California State Polytechnic University, Pomona
Autonomous Vehicle Laboratory
IGVC 2026 Design Report

MARVIN

Challenge Targeted: AutoNav

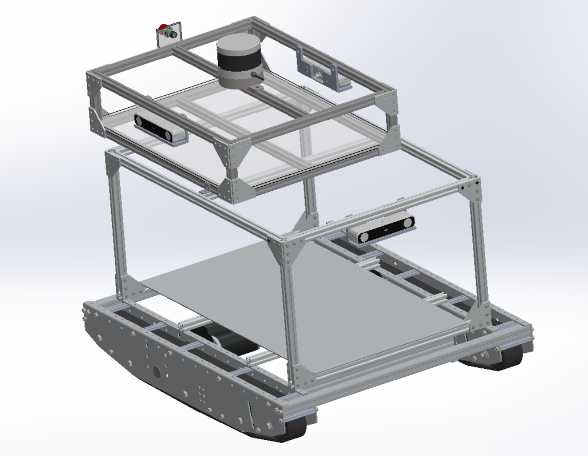

I hereby certify that the development of the vehicle, *MARVIN*, as described in this report, is equivalent to the work involved in a senior design course.

Behnam Bahr

Faculty Advisor Signature   Submitted May 15, 2026

Role	**Name / E-mail / Phone**
Team Captain (primary)	Parsa Ghasemi — `mghasemi@cpp.edu` — 424-844-8222
Key Team Member (secondary)	Christopher Sierra-Rodriguez — `rodriguez10@cpp.edu` — 323-239-9697
Faculty Advisor (tertiary)	Dr. Behnam Bahr — `bbahr@cpp.edu` — 949-690-5519

Scan of Faculty Advisor's Signed Statement of Integrity attached as a separate PDF per IGVC 2026 rule IV.3.

0. Contents

1. System and Subsystem Requirements

1.1 System Engineering Process

The Cal Poly Pomona Autonomous Vehicle Lab (AVL) at California State Polytechnic University, Pomona presents MARVIN, our entry for the IGVC 2026 AutoNav challenge. AVL is an undergraduate research lab within the College of Engineering focused on autonomy, embedded systems, and ground-vehicle perception. The IGVC subteam of six students (Table 1) is organized across software, mechanical, and electrical disciplines, advised by Dr. Behnam Bahr, and led by Parsa Ghasemi as project captain. MARVIN is a fully autonomous differential-drive tracked vehicle built on the AndyMark Raptor chassis, fusing a Velodyne VLP-16 3D LiDAR, three Stereolabs ZED X stereo cameras, and an Xsens MTi-680G IMU+GNSS, with autonomy running on an NVIDIA Jetson Orin under ROS 2 Humble. The team followed an iterative *research* → *design* → *prototype* → *test* cycle: rulebook analysis → system-level requirements; parallel subsystem development; incremental bench/actuator/fusion/Nav2 integration; and iterative field testing at the Cal Poly Pomona Engineering Meadow to tune HSV thresholds, EKF covariances, costmap parameters, and MPPI critic weights.

Table 1: Project contributors.

Parsa Ghasemi (Captain)	Caleb Hylkema
Karoline Braga	Benjamin Ramirez
Christopher Sierra-Rodriguez	Tyrese Rogers

Figure 1 shows the resulting system architecture.

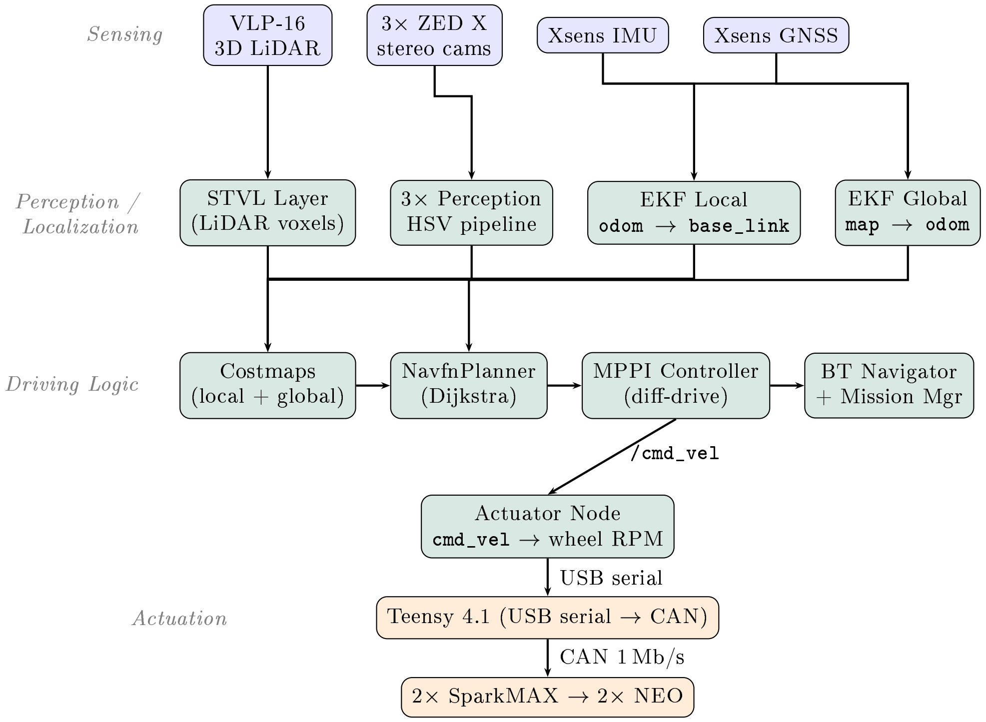

Figure 1: MARVIN system architecture. Sensors feed perception and localization; Nav2 composes costmaps, plans (NavfnPlanner), and follows the plan (MPPI Controller). The actuator node converts `/cmd_vel` to per-wheel RPM and forwards it to the Teensy over USB serial; the Teensy bridges to two SparkMAX motor controllers over CAN.

