RobotX 2026 Technical Design Report

Nanyang Technological University (Archimedes Autonomous Vehicles)

Alankritha Thiagarajan, Aloysius Chia, Chan Jin You, Chen Longting, Chen Yi Heng, Dinh Quang Anh, Ghosh Shoumil, Harshan Elangovan, He Zihang, Hu Tingli, Ishita Singhal, Jang Seo Jun, Kennard Garrison Johannes, Lee Yun Tung, Li Mu En, Li Tz Lung, Li Tz Yuan, Marco Louis Rudyanto, Mayank Shandilya, Murugappan Venkatesh, Nicholas Chandra, Racshanyaa Jagadish, Rafid Akil, Ryan Wong, Saanvi Khemka, Sujith Harirajan, Sushruth Muthukrishnan, Terry Teh, Vansh Jain, Vincent Gaudeo Lie, Vishnu Venkatesan, Vu Quang Hien, Weng Shizhuo, Wu Weizheng, and Yang Ruihua

Abstract— The RobotX 2026 Storm Response mission requires surface, aerial, and underwater vehicles to turn local observations into coordinated physical actions. Archimedes Autonomous Vehicles (AAV) is developing a three-vehicle system around repeatable Core tasks, with higher-tier cooperation introduced as its dependencies are verified. USV 2.0 combines camera and LiDAR perception with improved maneuverability for passage, docking, and firefighting. UAV 3.0 provides overhead cue reading and resource delivery, while UUV 1.0 performs underwater survey and repair. Each vehicle retains onboard feedback control; shared observations carry identity, capture time, and uncertainty. Protected power distribution and independent actuator isolation support safe integration. A venue-specific Gazebo environment exercises the intended autonomy and autopilot interfaces before coordinated field trials. Development has included 10 days of on-water testing. This report explains the mission-driven design choices, their principal tradeoffs, and the verification work needed to carry the system into competition.

I. Competition Strategy

A. Coordinated Mission Priorities

RobotX 2026 frames Safe Passage, Infrastructure Survey and Repair, Coordinated Logistics, and incident response as parts of a continuous port-recovery operation. The System of Systems (SoS) requirement makes cooperation a design constraint from the outset [1]. AAV assigns surface transit and manipulation to the unmanned surface vehicle (USV), overhead observation and delivery to the unmanned aerial vehicle (UAV), and submerged inspection and repair to the unmanned underwater vehicle (UUV). These roles exploit each domain's access and viewpoint without making a vehicle dependent on shore for its motion control.

Our 2024 campaign emphasized broad task exploration under limited integration time [2]. For 2026, we prioritize repeatable Core performance in the first three mission families, then introduce Advanced and Disruptive branches when their additional sensing, communication, and physical-action dependencies are ready.

The central tradeoff is between additional task capability and time available to test the complete mission. A new detector, radio path, or control mode is useful only when the downstream vehicle can act on it safely. We therefore reuse established autopilots and navigation infrastructure, concentrating custom development on maritime perception, cooperative task transitions, precision maneuvering, and electrical integration.

B. Readiness Before Complexity

The domain-specific Proof of Readiness requirements establish the first integration milestones [3]. Beyond those demonstrations, a competition capability must have a defined start condition, observable completion, cancellation behavior, and recovery path. The team develops all three vehicles in parallel, while admitting more demanding task branches only after their shared interfaces have been exercised together.

C. Evidence, Execution, and Control

Sensors produce observations, world models retain objects and constraints, executives choose objectives, and autopilots close feedback loops. A RoboBuoy detection becomes a route cue after its position, identity, course membership, and light history agree. Delivery requires aircraft acceptance within its energy reserve and existing commitments. These boundaries support independent algorithm development. Reusing the BlueBoat and BlueROV2 platforms limits hull-development work, reserving integration time for the four-thruster surface extension, aerial handoffs, and underwater task evidence. Higher-tier delivery remains conditional on these complete chains passing regression; adding a capability does not justify weakening a verified fallback.

