards the base. • In the event the vehicle tips over, Wireless E-stop will be used to power down the vehicle
Chassis hang-up on ramp	Loss of maneuverability	• The geometry if the vehicle is such that there is sufficient ground clearance for the maximum angle encountered during the course
Machine vision impairments from weather or lighting	Inaccurate readings of environment in bird's eye view and cost map	• Data set generated from the simulator contains a variety of environment lighting and weather conditions. Additionally, models are pre-trained on real world automotive datasets
Inability to pass between obstacles	Loss of maneuverability	• Skid-steered system allows for full motive control with no perceived configuration where the vehicle will be unable to maneuver between obstacles.
Loss of power	Inability to continue with the course run	• Battery capacity sized for excess run time to allow for multiple runs • Vehicle charging will be taking place as frequently as possible.
Object becomes lodged in wheels	Loss of maneuverability	• Wheels have been given sufficient clearance to ensure any expected debris (small rocks, etc.) can clear the vehicle chassis

8.2 Vehicle Failure Points and Resolutions

Careful consideration has been given to address system failure points. Table 8-2 summarizes the major concerns addressed in the vehicle design.

Table 8-2 - Vehicle failure points and mitigation

Point of Failure	Repercussions	Mitigation Technique
Battery harnessing breaks or is intermittent	• Possible arcing • Possible rail transients • Loss of total power to the electrical system	• Hot swap controller integrated into power distribution board to protect electrical system
CAN bus cable breaks or is intermittent	• Communication for navigation may be offline or give poor values	• Integrated firmware detects and moves vehicle to safe state
Loss of signal to wireless e-stop	• Wireless e-stop will not be able to stop the vehicle	• E-stop transmitter will periodically ping receiver to ensure connection, and if ping fails e-stop is activated on the vehicle • Ensure optimal placement of 915 MHz antennas on vehicle and transmitter
EMI from motors interferes with e-stop signal	• E-stop signal unintentionally triggers and stops the vehicle	• Redundancy in e-stop circuitry eliminates requirement for off-board e-stop signal wire • Control board containing main e-stop circuit is placed far from motors
Payload becomes dislodged	• Damage to internal components	• Dual redundant methods of securing the payload

8.3 Battery Thermal Event Protection

The 48V 50Ah Li-ion Battery with NMC Lithium-Ion Deep Cycle [2] chemistry is a high-performance battery. This battery features an integrated thermal protection system that is designed to prevent any thermal runaway events that may occur during operation. This system monitors the current of the pack and limits it to 50A. The sensors in the thermal protection system continuously monitor the battery temperature and detect any increase in temperature above a pre-set threshold. When the temperature exceeds the threshold, the control circuitry initiates a response to prevent thermal runaway.

9 Simulation

Paradigm has created a customized version of the CARLA simulator, based in Unreal Engine, specifically for simulating the Auto-Nav challenge. A unique Unreal Engine level was designed to replicate the competition setting as closely as possible. As depicted in Figure 9-1, the IGVC parking lot has been recreated with approximately +/- 0.5-foot accuracy. This includes similar placements of trees, greenery, lane lines, ramp, and parking spots. The IGVC level can be rearranged to generate various competition layouts for ROS navigation testing and computer vision data collection.

Figure 9-1 - Paradigm custom IGVC level in Unreal Engine

9.1 Simulations in Virtual Environments

Paradigm utilizes the CARLA simulator for data generation and as a robust platform for ROS navigation testing. The team leverages CARLA's ROS2 Bridge, a powerful tool that seamlessly facilitates the exchange of sensor and control data between the simulator and the ROS2 environment. This setup allows for live streaming of various sensor data directly into Paradigm's navigation pipeline, including camera images, IMU readings, and GPS coordinates. The real-time data feed simulates all inputs the navigation system would normally receive in an actual run, enabling the software team to test and debug the pipeline in a realistic but controlled environment.