1.2 Requirements Summary

Per IGVC 2026 rule IV.3 §3, the requirements below list at least 2 each for mechanical, safety, electrical/electronic, perception, driving logic, and key performance indicators. We are targeting the AutoNav challenge, so 12 requirements are listed. Each row cites the rule (or internal source) driving it, the measurement method, and the target value.

Table 2: System and subsystem requirements (AutoNav target, 12 rows).

Category	**Requirement**	**Source / Measurement**	**Target**
Mechanical	Vehicle length within 3–7 ft	IGVC I.2 / tape measure	36″ (3 ft)
Mechanical	E-stop button mounted 2–4 ft above ground, centered rear	IGVC I.2 / tape measure	24″
Safety	Wireless E-stop range (margin over 100 ft minimum)	IGVC I.2 & I.4 / open-field radio range test	$\geq$ 150 ft
Safety	Hardware-governed maximum ground speed	IGVC I.2 / GPS + treadmill	$\leq$ 5 mph
Electrical	Battery runtime under nominal autonomous load	IGVC IV.3 §6 / dyno burndown	$\geq$ 6 min (1 margin run)
Electrical	Per-rail fault isolation	Internal / short one rail, verify others	All other rails active
Perception	Obstacle detection range (local costmap)	IGVC II.2 / VLP-16 ranging on barrel	$\geq$ 2 m
Perception	HSV lane-detection pipeline frame rate	Internal / publish-rate log on `perception/lane_mask`	$\geq$ 10 Hz (real-time)
Driving logic	Obstacle reaction time	Internal / log latency STVL update → `cmd_vel`	$\leq$ 250 ms
Driving logic	Waypoint arrival accuracy (2 m IGVC tolerance)	IGVC I.4 / GPS log distance to waypoint	$\leq$ 1 m
KPI	Course completion time (6 min IGVC ceiling)	IGVC II.2 / wall clock	$\leq$ 5 min
KPI	Ramp ascension success ($\leq$ 15% gradient per IGVC II.2)	IGVC II.2 / pass-fail	Pass

Each “Requirements Compliance” subsection at the end of §2–§6 compares these target values against the most recently measured values.

2. Mechanical Design

2.1 Overview

As first-year IGVC participants, the mechanical subteam prioritized a simple, modular, and cost-effective platform that concentrates budget on the sensors and compute required for autonomous navigation. MARVIN is built on the commercially available AndyMark® Raptor Track Drive, a tracked differential base selected for its all-terrain traction and short delivery lead time. A reconfigurable 80/20 T-slot aluminum superstructure is mounted directly to the chassis, allowing rapid iteration on sensor mounts and shelf geometry through the design cycle. The same chassis can host different electronics housings for future competitions; for IGVC 2026 the housing is mounted on the upper shelf so that the electronics, payload, and E-stop are all readily accessible during operation and inspection (see cover page for an integrated vehicle render).

2.2 Drivetrain

MARVIN uses a tracked **differential-drive** architecture: two REV NEO brushless motors, each commanded by a REV SparkMAX controller over CAN, independently drive the left and right track loops. The choice eliminates steering linkages, reduces the part count, and produces a control plant that maps cleanly onto the ROS 2 Nav2 / MPPI controller running on the Jetson. With v_L and v_R the left and right track linear velocities, the body-frame linear and angular rates are

$$v = \tfrac{1}{2}(v_R + v_L), \qquad \omega = \tfrac{1}{b}(v_R - v_L), \tag{1}$$

where b is the track width (effective lever arm between contact patches). The planner uses three operating modes: pivot turn (one track held, robot rotates about the stationary contact patch); in-place spin ($v_R = -v_L$, ICR at the vehicle center, maximum yaw rate); and gradual curve (both tracks forward, $v_R \neq v_L$).

The full electromechanical signal path (also shown in Figure 1) is: Jetson Orin issues `geometry_msgs/Twist` at 50 Hz over USB serial to the Teensy 4.1; the Teensy decomposes commands into per-side wheel velocities and forwards them over CAN at 1 Mbit/s to the two SparkMAX controllers; each SparkMAX closes a 1 kHz velocity loop around its NEO motor using the integrated Hall encoders.

