Indian Institute of Technology Madras

Intelligent Ground Vehicle Competition 2026

Varaha

Team Captain: *Alan Royce Gabriel*, bs22b001@smail.iitm.ac.in, +91-934-257-9262
Key Team Member: *Ashwaat Tarun T S*, ashwaattarun@gmail.com, +91-988-404-9158
Faculty Advisor: *Dr. Sathyan Subbiah*, sathyans@iitm.ac.in, +91-875-457-1890

Targeted Challenge(s): Both AutoNav and Self Drive

I hereby certify that the development of the vehicle, Varaha, as described in this report, is equivalent to the work involved in a senior design course. This report has been prepared by the students of Team Abhiyaan under my guidance.

SATHYAN SUBBIAH, Ph.D.
Professor
Department of Mechanical Engineering
Indian Institute of Technology Madras
Chennai 600 036, India

Team Leads

Alan Royce Gabriel Yash Purswani

Team Members

Mechanical	Electrical	Software	Business and Design
Advait Abhijeet Kadam	Ashwaat Tarun T S	Eklavya Gade	Dikshansh Raipure
Joel Davis Christopher K	Lalith Akash M	Govind S Ashan	Yuvraj Singh
Pamarthi Shanmukha Akshay	Pranav R	Madhav Pradeep	
Sreya Devaraj	Shitheeksh H V	Vishal S	
	Vishal V		

Contents

1 System and Subsystem Requirements

1.1 System engineering process

Requirements were derived top-down from the IGVC 2026 rulebook (IV.3 §3) and locked at the team's internal Critical Design Review (CDR), following an earlier Preliminary Design Review (PDR). The team follows a V-model (Fig. 1): every requirement closes only when a measured value (bench, sim, or field) sits inside its target band.

Rulebook walkthrough and trade-offs. The PDR maps each course feature (lane width, ramp grade, switchback footprint, stop-line band, pedestrian halt) to one or more measurable requirements; conflicting requirements (low Center of Gravity vs. payload-bay height; sensor-mast height vs. envelope cap) are logged as trade-off entries before design freeze at CDR. No innovation enters the register unless it closes a scoring course feature.

Ownership and change control. Sub-team leads (Mechanical, Electrical, Software) own their requirements end-to-end and present closure evidence at PDR / CDR to the team captain and faculty advisor. Unmet requirements roll onto a P0/P1/P2 risk register reviewed weekly until resolved or escalated to the captain.

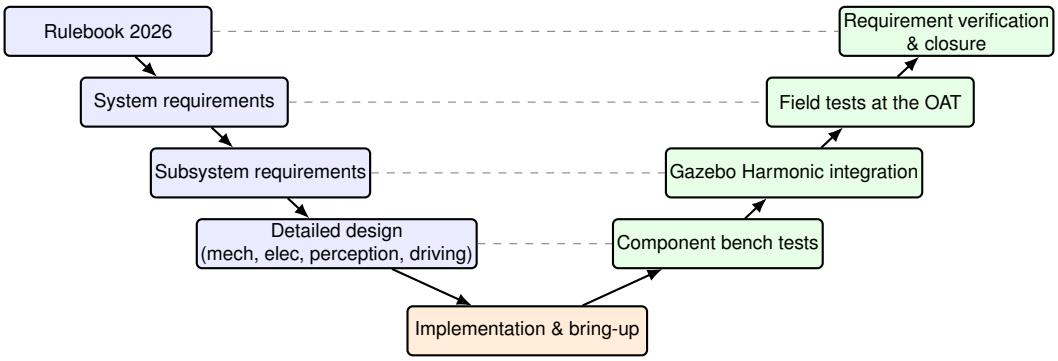

Figure 1: Varaha V-model. Dashed lines indicate level-by-level traceability.

1.2 Requirements register

Tab. 1 lists the common requirements and AutoNav side; Self Drive follows. Each row's **Status** cell hyperlinks to the corresponding section's target-vs-measured table.

Table 1: Varaha requirements register.

ID	Dom.	Rulebook	Requirement / driver	Target	Measurement	Status
Common (mechanical, safety, electrical)						
M1	Mech	I.2 p.5	Envelope inside 3–7 ft L, 2–4 ft W, $\leq$6 ft H	L $\leq$ 7 ft, W $\leq$ 4 ft, H $\leq$ 6 ft	Tape on assembled chassis	PASS
M2	Mech	I.3 p.5	20 lb cinder-block payload retained on the course	No slip or drop during loaded runs	Loaded-run observation	PASS
S1	Safety	I.2 p.5	Mech E-stop: red mushroom, $\geq$1 in $\varnothing$, 2–4 ft above ground, hardware cut	$\varnothing \geq$1 in; height within band; NC in coil leg	Multimeter on coil leg	PASS
S2	Safety	I.2 p.5	Wireless E-stop range $\geq$100 ft, hardware, judge-held	error rate $<$ 1 % at 100 ft LOS	Open-ground walk	PASS
E1	Elec	I.4 p.6	Safety light: solid powered, flashing autonomous, solid manual	Clean state-machine transitions on every input combination	Bench input cycle	PASS
E2	Elec	II.2 p.9	Battery energy supports a 6 min run with comfortable margin	$\geq$1 h continuous loaded run	Loaded soak	PASS
E3	Elec	I.2 p.5 / II.2 / III.3	Hardware-locked speed band 1–5 mph	Software cannot command $>$5 mph; MCU rejects $>$5 mph cmds	Bench + 44-ft straight	PASS

ID	Dom.	Rulebook	Requirement / driver	Target	Measurement	Status
AutoNav — II.2–II.5						
$P1_A$	Perc.	II.2 p.9	Lane following on the 10–20 ft asphalt course (white lanes only)	Lane mask precision/recall $\geq$0.90 on courtyard clip	Sim courtyard clip	PASS
$P2_A$	Perc.	II.2 p.9	Detect 2 ft ⌀ white-circle pothole	Cluster diameter $2\,\text{ft} \pm 24\,\text{in}$; recall $\geq$0.95	Sim cluster diameter	PASS
$P3_A$	Perc.	II.2 p.9	Detect solid obstacles (barrels, drums, natural obstacles)	100 % detection at $\geq$5 ft from sensor mast	Sim cloud review	PASS
$D1_A$	Drive	II.2 p.9	Reach all waypoints in 6 min; avg $\geq$1 mph from 44 ft	$\leq$6 min lap; $\geq$1 mph in first 30 s	Sim lap timing	PASS
$D2_A$	Drive	II.2 p.9	$\geq$5 ft clearance from boundaries and obstacles	STVL inflation $\approx$22 in (1.8 ft)	Sim collision monitor	PASS
$D3_A$	Drive	II.2 p.9	Climb $\leq$15 % ramp grade	elevation map ramp band 8–35°	Sim ramp climb	PASS
$K1_A$	KPI	II.5 p.13	Adjusted-time score (rule ranks by adjusted time)	Finish well under the 6 min cap; 0 student E-stops	Sim regression batch	PASS
$K2_A$	KPI	II.4 p.12	Penalty avoidance: boundary, obstacle-strike, careless driving	0 boundary, 0 hits across runs	Sim collision monitor	PASS

ID	Dom.	Rulebook	Requirement / driver	Target	Measurement	Status
Self Drive — III.3–III.6 + Appendix B/C						
$P1_S$	Perc.	III.6 / F.II.1	Real STOP vs fake (SOUP, IGVC, random) signs	STOP accept rate 100 %; fake accept rate 0 %	Sim sign-set	PASS
$P2_S$	Perc.	App. B / F.V.1	Mannequin (71.7 in, orange vest) and tyre detection	Detection at the 5–7 ft halt band	Sim mannequin gate	PASS
$P3_S$	Perc.	App. B / F.III.1	White and yellow lane detection on the Self Drive course	Both colours $\geq$0.90 P/R on courtyard clip	Sim courtyard clip	PASS
$D1_S$	Drive	App. B / F.III.1, V.1	Stop within 30 cm of stop line; 5–7 ft from pedestrian	30 cm to stop line; 5–7 ft pedestrian band	Sim distance log	PASS
$D2_S$	Drive	App. B / F.V.3	Lane-change in 13–10 ft window on $\geq$10 m radius curve	Trigger $\leq$13 ft; complete $\geq$10 ft past obstacle	Sim re-entry pose	PASS
$K1_S$	KPI	III.3 p.19	16 functions $\times$ 100 pt (1600 pt)	16/16	Sim BT regression	PASS
$K2_S$	KPI	App. C	Full Course 21-function consecutive run	Pass all 21 in one run with low penalty	Sim full-course run	PASS

2 Mechanical Design

2.1 Frame, layout, and chassis

Varaha is a fully enclosed differential-drive platform sized to the IGVC 2026 envelope (Sect. I.2: 3-7 ft L, 2-4 ft W, $\leq$6 ft H), the 20 lb cinder-block payload (16$\times$8$\times$8 in), and the 5 ft inter-obstacle clearance that sets the switchback footprint. The frame is **30$\times$30 mm Al 6061 T-slot extrusion**.

Why extrusion (not welded steel, sheet metal, or CF). As a travelling team from India, freight removes 1-2 months from the pre-competition development schedule. T-slot extrusion lets us disassemble the chassis into panels + bars, fly everything as checked baggage, and reassemble on-site a transport mode neither welded steel, sheet metal, nor a CF tub can match. Material selection was effectively a single-option decision.

