Millersville University of Pennsylvania

Project Phoenix

Statement of Integrity

I certify this report and Project Phoenix are products of the work and intellect of the below listed team members.

5/14/2026

Dr. John Wright Jr., Lead Advisor

Date Certified and Submitted

John.Wright@millersville.edu

Joseph LaMontagne, President	jolamont@millersville.edu	Kate Hilton	mdhilton@millersville.edu
Sofia Griffiths, VP	kjgriffi@millersville.edu	Josiah Kerr	jskerr@millersville.edu
Hecmarys Cintron, Treasurer	hecintro@millersville.edu	Benjamin Wright	brwrigh1@millersville.edu
John Hershey, Secretary	jjhersh1@millersville.edu	Rhiley Gordon	rdgordon@millersville.edu
Tristan Rush, PR Chair	tjrush@millersville.edu	Jeff Pinegar, Advisor	JPinegar@phoenixcontact.com
Benjamin Cromwell	bpcromwe@millersville.edu	Ian Troop, Advisor	iptroop@gmail.com
Zane Weaver	zaweaver@millersville.edu	Dr. John Haughery, Advisor	john.haughery@millersville.edu

Table of Contents

2. System and Subsystem Requirements

Project Phoenix is the eighth robot that the Millersville University Mobile Robotics Research Team has submitted to IGVC. This year's design was developed around a multi-stage predictive/reactive path planning system. Project Phoenix also utilizes an in-house manufactured chassis and drivetrain, rather than the tradition of using a wheelchair chassis.

2.1 Mechanical Design Requirements

From a mechanical perspective, the requirements the team focused on achieving were keeping the dimensions at minimum specifications, while utilizing industrial components. This requirement allows for the maximum amount of mobility on course, especially in the chicane. Custom designing the chassis around the drivetrain gives a more robust and accurate motor mounting system, allowing the motors to exist on the same axis and thus have no interference relative to one another. This also allows the motors to be center-mounted for omni-directional steering.

2.2 Safety Design Requirements

Project Phoenix's design prioritizes personal safety for the electrical system by utilizing effective and accessible E-STOP methods that completely isolate hazardous elements and power. The 48V required for the motors is safely controlled, as the team worked closely with SEW Eurodrive and Phoenix Contact sponsors to calculate component load and ensure effective electrical safety measures.

The team ensured safety for the mechanical system by minimizing pinch points on the frame and lid and minimize the chance of injury from moving parts through guarding. All relevant safety devices are to be mounted where they are visible and accessible to anyone.

2.3 Electrical Design Requirements

The electrical team worked with a heavy focus on industrial components and direct upgrades compared to previous systems. This grew from the team's sponsorships with SEW Eurodrive, Sick, and Phoenix Contact, as the team had ample access to high level industrial materials for this project.

2.4 Perception Design Requirements

The main requirements for perception included having a high enough resolution to detect small objects and obstacles, while simultaneously not overloading the PLC's CPU and keeping the dataset visualizable. This meant that reading large amounts of data and scaling the resolution down was necessary.

2.5 Driving Logic Design Requirements

The driving logic design requirements focus primarily on vehicle safety, followed by optimal pathing. The highest priority for Project Phoenix is to avoid collisions with obstacles or crossing of lane lines. Beyond that, the logic should attempt to optimize pathing for minimal steering, as steering requires the robot to have less forward velocity, resulting in slower run times.

2.6 Key Performance Indicators

Our KPIs for Project Pheonix were pre-established milestones such as physical construction, mounting all systems required for motion, object detection and avoidance, lane detection and avoidance, and GPS systems were all deemed important indicators, with the final integration stage being the most critical.

3. Mechanical Design

Project Pheonix focuses on an innovative mechanical design, implementing the experience of all team members to create the most capable robot yet from the team. Last year, the team discovered several mechanical design flaws during the competition. This year, the team developed targeted improvements to fix those issues and optimize the robot's performance at IGVC.

3.1 Vehicle Chassis

Project Pheonix is built on a custom chassis of welded square aluminum tubing. Its total dimensions are 24 inches wide by 36 inches long, with a height of 24 inches at the panel and 26 inches at the front box. The LiDAR sensor sits at a sensing height of 27 inches on a custom-made sheet metal mount. The main body of aluminum tubing is encased in riveted sheet metal to protect our batteries and payload.

3.2 Drive System

Project Pheonix utilizes a drop center drive train powered by two SEW Movimot motors. The wheels are inset into the frame, which simplifies the control algorithm as it allows us to make the robot's footprint a perfect rectangle. The design also utilizes a larger wheel diameter in comparison to previous entries to cut down on traction loss situations. Four caster wheels allow for the rotation point to be at the center of the robot while always supporting weight.