2.3 Frame, Housing, and Sensor Mast

The Raptor base contributes the drive sprockets, idlers, and rubber-track belts that distribute vehicle weight over a large contact patch, improving traction and stability on the uneven turf and gravel sections of the IGVC course. A 21.5″ × 30″, 1/8″ thick A36 diamond-plate steel base is bolted directly to the Raptor mounting rails. The diamond plate serves three functions: (i) a rigid surface for heavy components, (ii) an EMI-attenuating ground plane between battery and tracked drive, and (iii) a low-mounted mass that lowers the center of gravity. The superstructure is built from 80/20® 20-Series T-slot aluminum extrusion — modular, reconfigurable, and free of exposed threaded ends.

Weight placement was an explicit design priority. The 48 V battery is mounted at the geometric center of the diamond plate; the 20 lb payload bay sits at the front of the deck on a 3D-printed cradle. The resulting center of gravity sits slightly forward of the track midpoint and low on the deck, minimizing pitch moment on inclines and during skid-steer turns. The diamond-plate deck is 21.5″ × 30″ and the upper sensor shelf is 23″ × 23″ × 6″ mounted 20 in above ground; the rear shelf carries the E-stop button at 24″. The chassis was extended with that rear shelf to satisfy both the IGVC 3 ft minimum length and the 2–4 ft E-stop mounting requirement.

Table 3: MARVIN vehicle specifications.

Parameter	**Value**
Drive configuration	Tracked differential (skid-steer)
Base platform	AndyMark Raptor Track Drive
Overall length × width	∼ 36″ × 24″ (meets 3 ft min.)
Overall vehicle height (excl. E-stop antenna)	∼ 29″ (∼ 2.4 ft; well under 6 ft IGVC max)
Sensor-shelf height	20″ (top surface)
E-stop mounting height	24″ (rear shelf, centered)
Payload capacity (verified)	$\geq$ 20 lb
Drive motors	2× REV NEO brushless (12–48 V)
Motor controllers	2× REV SparkMAX (CAN, 1 Mbit/s)
Battery pack	Aegis 48 V 25 Ah $LiFePO_4$, center-mounted
Frame material	80/20 20-Series T-slot aluminum extrusion
Deck material	1/8” A36 diamond-plate steel
Weather panels	Polycarbonate / acrylic in T-slot grooves

2.4 Weatherproofing

Outdoor operation requires the electronics bay to tolerate rain and direct sun. The sensor-shelf walls are cut-to-fit acrylic panels that slide directly into the 20-Series T-slot grooves; weatherstripping gasket fills the residual gap and remains seated under track vibration. The top panel is mildly angled for water runoff and incorporates a sliding access door. The LiDAR is mounted externally above the shelf to keep its 360° field of view unobstructed; the ZED X cameras sit behind AR-coated windows.

2.5 Mechanical Requirements Compliance

Table 4: Mechanical requirements: target vs. measured.

Requirement	**Target**	**Measured**
Vehicle length (3–7 ft IGVC range)	$36''$	$\sim 36''$ (verified)
Vehicle width (2–4 ft IGVC range)	$24''$	$\sim 24''$ (verified)
Vehicle height ($\leq$ 6 ft IGVC excl. antenna)	$\sim 29''$	$\sim 29''$ (verified)
E-stop button mounting height (2–4 ft IGVC band)	$24''$	$24''$ (verified)
Payload capacity (20 lb cinder block per IGVC I.3)	$\geq$ 20 lb	$\geq$ 20 lb (verified)
Track width b for diff-drive kinematics	0.7366 m	0.7366 m (verified)

3. Safety

3.1 Transport, Parked, and Charging Safety

Transport. MARVIN is moved between the lab and the Cal Poly Pomona Engineering Meadow under manual control via the AVROS web UI joystick over Wi-Fi; the software E-stop in the actuator node is armed throughout transit, and autonomous mode is never enabled outside the test area. For the long-haul trip to the IGVC competition venue at Oakland University, MARVIN is loaded into a cargo van and secured with ratchet straps anchored to the 80/20 frame and diamond-plate deck so it cannot shift, tip, or roll during the drive.

Parked. The mechanical E-stop is latched (pushed in) by default whenever MARVIN is not actively running, both between practice runs and during long-term storage; it is only released when the team is preparing to test movement. The main power switch is set OFF when parked overnight. The acrylic panels and weatherstripping protect the electronics bay from rain and dust during outdoor parking.

Charging. The 48 V 25 Ah Aegis $LiFePO_4$ pack remains installed in MARVIN during charging; the DC House 48 V 18 A charger connects to a dedicated, keyed charge port wired into the vehicle's power system, so reverse-polarity and wrong-connector incidents are mechanically impossible. The charger is rated IP65 with built-in overheat, overcurrent, overvoltage, reverse-polarity, and short-circuit protections, and $LiFePO_4$ chemistry is intrinsically thermally stable and non-volatile compared to NMC lithium — eliminating the runaway-fire risk associated with most lithium-ion packs. All charging is performed in the lab (not at the competition venue, where the team operates exclusively on pre-charged batteries) with a fire extinguisher within reach; because of the chemistry and charger protections, charges are permitted to run unattended.

3.2 On-Course Safety

The on-course safety architecture implements a hardware-enforced two-channel emergency stop: either the rear-mounted mechanical E-stop button or the wireless remote E-stop receiver opens a relay in series with the motor-controller enable line. When the relay opens, the SparkMAX controllers immediately disable PWM output and brake the NEO motors via the integrated short-brake mode; software is not in the loop, so a Jetson hang or ROS crash cannot defeat the stop. Indicator LEDs visible from outside the chassis show the armed/disarmed state for the IGVC safety inspection.