Layout, developed inward from the rule envelope. Drivetrain at the rear; Power-Distribution and Main PCBs alongside the drive zone; payload bay central with CoG over the wheelbase; battery between the rear wheels for low CoG; front castor in the nose. The mid-section is waisted to PCB-plus-harness width, eliminating unused volume. External envelope: 43$\times$31$\times$41 in (inside requirement M1).

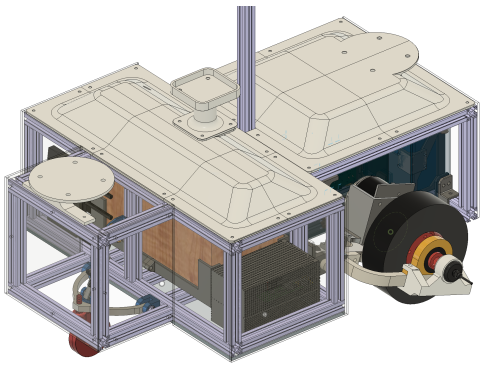

(a) Isometric view with top cover.

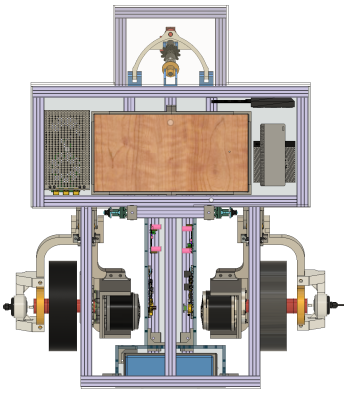

(b) Top view without cover.

Figure 2: Varaha frame and component layout.

2.2 Drive system

Varaha has two independently driven rear drive wheels, along with a single passive front castor on the nose. This ensures that all three points are always in contact with the ground. **MY1016Z3 24 V 350 W** brushed DC motor was selected by torque-first sizing.

Assumption ladder (SI throughout). Vehicle gross $m = 50\,\text{kg}$ (chassis, sensors, battery, 9.07 kg cinder block); rolling-resistance coefficient $C_{rr} = 0.03$ (rubber on coarse asphalt); 15 % ramp grade ($\sin\theta = 0.148$); target on-flat acceleration $a = 2\,\text{m/s}^2$; wheel diameter 10 in ($r = 0.127\,\text{m}$); 30 % headroom to account for thermal derating and other mechanical losses.

Worst case (flat acceleration). The pull-out from rest dominates over the steady ramp climb:

$$F_{flat} = m\,a + m\,g\,C_{rr} = 114.7\ \text{N}, \qquad \tau_{flat} = F_{flat}\,r/2 = 7.28\ \text{N·m per motor}$$
$$F_{ramp} = m\,g\,(\sin\theta + C_{rr}) = 87.3\ \text{N}, \qquad \tau_{ramp} = F_{ramp}\,r/2 = 5.54\ \text{N·m per motor}$$

The **MY1016Z3 (11 N·m continuous / 55 N·m stall)** covers $\tau_{flat} = 7.28\,\text{N·m}$ with $\sim$1.5$\times$ margin after headroom and absorbs the combined transient ($\tau_{flat} + \tau_{ramp} \approx 12.8\,\text{N·m}$) inside the stall ceiling. Top-speed power $P = F_{flat}\,v_{max} \approx 257\,\text{W}$ sits inside the 2$\times$350 W$= 700\,\text{W}$ installed.

Motor operating point. With an internal gear ratio of **9.78 : 1**, the MY1016Z3 has a **324 RPM** no-load output speed (= wheel no-load via the 1:1 direct coupling). At 5 mph the wheels turn at 168 RPM, which is $\sim$52 % of no-load, landing the motor in the **peak-power region** of the brushed-DC torque-speed curve.

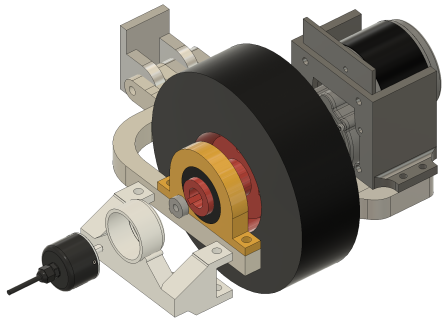

Figure 3: Direct coupler between motor output and wheel hub, with the encoder mounted on the same coupler shaft.

Direct-coupled drivetrain and encoder. A custom in-house coupler links the motor output shaft directly to the wheel hub (**gear ratio 1:1**); the encoder is mounted on this same coupler so it reads the very shaft that drives the wheel. This removes any gear pair from the drivetrain and from the feedback path, so neither inter-gear backlash nor a coupling-train failure can corrupt low-speed odometry.

Wheels. A tubeless tyre of 10 in diameter (solid rubber) was chosen, ensuring that there is no risk of puncturing on the asphalt course, and gives consistent traction across the ramp and over any surface imperfections.

2.3 Suspension

Passive suspension on both rear drive wheels and the front castor keeps all three contact points loaded over surface cracks, expansion joints, ramp lips, and painted-line ridges, preventing wheel slip and the resulting odometry steps.

Rear trailing arm (Al 6061). Trailing-arm was chosen over a double-wishbone (more lateral space, more precision joints) and over a spring-loaded fixed mount (camber variation introduces a systematic odometry scaling error on a differential-drive). The trailing arm holds camber constant through the full vertical travel and carries the spring + damper element that attenuates vibration into chassis and electronics. ANSYS Static Structural FEA confirms a minimum **safety factor of 2** on Al 6061 yield (Fig. 4).

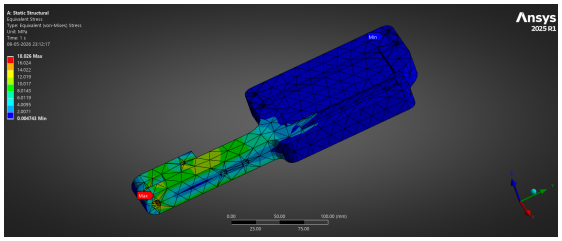

(a) Inner trailing arm.

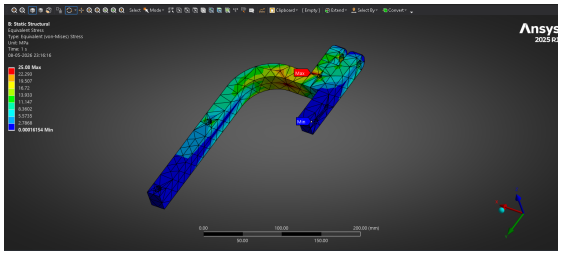

(b) Outer trailing arm.

Figure 4: ANSYS Static Structural von Mises plots on the rear trailing-arm pair; minimum SF 2 on Al 6061 yield.

Front castor trailing arm. Targets the *ramp entry/exit transition*, the largest one-off jolt on the IGVC course: a rigid front mount transmits a sharp impulse into the chassis. The castor trailing arm absorbs the lip and re-loads the front contact within a fraction of the wheel's roll-over time, keeping chassis pitch impulse small. ANSYS FEA on the castor arm also meets **Safety Factor** $\geq$ **2** on Al 6061 yield (Fig. 5).

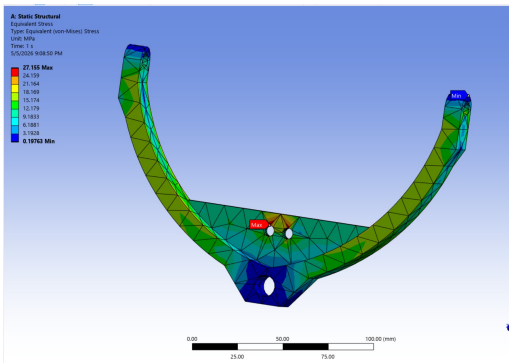

Figure 5: Castor trailing-arm FEA von Mises plot.

2.4 Sensor mast

A **930 mm** vertical Al-extrusion mast with **three-point bracing** (chassis-floor base + top brace to the upper frame + lateral strut along the longitudinal direction). The three-point arrangement was selected after benchmarking it against a **single-point** baseline (chassis-floor base only): an IMU at the mast tip showed markedly larger roll and Gyro X spikes under the single-point mount when driven over the same profile, while the three-point setup produced a smooth response (see §8.5, Fig. 26). The mechanical E-stop is fixed on the mast 2-4 ft above ground on a 3D-printed mount.

2.5 Weatherproofing, mounts, and thermal

Enclosure. Fully enclosed 3D-printed PETG top cover (lightweight, impact-resistant under wind/rain), polycarbonate / acrylic sides for visual inspectability without opening the chassis, acrylic bottom, rubber-gasket + rubber-washer-fastener seal at every panel-to-extrusion joint. IP-rated motors mount *outside* the sealed envelope so the drivetrain breathes, avoiding sealed-compartment complexity.

Serviceability. Each enclosure panel is fastened to the extrusion skeleton with just **four bolts + T-nuts** and removes in **under a minute**, so battery, PCBs, motor mounts, and the sensor mast base are all reachable for inspection or swap without specialised tools.

3D-printed mounts. Custom PLA+ for battery, payload, PCBs, and panel-joining brackets. Battery and payload mounts sit on the chassis extrusion base, so vertical gravitational load goes through the chassis the mounts resist only the lateral/longitudinal inertial forces ($F = ma$, an order of magnitude smaller than $F = mg$).

Thermal. Passive cooling; Jetson Orin has an integrated fan; brushed DC motors mount outside the chassis for passive air cooling; master PCBs are fitted with heatsinks. Adequate for the IGVC run duration; reduces the mass, complexity, and power draw of a forced-air system.