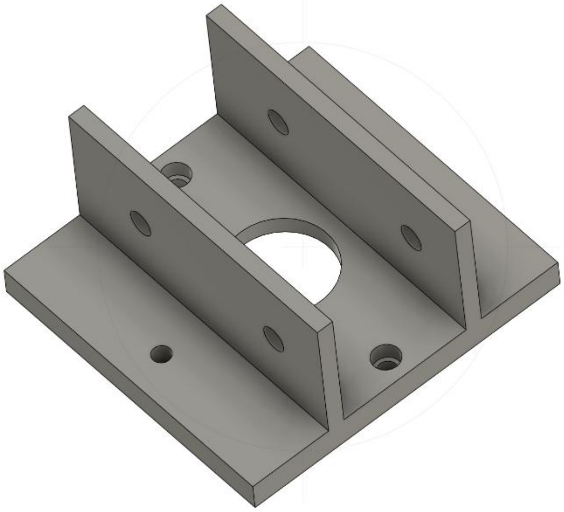

Figure 1: Motor mount brackets

The Motors are mounted on custom machined brackets that are bolted directly into the frame, which prevents slippage. A machined shaft mates to the motor inside the frame and bolts directly to the wheel.

3.3 Weather Proofing

The team's focus on industrial components allows a high level of weather proofing. While many of our components are open to the weather, they are each individually rated for the elements. external components, such as the motors, LiDAR, and panel enclosure are IP65 rated.

4. Safety System

The safety system design on Project Phoenix was made to be compatible and effective with the industrial-grade components used on the robot. It utilizes industrial safety components and practices to ensure the safety of all personnel and equipment while testing and operating.

4.1 Safety in Movement States

When transporting, if the robot is E-Stopped or in any state except MANUAL or AUTO, the power to the motors is isolated. The robot is in a freewheel state, so it can be safely moved without back feeding current or having the motors move autonomously when travelling. MANUAL mode allows the robot to remotely move so it can be safely positioned without straining team members by physically moving it.

While parked, the robot will indicate it is in the STOPPED state through the stack light mounted visibly on the top of the robot.

While charging, batteries are physically unhooked from the system, ensuring there is no unintended current flow. Furthermore, using 24V Dakota Lithium batteries, with their designed charger, ensures that no thermal runaway can occur. These batteries also contain a built-in control circuit that balances load between individual cells and prevents overcharging and shorting.

4.2 Safety while Operating

While running, the robot indicates its current state with the stack light LEDs. These states include running, danger due to physical object proximity, and ESTOP. The robot can be stopped with the ESTOP button which will lock the motors. The ESTOP button is visible and accessible as it is placed on the top of the robot. The robot can also be stopped by the remote ESTOP, which can be triggered remotely and from over 150 feet away. The robot is always monitored by at least two team members, who can either trigger the remote ESTOP or activate the mounted ESTOP button.

4.3 Electrical Safety

We use a single-channel live-wire safety system. This ensures that in the event of power loss to the ESTOPs or if they trigger, the robot will lock out the motors. If the ESTOP or Remote ESTOP is triggered, the Safety Relay will trigger, sending a signal to the PLC to turn the ESTOP indicator light on and send a signal to the Force-Guided Contactor to trigger. The Force-Guided Contactor is a device that will forcibly isolate power to our devices, even in situations where normal contactors fail. This effectively isolates all power to the motors. The team collaborated with its sponsors to calculate the load of all components and properly size the Force-Guided Contactor to shut off necessary components without fail.

5. Electrical Design

5.1: Wiring Schematic

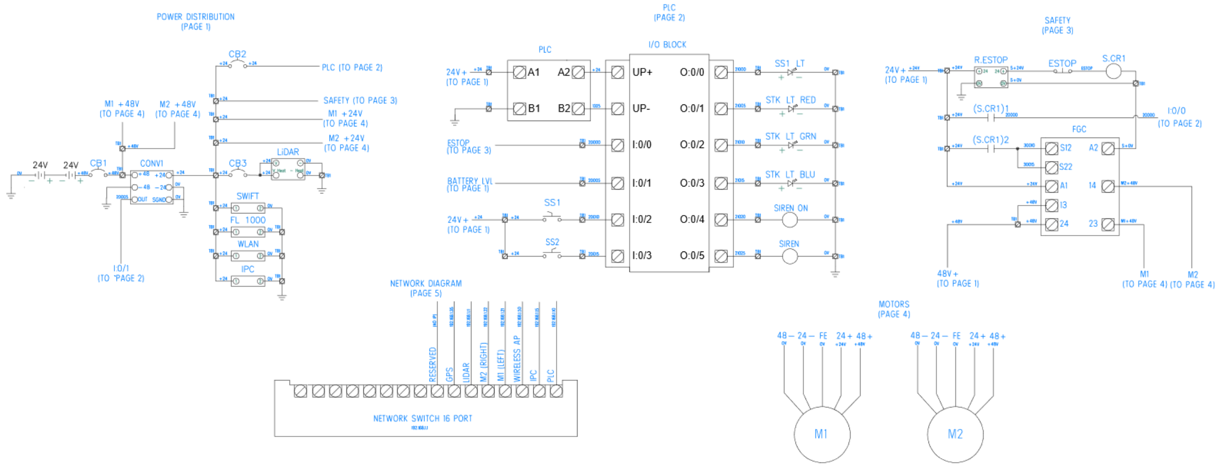

Figure 2: Wiring diagram

5.2: PLC I/O Table

Address	Description	Wire No.
I:0/0	Emergency Stop Signal	20000
I:0/1	Low Battery Level Signal	20005
I:0/2	Selector Switch AUTO	20010
I:0/3	Selector Switch MANUAL	20015
O:0/0	Selector Switch Light	21000