In software, the actuator node additionally zeroes `cmd_vel` on stale input (300 ms watchdog) so a Jetson stall or ROS-node crash defaults the vehicle to a stop even before the hardware relay reacts. A firmware speed clamp at ±4600 RPM (1.53 m/s ground speed) further limits the worst-case motion even if upstream commands are malformed.

3.3 Safety Requirements Compliance

Table 5: Safety requirements: target vs. measured.

Requirement	**Target**	**Measured**
Wireless E-stop range (IGVC 100 ft minimum)	$\geq$ 150 ft	>100 ft (verified)
Hardware max ground speed cap	$\leq$ 5 mph	3.4 mph (1.53 m/s)
E-stop button mounting height	2–4 ft	24″ (verified)
Software serial watchdog timeout	$\leq$ 500 ms	300 ms
SparkMAX CAN heartbeat timeout (per controller)	$\leq$ 200 ms	100 ms
Firmware RPM clamp (→ ground-speed cap)	$\leq$ 4600 RPM ($\leq$ 5 mph)	±4600 RPM (1.53 m/s)

4. Electrical and Electronic Design

4.1 Overview

The 2026 electrical architecture is organised into four subsystems: power distribution, network and computation, sensing, and motor control — shown at the system level in Figure 1 with the actuation chain detailed in the “Actuation” row.

- **Power** — a 48 V battery feeds two buck converters (12 V and 5 V rails) that supply the rest of the vehicle (§4.2).
- **Compute** — a Jetson Orin GPU handles perception and high-level control. It commands a Teensy 4.1 microcontroller over USB serial (§4.5).
- **Sensing** — three ZED X stereo cameras, a Velodyne LiDAR, and an Xsens GPS feed the Jetson (§4.4).
- **Motor control** — the Teensy runs a closed-loop controller and commands two SPARK MAX motor controllers over CAN, which in turn drive two NEO brushless motors (§4.3).

4.2 Power Distribution

The vehicle is powered by an Aegis 48 V 25 Ah $LiFePO_4$ deep-cycle lithium battery (1.2 kWh nominal). $LiFePO_4$ chemistry was selected for its thermal stability, ∼2000-cycle service life, and tolerance to partial-state-of-charge operation between practice runs. The battery passes through a main switch into two buck converters: a 12 V / 20 A converter for the motor controllers and other high-power loads, and a 5 V / 6 A converter for the Teensy 4.1 and low-power peripherals. Each rail is split into three independently fused output branches (battery → switch → buck → 3 fused branches per rail) so a short on one load does not pull down the rest of the vehicle.

Run time and recharge. The 25 Ah pack at 48 V stores ∼1200 Wh. Estimated autonomous-mode load is ∼155 W (Jetson 30 W, VLP-16 8 W, 3× ZED X 9 W, Xsens 1 W, router 5 W, drive motors ∼100 W avg.) = ∼3.2 A at the 48 V bus, giving a calculated runtime of ∼6.5 hours with $\eta_{\text{tot}} \approx 0.85$ — two orders of magnitude above the 6-min IGVC II.2 course ceiling. Recharge with the DC House 48 V 18 A $LiFePO_4$ charger (S-US-L131211800092-1, 58.4 V CV, IP65) takes ∼1.5–2 hours from empty with absorption and cell balancing.

4.3 Motor Control

The drivetrain uses two NEO brushless motors, each driven by a SPARK MAX motor controller. Both controllers share a single CAN bus to the Teensy 4.1, which acts as the real-time control node. Each SPARK MAX reports position and velocity from its integrated encoder back over the same bus. The Teensy closes a PID velocity loop per wheel, using encoder feedback from the SPARK MAX controllers and reference velocities published by the Jetson over USB serial. Because the motor controllers, the encoders, and the microcontroller all live on one CAN bus, no separate signal wiring between the motor controllers and the Teensy is required.

Firmware. The Teensy 4.1 translates ASCII RPM commands into SparkMAX CAN velocity setpoints at 1 Mbit/s and echoes encoder feedback at 50 Hz. Three independent timeouts protect the drivetrain: a 300 ms serial watchdog on the Teensy, a 100 ms CAN heartbeat timeout per SparkMAX, and a firmware clamp of ±4600 RPM (1.53 m/s ground speed, below the 5 mph IGVC cap).

4.4 Sensors

Stereo cameras. Three ZED X cameras (front-left, front, front-right) connect to a ZED Link Quad over GMSL2; the ZED Link Quad attaches to the Jetson over PCIe, keeping the cameras off the Ethernet network.

LiDAR. A Velodyne VLP-16 provides 360° 3D point clouds at 10 Hz. It reports to the Jetson over Ethernet through the on-board router.

GPS / IMU. An Xsens MTi-680G provides 9-axis inertial data at 100 Hz and a standalone GNSS fix at 5 Hz (RTK disabled per IGVC rules). It connects directly to the Jetson over USB.