2.6 Mechanical requirements: target vs. measured

Table 2: Mechanical requirements: target vs. measured.

ID	Requirement	Target	Measured
M1	Vehicle envelope inside rule I.2	L $\leq$ 7 ft, W $\leq$ 4 ft, H $\leq$ 6 ft	43$\times$31$\times$41 in: PASS
M2	Payload retention on the course	No slip or drop during loaded runs	PASS: ran with the 20 lb cinder block on multiple loaded sim and bench rolls; no slip or movement observed

3 Safety

Varaha layers multiple independent safety mechanisms across mechanical, electrical, wireless, and software domains. **No single failure can drive the vehicle.**

Power-path cuts (mechanical, master, contactor) (S1). A **1-in ⌀ red mushroom push-button mounted centre-rear at 36 in above ground** (inside the rulebook 2–4 ft band) wires its NC contact *in series* with the supply to the Power-Distribution PCB contactor coil. Pressing it mechanically breaks the coil current and drops an **Altran ALEV100** (24 V coil, 100 A) contactor, removing 24 V from both motor controllers while compute, sensors, and Main PCB stay energised so telemetry continues.

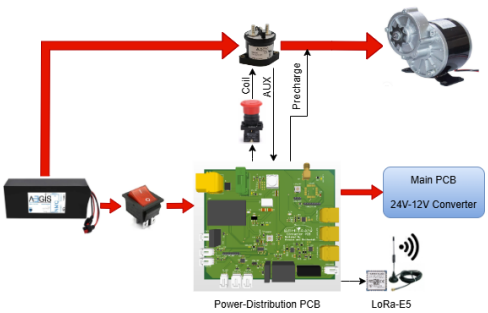

Figure 6: Power-path safety.

The MCU only *reads* the contactor's NO aux contact (3.3 V state); firmware refuses to enable motor PWM until that confirmation asserts, so a stuck coil, welded contact, or broken aux line all collapse to "motors off".

A separate panel-mounted **12 A master toggle** removes power from the entire vehicle for storage / charging / inspection; the MCU has no authority over the coil leg, so software cannot override the cut. At the battery terminal, a **MAXI bolt-down 70 A slow-blow fuse** provides last-line over-current protection against any sustained fault that escapes downstream fusing.

Wireless E-stop with redundancy ((S2)). The primary kill is a **LoRa-E5 (Wio-E5, STM32WLE5)** on the Power-Distribution PCB matched to a hand-held remote held by the judge; a secondary **ESP-CAM** module gives a redundant kill on a separate channel. Loss of either authenticated heartbeat drops the contactor coil through the same hardware leg as the mechanical E-stop. An open-ground range walk confirmed error rate $<$ 1 % at $>$150 m.

Software-side safety. Indicator lights: the Orin asserts an autonomous-mode signal over CAN; the Main PCB flashes the DRL strips while the signal is present and reverts to solid (idle / teleop) on signal loss. **MPPI trajectory validator:** every chosen trajectory passes through a final collision check against the costmap (§6). **Heartbeat watchdog:** the Main PCB emits a CAN heartbeat; dependent PCBs shut their outputs down on its loss.

Transport, parked, charging. *Transport:* fitted hard case with cut-foam saddles, contactor open, master switch off, battery disconnected at the XT90. *Parked:* master switch OFF, visible "battery live" indicator removed from the mast. *Charging:* BMS-protected charger in a non-flammable tray with smoke detection during overnight charges; wall side fused. A laminated pre-run safety checklist runs before every session.

3.1 Safety requirements: target vs. measured

Table 3: Safety requirements: target vs. measured.

ID	Requirement	Target	Measured
S1	Mech E-stop, $\geq$1 in $\varnothing$, 2–4 ft above ground, hardware-cut	1-in $\varnothing$ button, centre-rear at 36 in above ground, NC in coil leg; MAXI 70 A slow-blow fuse at battery	PASS: continuity opens on press; bench-verified
S2	Wireless E-stop range $\geq$100 ft, error rate $<$ 1 %	error rate $<$ 1 % at 100 ft LOS	PASS: $>$150 m on open-ground walk

4 Electrical/Electronic Design

Varaha's electrical architecture centres on a **25.6 V nominal / 29.2 V max, 30 Ah NMC(Lithium Nickel Manganese Cobalt Oxide)** pack feeding three custom PCBs - a **Power-Distribution PCB**, a **Main PCB**, and a **12 V Power-Protection PCB** - with an **STM32H5** as the on-board MCU.

All inter-board communication runs on **CAN** (differential signalling, built-in error detection, intrinsic common-mode rejection in the motor-rich environment): the Jetson Orin issues velocity setpoints and mode flags and receives encoder velocities, motor currents, battery voltage, and contactor state.

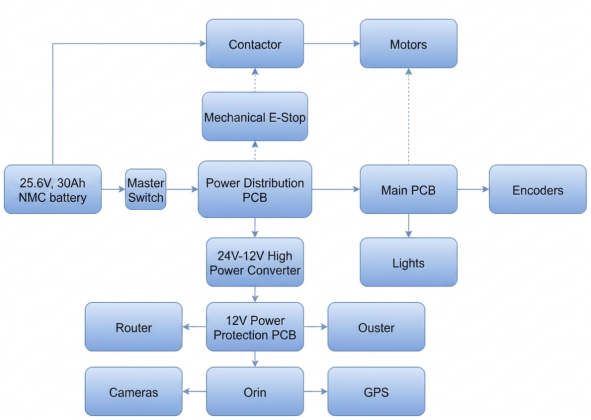

Figure 7: System power-flow.

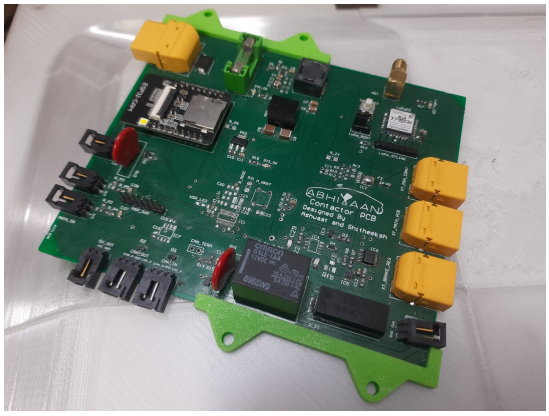

Figure 8: Power-Distribution PCB.

Power-Distribution PCB. The 25.6 V battery enters through an **ESD diode** and a cartridge fuse. A three-stage rail hierarchy follows - a **LMR51625** buck (24$\rightarrow$12 V), a 12$\rightarrow$5 V buck and a LDO converting 5$\rightarrow$3.3 V. The first buck stage (24$\rightarrow$12 V) was chosen for its wide 4.3–60 V input window, so that fluctuations in the 24V line caused by motor-commutation transients cannot disrupt regulation.

Board-level control. A **LoRa-E5** (STM32WLE5) module runs the wireless-E-stop firmware while an **MSPM0** microcontroller drives the contactor and pre-charge relays and handles CAN telemetry.

Contactor. The 12 V rail energises an **Altran ALEV100** (24 V coil, 100 A). Its NO auxiliary contact returns a 3.3 V state signal so the MCU only enables motor PWM after hardware confirmation - a stuck or welded contact fails to "motors off".

Pre-charge. An Omron 12 V relay places a **20 Ω, 7 W** wire-wound resistor in series with the 24 V supply for 500 ms before the contactor closes. With 4000 μF total bulk cap (2000 μF per Cytron MD30C), peak inrush is $24/20{\cdot}1.22 = 1.46$ A and decays exponentially well within the 7 W rating - eliminating the inrush sparks observed without any pre-charging at the E-stop and XT90 contacts (Fig. 9).

Over-voltage clamp. An **LM393** comparator engages a **2 Ω, 50 W** brake resistor at $V_{HIGH} = 31$ V with $V_{LOW} = 29.8$ V hysteresis, catching regenerative-braking spikes (especially with the battery disconnected). The clamp was modelled in **MATLAB Simulink** under repeated regen pulses (Fig. 10): without the clamp the bus accumulates charge to unsafe levels; with it active the bus is held inside the hysteresis band for the entire test.

Overvoltage Threshold rationale. The **Cytron MD30C is rated for a 30 V maximum bus voltage**; 29.8 V (off) sits just above the 29.2 V battery-full level so normal operation never trips the clamp, and 31 V (on) engages the brake resistor before the MD30C's absolute-maximum input is reached.

Figure 9: Sparks in XT90 connector without pre-charging.

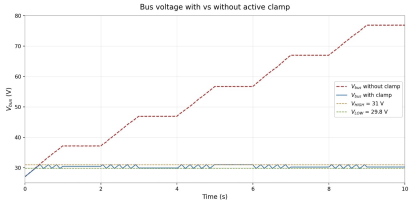

Figure 10: Bus voltage with/without clamp under repeated regen pulses.

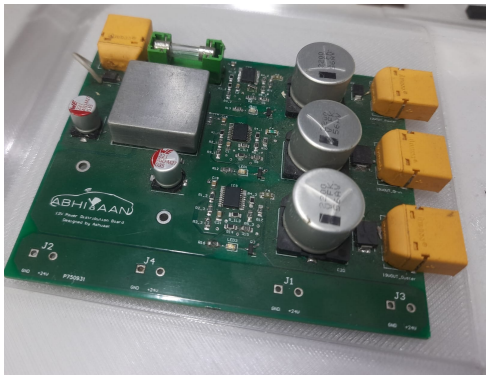

Figure 11: 12 V Power-Protection PCB.