O:0/1	Stack Light Red	21005
O:0/2	Stack Light Green	21010
O:0/3	Stack Light Blue	21015
O:0/4	Siren Trigger Signal	21020
O:0/5	Siren Power	21025

5.3 Component Overview

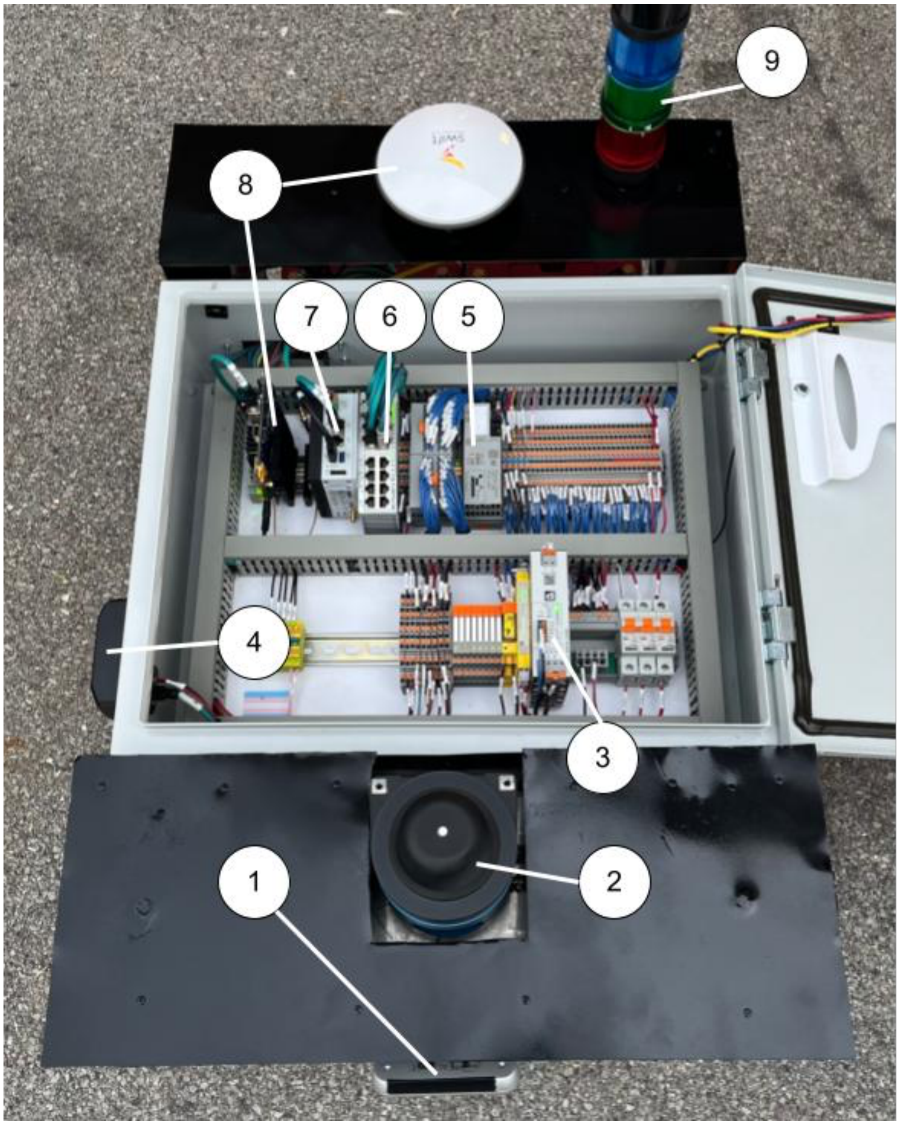

Figure 3: The inside of the control Panel

Dakota Lithium Batteries (not shown) – The device uses two 24V Dakota Lithium batteries to provide power to the system. They were chosen for their capacity, rechargeability, and reliability witnessed in previous years.

SEW Eurodrive Motors (not shown) – The robot utilizes two 48V high-torque motors to drive the robot. These Industrial motors were chosen for the high amount of torque they produce at 48V, and to showcase the components of one of our sponsors, SEW Eurodrive.

RealSense Camera (1) - This is a stereoscopic 3D vision camera. It was chosen for its cost-effectiveness, vision capabilities, and compatibility with the IPC.