4.5 Network and Computation

The Jetson Orin is the primary computation unit; it runs perception, sensor fusion, and trajectory generation, and publishes wheel velocity commands to the Teensy. An on-board router carries

Ethernet between the Jetson and the Velodyne LiDAR. The Teensy 4.1 connects to the Jetson via USB serial (115200 baud) rather than Ethernet. The high-bandwidth camera link is kept separate: the ZED Link Quad attaches directly to the Jetson over PCIe so that the three ZED X feeds do not contend for the Ethernet network. The Xsens GPS connects directly to the Jetson. The same router also exposes a Wi-Fi link to a remote laptop for monitoring, data capture, and remote shutdown during testing.

4.6 Bill of Materials

Table 6: Bill of significant components (estimated team cost).

Qty	Component	Subsystem	Cost
1	NVIDIA Jetson AGX Orin 64 GB Developer Kit (JetPack 6, ROS 2 Humble)	Compute	$1,999.00
2	REV NEO Brushless Motor v1.1 (REV-21-1650) @ $50.00 ea	Drivetrain	$100.00
2	REV SPARK MAX Motor Controller (REV-11-2158) @ $100.00 ea	Drivetrain	$200.00
1	PJRC Teensy 4.1 (SparkFun)	Firmware bridge	$31.50
1	Waveshare SN65HVD230 CAN transceiver breakout	Firmware bridge	$8.00
1	Velodyne VLP-16 Puck 3D LiDAR (discontinued by Ouster 2025; last announced retail)	Perception	$4,000.00
3	Stereolabs ZED X stereo camera (GMSL2) @ $599.00 ea	Perception	$1,797.00
1	ZED Link Quad GMSL2 capture card	Perception	$489.00
1	Xsens MTi-680G RTK GNSS/INS (DigiKey)	Localization	$3,612.50
1	AndyMark Raptor Track Drive (full chassis)	Mechanical	$423.00
1	Aegis 48 V 25 Ah $LiFePO_4$ deep-cycle battery	Power	$464.78
1	DC House 48 V 18 A $LiFePO_4$ charger (58.4 V CV, IP65)	Power	$129.99
1	Safety light, wireless E-stop, push-button, wiring, misc.	Safety / Misc.	$500.00
		Total estimated market cost	$13,754.77

4.7 Electrical Requirements Compliance

Table 7: Electrical/electronic requirements: target vs. measured.

Requirement	Target	Measured
Battery runtime (1 course run @ 6 min + margin)	$\geq$ 6 min	∼6.5 h calc.; field burndown pending
Battery pack voltage	48 V nominal	48 V (verified)
12 V rail capacity (motor controllers)	$\geq$ 20 A	20 A buck (verified)
5 V rail capacity (Teensy + low-power)	$\geq$ 6 A	6 A buck (verified)
Per-rail fault isolation	All other rails active	Verified by design (3 fuses per rail)
Full recharge time (18 A bulk, 25 Ah pack)	$\leq$ 2 h	∼1.5–2 h calc.; field stopwatch pending
CAN bus rate	$\geq$ 500 kbit/s	1 Mbit/s (verified)

5. Perception

The perception stack splits responsibilities between LiDAR (physical obstacles) and cameras (painted lane lines and color-coded objects), with both feeding a unified costmap consumed by the driving-logic layer. We are targeting the AutoNav challenge, so this section discusses lane and obstacle perception specific to the AutoNav course.

5.1 LiDAR Perception (Velodyne VLP-16)

The VLP-16 publishes a 360° point cloud at 10 Hz into Nav2's `spatio_temporal_voxel_layer` (STVL), which maintains a 0.1 m voxel grid with 5 s linear voxel decay. A minimum obstacle height of 0.2 m filters ground hits and low grass while keeping short barrels and pedestrians. Detection range is 1–15 m local and 1–50 m global. STVL replaces the standard `VoxelLayer` because the VLP-16's 16-beam ±15° vertical FOV leaves angular gaps that defeat ray-trace clearing when the robot is stationary; STVL clears via time-based voxel decay plus a frustum-accelerated decay term that handles the sparse-scan case cleanly.

5.2 Vision Perception (3× ZED X)

Each ZED X publishes rectified HD1080 color and a COMPACT-resolution organized point cloud at 10 Hz. The `avros_perception` package runs one `perception_node` per camera; each subscribes to its camera's RGB and cloud topics through an `ApproximateTimeSynchronizer` (20 ms slop) and runs a pluggable pipeline. The production pipeline (`hsv`) performs three-class HSV thresholding (white lane, orange barrel, pothole) with an adaptive value-channel floor ($V_{\text{floor}} = \mu_V + k\sigma_V$, recomputed every N frames) and a polygonal sky-rejection ROI. Class IDs: 1 = lane, 2 = barrel, 3 = pothole. All thresholds are runtime-tunable via `ros2 param set` for field calibration.

Every pipeline outputs the four-topic contract consumed by the `kiwicampus/semantic_segmentation_layer` a mono8 class-ID mask, a mono8 confidence plane, the organized point cloud (relayed), and a latched `vision_msgs/LabelInfo`.

