Turbo Blue

May 15th, 2025

32nd Intelligent Ground Vehicle Competition

Lawrence Technological University

Team Captains

Angela Ladd aladd@ltu.edu

Andrew Turner aturner2@ltu.edu

Team Members

Constantine Gastouniotis cgastouni@ltu.edu

Andreu Lombarte alombarte@ltu.edu

Oliver Pineda opineda@ltu.edu

Faculty Advisor

Dr. Gaurav Singh

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

The Intelligent Ground Vehicle Competition (IGVC) represents one of the most rigorous and multidisciplinary challenges in autonomous robotics at the undergraduate level. It demands excellence not only in vehicle design, but also in systems integration, safety, and software development under strict real-world constraints. For the 2025 competition, the Blue Devil Robotics team from Lawrence Technological University proudly presents Turbo Blue, a fully custom-built autonomous ground vehicle designed to compete in both the Auto-Nav and Self-Drive competitions.

This dual-track objective significantly shaped our engineering approach. Turbo Blue was developed to meet the distinct technical demands of both challenges: precision lane following and GPS-based waypoint navigation for Auto-Nav, and dynamic perception tasks such as traffic sign recognition and line detection for Self-Drive. The robot's architecture was deliberately structured for modularity, real-time response, and fault tolerance—capable of adapting between mission profiles with minimal reconfiguration.

Turbo Blue features a rugged, low-profile aluminum chassis, custom-welded motor mounts, and a swing-arm suspension system tailored for uneven outdoor terrain. Electrically, the vehicle employs a centralized USB hub for sensor connection, fused power distribution boards, and both hard-wired and wireless E-stop systems. The compute stack runs on a 12th Gen Intel(R) Core(TM) i7-12700H, managing sensor fusion and motion planning through ROS2, with full support for behavior trees, EKF-based localization, and dynamic path generation.

Given early hardware limitations, our team adopted a simulation-first development strategy, utilizing a high-fidelity Gazebo environment to validate lane detection, GPS waypoint navigation, and obstacle avoidance long before physical integration began. This decision proved instrumental in accelerating the software implementation and minimizing the risk of running out of time.

Every aspect of this project adheres to industry-standard engineering practices. We applied ISO 2768 for mechanical tolerancing, IPC-A-610 for electronic assembly, IEEE 315 for schematic documentation, and ISO 12100 for safety risk assessments. Our development process included rigorous subsystem testing, integration trials, and compliance verification to ensure Turbo Blue is competition-ready.

This report outlines the complete technical development of Turbo Blue, including our design methodology, subsystem architectures, innovations, test results, safety measures, and lessons learned. We respectfully submit this final report to the IGVC judges as a comprehensive representation of our team's work, professionalism, and commitment to the spirit of the competition.

2. Team Organization

The Blue Devil Robotics team is composed entirely of senior students majoring in Robotics Engineering at Lawrence Technological University**.** This capstone team was assembled to address the complex integration challenges of autonomous vehicle design and to compete in both the Auto-Nav and Self-Drive categories at IGVC 2025. A breakdown of the cumulative work hours of each team member can be viewed below is Table 1.

Table 1: Organization of Blue Devil Team

Name	Department	Class Year	Role	Hours Contributed

Constantine Gastouniotis	Robotics Engineering	Senior	Electrical Lead	290
Angela Ladd	Robotics Engineering	Senior	Documentation Lead	230
Andreu Lombarte	Robotics Engineering	Senior	Mechanical Lead	220
Oliver Pineda	Robotics Engineering	Senior	Safety Officer	40
Andrew Turner	Robotics Engineering	Senior	Programming Lead	230

3. Design Assumptions and Process

The design process for Turbo Blue began with a comprehensive review of the IGVC rules, past competition courses, and system requirements for both Auto-Nav and Self-Drive events. Early in the semester, the team identified specific technical constraints and environmental challenges likely to be encountered during the competition. These included variable lighting conditions, inconsistent GPS availability, and the need to detect both static and dynamic obstacles across open outdoor terrain.

One of the core assumptions was that robust autonomy could not be achieved without early software validation, even in the absence of physical hardware. This led to the adoption of a simulation-first workflow, where the entire perception and navigation stack would be developed and tested in a virtual environment before deployment on Turbo Blue. The team assumed that key sensors—including the LiDAR, GPS, IMU, and camera—could be effectively emulated in simulation with sufficient fidelity to tune algorithms and troubleshoot integration challenges in advance. ROS2 was chosen for the software environment under the assumption that its improved timing accuracy and node lifecycle management would be beneficial in supporting reliable autonomous behavior across multiple mission modes. This approach proved valuable for parallelizing development tasks across electrical, mechanical, and software domains without requiring immediate hardware access.