12 V Power-Protection PCB. A galvanically isolated **Mean Well SD-150B-12** (24→12 V, 150W) DC-DC converter feeds this board through an input fuse and ESD diode.

The 12 V rail then forks into two independently-protected paths:

- **12 V path** - eFuse + PCB-mount fuse + bulk cap, supplying on-board routers and the monitor.
- **19 V path** - a 200 W boost converter splits onto two parallel, independently-fused lines, one for the Jetson Orin and one for the Ouster LiDAR, so a fault on one load cannot collapse the other.

Main PCB. Centralises motor control, encoder interfacing, **safety-light driver (E1)**, and current sensing. Its power path is the same as the Contactor PCB, making the PCB resistant to fluctuations in the main 24V power line.

Encoder isolation. The 12 V output reaches the encoders through a **galvanically isolated DC-DC converter**, breaking the ground loop that historically injected commutation noise into the encoder pulse train.

Motor Phase-current sensing. TMCS1126 Hall-effect sensors provide galvanic isolation between motor phase and MCU analogue front-end. The phase current is used to control torque to the motor with a PID loop.

Figure 12: Main PCB.

CAN. A **TCAN1472** transceiver runs **CAN-FD at 250 kbit/s arbitration / 1 Mbit/s data**. CANSIC suppresses stub-reflection ringing. Split termination (62 Ω + 62 Ω to a midpoint cap) used

for common-mode filtering.

Drive motors and cascaded control. Two **MY1016Z3 24 V 350 W brushed DC** motors (sizing in §2) are driven by **Cytron MD30C** controllers, with PWM and direction signals galvanically isolated from the microcontroller. The H5 closes a current loop on the TMCS1126 phase reading and a velocity loop on the encoder rate; closed-loop tracking RMS 96.4 mm/s (≈ 6.8 % of mean cmd vel), characterised in §8.5. RPM limits are enforced at the firmware level so that the 5mph speed limit (**E3**) is adhered to.

Encoder pickup. Quadrature encoders (**Autonics E50S8-1024-3-T-1**, 1024 PPR) run at 12 V (not 5 V) to maximise edge-noise margin; pulses are opto-isolated, level-shifted to 3.3 V, and captured on the STM32H5's hardware-timer inputs at sub-microsecond resolution.

E-stop Remote PCB. A hand-held controller (Fig. 13) with **two independent transmit paths** - a primary **LoRa-E5** (sub-GHz long-range) and a backup **ESP32**. An on-board OLED shows the bot status and the vehicle's main battery voltage; the remote charges its onboard LiPo battery from any USB Type-C source.

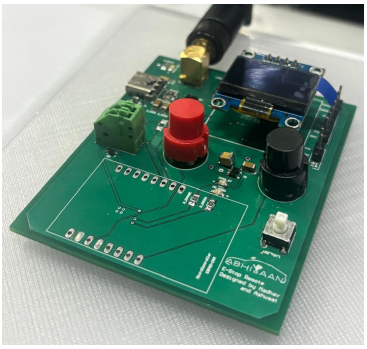

Figure 13: E-stop Remote PCB.

4.1 Power budget and runtime (E2)

The 25.6 V / 30 Ah NMC pack stores **768 Wh nominal**; the BMS cuts at 20 % SoC, leaving ≈**614 Wh usable**. Tab. 4 sums worst-case sub-system draws.

Table 4: Worst-case power budget.

Subsystem	Worst-case draw
Drive motors (2× MY1016Z3, peak from battery)	400 W
Jetson AGX Orin (max)	65 W
Ouster OS1-32 LiDAR	16 W
PCBs, sensors, comms, indicator lights (combined)	10 W
Total worst case	**491 W**

Worst-case runtime. 614 Wh / 491 W ≈ 1.25 h (≈ 75 min) under continuous worst-case load (both motors at peak simultaneously). With realistic motor duty cycles - cruise at the IGVC operating band with intermittent acceleration - field-measured runtime is >**2 h continuous loaded run**, comfortably above the 6 min IGVC AutoNav cap.

Recharge. Charging at **5 A** (128 W charging power at 25.6 V mean) takes ≈ 6 h to refill the pack from the undervolage cut-off level.

4.2 Sensor and compute suite

Table 5: Sensor and compute suite.

Device	Part / specification
Compute	NVIDIA Jetson AGX Orin
LiDAR	Ouster OS1, 32-channel, 10 Hz (with internal IMU used for pose fusion)
Wide-angle cameras	2× IMX335
Fish-eye camera	GC2093
GNSS	Unicore UM982 (multi-constellation, dual-antenna heading, Galileo HAS L6)
Encoders	2× Autonics E50S8-1024-3-T-1, 1024 PPR, 12 V
Motor controller	2× Cytron MD30C
Wireless E-stop	LoRa-E5 primary, ESP32 ESP-NOW secondary

4.3 Electrical requirements: target vs. measured

Table 6: Electrical requirements: target vs. measured.

ID	Requirement	Target	Measured
E1	Safety light: solid powered, flashing autonomous, solid manual	Switches successfully between idle and autonomous mode	PASS, bench-verified
E2	Battery energy supports a 6 min run with comfortable margin	768 Wh nominal pack; $\geq$1 h continuous run target	PASS: $>$2 h continuous loaded run measured
E3	Hardware-locked speed band 1–5 mph	Speed does not cross 5mph even if software commands it	PASS, bench-verified at controller level; PASS on 44-ft straight

5 Perception

Perception produces a road / ramp / obstacle classification, lane and pothole geometry, and a metric pose, via a two-stage cloud pipeline (Figs. 14 and 15) and a parallel pose-fusion path (Fig. 17). **Four ROS 2 nodes are shared across both challenges**; **Self Drive** adds an `object_detector` (YOLO26 + PaddleOCR) for stop signs, potholes, mannequin, tyres, and barrels.

5.1 Vision

Sensors. The **Ouster OS1-32** runs at **10 Hz**, 360° horizontal ($\approx$10° lost at the rear by the camera-mast pole), $\sim$84° vertical ($\pm$42° about horizontal), 32 channels; datasheet range 90–170 m, software-limited to **20 m on the obstacle-marking layer** for compute-budget reasons. Cameras run at **640$\times$480 / 30 fps**: the wide-angle **IMX335** has $\sim$85° horizontal FOV, the **GC2093** fisheye $\sim$170° horizontal. **Calibration:** camera intrinsics are obtained with a checkerboard target; camera-to-LiDAR extrinsics are calibrated by recording ROS 2 bags of a target visible to both sensors and running them through `ros2_calib`.

Elevation-map segmentation (`road_obstacle_segmentation`). A three-pass CUDA kernel builds a **ground elevation map** (0.15 m cells) over a 50 m $\times$ 50 m window using `atomicMinFloat` / `atomicMaxFloat` per cell. Cells classify by **vertical span** ($>$0.25 m $\rightarrow$ obstacle) and a **4-neighbour Sobel slope** on $\min z$ ($<$8° road, 8°–35° ramp, else obstacle). A **0.10 m height-jump rule** reclassifies a flat cell above a flat neighbour as obstacle (catching the floating tip of an obstacle whose base was missed).

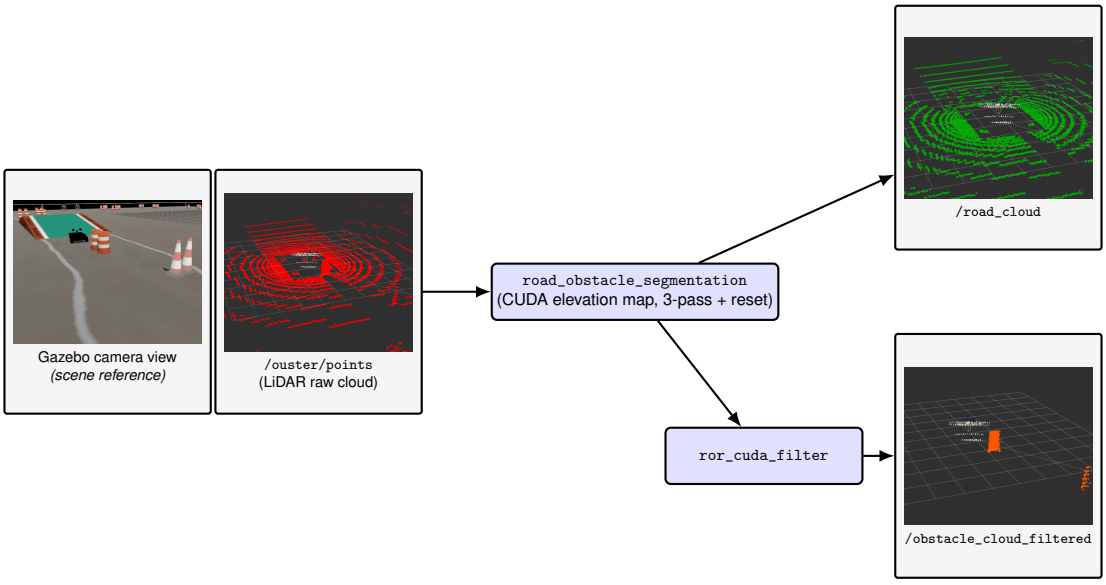

Figure 14: Stage 1: Segmentation

Lane and pothole projection (`ground_densifier_cuda`). The densifier raster-scans the ground at **0.05 m**, lifts each cell to its elevation-map z, and projects it into every camera (`plumb_bob`, fisheye, and `rational_polynomial` on the GPU). A **polar shadow map** ($\pm 5^\circ$ per obstacle) prevents colour-painting behind occluders; a **gravity-constrained RANSAC plane** (axis $\hat{z}$, 3° tolerance) fills unreachable corners. AutoNav applies an **RGB white threshold**; Self Drive adds yellow on the same kernel. Masked points are clustered, then PCA-split into linear (lane) and isotropic (pothole) groups. Reliable detection range is $\sim$**3–4 m ahead of the bot** for both lanes and potholes (limited by camera ground-projection horizon); solid obstacles are detected at the LiDAR's software limit of **20 m**. Lane publication is **climb-only** (positive IMU pitch $> 8^\circ$): on climbs the cameras see over the ramp's far edge and project distant ground pixels as phantom near-lanes (see §8.1).