D. Task Allocation and Cooperative Decisions

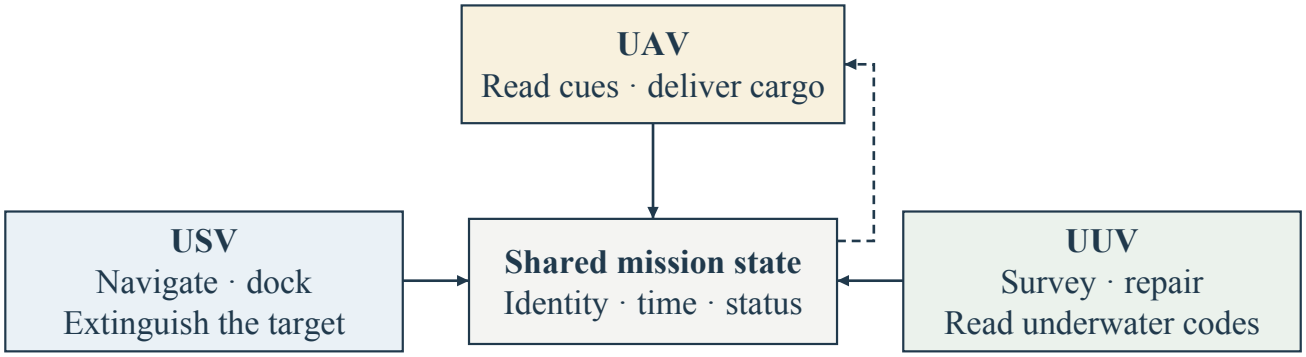


Fig. 1. Vehicle roles in the RobotX mission. Observations and requests update shared mission state; accepted requests initiate local execution.

TABLE I. Mission requirements and the AAV design response.

Mission	Core development objective	Higher-tier design consequence
Safe Passage	USV reads side beacons, completes entry/exit orbits, and follows the permitted route.	UAV reads top cues; Disruptive-tier updates can invalidate the USV route.
Infrastructure Survey	UUV reports the ordered pre-repair pipeline state; repair adds Core points.	Repair and timed code reading trigger UAV resource delivery.
Coordinated Logistics	USV docks in the safe bay, reports docking, and extinguishes the activated target.	The observed code identifies the requested aerial delivery.
Incident response	The addressed vehicle acknowledges assistance, transits, and waits for clearance.	Persistent keep-outs and moving-object separation constrain ongoing work.

Fig. 1 and Table I link the competition requirements to the capabilities we are building [1]. In Safe Passage, a buoy's physical identity persists while its light state changes. The USV must distinguish flashing and steady blue cues to execute the correct entry and exit orbits. At higher tiers, the UAV's overhead observations become part of the route decision, so freshness matters as much as classification.

Task 2 gives cooperation a concrete sequence: locate the start region, survey the pipeline in order, perform the admitted repair, observe the resulting code, and request the corresponding delivery. Visual endpoint scouting is the first search method; acoustics provides a fallback cue toward the active green endpoint. The UUV captures a light sequence locally when it occurs, rather than waiting for an acoustic acknowledgment [4].

For Task 3, the USV first establishes a dock-relative pose and reports that it is fully docked. It then shoots water at the activated target and observes completion before advancing to delivery. A waypoint arrival alone cannot establish docking, extinguishment, repair, or payload release. Each transition therefore uses evidence appropriate to the physical task.

Aircraft observation and delivery duties are scheduled explicitly because they compete for flight time and sensor coverage. Task 4 can interrupt the addressed vehicle, and its restrictions remain active until cleared. The separate later-round Distractors challenge adds vessel identification and, at its highest tier, underwater tagging. Those extensions reuse the same identity and reporting interfaces without changing the initial Core development priority.

The USV executive composes reusable navigation and action primitives with prerequisites, completion evidence, and cancellation. Nav2 supplies behavior-tree navigation; the higher-level executive uses explicit task phases. After an interruption, observed physical state determines whether an action can resume.