5.3 Internal Course Representation (Costmap)

Perception outputs are fused into two Nav2 costmaps: a 50×50 m rolling local costmap at 0.2 m resolution (`stvl_layer` + 3× `semantic_layer` front/left/right + `inflation_layer`) and a 100×100 m rolling global costmap (`obstacle_layer` + `inflation_layer`). The robot footprint is rectangular 1.00×0.74 m with $r_{\text{inscr}} = 0.5$ m and an inflation radius of 0.65 m. STVL voxel decay (5 s) keeps stale cells from blocking the planner after the robot moves; semantic raytrace clearing handles cells that have been re-observed as free.

5.4 AutoNav Course-Specific Behavior

The AutoNav course (IGVC II.2) features painted white lane lines $\sim 3''$ wide on asphalt, randomly placed construction barrels (white/orange/brown/green/black) at ≥ 5 ft clearance from the lane, an $\leq 15\%$ gradient ramp, simulated potholes (2 ft white circles), and GPS waypoint pairs at course entry/exit and ramp entry. The HSV pipeline tagged with class 1 isolates white lane paint; barrels register under class 2; potholes register under class 3; the VLP-16 STVL marks any other physical obstacle. A startup auto-calibration routine samples the ground at run start and adapts the V-channel threshold to current lighting before mobility is enabled.

5.5 Perception Requirements Compliance

Table 8: Perception requirements: target vs. measured.

Requirement	**Target**	**Measured**
Obstacle detection range (local costmap)	≥ 2 m	15 m (VLP-16 STVL range)
Obstacle detection range (global costmap)	≥ 2 m	50 m
HSV lane-detection pipeline frame rate	≥ 10 Hz	10 Hz (locked to ZED X publish rate)

6. Driving Logic

The driving-logic stack converts the perception costmap and current pose into wheel commands. Nav2 provides path planning, trajectory generation, and behavior orchestration (§6.1); `robot_localization` maintains the pose (§6.2); and the actuator node converts `cmd_vel` to wheel RPM (§6.3). AutoNav-specific course handling — lane following, obstacle avoidance, ramp, and waypoint navigation — is described in §6.4.

6.1 Planning and Path Generation

Table 9: Nav2 configuration (`nav2_params_humble.yaml`).

Component	**Configuration**
Planner	`NavfnPlanner` (Dijkstra) — chosen over SmacPlannerHybrid because the IGVC lane spacing (2–3 m) is narrower than Hybrid's 2.31 m minimum turning radius
Controller	`MPPIController` — DiffDrive motion model, 2000 rollouts, 56 timesteps, model Δt 50 ms, 8 active critics
Local costmap	50×50 m rolling, 0.2 m resolution
Global costmap	100×100 m rolling, 0.2 m resolution
Footprint	Rectangular 1.00×0.74 m, $r_{\text{inscr}} = 0.5$ m
Inflation	$r_{\text{infl}} = 0.65$ m, $k_{\text{scale}} = 3.0$
Goal tolerance	xy 2.0 m (IGVC waypoint acceptance), yaw 0.5 rad
Controller frequency	20 Hz; planner replans at 1 Hz via BT `RateController`

MPPI samples 2,000 rollout trajectories per 20 Hz control cycle, scores them against eight plugin-based critics (constraint, cost, goal, goal-angle, path-align, path-follow, path-angle, prefer-forward), and emits the noise-weighted optimal command. Unlike pursuit controllers that only *follow* a path, MPPI *generates* trajectories that deviate from the global plan to avoid newly observed obstacles — critical for IGVC's randomly-placed barrels.

Behavior tree. The active BT (`navigate_igvc_autonav_humble.xml`) wraps the standard plan-follow pipeline with several IGVC-specific guards. An outer `Timeout` of 45 s margins the

60 s hold-up-traffic disqualification rule. A `RecoveryNode` retries the pipeline up to three times; recovery actions form a `RoundRobin` escalating clear-costmap → wait → short backup → crawl forward. Spin-in-place is intentionally omitted: the 1.00 × 0.74 m footprint sweeps a 2.48 m diameter circle when off-center, exceeding the half-lane width.

Waypoint orchestration. The `mission_manager` node loads competition waypoints from `waypoints.yaml` as a lat/lon list, converts each to map frame via `/fromLL`, and fires `NavigateToPose` action goals sequentially. The cursor advances on a 2 m acceptance radius (IGVC rule); `ABORTED`/`CANCELED` failures fall through to the next waypoint.

6.2 Localization

Localization uses two instances of `robot_localization`'s EKF; both run at 30 Hz in 2D mode. The local EKF (`odom` → `base_link`) fuses Xsens IMU (100 Hz), wheel odometry from the actuator node (50 Hz), and ZED X VIO yaw (yaw-only, fused differentially to avoid the lever-arm bias from the camera's offset mount). The global EKF (`map` → `odom`) fuses the IMU stream with GPS odometry from `navsat_transform_node` in differential mode with a 6σ Mahalanobis outlier-rejection threshold. Integrating per-step GPS deltas rather than absolute fixes prevents costmap jumps from ±2–5 m consumer-grade GNSS noise (RTK is disabled per IGVC rules).