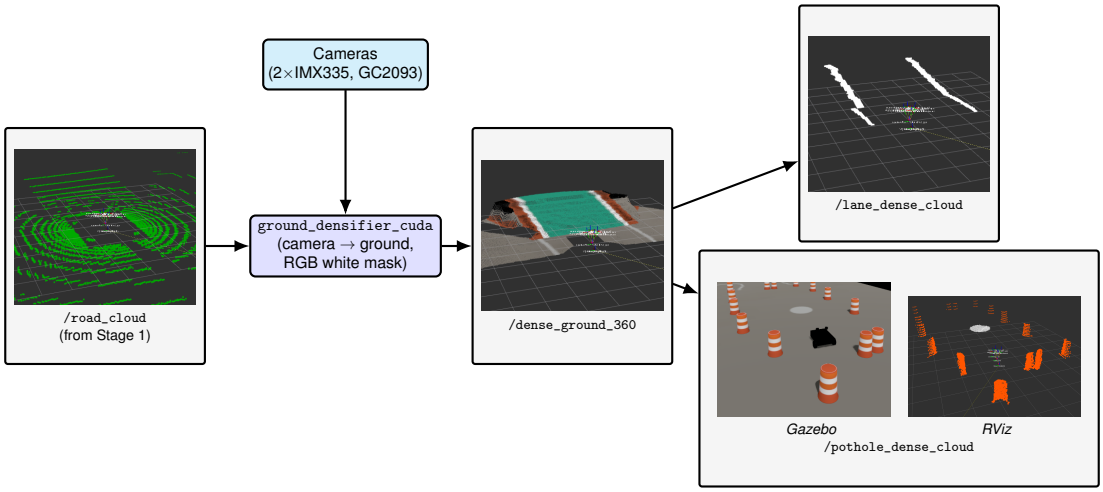

Figure 15: Stage 2: Densification.

Object and sign detection (Self Drive). The same four nodes above run alongside an `object_detector` pipeline: **YOLO26 medium** on **TensorRT FP16** detects road signs, potholes, barrels, mannequins, and tyres at **50–60 FPS** with OCR running concurrently. When YOLO confidence on a road sign exceeds 0.5, the cropped sign region is passed to **PaddleOCR** (every fifth frame to bound load), which reads the text. A **regex whitelist** over the recognised text classifies the sign as STOP or UNKNOWN; only a real STOP triggers a stop, while SOUP, IGVC, and random-letter fakes are pass-throughs. A single-stage YOLO26 was tried first but confused octagonal red signs (STOP vs. SOUP); the detect-then-read split preserves precision and frame rate.

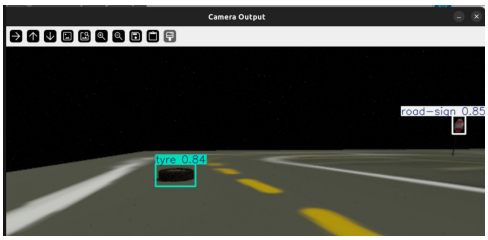

Figure 16: YOLO26 detections in simulation: a tyre and a road sign correctly identified with confidence scores.

Training. 500 images / 200 epochs / 480×480 input on a combined Roboflow + Kaggle + IGVC (from YouTube clips) + field-testing dataset, annotated in Label Studio with Albumentations augmentation (rain, noise, sun-flare, shadow, blur).

Table 7: YOLO26 medium validation metrics across all object classes.

Class	Img	Inst	P	R	mAP@0.5
All	103	272	0.948	0.916	0.962
Barrel	18	70	0.971	0.957	0.982
Mannequin	42	127	0.953	0.906	0.964
Pothole	10	13	0.904	0.923	0.936
Road sign	28	28	0.910	1.000	0.992
Tyre	24	34	1.000	0.796	0.937

5.2 Internal representation

Pose (Fig. 17). A **GTSAM iSAM2 factor graph** fuses wheel odometry (50 Hz), the Ouster IMU (100 Hz, Madgwick yaw), GNSS (10 Hz, UM982 with **Galileo HAS**) and the UM982 dual-antenna GNSS compass reading. A **fast thread** dead-reckons on every wheel callback and publishes `/odom` at 50 Hz so the controller never starves; a **slow thread** inserts wheel/IMU `BetweenFactorPose2` and a GPS `PriorFactorPose2` (plus compass-yaw prior when sub-degree heading is available) on each GPS arrival, solves the factor graph to get the smoothed global pose, and atomically swaps in the new `map`→`odom` transform. Per-fix UM982 covariances feed the graph directly, and high-σ fixes are **auto-deweighted**, with no hard cutoff.

Time alignment. The Orin runs as **PTP grandmaster** on the Ethernet bus with the Ouster as PTP slave, aligning `/ouster/points` and `/ouster/imu` stamps to a common clock.

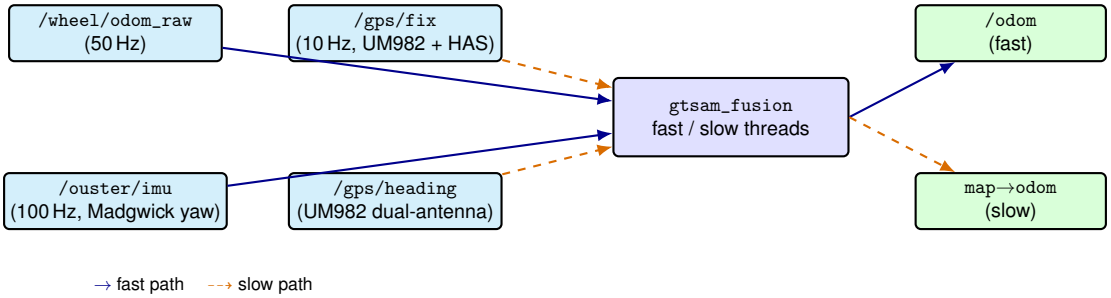

Figure 17: Pose-fusion path.

Costmap. Lane, pothole, and filtered-obstacle clouds fuse into a Nav2 **Spatio-Temporal Voxel Layer (STVL)** at **10 Hz mark / 5 Hz publish** (Fig. 18). Obstacle voxels decay with a **1 s half-life**, lane voxels with **15 s**, so that transient field clutter does not freeze the plan while painted-lane evidence persists. STVL's 3D representation marks ramps traversable while inflating overhead obstacles. **Cross-run caching is forbidden by IGVC**; we re-spawn the costmap every run.

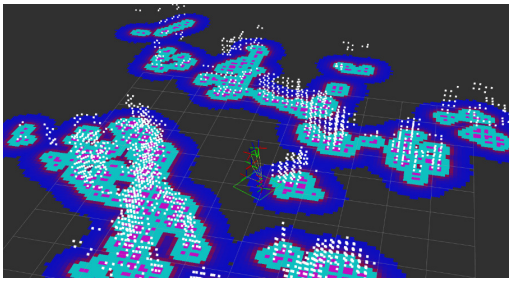

Figure 18: STVL costmap snapshot.

5.3 Perception requirements: target vs. measured

Table 8: Perception requirements: target vs. measured.

ID	Requirement	Target	Measured
$P1_A$	White lane following on the 10–20 ft asphalt course	Detect white lanes; reliable to ≥ 3 m ahead	RGB white threshold + densifier projection; reliable lane ∼3–4 m ahead (sim)
$P2_A$	2 ft ∅ white-circle pothole	Cluster diameter 1.0 ± 0.6 m; reliable to ≥ 3 m ahead	1.0 ± 0.6 m via PCA bounding-box; reliable pothole ∼3–4 m ahead (sim)
$P3_A$	Detect solid obstacles (barrels, drums, natural obstacles)	Detection out to the LiDAR software limit	Obstacles detected up to 20 m range (sim)
$P1_S$	Real STOP vs SOUP/IGVC/random fakes	STOP accept rate 100 %; fake accept rate 0 %	2-stage YOLO26 + PaddleOCR + regex: 100 % STOP accept, 0 % fake accept on sim sign-set
$P2_S$	Mannequin and tyre detection	Detected at the rule stop bands; mAP@0.5 ≥ 0.90	YOLO26; mannequin mAP@0.5 $= 0.964$, tyre $= 0.937$ (Tab. 7)
$P3_S$	White and yellow lane detection (Self Drive course)	Both colours classified ≥ 0.90 P/R on courtyard clip	RGB white + yellow on same densifier kernel; courtyard clip ≥ 0.90 P/R (sim)

6 Driving Logic

Both courses share a single autonomy stack on the Jetson AGX Orin: the same GTSAM pose, the same **STVL** costmap, the same Nav2 **SmacPlannerLattice** global planner, the same Nav2 **MPPI** local controller, and the same software safety chain that clamps every command before it reaches the motors. AutoNav and Self Drive differ only above this layer, in goal sequencing and which perception channels feed the costmap.