II. Design Strategy

A. USV 2.0: Maneuverability and Electrical Integration

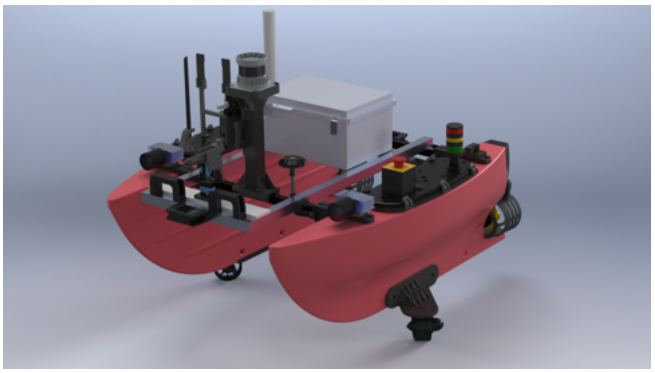


Fig. 2. USV 2.0 CAD model. The catamaran provides a common platform for surface sensing, propulsion, task hardware, and serviceable electronics.

The BlueBoat-based USV (Fig. 2) supports both open-water transit and close-range work at the docking station. Two stern thrusters provide efficient forward motion, while two diagonal T200 thrusters add the lateral force needed to align with a bay and resist disturbances during firefighting. The added complexity addresses a differential-thrust limitation: removing lateral error otherwise requires a heading change and more maneuvering room.

The revised thruster mounts use narrower ribs to reduce blockage while retaining a triangulated support. A matched computational fluid dynamics (CFD) comparison at 3 m/s predicted an approximately 11.8% reduction in mount drag and a 2.05% reduction in total modeled drag. The total change is $(67.453331 - 66.072583)/67.453331 \times 100\% = 2.05\%$. These are design-comparison results; the comparison method is summarized in Appendix G.

Sensor placement affects maneuvering. Forward cameras provide task detail and the LiDAR supplies surrounding geometry, but the hull, payload, and mast obscure parts of its view. Measured exclusion masks must retain nearby hazards during lateral motion.

1) Power, Electrical Integration, and Safety: The electrical architecture separates protected distribution, actuator isolation, and sensitive acoustic acquisition. The nominal 14.8-V battery system feeds independently protected branches. A regulated 12-V service rail supplies appropriate downstream loads; it is distinct from raw battery voltage. The power distribution board (PDB) organizes branch protection and connections for fault diagnosis. Protection is sized for the wiring and credible load envelope, including startup and simultaneous loads, rather than only nominal device ratings.

Compute, sensing, communications, logging, and safety indication remain powered when the actuator bus is isolated. Operators can therefore inspect faults and recover logs after a stop. High-current propulsion and converter returns are kept away from hydrophone signal-return paths to reduce conducted interference.

The Jetson handles perception and mission computation; ArduRover runs on a Raspberry Pi with Navigator flight controller and I/O hardware. A Pi failure can remove low-level control, so this division does not replace independent actuator isolation.

A hard-wired onboard stop and an independent handheld wireless stop control a hardware permission latch and external contactor. The isolation path removes power from propulsion and task actuators, including auxiliary feeds. Software remains responsible for orderly behavior, but cannot substitute a zero-thrust command for physical disconnection. The verification target is the handbook's wireless disconnection time of less than 2 s, including loss of transmitter power [5].

Restoring a radio connection does not automatically re-arm the vehicle. Physical stops and hard faults require local reset; remote-only recovery requires a deliberate qualified reset followed by separate arming. This separation prevents a repaired communication path from unexpectedly restarting a propeller or pump. The safety prototype uses secured modules and a soldered carrier. Power-up and brownout begin with actuator permission removed. Architecture drawings and the PDB, safety-board, and acoustic-board views are grouped in Appendices B and C.

Electrical work begins with controlled supply and output-inhibited integration. Load transitions, brownouts, loss of the wireless-stop transmitter, and reset sequences are checked before propulsion trials. Acoustic tests progress from electrical injection to wet calibration with surveyed source positions, then to operation with thrusters and communication hardware active. These tests separate electrical and acoustic effects.