The use of CARLA for ROS testing provides significant benefits to the team. It enables testing and debugging before the vehicle is fully assembled, accelerating the development cycle. This method also allows the team to simulate and test their navigation algorithms against various environmental conditions and course layouts without any of the risks or costs associated with physical testing. This would be difficult, if not impossible, in the real world where conditions are constantly changing. This approach also allows Paradigm to ensure that its navigation system is robust and reliable, capable of handling the wide range of scenarios the vehicle might encounter during the competition.

9.2 Theoretical Concepts in Simulations

Paradigm employs the CARLA [3] Python API to generate a diverse dataset for training computer vision models. To create the IGVC dataset, the B.O.A.T. is integrated into the CARLA level. The team has ensured that the simulation's camera positions and intrinsic parameters match those of the real cameras mounted on the vehicle.

To cultivate a diverse dataset, the B.O.A.T. is randomly placed at different positions within the course, with varying orientations. This process is repeated approximately 20,000 times. At each frame, corresponding camera input images are collected and BEV ground truth label images. In addition to the variations in vehicle placement and orientation, other dynamic factors are also manipulated to enhance the diversity of the dataset. These include changes in the sun's position (reflecting the time of day), weather conditions, track layout, ground material, obstacle materials, and the B.O.A.T.'s specific position. With these measures, the team aims to ensure the dataset's robustness, capturing a wide range of scenarios the autonomous vehicle might encounter during the competition.

To train Paradigm's CVT models, they first are pe-trained nuScenes dataset [4], a large-scale real-world dataset for autonomous vehicles, where the teams model achieves near state-of-the-art BEV prediction performance. Pre-training the models allows the model to generalize a BEV prediction strategy and familiarize itself with many out-of-sample scenarios, enhancing its robustness and transferability from simulation to real-world scenarios. Subsequently, these models are fine-tuned using Paradigm's custom generated IGVC dataset. Initial testing has demonstrated the CVT models can operate in parallel at 43 FPS on the robot's RTX 4080 GPU and achieve an Intersection Over Union (IoU) score of 64%.

10 Testing

Various components, systems, and subsystems of the vehicle have undergone testing. Key highlights include:

- Firmware: Extensive unit testing and validation on the ESP32-S3 microcontroller confirmed the functionality of the finite state machine and safety lights. Further testing validated the CAN and PWM communication protocols on the control board.
- Navigation Software: Simulator testing is actively conducted to validate the navigation software's performance, including sensor noise and fault events.
- Motor and Drivetrain: Bench-top tests were carried out on the motors and skid steer drive train modules to ensure their functionality and reliability.
- Computer Vision: Real-world lane line tests were performed to verify the accuracy and effectiveness of computer vision models.
- Circuit Boards: Initial bench-top power rail and circuit testing has been completed on all circuit boards, and a motor has been driven by the boards. Remaining testing will focus on testing the complete assembled system at increased electrical load.

11 Performance

Performance evaluations on various subsystems have yielded the following key results:

- Computer Vision: The CVT models run efficiently at 43 FPS on the RTX 4080 GPU, with a solid IoU score of 64% indicating accurate object detection.
- Emergency-Stop Circuit Boards: These components have been validated to function effectively up to 800 meters, ensuring safety.
- Motors: These have been tested to operate at 2,000 RPM and when paired with the planetary gearbox runs at 200 RPM, confirming the powertrain performance and motor specifications.

Testing will persist to further refine subsystem performance.

12 References

[1] B. Zhou and P. Krähenbühl, "Cross-view Transformers for real-time Map-view Semantic Segmentation," *Computer Vision and Pattern Recognition,* 2022.

[2] "48V 50Ah Li-ion Battery - NMC Lithium Ion Deep Cycle," AEGIS, [Online]. Available: https://www.aegisbattery.com/collections/48-volt-nmc-li-ion-batteries/products/48v-50ah-li-ion-battery-pvc-1?variant=30306983673954. [Accessed 15 May 2023].

[3] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez and V. Koltun, "CARLA: An Open Urban Driving Simulator," *Conference on Robot Learning,* 2017.