6.1 AutoNav

Mission control combines two goal sources: a GPS-waypoint sequencer that drives the bot toward each rulebook waypoint, and a lane-following node that fills in the gaps with intermediate goals to keep it centred between the visible lane clusters.

GPS waypoint following (`waypoint_follower_go`). Each rulebook waypoint is a WGS84 `(lat, lon)`. A 5 Hz timer reads `/gps/fix`, computes the flat-earth offset from current fix to target, adds it to the current map-frame pose, and dispatches a `NavigateToPose` goal, always reflecting the latest GPS solution. Goals are re-issued only on $>$0.2 m drift; arrival is declared at $<$3.5 m to the goal. Under GNSS dropout the controller continues on the last accepted goal using the GTSAM-fused odom.

Lane following (`lane_follower_go`). A 5 Hz C++ node that issues `NavigateToPose` goals to Nav2 from the two lane-cluster streams. The goal is the **midpoint of the forward tips** of the left and right clusters (or a $w/2$ offset from the visible tip when one drops out), then **spiral-searched** through the global costmap for the nearest free cell whose straight-line path does not cross the dense lane cloud, so that a snapped goal cannot tell Nav2 to cut across a painted line.

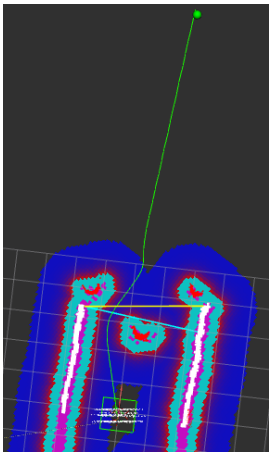

Figure 19: Lane-following.

Global planner: SmacPlannerLattice. An A*-based state-lattice planner that searches over pre-computed motion primitives, so every returned path is **kinematically feasible by construction** rather than post-smoothed. This is one of the fastest Nav2 global planners that also respects the robot's kinematics. A custom lattice file encodes Varaha's diff-drive-plus-front-castor kinematics so $r_{\min}$ is honoured from the first expansion.

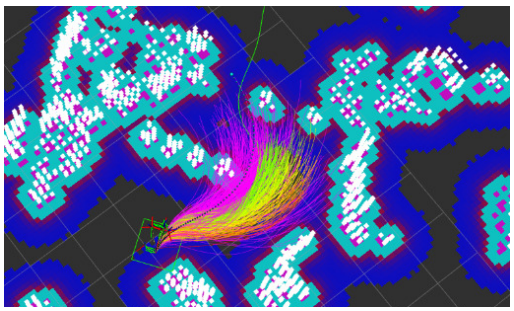

Figure 20: MPPI sampled trajectories (colored) and the chosen softmax-weighted average (dotted).

Local controller: Nav2 MPPI. A sampling-based predictive controller. Each cycle it draws a batch of **Gaussian-perturbed** control sequences around the previous optimum, forward-simulates each as a candidate trajectory under the diff-drive motion model, scores every candidate against a critic stack (path-align, obstacle, goal, goal-angle, prefer-forward), and emits a **softmax-weighted average** over safe trajectories rather than a single best path. The integral over many cost-weighted samples makes the controller **robust to short-lived sensor glitches**: a single noisy sample cannot snap the command.

Course features. Ramps are marked traversable by the elevation-map $8°$–$35°$ **slope band** so MPPI commits to them rather than routing around them. **Switchbacks and traps** are handled by costmap inflation alone. The planner backs out of dead-ends because the lattice expansion exhausts valid neighbours; no special detector or escape behaviour fires. **Potholes** inflate identically to a barrel; no separate behaviour.

6.2 Self Drive

Self Drive runs the 21-step Full Course as a **BehaviorTree.CPP v3** tree ticked at 20 Hz. Behaviours emit a global Path dispatched to MPPI. The diff-drive-plus-front-castor chassis forbids in-place rotation, so every planned path respects $r_{\min} \approx 1.5$ m.

Lane keeping (Q.1, F.III.1, F.VI.1). `AlignToWhiteLanes` keeps the centreline as the midline between the white right-edge and yellow centreline cluster fits, or as a $\pm w_{\text{lane}}/2$ offset from whichever single line is visible, and emits a 2.0 m smooth-step heading-and-lateral blend to MPPI. On curved roads `lane_follower_node` adds a sliding-window polyline whose median lateral re-centres each window, naturally tracing the rulebook's 10 m **minimum-radius inside curve** without an explicit curvature estimator.

Stop signs (F.II.1, F.III.1). YOLO26 + PaddleOCR classifies signs as STOP or UNKNOWN; forgeries (SOUP, IGVC, off-coloured red octagons) collapse to UNKNOWN. `DetectStopBar` runs a 2-D PCA on the white cloud and admits a stop bar only when linearity $(\lambda_1 - \lambda_2)/\lambda_1 \geq 0.6$ *and* perpendicularity $|a_y| \geq 0.85$. The BT cruises until the projected bumper-to-line distance is ≤ 0.30 m (the F.III.1 tolerance band), dwells 3 s, then yields (Fig. 21).

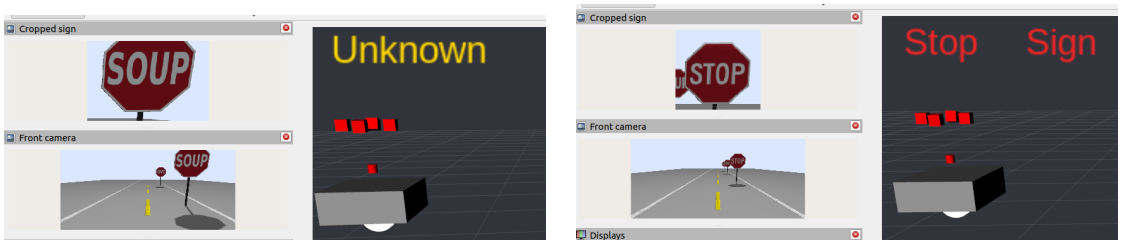

(a) Fake (SOUP) rejected. **(b)** Real STOP triggers a halt.

Figure 21: Stop-sign classifier: fakes collapse to UNKNOWN; only real STOP fires the halt.

Intersection turns (Q.3 left, Q.4 right; F.III.2–3). Both turns use the **same two-cluster pattern**: cluster the perpendicular lane returns into the cross-street's edges, fix an intermediate goal at the intersection centre (line through the two cluster centroids, intersected with the bot's heading vector), and place the final goal at the same lateral offset δ the bot held from its right edge, rotated $-\pi/2$ **for right** (single edge transfer) or $+\pi/2$ **for left** (cross both lanes, target the far edge). SmacPlannerLattice connects the two waypoints with kinematically feasible primitives, so $r_{\min}$ is honoured by the lattice itself rather than a hard-coded arc, and the geometry adapts to whatever lane width the cross-street presents on a given run. The "One Way" sign called out in F.III.2 is visual confirmation only. The turn direction is selected by the BT branch, not perception (Figs. 22 and 23).

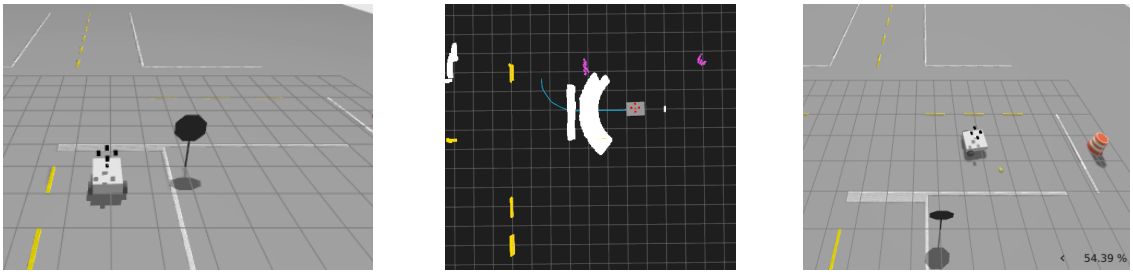

(a) Start. **(b)** Path generated. **(c)** Reached.

Figure 22: Right-turn intersection: start → planned path → reached.

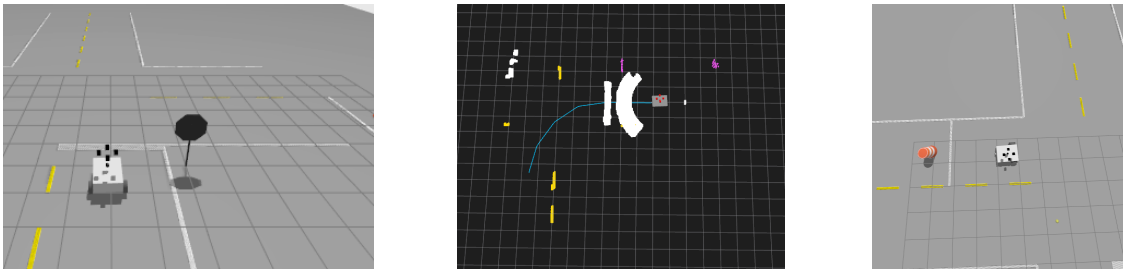

(a) Start. **(b)** Path generated. **(c)** Reached.