LiDAR (2) – The robot uses a 2D 360° SICK NAV310 LiDAR to detect physical barriers around itself. This LiDAR was chosen specifically for the flexibility of a 360° LiDAR to fully customize the angle the robot perceives. The team did not need a 3D LiDAR as all obstacles are in the 2D plane. One of our sponsors, SICK Sensor Intelligence, provided a suitable sensor. The team utilized this in conjunction with the RealSense camera for an added safety factor, as the LiDAR has a higher effective range and reliability.

Phoenix Contact 24V Power Supply (3) – This power supply was chosen to convert 48VDC from the batteries to 24VDC for the controls wiring, and components. This was due to its industrial focus, power capabilities, and to showcase Phoenix Contact's products.

WLAN Module (4) – The design utilizes the WLAN 1120 from Phoenix Contact to remotely access the robot's software. This module was chosen for the industrial focus, communication compatibility, industrial security, and to support the sponsor.

Axioline F 2152 Phoenix Contact PLC (5) – The robot uses this as the main logic processor. The PLC runs the inputs, outputs, and decision-making. This PLC was chosen for its industrial focus, as the team opted to use a PLC over the onboard computer for our drive control. This decision was made due to the PLCs ability to integrate with industrial components well. The PLC is also very flexible, with interchangeable I/O cards if more inputs or outputs are required. The team was also generously given this component from one of its sponsors, Phoenix Contact, to highlight their PLC.

Network Switch (6) – The design uses a 16-channel FL1000 industrial network switch from Phoenix Contact to communicate between devices. It was chosen for its industrial focus, compatibility, number of network ports, and to showcase one of our sponsor's products.

IPC (7) - This is a small industrial PC to run stereoscopic camera post-processing. This was chosen for its compatibility with other components, cooling factors, cost effectiveness, and to show off Phoenix Contact's IPC.

SWIFT (8) - The Piksi Multi GNSS Module is our GPS module. It was chosen for ethernet and power compatibility, as well as for its performance.

Stack Light (9) - The Phoenix Contact Stack Light warns personnel of hazards as well as indicate the state of the robot. This was chosen as it interfaces well with the PLC, has many colors and effects for many different indication states, and shows off Phoenix Contact's Stack Light.

5.4 Power Distribution & Capacity

The system is powered by two 24V Dakota Lithium batteries, wired in series to provide 48V. The batteries have a capacity of 60AH. Because of this high capacity, the robot can run for weeks of testing without needing a recharge. The batteries can be recharged fully over the course of 12 hours. Our motors are driven by the 48V power from the batteries. All other components are driven on a 24V output from our power supply.

6. Perception

Project Phoenix Utilizes three sensors for perception across three main systems. The Swift navigation system is utilized for global positioning. The Sick NAV310 LiDAR is utilized to detect obstacles such as barricades and barrels. The LiDAR was chosen for this job as it is, within the preconditions of the competition, guaranteed to detect obstacles. The detection of potholes and lane lines is left to the onboard Realsense camera, utilizing stereoscopic vision to give us proper depth perception and tracing.

6.1 Positioning

The swift navigation system consists of a Piksi Multi GNSS Module, which provides Project Phoenix with a GPS waypoint. The system provides an accuracy of up to 1cm, leaving no room for doubt on where the robot's location is. The system is also capable of refreshing its data every 50ms, meaning a 20Hz refresh rate. This functionally means new data is available on demand and won't be stale when utilized. The Piksi Multi GNSS Module can boot fully and acquire signal within 10 seconds allowing for rapid startups.

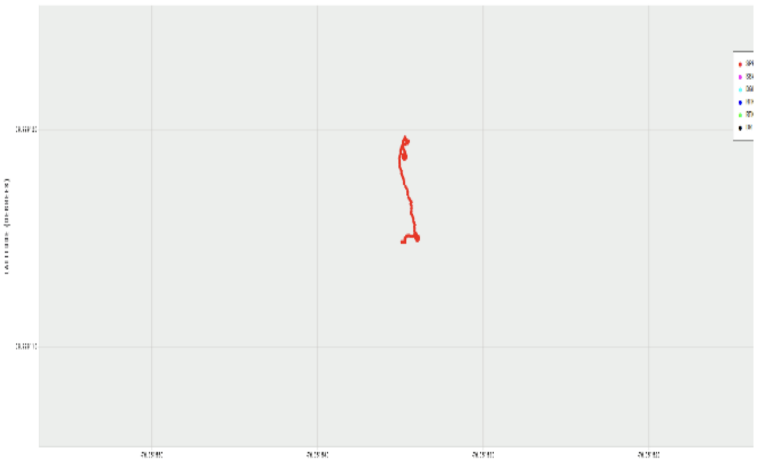

Figure 4: Swift navigation system mapping out a live coordinate solution (39.999098, -76.351486).