From an architectural standpoint, the team assumed that the system would need to support frequent iteration and reconfiguration. This led to the early decision to pursue a fully modular electrical layout, using a centralized USB hub and breakout wiring system to simplify integration and troubleshooting. Mechanically, all sensor mounts were designed to be removable and repositionable to accommodate ongoing tuning and environmental adjustments. In parallel, the team prioritized implementing a suspension system after reviewing documentation and footage from previous LTU teams. It was evident that a lack of rigidity and vibration isolation had contributed to inconsistent sensor readings and mechanical wear. To address this, the chassis was designed with suspension components to better absorb terrain-induced shock, increase system stability, and protect sensitive electronics such as the LiDAR and vision system during prolonged outdoor use.

The design process itself followed a phased development model. The first phase involved research, benchmarking, and constraint identification. The second focused on CAD, electrical schematics, and simulation environment setup. The third phase, which aligned with the arrival of hardware, focused on physical integration and initial lab testing. The final phase, ongoing during this report's submission, centers on testing, calibration, and mission tuning. Each phase included formal design reviews and documentation checkpoints to ensure alignment between team members and subsystems.

By grounding the project in clearly defined design assumptions and a structured engineering process, the team was able to make measurable progress despite logistical constraints and shifting timelines. This deliberate approach resulted in a vehicle that reflects both technical rigor and practical adaptability.

4. System Architecture

Turbo Blue was developed around a modular and tightly integrated system architecture that combines mechanical durability, robust electrical design, and ROS2-based autonomous software. The architecture was designed from the ground up to support both IGVC competition categories: Auto-Nav and Self-Drive. Each subsystem was developed independently, but with a focus on seamless integration, allowing the team to iterate rapidly and debug efficiently during development and testing.

The mechanical structure consists of a custom-fabricated chassis made from 6061-T6 aluminum, reinforced with machined brackets and welded joints. The chassis was designed using SolidWorks and manufactured in-house, with tolerances and fits defined according to ISO 2768 and ASME Y14.5 standards. A swing-arm suspension was incorporated to provide vibration isolation and terrain adaptability, minimizing disturbance to sensors during navigation. Sensor mounts were 3D printed in PLA+ to allow for quick modifications and testing across multiple mounting configurations.

The electrical system was built to emphasize safety, modularity, and clean cable routing. Two 12V batteries in series are used to supply power to all necessary components of the robot. Rather than using an embedded processor, Turbo Blue’s compute stack is powered by a standard onboard laptop, which hosts all perception, localization, and navigation nodes within the ROS2 Humble environment. The laptop interfaces with sensors through a high-speed USB 3.0 hub mounted internally, streamlining integration and reducing the need for separate sensor drivers or external processing hardware. The electrical system includes an Atlas GPS module, a Waveshare 10 DOF IMU, a SICK 2D LiDAR for 2D obstacle detection, and a Lucid Vision Labs camera for visual lane tracking and environmental perception. An Arduino Mega handles encoder data and safety-critical tasks such as E-stop logic and LED signaling, isolated from the main computing environment to prevent software faults from impacting system shutdown.

Safety was a top design priority throughout the system’s architecture. The system includes fuse protection on all major power rails and isolated wiring for high-current components. A status LED on the rear of the vehicle indicates system state, solid when powered and flashing when in autonomous mode which is offering visual feedback for operators and safety officials.

Software integration is handled entirely within ROS2, using a layered node structure and launch files for modular control. Sensor data from the LiDAR, GPS, and camera feed into the perception pipeline, which constructs a real-time costmap representing obstacles and lane lines. Localization is achieved by fusing encoder, GPS, and IMU data through an Extended Kalman Filter using the robot_localization package. The global planner uses an A* algorithm to generate paths to navigation goals, while the local planner implements the Dynamic Window Approach (DWB) to compute safe velocity commands based on the robot’s position and the dynamic costmap. These velocity commands are passed through the control node and relayed to the motor drivers via the Arduino using the micro-ROS library. The core software structure of Turbo Blue is depicted in Figure 1, providing an overview of its system architecture from a high-level point of view.