Figure 23: Left-turn intersection: start → planned path → reached.

Pedestrians (F.V.1, F.V.2). A forward-cone monitor on `/obstacle_cloud` halts the bot at 1.83 m, centred in the rulebook 5–7 ft halt band so a single-tick perception lag at 4.5 mph leaves ± 0.30 m margin on each side. The LiDAR cone return and the YOLO + HSV-orange semantic channel race; whichever fires first wins, keeping perception-to-brake latency well inside the 0.75 s budget the dynamic case (F.V.2) imposes at 4.5 mph.

Lane change (F.V.3, F.V.4, F.I.2, F.VII.1). `DetectObstacleAndLaneChange` triggers at cone range $\leq$ 3.0 m and emits a $\tanh$ **S-curve** in the shape of an outbound shift, then a 4 m straight, then a return shift. MPPI sees no curvature jump at segment boundaries, with re-entry inside the 13–10 ft rulebook band. The same primitive serves barrels, tyres (F.I.2), the pedestrian-blocking-vehicle case (F.V.3, with a 0.3 m/s slow-down), and potholes.

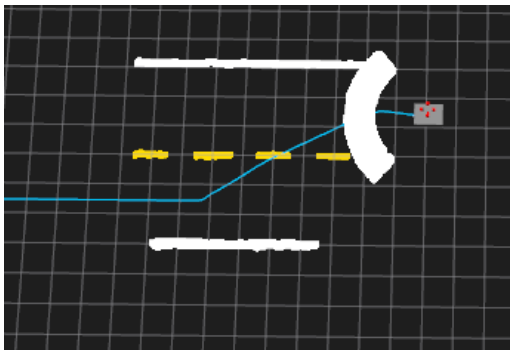

Figure 24: Lane-change S-curve generated past an obstacle.

Parking (F.IV.1–F.IV.3). Pull-out is a scripted open-loop sequence on `/cmd_vel_raw`: 1.5 m straight, a quarter-circle arc at $r_{\min}$, then a 3-ft barrel halt.

Pull-in (F.IV.2) is a quarter-circle entry arc plus a 0.6 m centring straight, wrapped in a BT `Parallel` alongside a barrel-stop monitor so a 0.1–0.2 m goal overshoot never costs a line-cross penalty. Parallel parking (F.IV.3) is executed as a three-arc 3-point manoeuvre (forward arc, controlled reverse at 0.3 m/s, forward straightening) inside the parking-zone boundary; the low reverse speed avoids the front-castor shimmy seen at full reverse velocity, and every sub-segment stays inside the lattice's $r_{\min}$ envelope so no white line is crossed.

6.3 Driving-logic requirements: target vs. measured

Table 9: Driving-logic requirements: target vs. measured.

ID	Requirement	Target	Measured
$D1_A$	Reach all waypoints in 6 min; avg $\geq$1 mph from 44 ft	$\leq$6 min lap; $\geq$1 mph in first 30 s	1 min 42 s best / ~1 min 50 s median across the AutoNav regression batch (sim)
$D2_A$	$\geq$5 ft clearance; ramp $\leq$15 %	STVL inflation ≈22 in; ramp band 8–35°	0 boundary exits, 0 obstacle hits across the regression batch (sim)
$D1_S$	Stop $\leq$30 cm of stop line; pedestrian 5–7 ft	30 cm to stop line; 5–7 ft gates (rule band)	Stop-bar distance $\leq$30 cm and pedestrian halt within 5–7 ft on every SD regression run (sim)
$D2_S$	Lane-change 13–10 ft on $\geq$10 m radius curve	Trigger $\leq$13 ft; complete $\geq$10 ft past obstacle	Trigger at $\leq$3.0 m cone range; re-entry pose inside the 13–10 ft band on every lane-change run (sim)

Discussion. All four driving-logic requirements close inside their target bands across the deterministic Gazebo regression.

D1$_A$: the 1 min 42 s best / ∼1 min 50 s median lap leaves ∼70 % margin under the 6 min cap. This leaves comfortable headroom for an off-nominal recovery.

D1$_S$ and **D2$_S$**: every sim run lands inside the rulebook bands, and the lane-change `tanh` S-curve preserves MPPI's continuous-curvature assumption end-to-end. The remaining unknown is field transfer of these sim numbers (discussed in §8.5).

7 Key Performance Indicators

Tab. 10 lists 3 AutoNav (A1–A3) and 6 Self Drive (S1–S6) KPIs. Targets are rulebook-driven (or tighter for margin); measured values are tagged *(sim)*, *(bench)*, or *(field)* by test rig.

Table 10: KPIs: target, measured, rule-and-measurement basis.

KPI	Description	Target	Measured	Rule + measurement method
AutoNav				
A1	End-to-end course-completion time	<5:00 (margin under 6 min cap)	1 min 42 s best / ∼1 min 50 s median (sim); 23 s best on ∼100 ft OAT segment (§8.5)	II.5 ranks by adjusted time; lap timer on /odom + GPS-stamped start/finish gates
A2	Boundary exits and obstacle hits	0 boundary, 0 hits	0/0 across the AutoNav regression batch (sim)	II.4 penalty schedule; collision monitor + lane-exit log on every run
A3	Average speed over first 30 s	≥ 1 mph (II.2 DQ floor)	∼3.6 mph cruise (5.3 ft/s), 3.6× floor (sim)	II.2 disqualifies below 1 mph; /odom integration over the first 30 s
Self Drive				
S1	Functions passed (out of 16)	16/16	16/16 (sim)	III.3, 100 pt each; BT regression PASS/FAIL per function
S2	Full Course score (21 consecutive functions)	≥ 1900/2100 (margin under cap)	Full course completed in ∼4 min (sim)	App. C scoring; full-course end-to-end run with score tracking
S3	Distance from front bumper to stop line at full stop	≤ 30 cm	≤ 30 cm on every SD regression run (sim)	App. B intersection rule; /odom distance log at PCA stop-bar acquire
S4	Pedestrian halt band relative to mannequin	5–7 ft (rule band)	5–7 ft on every SD regression run (sim)	App. B pedestrian rule; LiDAR-cone distance log at halt
S5	Stop / fake-stop sign FP rate (OCR whitelist)	0 % FP on SOUP/IGVC fake-sign set	0 % FP on sim sign set	III.6 forgery rule; OCR-whitelist outcome on each sign in the regression
S6	Lane-change completion past obstacle on ≥ 10 m curve	≥ 10 ft past obstacle	≥ 10 ft on every lane-change regression run (sim)	App. B lane-change rule; /odom re-entry pose log

Discussion.

AutoNav: A1 leaves ∼70 % margin under the 6 min cap; A2's zero-boundary / zero-obstacle outcome confirms STVL's ∼1.8 ft inflation; A3 sits 3.6× above the 1 mph DQ floor.

Self Drive: 16/16 functions, ∼4 min full course, stop-line / pedestrian / lane-change re-entry all inside rule bands, OCR whitelist 0 % FP.

Sim-to-field gap: deltas characterised in §8.5; mitigations are iSAM2 covariance deweighting, climb-only lane gating, and YOLO26 field-frame top-up.

8 Analysis of Complete Vehicle

8.1 Lessons learned during construction and integration

- *Recalibration is a planned assembly step.* The chassis disassembles for checked-baggage transport, so every reassembly invalidates the camera–LiDAR extrinsics. The first reassembly cost half day; a pre-cut checkerboard fixture plus a single-launch SOP now re-runs intrinsics + extrinsics in $\sim$30 min, scheduled as task one after every assembly cycle.
- *One CAN dictionary, enforced by CI.* Independent firmware revs on the Main and Power-Distribution PCBs created subtle ID conflicts that only surfaced at integration. A single `can_dict.yaml` checked into both repos is the source of truth; CI rejects firmware emitting an undeclared ID.
- *Pre-routed labelled harnesses halve assembly.* First fully-weatherproofed run: $\sim$6 h. With pre-routed harnesses, connector ID tags, and harness diagrams shipped alongside each PCB schematic, the second took $\sim$1.5 h.

8.2 Component failures during testing

Table 11: Component failures, fix, validation.

Dom.	Failure $\rightarrow$ cause $\rightarrow$ fix	Validation
Mech	PLA+ cradle crept under cinder-block soak $\rightarrow$ low infill $\rightarrow$ infill + wall count raised.	Re-soak: no creep; PASS.
Mech	Mast deflection on single-point bracing $\rightarrow$ unsupported longitudinal moment $\rightarrow$ three-point bracing added.	IMU compare (Fig. 26); PASS.
Elec	12 V ripple under load $\rightarrow$ low bulk cap $\rightarrow$ cap added.	Scope: within encoder noise margin; 24 h soak PASS.
Elec	Encoder counts dropped $\rightarrow$ cable near motor lead $\rightarrow$ re-routed and shielded.	0 dropped on bench-and-roll; Field-load run PASS.
SW	GTSAM iSAM2 pose knocked off-course by occasional several-metre GNSS jumps under partial occlusion $\rightarrow$ standalone UM982 fix is bursty / multipath-dependent $\rightarrow$ per-fix covariance passed to the factor graph; high-σ fixes auto-deweighted, no hard cutoff.	Replay of field recorded ROS bag: pose stays smooth through the jumps; PASS.
SW	Phantom near-lanes appeared while climbing the ramp $\rightarrow$ cameras saw over the ramp's far edge; densifier projected distant ground pixels onto the road cloud $\rightarrow$ lane publication gated to positive (climbing) IMU pitch $> 8^\circ$ only.	Replay of ramp-climb bag: phantom lanes suppressed; PASS.