[4] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Lion, Q. Xu, A. Krishnan, Y. Pan, G. Baldan and O. Beijbom, "nuScenes: A multimodal dataset for autonomous driving," *Computer Vision and Pattern Recognition,* 2019.

Appendix A – Control Board Schematic

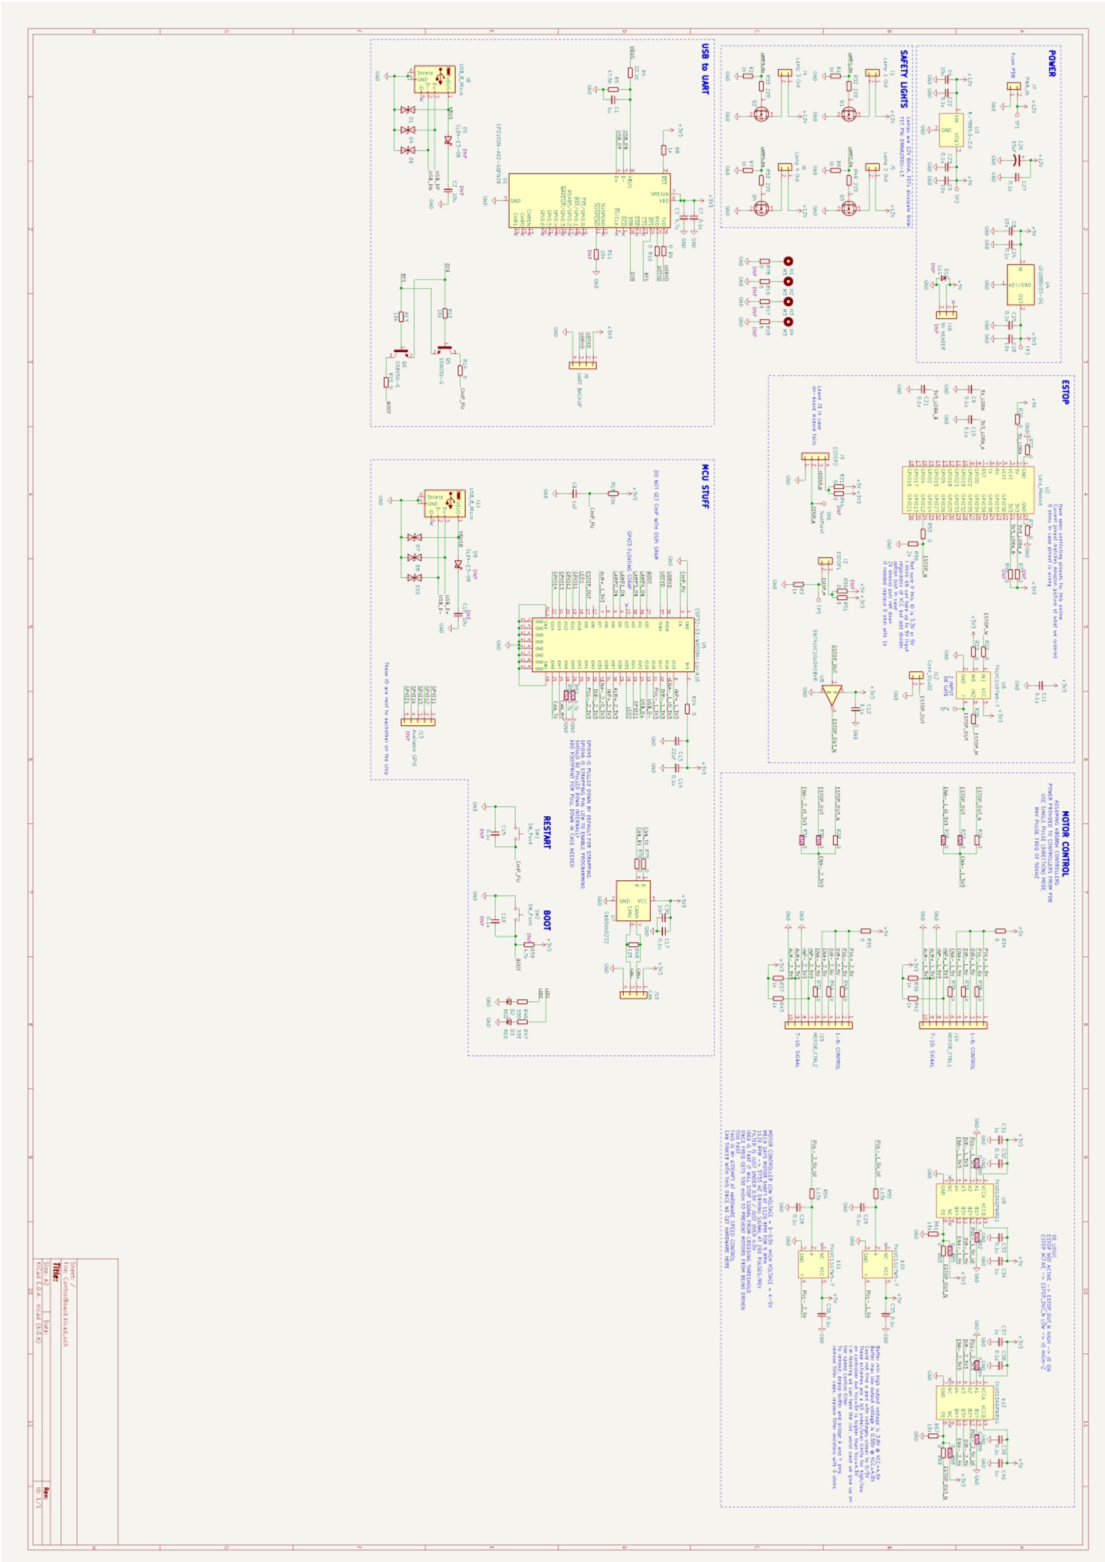