Data is extracted from the system through ethernet via dedicated python middleware. The script extracts raw SBP strings and alters them to be digestible for the PLC. This digestible form consists of a structure containing the following variables: Latitude, Longitude, Heading, North Velocity, and East Velocity. These variables are all in LREAL format, which is equal to a double in non-industrial programming. Offloading this work to a python middleman allows for the PLC to focus on other tasks such as motor control and safety systems.

6.2 Object Detection

Objects are detected through the NAV310 LiDAR donated to the team from Sick Sensor Intelligence. This LiDAR can scan up to 360° with a variable angular resolution from 0.0625° to 1.825°. It also boasts variable scan frequency of 5Hz to 20Hz. Our application utilizes a 180° scan at a half degree resolution with 20Hz frequency, meaning 360 sample points. The LiDAR has a functional range of 250m, while maintaining an accuracy of ± 10 mm. It has an IP65 rating, meaning it will stand up to rainy conditions and dusty environments.

The scan data from the LiDAR is sent via telegram packet to the PLC directly. The PLC issues a "sera LMDscandata" byte command via CoLa B, and receives an "SSn LMDscandata" response, containing a header and scan data. This data is presented sequentially in big endian, with every 2 bytes being equal to a half degree poll. The angle is implied, and the 2 bytes supplied are the distance measured at that point.

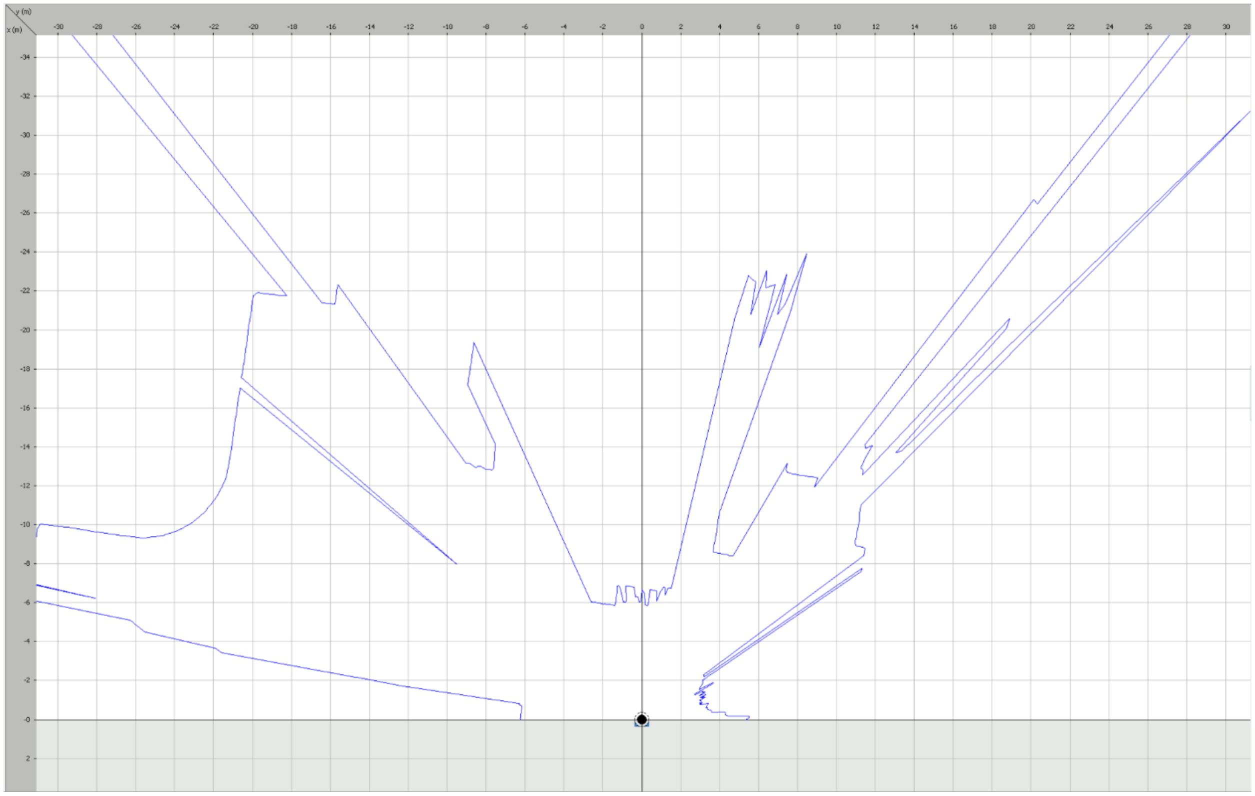

Figure 5: A sample LiDAR scan for our application.

6.3 Lane and Pothole Detection

Lane Detection is completed using an Intel RealSense Stereoscopic camera. This sensor has an official effective range of 3 meters, but in testing we have demonstrated a relatively accurate (>85% accuracy) range of at least 5 meters. The camera has been modified to increase its IP rating, ensuring it can withstand rain. It has also been mounted in a way that allows its angle to be adjusted as needed, ensuring vision can be tuned to specific needs.