B. USV Perception and Acoustic Search

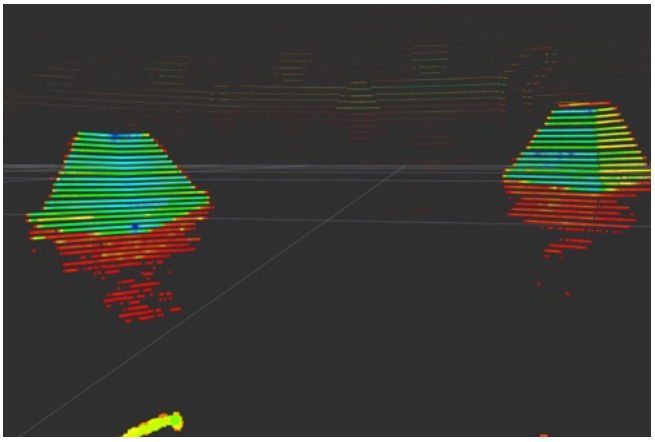


Fig. 3. LiDAR point-cloud view of RoboBuoys. The spatial returns support geometric clustering and obstacle checks; task identity and light state come from complementary visual observations.

1) YOLO Detection and Geometric Association: Maritime perception must distinguish the space a vehicle can occupy from the task meaning of what it sees. The Ouster LiDAR supplies geometric observations, while two independently calibrated Basler cameras supply color, appearance, and temporal light-state evidence. Fig. 3 shows spatial returns used by this processing path. An unclassified return remains relevant to collision avoidance.

The current camera runtime includes an Ultralytics YOLO-compatible detector that returns bounding boxes, classes, and confidence scores [6]. These are detections, not pixel-wise segmentation or metric range. The selected field design detects structures and components first, then estimates color and temporal signal states within the relevant regions. This avoids creating a new physical buoy whenever its light changes.

The selected pipeline uses PCL Euclidean clustering with a KdTree, then projects spatial clusters into each calibrated camera image for detection-to-cluster association. This associates appearance with measured geometry using calibrated sensor poses [7], [8]. The approach follows established maritime sensor-fusion practice [9], [10]. Camera overlap improves visibility, but the shared LiDAR observation is counted once; the two cameras are not assumed to form a calibrated stereo pair.

Tracks retain physical identity separately from mutable side-facing and top-facing light states. A single dark frame does not establish an OFF beacon, and a single blue frame cannot distinguish flashing entry from steady exit. Source-time observations support temporal interpretation. The mission planner uses these task objects, while a rolling geometric layer represents nearby occupied space [11]. For firefighting, a YOLO box nominates a window region; visible aperture corners and the front dock plane determine the usable opening. Background returns through the hole cannot be used as its firing range.

One USV world-model writer associates source-stamped aerial observations with surface tracks. Adjacent-course objects remain collision hazards but cannot become mission targets. Persistent semantic memory supports reacquisition; immediate clearance uses current geometry. Historical observations retain their age through communication delays and position resets.

2) Acoustics: When visual search fails to locate the Task 2 start region, passive acoustics supplies a directional cue toward the active green endpoint; it cannot identify the silent endpoint [4]. Four Aquarian AS-1 hydrophones feed a custom analog front end (AFE) with input protection, approximately 20-dB preamplification, and sixth-order low-pass filtering [12], [13]. Four USB-1808X channels are sampled simultaneously at 200 kS/s per channel, giving a 100-kHz Nyquist frequency [14].

The receiver covers the handbook's 25–40-kHz pinger band [15]. Burst detection and physically bounded generalized cross-correlation provide candidate time differences of arrival [16]. Calibration and geometric consistency checks are needed to reject ambiguous tonal and reflected arrivals. A bearing is associated with the vehicle pose at capture time; several separated observations can constrain the start region. A single bearing does not determine range. Appendix C-B summarizes the electrical evidence and calibration work.