Appendix B – Power Distribution Board Schematic

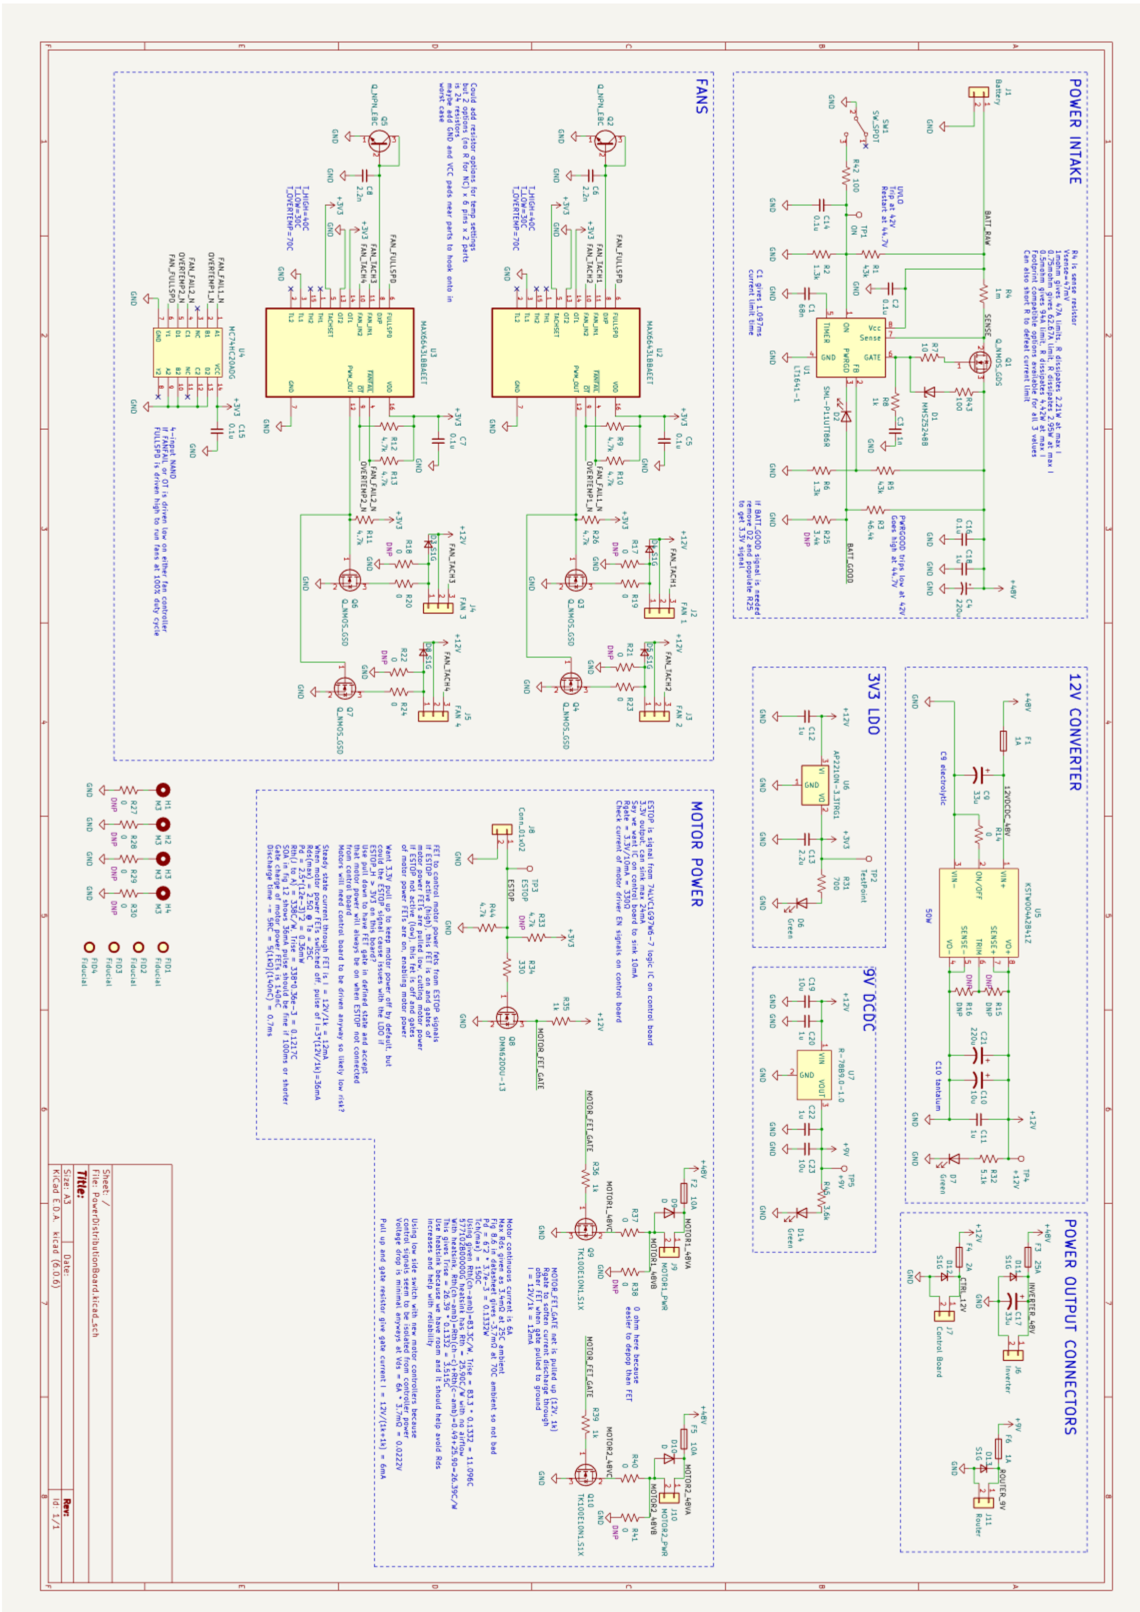

Appendix C – Emergency-Stop Transmitter Board