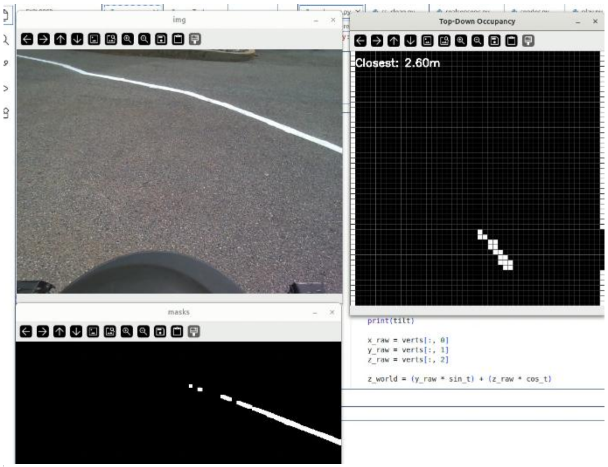

```
print(tilt)

x_raw = verts[:, 0]
y_raw = verts[:, 1]
z_raw = verts[:, 2]

z_world = (y_raw * sin_t) + (z_raw * cos_t)
```

Figure 6: The raw color image from the RealSense camera, the white mask, and the final camera occupancy grid output

A python middleman takes the depth and color frames from the RealSense camera, applies a keystoning transformation to account for mounting angle, and merges them to generate a pointcloud. Once the pointcloud is generated, all data is thrown away that exists more than

100mm off the ground. This helps to throw away noise and information otherwise not wanted or needed. A color filter is applied, and the resulting pointcloud is blurred then sharpened to further reduce noise. Finally, the grid is flattened and shrunk down to a 50x50 occupancy grid with each cell being 200mm x 200mm. This is sent to the PLC via a TCP socket connection.

6.4 Internal Awareness

The core of awareness is held in the AXL F 2152 PLC, donated from Phoenix Contact. The PLC stores a 50x50 grid of cells, each cell possessing a LiDAR occupied and camera occupied bit. Each grid cell is 200mm x 200mm, giving the robot 10 square meters of vision in front of it. This allows the robot to have a full understanding of all obstacles ahead of it. Additionally, line data is predicted to identify lane lines beside the robot on either side to assist with navigation. The body of the robot exists at coordinate (0, 25), meaning half the dataset shows obstacles on the right and half on the left. The robot is approximately 4 grid squares wide, inflated for a margin of safety.

This grid is built by first receiving the camera grid from a vision script. This script gets directly cast to the final occupancy grid, modifying all camera occupied grid locations. Then, LiDAR data gets converted from theta magnitude to a cartesian array. This allows for the data to be merged into the LiDAR occupied bits.

6.5 Real world Application of Perception

Testing with real datasets has revealed the accuracy of various sensor systems to be on par with or better than advertised. The NAV310 sick LiDAR was the easiest to verify, as we placed objects at known distances of 1, 2, 5, and 10 meters. The LiDAR accurately gave these measurements down to under a centimeter, beyond what we could accurately measure for. The Swift navigation system also proved to be more accurate than testable, as centimeter level precision is survey grade accuracy. This will be tested using IGVC’s official survey points, but all testing completed thus far has shown it to be more accurate than any GPS system the team currently possesses.

The vision system testing led to some interesting results. Intel advertises the RealSense camera to have an operating range of 0.3-3.0 meters projected from the camera, and measurements within these margins have proven to be accurate to ± 20mm. In practice, the team verified an operating range of up to five meters with ±30mm.

7. Driving Logic

The navigation system for Project Phoenix is tailored to the Autonav course. At its core, the architecture relies on the priorly mentioned grid for perception and awareness. This map is critical to the Reactive, Predictive, and Positioning system, as it is used to determine the safest location to target.

The highest priority in the driving logic is the Reactive system, referred to internally as "DANGER" mode. The algorithm scans the first three rows of the occupancy grid, from column 22 to 28, so that the system can detect the obstacles close to the chassis. If a hazard is detected, the PLC cancels the predictive navigation to avoid a collision and strobes the blue stack light indicator. Then, it reduces the drive scale and the acceleration, which globally slows down the robot speed and makes the turns more precise. Next, it issues a stop forward command and executes an in-place turn toward the side with the least number of obstacles. This fail-safe ensures the robot can autonomously extract itself from tight spaces or sudden obstructions before proceeding. The predictive path planning mode minimizes the usage of this system, which allows us to maintain higher speeds for longer periods.

If no hazards are of concern, the system transitions from reactive to predictive navigation. The PLC performs a row-by-row scan of the remaining grid to identify the closest left and right obstacles. If obstacles are not identified within the safe path ahead of the robot, it continues to the next row. If obstacles are detected within the safe path, the PLC calculates a central target coordinate within identified gaps and passes that off to a system that calculates the desired forward command and turn command. This ensures good entries into obstacles like chicanes and ensures the avoidance of obstacles such as single barrels or potholes scattered around the course.