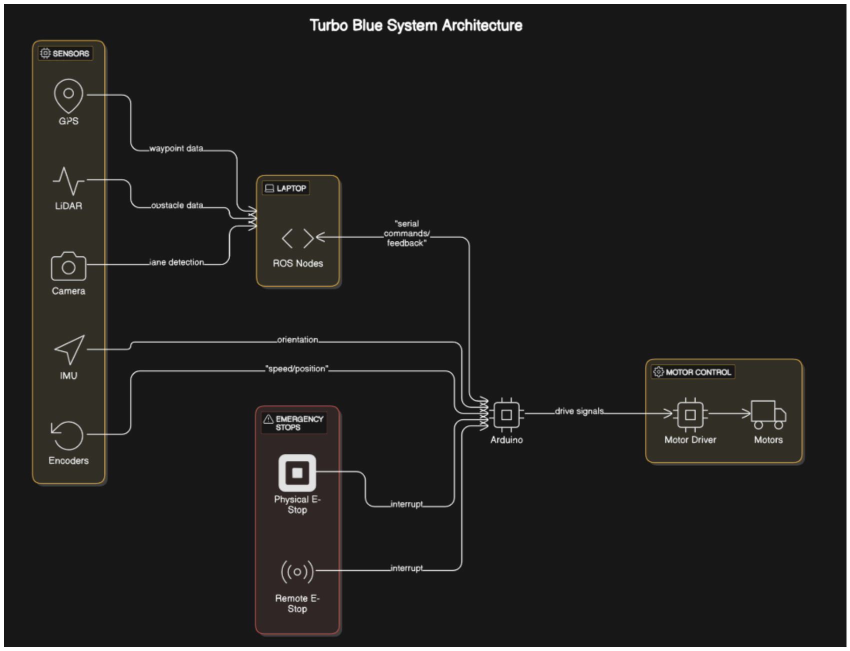

Figure 1. Turbo Blue System Architecture.

All components are coordinated through this centralized system architecture, ensuring that each subsystem can function independently while contributing to the overall autonomy pipeline. The result is a flexible and resilient platform, capable of adapting to IGVC's diverse challenges while maintaining the modularity required for future expansion and research.

5. Effective Innovations

Turbo Blue was engineered with simplicity, adaptability, and reliability at the forefront. Throughout the development process, the team focused not only on meeting IGVC requirements, but also on introducing practical, high-impact innovations that reduced system complexity and improved real-world performance. These innovations emerged from hands-on testing, system bottlenecks, and the lessons learned from managing a fully custom robotic platform with limited time and resources.

One of the most impactful innovations was the decision to directly connect all of the incorporated sensors to the onboard computer via native ports. This eliminates the need for an intermediate USB or Power-over-Ethernet hub. During early prototyping, the team experimented with centralized hubs for sensor integration, but quickly discovered that such configurations introduced added points of failure, USB enumeration conflicts, and unnecessary cable management issues. By shifting to a direct-connection strategy, the team dramatically simplified both the physical wiring and the software-side device mapping. This also made in-field troubleshooting much faster, allowing the team to easily swap components or reset connections without interrupting the system as a whole.

Another standout innovation was the use of a dual-boot setup on the onboard computer to support both hardware-focused development in Windows and ROS2-based autonomy development in Ubuntu. While virtual machines were initially considered, they proved unstable when handling camera and LiDAR drivers. Transitioning to dual-boot allowed the team to run ROS2 natively, ensuring reliable performance, low latency, and full compatibility with all required autonomy tools.

On the mechanical side, Turbo Blue's custom swing-arm suspension stands out as a key differentiator from most IGVC vehicles. Previous teams' driving tests revealed that uneven terrain caused excessive camera vibration and introduced noise into the LiDAR scan data. The design uses a single-axis pivot with dampened bushings, offering smoother ride dynamics while maintaining structural simplicity. This innovation significantly reduced motion artifacts in sensor data and made downstream filtering more effective. A close up of the CAD can be seen in Figure 2.