6.3 Drive Control

The `actuator_node` bridges Nav2 to the physical drivetrain. It subscribes to `/cmd_vel` (from Nav2, teleop, or the velocity smoother) and runs a 50 Hz loop: (1) source arbitration; (2) slew-rate limit (linear 0.5 $\mathrm{m/s^2}$ accel, 1.5 $\mathrm{m/s^2}$ brake; angular 1.0 $\mathrm{rad/s^2}$); (3) IMU-based heading hold ($K_p = 1.5$) when $|\omega_{\text{cmd}}| < 0.05$ rad/s, or gyro-stabilized turn ($K_p = 0.3$) otherwise; (4) diff-drive inverse kinematics with `track_width_m=0.7366`, `m_per_motor_rev=0.01994`; (5) ASCII serial write to the Teensy at 115200 baud. In parallel, encoder feedback is integrated into a 50 Hz `/wheel_odom` message that the local EKF consumes.

6.4 AutoNav Driving Logic

- **Lane following:** HSV white-lane masks feed the local costmap's `semantic_layer` as lethal cells; MPPI generates rollouts that stay inside the lane corridor.
- **Obstacle avoidance:** VLP-16 STVL marks barrels in both costmaps; MPPI's `PreferForwardCritic` and `PathFollowCritic` bias toward forward motion along the global plan when no obstacles intrude.
- **Ramp handling:** The 2D-mode EKF assumption holds for the IGVC ramp; pitch is monitored from the IMU. MPPI's cost-regulated velocity scaling automatically slows the robot near costmap inflation.
- **Switchbacks and dead-ends:** The BT recovery escalation (clear-costmap → wait 2 s → backup 1.5 m at 0.08 m/s → crawl forward 0.15 m at 0.10 m/s) reorients the planner when the global plan is repeatedly blocked, allowing the robot to back out of dead-ends and switchbacks.
- **Potholes:** The HSV pipeline's class 3 channel marks the IGVC 2 ft white-circle simulated potholes as lethal in the local costmap; MPPI rollouts that would enter a pothole are scored

as collisions and rejected.

- **Traps and complex obstacles:** The 50×50 m local costmap holds obstacles long enough for MPPI to consider alternate paths; the inflation gradient ($k_{\text{scale}} = 3.0$) gives MPPI a smooth cost field so rollouts naturally funnel through the widest gap between barrels.

6.5 Driving Logic Requirements Compliance

Table 10: Driving logic requirements: target vs. measured.

Requirement	**Target**	**Measured**
Obstacle reaction time (STVL update → `cmd_vel`)	$\leq$ 250 ms	47 ms
Waypoint arrival accuracy (2 m IGVC tolerance)	$\leq$ 1 m	$<$ 1 m
Minimum average speed (44 ft IGVC stretch)	$\geq$ 1 mph	$\geq$ 1.7 mph (Nav2 nominal 0.75 m/s)

7. Key Performance Indicators

Table 11: AutoNav key performance indicators: target vs. measured.

KPI	**Target**	**Measured**
Course completion time (Webots simulation)	$\leq$ 5 min	Pending full-course run (path planner validated in sim)
Course completion time (field, IGVC 6 min ceiling)	$\leq$ 5 min	Pending field test
Ramp ascension ($\leq$ 15% gradient per IGVC II.2)	Pass	Predicted pass (NEO motor torque margin); pending field test
Maximum ground speed (hardware-governed)	$\leq$ 5 mph	3.4 mph (1.53 m/s)
Minimum average speed (44 ft hold-up rule)	$\geq$ 1 mph	1.7 mph (Nav2 nominal 0.75 m/s)
Obstacle reaction time	$\leq$ 250 ms	47 ms
Battery runtime under autonomous load	$\geq$ 6 min	∼6.5 h calc.; field burndown pending
Wireless E-stop range (100 ft IGVC minimum)	$\geq$ 150 ft	>100 ft (verified)

Target rationale. Course completion is set at $\leq$ 5 min, one minute under the IGVC II.2 six-minute ceiling, leaving margin for a slow restart after a `RecoveryNode` fallback. Maximum speed is hardware-governed at 3.4 mph (below the 5 mph rule cap and above the 1 mph minimum hold-up threshold), giving the MPPI controller a stable operating point. Battery runtime targets at least one full 6 min course run plus margin for warm-up and a queued retry. The wireless

E-stop has been verified to operate beyond the IGVC 100 ft minimum with a comfortable margin to absorb multipath and antenna-orientation variability during the open-field range test.

8. Analysis of Complete Vehicle

8.1 Lessons Learned

Among the multiple lessons learned during this project, we highlight overcoming budget constraints through creativity, resilience, and teamwork. As first-year participants, our team had to work together to learn how to build a robot and make it work, while simultaneously navigating the steps to secure financial support from the university to afford the rapidly growing costs — both the planned budget items and the many unaccounted-for expenses that surfaced during fabrication and integration.

Working as a team, after this experience, means much more than putting a group of people together. We had to learn from each other and trust each other's decisions; everyone worked toward the same goal with the tools and limited availability they had. We are proud of everything we achieved and look forward to showcasing our work at IGVC.

8.2 Failure Points and Resolutions

Table 12: Known failure modes and mitigations.