If the robot enters the waypoint navigation section, referred to as “No Mans Land”, GPS mode activates. This system employs a modified Pure Pursuit logic to reach specific coordinates. By calculating the bearing and distance to a waypoint using the Haversine formula, the controller determines the necessary curvature to stay on course. This is used to determine desired forward and turn commands, and if there are no obstructions along this vector, the robot executes the maneuver.

Due to high levels of magnetic interference from underground powerlines and the motors on the chassis, we have elected not to use a magnetometer this year. Instead, we perform a Course-Over-Ground (COG) calculation utilizing prior locations relative to the current across distance and velocity. When the reported location is within half a meter of a given waypoint, an audible siren is triggered, and the target waypoint is updated.

The final stage of the pipeline converts the forward and turn commands from a range of [0.0, 1.0] to a range of [-300, 300], representing the raw motor speeds. Due to the utilization of differential steering, steering and speed control are statically linked. This allows us to greatly simplify our control algorithm.

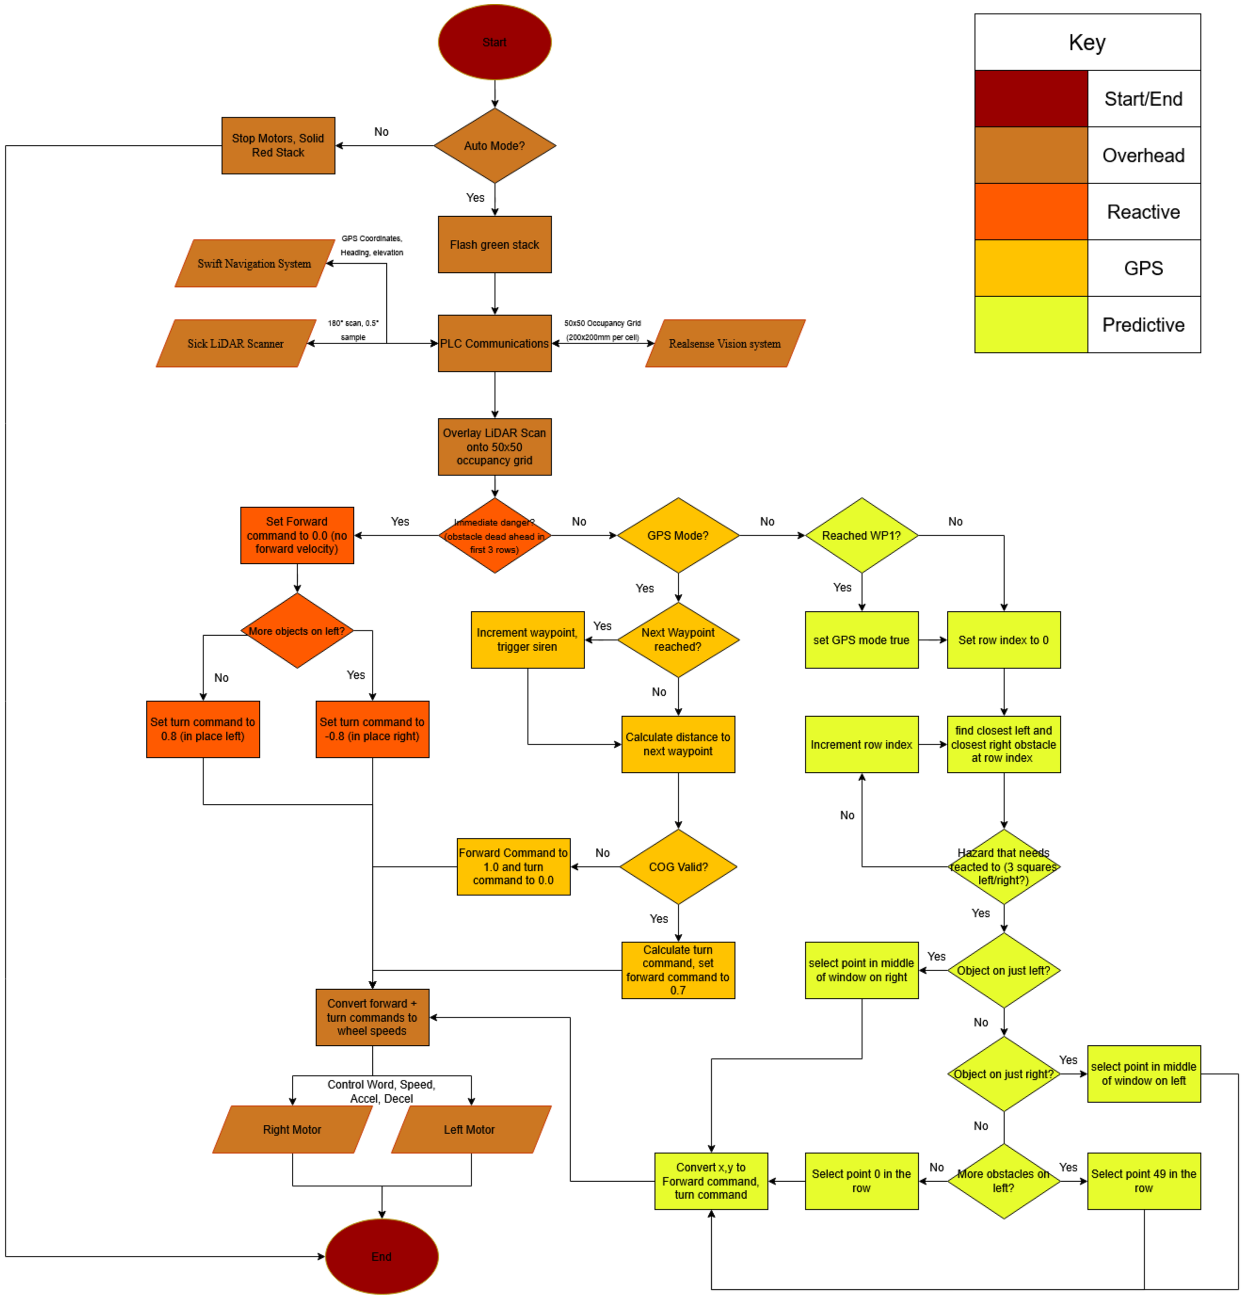

Figure 7: Flowchart of Project Phoenix's drive logic

8. Key Performance Indicators


Figure 8: An early development of the robot

This year's focus was on a new driving algorithm and a custom chassis designed to fulfill specific needs, which is reflected in our KPIs. The team's first main KPI was having a custom, ground up developed chassis. This included the initial design requirements for size constraints and a design that was built around the control system.

The second KPI included having all necessary components mounted for drive testing. This included the control cabinet, batteries, motors, wheels, and castors. While far from complete, achieving this state allowed the team to transition from simulated tests of code to functional real-world testing on at the Team's home course.


Figure 9: preliminary drive testing

The final milestone for Project Pheonix was integration of all control systems such that we get to a state where the robot can qualify for the AutoNav competition. This involves having a drive algorithm, integrating both Camera and LiDAR data into that algorithm, and GPS integration. Once this KPI is achieved, the project is almost complete, leaving only extended testing and precision tuning.

9. Analysis of Complete Vehicle

The completed vehicle of Project Phoenix represents strength in industrial reliability and efficiency. By utilizing an AXL F 2152 PLC, the robot can take advantage of discrete IO, industrial communications, and robust fail-safes.

9.1 Operational Performance