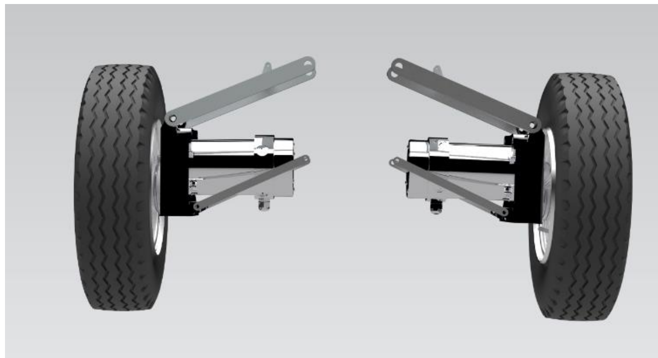

Figure 2. CAD of Suspension

Finally, the team introduced a real-world GPS playback testing method to continue development even when outdoor testing was not feasible. By capturing GPS data during prior test runs and feeding it back into the localization stack within Gazebo, the software team was able to tune the EKF, planner behaviors, and mission timing in a controlled and repeatable environment. This process enabled debugging of navigation errors and tuning of motion control parameters without requiring open-sky conditions or clear weather, giving the team an edge in terms of testing hours and system reliability. Each of these innovations was rooted in the practical needs of building a real-world autonomous system. Rather than over-engineering or relying on overly complex solutions, the team focused on minimalism, robustness, and clarity—leading to a vehicle that is easier to debug, maintain, and scale. These decisions reflect a strong understanding of field robotics constraints and a proactive approach to solving real integration challenges.

6. Mechanical Design

The mechanical design of Turbo Blue reflects a balance between structural integrity, modularity, and ease of maintenance. As a fully custom ground vehicle, Turbo Blue was designed from scratch to accommodate both Auto-Nav and Self-Drive challenge requirements while remaining compact, rugged, and accessible for field repair and modification.

The design process began with preliminary sketches and design constraints based on IGVC vehicle dimension and safety rules. From there, the mechanical team moved into a CAD phase using SolidWorks. Here, the full frame, mounting plates, sensor brackets, and drive system were iteratively modeled. The team followed ISO 2768 and ASME Y14.5 standards for tolerancing and dimensioning. Weekly internal reviews were held to evaluate trade-offs between weight, structural support, and mounting flexibility, especially for sensor placement and electronics access.

Once the CAD model was finalized, 2D drawings were generated for fabrication. Most structural components were manufactured in-house using band saws, belt sanders, drill presses, and welding tools. All metal components were constructed from 6061-T6 aluminum for its favorable strength-to-weight ratio and corrosion resistance. After machining, surfaces were deburred, sandblasted, and cleaned prior to assembly.

Turbo Blue’s chassis consists of a welded aluminum frame with integrated vertical and horizontal supports. A front-mounted sensor tower provides an unobstructed field of view for the LiDAR and camera systems. The drive system includes two NPC-T74 DC motors connected to 16-inch wheels via welded mounts and precision-aligned shaft collars. The motors were selected for their high torque and compatibility with IGVC power constraints. Sensor mounts, electronics trays, and motor guards were 3D printed using PLA+ and mechanically reinforced where necessary with aluminum backing.

Frame dimensions were chosen to optimize both stability and sensor visibility while remaining within the IGVC course width limits. The center of mass will be kept low utilizing the batteries and payload in the center of our base. Access to internal systems is maintained through removable weatherproofing panels, enabling fast servicing and cable management. Mounting holes were pre-drilled to allow repositioning of sensors and to accommodate future upgrades without needing to fabricate a new chassis.

A major mechanical feature of Turbo Blue is the custom swing-arm suspension, designed to isolate vibration from the motor mounts and wheels. The swing-arm uses a single-axis pivot and passive damping to absorb ground-level shocks, reducing motion distortion in the LiDAR and camera feeds. The decision to use a swing-arm was made after early drive tests revealed significant instability in image capture and localization due to mechanical jolt. The suspension system was iteratively tuned using simulated loading and refined during outdoor test runs to improve ride quality and sensor consistency.

Weather protection was implemented through a combination of enclosure design and component sealing. All external cable pass-throughs use waterproof bulkhead connectors, and exposed wiring is covered in braided sleeving with heat-shrink protection. While the vehicle is not fully submersible, it is capable of operating safely in light rain and on wet surfaces, as required by the IGVC competition environment.

7. Description of electronic and power design

Turbo Blue’s electrical and power systems were engineered with an emphasis on modularity, safety, and serviceability. Every component was selected to support reliable operation during extended field runs while remaining compliant with IGVC rules and safety protocols. The vehicle’s electronics are organized into clearly defined subsystems—compute, sensors, power delivery, motor control, and safety—each integrated through a custom-wired architecture that supports fast troubleshooting and robust performance.