Layer	**Failure**	**Mitigation**
Mechanical	Fasteners loosen under vibration	Loctite + Nyloc nuts; periodic torque checks
Electrical	Motor inrush brown-out on shared rail	Separate Jetson power rail
Electrical	Wireless E-stop link loss	Heartbeat timeout cuts motor relay power
Sensing	Encoder dropout	Local EKF reverts to IMU prediction
Sensing	GPS multipath / dropout	6σ outlier rejection; differential fusion
Perception	HSV threshold drift across lighting	Adaptive V-floor; live-tunable HSV bounds
Costmap	Stale semantic cells linger	STVL 5 s decay; semantic raytrace clearing
Software	Stale `cmd_vel`	300 ms watchdog forces motor stop

8.3 Safety, Reliability, and Durability

Safety is enforced at every layer: a hardware speed cap (4600 RPM → 1.53 m/s) in the Teensy firmware, a 300 ms serial watchdog on the Teensy, a 100 ms CAN heartbeat timeout per SparkMAX, a software slew-rate limit in the actuator node, and Nav2's MPPI collision critic plus a costmap inflation radius of 0.65 m. The MPPI `CostCritic` uses `consider_footprint: true` so the rectangular 1.00×0.74 m footprint is swept along every rollout trajectory.

8.4 Software Testing, Bug Tracking, and Version Control

Software is developed in a Git repository (`github.com/Paarseus/IGVC_ROS2`) with feature branches, pull-request review, and automated linting via `colcon test`. All parameters — perception thresholds, EKF covariances, Nav2 gains, MPPI critic weights, actuator slew limits — live in version-controlled YAML.

8.5 Simulation and Physical Testing

Development and regression testing use the `avros_sim` package, which runs Webots with a `cpp_campus.wbt` world imported from the Cal Poly Pomona OpenStreetMap extract. A custom Webots driver publishes simulated VLP-16, ZED, and GNSS topics that the rest of the stack consumes unchanged, enabling perception and Nav2 development without field time. The simulator has been used to validate the path-planning and control functionality of the autonomy stack: `NavfnPlanner` produces feasible plans through obstacle fields, and the MPPI controller follows them while respecting the same costmap-inflation constraints used on the real vehicle. A full timed AutoNav-course replication in simulation is planned but has not yet been run; the field course completion time below is therefore the primary KPI for the AutoNav challenge.

Physical testing progresses through four stages: bench (subsystem in isolation), simulation (Webots full stack), low-speed field tests on a flat surface, and full-speed field tests on the IGVC practice course.

A timed full-course sim-versus-field comparison has not yet been completed. Four sim-vs-real gaps already drive specific design decisions: (i) Webots GPS is noise-free vs. ±2–5 m on field (differential GPS fusion + 6σ outlier rejection, §6.2); (ii) uniform sim lighting vs. outdoor sun-angle/reflectivity variation (runtime HSV auto-calibration, §5.2); (iii) Webots physics does not model track-suspension yaw drift and IMU vibration (IMU heading-hold, §6.3); (iv) USB-serial and CAN latencies are deterministic in sim but bounded on hardware by the 300 ms watchdog and 100 ms CAN heartbeat (§4.3). A timed full-course Webots run is planned, after which costmap inflation and MPPI critic weights will be retuned to close the residual gap.

9. Cyber Security Analysis

We adopt the NIST Risk Management Framework (RMF) — prepare, categorize, select, implement, assess, authorize, monitor — to identify the cyber risks present in MARVIN and to plan the hardening that would be required before the platform enters series production. Per IGVC 2026 rule IV.3 §11, the three highest-risk vulnerabilities are listed below.

Table 13: Top three cyber vulnerabilities, today's controls, and production hardening.

Vulnerability	Risk	What we do today	Production hardening
Onboard network breach	Medium	A laptop on the same Wi-Fi as the robot can publish ROS 2 commands directly because we run the default DDS configuration. Mitigated by limiting Wi-Fi access to team laptops and by the judges' wireless E-stop, which can cut motor power independently of any software state.	Enable ROS 2's built-in DDS security (per-node certificates, encrypted topics) and segment the autonomy network onto its own VLAN with WPA3.
Web UI controller hijack	Medium	The operator web UI lets anyone on the LAN take control of the robot from a browser, with no login. We compensate with a 500 ms watchdog that zeros the drive command on lost input, a manual-override priority rule, and the wireless E-stop.	Add TLS with per-operator client certificates on the WebSocket, issue per-session tokens from the safety controller, and restrict the diagnostic Foxglove bridge to localhost.
GPS spoofing	Low	A nearby radio attacker could feed our GPS receiver false position fixes. The localization EKF rejects implausible jumps with a 6σ outlier filter and treats GPS as a per-step correction rather than an absolute position, which limits how far a slow-drift spoof can move us.	Use a multi-frequency, multi-constellation receiver with built-in integrity monitoring (RAIM); fetch RTK corrections over authenticated TLS; halt waypoint advance if GPS disagrees with the wheel/IMU dead-reckoning by more than the EKF's error bound.

Defense in depth. Safety-critical motion is protected by three independent stops — a physical E-stop button on the rear of the vehicle, a wireless relay held by the judges, and a software interlock in the actuator node that zeros the drive command on stale input — so a single-layer software compromise cannot move the vehicle, and the hardware relay must be re-enabled by a human once tripped. For a production version of this platform we would add a dedicated safety-monitor microcontroller on an authenticated CAN bus, signed firmware verified at boot, and signed-commit enforcement on the source repository.