The vehicle achieves a measured top speed of 4.93 mph. Its differential drive configuration allows for zero-point turns, providing the agility necessary for tight chicanes. Power is supplied by two 24V 60Ah LiFePO4 batteries in series. Testing shows a massive power capacity, with only one full discharge recorded over weeks of daily operation. The robot is capable of running for the entire duration of the competition without recharging, maximizing our onsite testing time.

9.2 System Latency and Safety

The software architecture is designed for deterministic execution. Each PLC routine runs under 50ms, or a watchdog timer is triggered. This guarantees a total critical time of 150ms. At max speed, the robot processes and reacts to an obstacle within 350mm of travel by quickly braking and applying steering torque. To ensure waterproofing, the system has been field-validated in active rain conditions to confirm the efficacy of the IP65 weatherproofing.

10. Cyber Security Analysis

Many modern vehicles are increasingly vulnerable to cyber-attacks due to their increased dependency on software and communications. Each proprietary software adds another potential vulnerability. We mitigate this by utilizing industrial systems, each with their own unique security systems, passkeys, and fail-safes. Project Phoenix still has security vulnerabilities that would need to be addressed if the robot were to go into production.

All communications are done over an ethernet network that has a wireless access point utilizing WPA3 encryption. This is very hardened, however the switch behind it does not filter mac addresses. If a malicious actor acquired the unencrypted password to the network by social engineering, they could gain access. Further, TCP communication is not completed utilizing TLS encryption. Adding this feature would greatly harden communication between devices and protect against a wider range of attacks.

Another risk to the robot is GPS spoofing, which would affect any device utilizing the GNSS locality. To address this, the robot could ensure calculated speed never passes a certain threshold to detect anomalies in GPS systems. Additionally, the robot could utilize more constellations across broader spectrums in addition to cell towers to attempt to ignore spoofed signals.

The most significant risk to the robot is if a malicious attacker gains access to the physical robot, they could damage internal components. To mitigate this, the existing lock on the panel should be replaced with a proper keyed one, and an alarm should be fitted to go off if unauthorized access is attempted.

11. Acknowledgements

The Millersville University Mobile Robotics Research Team would like to formally thank its sponsors of Phoenix Contact, SEW Eurodrive, Sick Sensor Intelligence, and The ATMAE Board of Accreditation. Without these sponsors, Project Phoenix would not be possible, and the team would not have been able to travel to the competition.