Turbo Blue is powered by two 12V batteries connected in series. The Cryton Technologies SmartDriveDuo-60 motor driver is used to provide easy control of the motors while also purposing as a 5V supply for the Arduino subsystem. The whole electrical system has an estimated full-power run time of 2.5 hours. This provides plenty of time for testing and competition runs. Due to the modularity designed into Turbo Blue, replacing the batteries is a quick process allowing the team to preform testing on full power without losing heaps of time fidgeting with the power system. Recharge time for the battery system is approximately 4.5 hours per battery using a standard 12V charger rated at 7 amps.

Turbo Blue’s systems are centered around a standard onboard laptop, chosen for its compatibility with ROS2, ease of driver support, and ability to run simulation and autonomy stacks in real time. The LiDAR, camera, and GPS are connected directly to the computer via native USB 3.0 ports. This direct connection strategy was adopted to eliminate the need for an external USB or PoE hub, reducing wiring complexity, and improving reliability. The vehicle’s motor controllers are interfaced via an Arduino Mega, which also handles digital signals for the safety light and emergency relay logic. The complete schematic can be viewed in Figure 3.

Turbo Blue features both mechanical and wireless emergency stop (E-stop) systems, each capable of triggering a relay that instantly disables motor power. The mechanical E-stop is mounted on the rear of the vehicle for quick physical access, while the wireless system allows for remote shutdown from the pit area. Both systems are electrically isolated from the compute stack and cannot be overridden by software, ensuring that autonomous logic cannot prevent emergency intervention. A rear-facing LED indicator signals system status to judges and operators: solid when powered on, flashing in autonomous mode, and off when disarmed. These safety mechanisms were tested under both lab and field conditions to ensure consistent performance under realistic scenarios.

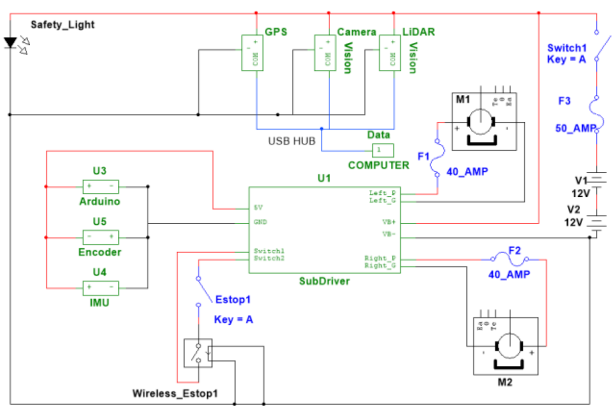

Figure 3. Electrical Diagram Schematic.

8. Description of Software

Turbo Blue’s autonomy software stack is built entirely in ROS2 Humble and designed to run in real-time on a standard onboard computer. The system is organized into independent nodes responsible for sensor processing, localization, mapping, planning, and control. All components were developed and tested using a simulation-first strategy, enabling high modularity and consistent behavior across virtual and real-world environments.

Turbo Blue perceives its environment through four primary sensor modalities: LiDAR, GPS, IMU, and a forward-facing RGB camera. The LiDAR provides a 2D scan of the vehicle’s surroundings, which is filtered and transformed into a local obstacle costmap using ROS2's costmap_2d and voxel_layer plugins. The Lucid Vision camera captures a real-time image stream processed using OpenCV-based filtering. Color thresholding, edge detection, and contour analysis are used to detect painted lane lines and other relevant high-contrast features, which are then projected into a top-down occupancy map for planning. Sensor data from the GPS and IMU are fed into the robot_localization package, which fuses inputs through an Extended Kalman Filter (EKF) to generate a stable, drift-reduced estimate of the robot's position and orientation. Encoder data from the motor system is also fused into this estimate to improve short-term accuracy during GPS dropouts.

While Turbo Blue does not use semantic object classification (e.g., for pedestrians or signage) in Auto-Nav mode, its perception stack is optimized for geometric classification. The LiDAR identifies obstacles based on range discontinuities and angular resolution, which are used to create a rolling costmap centered on the robot. Lane information from the camera is treated as a soft constraint in planning; it is used to infer the intended travel corridor, even when GPS goals are not aligned with painted lanes. All detected features are stored in local costmaps that update at ~10 Hz and are visualized in RViz for debugging. These costmaps combine current scene data with time-decayed history, effectively generating a short-term world model that reflects the robot's dynamic environment.