8.3 Software process

Two git repositories (one per challenge), branch protection, mandatory PR review, and CI gating. CI runs `ament` unit tests, static analysis, a `gtest` suite for the GTSAM iSAM2 fusion node, and a deterministic-seed Gazebo regression on every push. Tasks are tracked on Notion with P0/P1/P2 tags. ROS 2 bags attach to their tickets and replay in subsequent CI runs.

8.4 Simulation-based testing

Two **Gazebo Harmonic** worlds. The **AutoNav** world replicates the rulebook geometry (10–20 ft lanes, ramp, switchback, centre-island, 2-ft pothole, GPS waypoints); the **Self Drive** world replicates the parking-lot test track, all 16 function-test stations, the 21-step Full Course, and the SOUP/IGVC fake stop-sign set. Each world is exercised through a deterministic seed scheduler with KPI scripts capturing pass/fail per scenario.

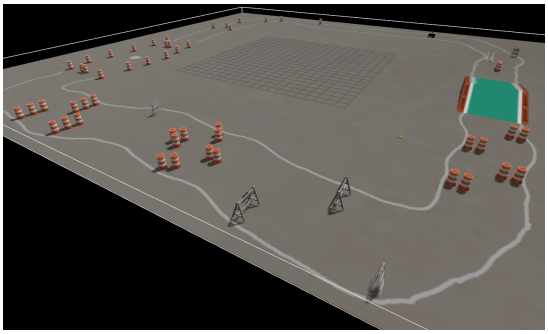

Figure 25: Gazebo Harmonic AutoNav world snapshot.

As of the snapshot, **16/16 Self Drive functions pass** (KPI S1), **AutoNav penalty count is 0/0** (KPI A2), **1 min 42 s best AutoNav lap** (KPI A1), and the Self Drive Full Course completes in ∼4 min (KPI S2).

8.5 Physical testing and sim-vs-actual deltas

Physical testing uses the team's indoor workspace for Nav2 tuning and bench rigs, and the IIT Madras **Open Air Theatre (OAT)** for GNSS characterisation and a **scaled ∼100 ft lane segment** (no full-course assets available); best autonomous traversal **23 s**, consistent with sim at the same length.

Camera-mount stability (3-point vs single-point). A high-rate IMU at the mast tip was logged under the same drive profile in both bracing configurations; the 3-point arrangement (Fig. 26) produced a markedly smoother response and is on the final vehicle.

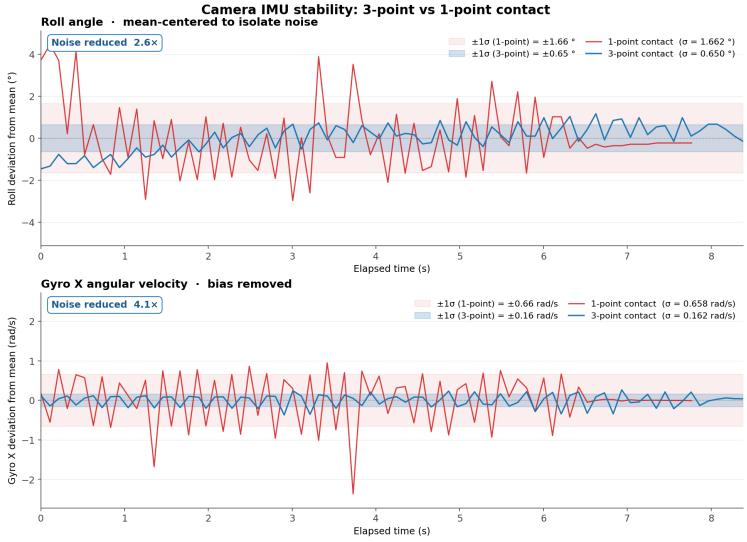

(a) Roll and Gyro X under the same drive profile.

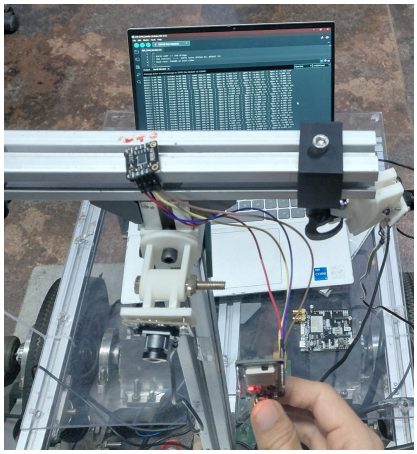

(b) Mast-tip IMU bench setup.

Figure 26: Mast bracing comparison: the 3-point arrangement smooths the IMU response versus the single-point baseline.

Closed-loop velocity tracking. The cascaded current/velocity PID was exercised end-to-end on the assembled vehicle with a joystick-driven erratic profile spanning the IGVC 1–5 mph band; `cmd_vel` and encoder-derived measured velocity were logged live in **PlotJuggler**.

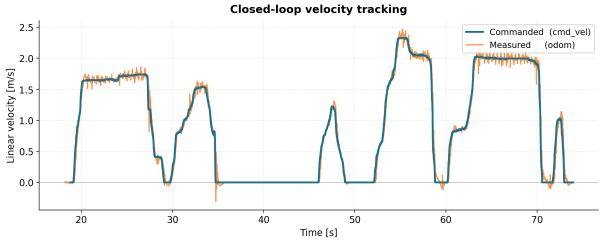

Figure 27: Commanded vs measured velocity over an erratic joystick profile.

The loop tracked with **RMS error 96.4 mm/s** (≈ 6.8 % of mean cmd vel) and **MAE 62.9 mm/s**; the ratio holds across the 1–5 mph IGVC band, so the loop does not degrade at the high end.

Sim-vs-real deltas.

- *GNSS scatter is non-Gaussian in real life.* Sim injects white-Gaussian noise; the standalone UM982 fix is bursty and multipath-dependent under partial occlusion. iSAM2 covariance deweighting handles this; the lane-anchor heading fallback (§6) catches the residual.
- *Ramp-lip jerk is harsher in field than sim.* The castor trailing arm (§2) absorbs most of the impulse.

9 Cyber Security Analysis

For a deployed Varaha, the three highest-residual-risk attack vectors are organised by where the attacker stands: **(T1)** wireless RF, **(T2)** the on-board network, and **(T3)** physical access. Tab. 12 summarises each vector and its headline defence. NIST SP 800-53 r5 control IDs (`IA-5`, `SC-7/8/12/13`, `SI-7`, `AC-3/4/17`, `PE-3`, `CM-5`) are cited as reference vocabulary.

Table 12: Threat-vector summary.

ID	Vector	Attack	Headline defence
T1	Wireless RF	Sniffs / forges / replays the LoRa wireless E-stop packet	PSK + AES-256-GCM + rolling counter
T2	On-board network	SSHs into the Orin or injects ROS 2 traffic onto the autonomy bus	Tailscale-only + key-only SSH + Zenoh mTLS + key-expression ACLs
T3	Physical access	Reads / rewrites MCU firmware, plugs a USB HID, attaches a display	Debug Authentication + Secure Boot + display blocked + `usbguard`

T1: Wireless E-stop spoofing. Plain LoRa is unauthenticated, so an adversary could sniff, forge, or replay the safety packet. Three independent firmware layers: a **hardware-rooted 256-bit key** held in a secure element on each LoRa-E5 (no flash-resident secret to extract); **AES-256-GCM** authenticated encryption (16-byte tag), so any bit-flip causes a hard reject; and a **monotonic rolling counter** inside the encrypted payload that defeats replay. An authenticated dead-man heartbeat asserts E-stop on key loss.

T2: Network attack. An adversary with access to the bot's onboard Wi-Fi could SSH into the Orin or inject malicious messages onto any ROS topic. SSH is bound to `tailscale0` (WireGuard) with key-only auth and UFW dropping every other port. ROS 2 / Zenoh runs mTLS-authenticated sessions over Tailscale only, with multicast disabled and a peer allowlist. Secure Boot + dm-verity blocks a tampered rootfs from presenting a valid mTLS identity.

T3: Physical access. The Main PCB STM32H5 uses **ST Debug Authentication**: per-unit HSM-generated keys burned into OTP at the contract manufacturer, so any SWD session must present an HSM-signed certificate before the debug interface unlocks. On the Jetson Orin, no graphical session runs (`getty` disabled), so HDMI/DisplayPort yields a blank screen; `usbguard` denies every USB HID, Ethernet, and mass-storage class, allowing only the cameras and the STM32 USB-CDC link. This means that a plugged-in keyboard is torn down before any `evdev` event reaches userspace.

SSD swap: Jetson Secure Boot verifies the bootloader, kernel, and initrd against keys burned into the Orin's OTP fuses, so a swapped SSD with unsigned firmware fails signature check and never boots. the drive is encrypted with a key sealed to the SoC's fuse-bound identity, so the original SSD cannot be read or modified off-board.

Open issue: Contactor PCB LoRa-E5. The LoRa-E5 does not support Debug Authentication, so its SWD pad is permanently unlockable, making it the weakest link in the firmware-tamper chain. The change required to secure this PCB is demoting the LoRa-E5 to a **UART-bridged dumb radio behind a second STM32H5**: AES-GCM and replay protection move into the locked H5, so flashing the E5 yields only a radio relaying ciphertext to a closed port.