Lane following is performed by processing RGB camera frames using OpenCV. The image is filtered for IGVC-standard paint colors (white and yellow), followed by Canny edge detection and Hough line transforms to determine lane boundaries. These features are converted into a lane corridor and used to bias the local planner's costmap. Obstacle avoidance is handled through LiDAR-based occupancy grids, which inflate around known obstacles and enforce clearance margins during navigation. When a potential collision is detected, the planner reroutes dynamically within the defined corridor, either around the obstacle or to a fail-safe stop depending on the available options.

GPS waypoints are defined in latitude/longitude and converted to the local map frame using a navsat_transform_node. This allows the robot to localize globally while planning trajectories within a locally consistent frame. GPS accuracy is improved by the EKF using IMU data to interpolate orientation and acceleration. The result is a filtered pose that feeds into both global and local planning pipelines.

Turbo Blue builds a short-term world model by combining real-time sensor data with a rolling costmap window. The costmap retains obstacle information over time using exponential decay, allowing the robot to retain awareness of transient obstacles (e.g., cones or barrels) that may momentarily exit sensor range. Lane features and static structures are fused into this map, providing context for trajectory generation even under partial occlusion.

Turbo Blue supports two operating modes: Auto-Nav (GPS-based waypoint navigation) and Self-Drive (visual stop signs, traffic control, and lane logic). The operating mode is selected via a launch-time parameter or physical toggle switch. In Auto-Nav mode, the global planner receives a list of GPS waypoints and navigates through the course using LiDAR for obstacle avoidance. In Self-Drive mode, the robot switches to vision-priority behavior, including lane following, stop sign detection, and parking logic. Switching between modes triggers a reconfiguration of costmap layers and behavior tree parameters without restarting the system.

The global path is generated using an A***** planner, which creates an efficient path through known free space from the robot's current location to the next waypoint. This path is passed to the DWB Local Planner, which evaluates velocity command options based on costmap data, robot kinematics, and dynamic constraints. Commands are published as linear and angular velocities to the motor interface. Motion is monitored using encoder feedback, which is compared against expected path progress to trigger alerts in the event of wheel slip or actuator lag.

9. Cyber Security Analysis using RMF

To ensure Turbo Blue's operational security during IGVC, particularly in the pit area where wireless interference and physical access risks are highest, the team performed a cybersecurity analysis using the NIST Risk Management Framework (RMF). The RMF provides a structured process to identify, assess, and mitigate risks to critical systems. Our implementation of the RMF focused on scaled, practical controls tailored to a field-deployable robotics system. The steps included categorizing key assets such as the onboard compute system and its connected sensors, selecting relevant security controls from NIST SP 800-53, implementing those controls in software and hardware, assessing their effectiveness through in-lab testing, and monitoring system behavior through status indicators and logging tools.

In modeling potential threats, the team considered scenarios such as spoofing ROS2 messages over the network or connecting unauthorized devices to the robot's onboard Wi-Fi during pit setup. The impact of such threats could include node crashes, invalid motion commands, GPS or costmap corruption, or an unintentional E-stop override. These risks were deemed high-impact due to their ability to compromise both safety and competition scoring.

To counter these threats, the team implemented a layered set of cyber controls. USB ports on the onboard computer were configured as read-only through system-level udev rules, preventing execution of files from inserted devices. Network security was further enhanced by applying firewall rules that restrict ROS2 DDS communication to known IP ranges only, and by disabling broadcast discovery to prevent rogue topic injection. Additionally, the E-stop system was physically isolated from the ROS environment to ensure that software interference could not accidentally engage or disengage critical safety features.

All controls were validated through simulations during development. These included attempting BIOS access during boot, spoofing ROS2 topics, and monitoring system response to unexpected data packets. In each case, the implemented safeguards prevented damage or unauthorized access, and logs were generated to assist in identifying the source of the event. These cybersecurity measures ensure that Turbo Blue is resilient to both accidental and malicious disruptions and remains secure throughout the competition environment.

10. Analysis of Complete Vehicle

Throughout the development of Turbo Blue, the team encountered and overcame numerous engineering, integration, and logistical challenges. One of the most important lessons learned during construction was the value of modular design, not just in theory, but in practice. By designing each subsystem to be independently testable, swappable, and diagnosable, the team was able to make consistent progress even when specific components were delayed or malfunctioning. This approach also allowed parallel work across team members, which became critical as project deadlines approached. System integration taught the team the importance of proper documentation, interface standardization, and consistent test procedures, particularly when bridging mechanical, electrical, and software domains.

During field and bench testing, several hardware vulnerabilities were identified that could have critically impacted competition success. The most significant issues included intermittent USB connection failures, LiDAR dropout due to power fluctuations, and mechanical stress at sensor mounting points. These were mitigated by switching to direct USB connections with reinforced cabling and redesigning sensor mounts using reinforced PLA+ and mechanical braces. Power isolation between the motor and compute systems also prevented brownouts when the motors surged under load. Additionally, the team discovered that signal noise from motor actuation occasionally affected GPS readings, which was resolved by introducing improved grounding and shielding practices.

Safety, reliability, and durability were integral to every design decision. Electrical fusing, emergency stop redundancy, waterproofing, and protected cable routing were all implemented based on worst-case environmental scenarios. Physically, Turbo Blue was tested against impact, vibration, and weather conditions using lab simulations and controlled outdoor testing. When unexpected problems emerged—such as sensor desync under low lighting, or camera jitter on rough terrain—the team responded with iterative hardware and software changes, including implementing lane-detection confidence thresholds and introducing suspension dampening via the custom swing-arm design.

Key hardware failure points for Turbo Blue include encoder slippage, sensor desynchronization, and power instability. The EKF continuously compares encoder, GPS, and IMU data. If encoder slippage occurs, the system automatically reweights GPS and IMU inputs to maintain accurate localization without stopping. In cases of sensor dropout—such as GPS or LiDAR loss—the autonomy stack degrades gracefully, relying on remaining sensors and triggering a soft stop only if localization confidence falls too low. The onboard computer is powered through a regulated converter with brownout protection, and all critical subsystems are electrically isolated. Real-time monitoring nodes detect faults and trigger fallback behaviors autonomously, ensuring safe operation without manual intervention.

Turbo Blue's predicted performance metrics were evaluated using both simulation and real-world testing. Its acceleration and reaction time to dynamic changes in the costmap are within 300ms, making it responsive enough for obstacle-heavy environments. In testing the average speed is 4.8 MPH. While Self-Drive ramp climbing is not applicable, the vehicle maintains traction and balance across uneven terrain and can handle a ramp up to 26 degrees. GPS waypoint navigation achieves an average accuracy of 0.5 to 0.8 meters RMS, depending on sky visibility and IMU stability. Battery life across both 12V batteries averages 2.5 to 3 hours under autonomous load, with the compute system drawing a consistent base load while motor usage scales depending on course layout. Obstacles are detected up to 10 meters away by the LiDAR, with effective avoidance behavior starting at 4-6 meters. In simulations and field testing, Turbo Blue demonstrated consistent pathing through switchbacks and simple traps, with difficulty increasing in environments with high visual clutter. Lane detection remains robust on clear markings but degrades in direct sunlight or in shadowed transitions—an area the team continues to refine.

Software testing followed a structured Git-based version control workflow with milestones, branches, and merge approvals. Bugs were tracked using a shared issue board, and critical system behaviors were logged through custom ROS2 nodes that monitored node health, perception confidence, and planner decisions. Software-in-the-loop (SIL) testing was performed in Gazebo with simulated GPS, LiDAR, and camera inputs modeled on an IGVC-style course. These virtual runs were used to test new algorithms, evaluate planner tuning, and simulate GPS loss, LiDAR occlusion, and path obstructions. Compared to predictions, real-world testing revealed greater variability in camera and GPS reliability, as well as increased difficulty in achieving clean obstacle costmaps on tall grass and uneven terrain. However, core navigation, lane following, and emergency stop behaviors performed consistently in both environments.

Physically, Turbo Blue has undergone rigorous testing including bench-level subsystem validation, indoor odometry and localization trials, and outdoor autonomy runs on varied terrain. These real-world sessions were used to measure the robot's ability to transition between waypoints, navigate tight turns, and recover from unexpected obstacles. The system consistently demonstrated reliable autonomous behavior over multiple full-course simulations, validating both the design and the integration work completed by the team. Although small deviations between simulation and reality were noted—particularly in perception timing and localization drift—Turbo Blue remains stable, responsive, and fully operational under the competition's expected conditions.

11. Unique Software

Turbo Blue was designed to compete in both the Auto-Nav and Self-Drive portions of IGVC 2025, each requiring its own tailored sensor configurations, autonomy logic, and software behaviors. To support both modes within a single platform, the team developed modular software stacks and parameterized behavior tree logic, enabling the robot to switch between navigation priorities with minimal reconfiguration.

In Auto-Nav mode, Turbo Blue focuses on GPS waypoint navigation, obstacle avoidance, and basic lane tracking. The system leverages the GPS module combined with an onboard IMU, fused through an Extended Kalman Filter to provide reliable localization in real-world conditions. A global A* planner is used to generate routes from the robot's current position to predefined GPS waypoints, while a dynamic local planner (DWB) responds to changes in the costmap. Lane tracking in Auto-Nav is performed using vision-based color filtering and edge detection on the camera feed, identifying white and yellow markings typically found on IGVC courses. The LiDAR generates a 2D rolling obstacle map used to adjust path selection in real time. A reduced field of view and streamlined costmap configuration are used in this mode to optimize planner speed and ensure responsiveness around tight switchbacks or confined course corridors.

In Self-Drive mode, Turbo Blue operates with a greater emphasis on real-time visual interpretation and lane-dominant decision-making. While the GPS and obstacle planning stack remains active, priority is given to camera-based detection and contextual behavior. This information is passed through a custom perception node that assigns context tags to each visual structure, which in turn influence trajectory generation through a behavior tree-based logic controller. Motion planning is constrained to center-lane biasing, and stopping behavior is triggered by the presence of stop lines or dead ends. While Turbo Blue does not yet support semantic object detection, the lane-following system is robust across variable lighting and moderate occlusion due to its use of multiple filtering and edge detection strategies. A real-time lane confidence score is published for diagnostic use and is also used by the planner to reduce speed or increase caution in ambiguous visual environments.

Both modes share core infrastructure—including localization, motion control, and safety systems—but adjust costmap parameters, planner tolerances, and visual prioritization logic depending on the active mode. The operating mode is selected via a runtime parameter or manual toggle before launch. This dual-mode architecture allows Turbo Blue to adapt to different competition requirements without swapping hardware or restarting the core system, making it both versatile and field-efficient for back-to-back runs in Auto-Nav and Self-Drive categories.

12. Initial Performance

As of this report, Turbo Blue has undergone extensive testing in simulation, controlled indoor environments, and outdoor field conditions. The vehicle has demonstrated reliable core functionality across all major subsystems, including autonomous navigation, obstacle avoidance, and lane tracking. While full-scale field testing under competition-like conditions is still ongoing, early performance metrics indicate that the system is well-aligned with IGVC requirements and is on track for successful competition deployment.

In terms of navigation, Turbo Blue consistently reaches GPS-defined waypoints with an average accuracy of 0.5 to 0.8 meters RMS, depending on sky visibility and surrounding structures. The localization system, which fuses data from GPS, IMU, and wheel encoders through an Extended Kalman Filter, has proven stable across both open-field and urban-simulated environments. Obstacle detection using the LiDAR is reliable up to 10 meters, with real-time avoidance behavior activating around the 4 to 6 meter range. The global and local planners maintain smooth trajectory execution and recover effectively from temporary path obstructions.

Lane detection performance has improved significantly over multiple iterations of the vision pipeline. Under ideal lighting conditions, Turbo Blue's camera-based detection can identify and track standard white and yellow lane markings with sub-15 cm deviation. Current RMS error while tracking a lane is approximately ±11.8 cm, with system robustness maintained through adaptive color thresholds and edge filtering. However, performance can degrade in direct sunlight or under heavy shadow, particularly when lane paint is worn or covered by debris—this remains an area of active tuning.

Simulation-based testing using Gazebo has provided valuable insight into the system's behavior under controlled failure conditions. In these tests, the robot successfully rerouted around artificial obstacles, recovered from simulated GPS dropouts, and executed waypoint-to-waypoint navigation without manual intervention. Scene replay from previous GPS logs also enabled stress testing of the planning and localization stack. These virtual outcomes have been validated through selective real-world trials, which confirm close alignment between simulated predictions and physical performance.

While Turbo Blue is not without remaining challenges—particularly in complex visual environments and under extreme lighting shifts—it has demonstrated consistent and predictable behavior in all core areas. Continued refinement of the perception pipeline, expansion of Self-Drive behavior logic, and full-course outdoor testing will be the focus of the final development phase leading up to competition. Based on current performance, the vehicle is expected to complete the Auto-Nav course successfully and demonstrate meaningful capabilities in the Self-Drive competition, with real-time reliability and safety protocols already well-established.