C. Communications and Operator Supervision

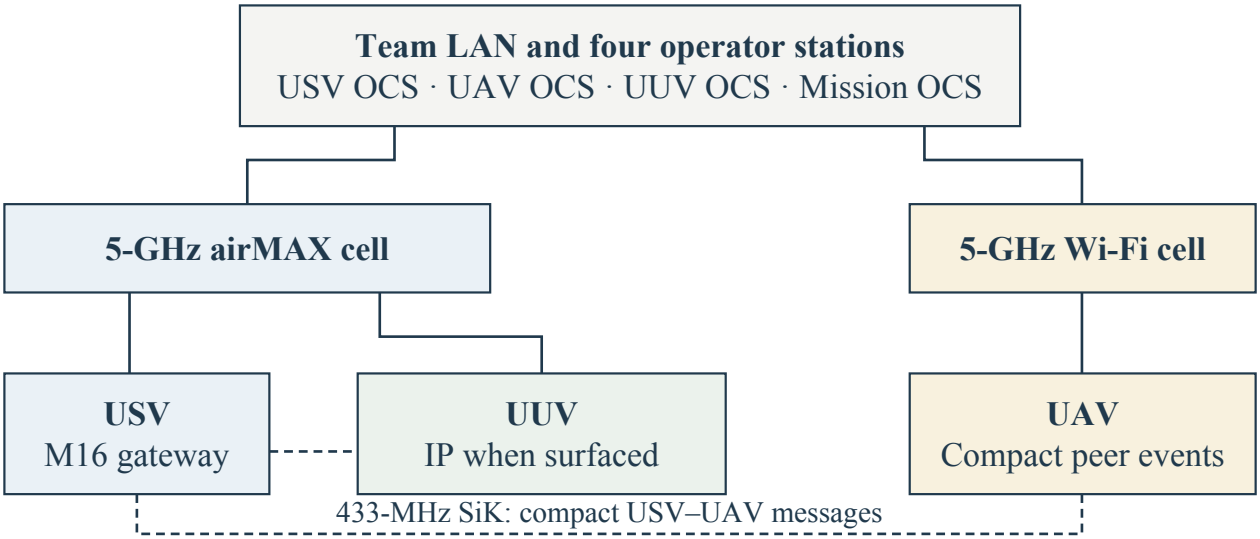


Fig. 4. Fleet communication paths. Solid lines denote IP connectivity; dashed lines denote compact peer links. The short USV–UUV dashed link is acoustic. Manual control and emergency-stop paths remain independent.

The communication design reflects the information each mission needs. A complete buoy map or camera stream requires an IP link, whereas a delivery request or repair event can be represented as a compact message. Fig. 4 separates these uses so that low-bandwidth links are not burdened with raw sensor data.

The USV and surfaced UUV are independent stations on a shore-based 5-GHz airMAX cell. The UAV uses a separate standard 5-GHz Wi-Fi cell connected to the same team LAN. This arrangement accommodates the vehicles' different radio interfaces while preserving a common route for mission products, telemetry, GNSS corrections, and selected operator video. Bulk logs and raw sensor streams are retained onboard unless a transfer is needed.

A paired 433-MHz SiK link carries admitted compact transactions between the USV and UAV [17]. It supports cooperation when a small event is sufficient, but does not replace IP delivery of full buoy fields or frequent aerial state updates. Underwater, an M16 Standard modem on the UUV communicates with an M16 Extended gateway on the USV. The gateway forwards selected events while preserving their origin and observation time.

The distinction between receiving a message and completing a task is central to reliable cooperation. For example, receiving a resource request confirms delivery of the request, local acceptance confirms that the UAV can undertake it, and observed payload release establishes the physical outcome. Retries preserve the original event identity so that a delayed duplicate cannot trigger a second irreversible action. Freshness limits are evaluated against capture time rather than the time of retransmission.

Four shore laptops provide vehicle-specific operator control stations (OCSs) and a fleet-level Mission OCS. Operators see actual mode, navigation status, link health, and report state. Unavailable or stale information is displayed as such, allowing a vehicle fault to be distinguished from a broken display path.

Mission OCS alone translates eligible reports into the organizer's MQTT/Protocol Buffers interface [18]. A separate Ethernet interface reaches RoboCommand without bridging the organizer network into the team LAN. Peer exchange does not require a Mission OCS application relay, and vehicle feedback loops remain onboard. This division keeps official reporting consistent while limiting the effect of a shore application failure. Appendix E gives the supplementary information-flow view and the communication failure cases used in testing.

D. Navigation and Control

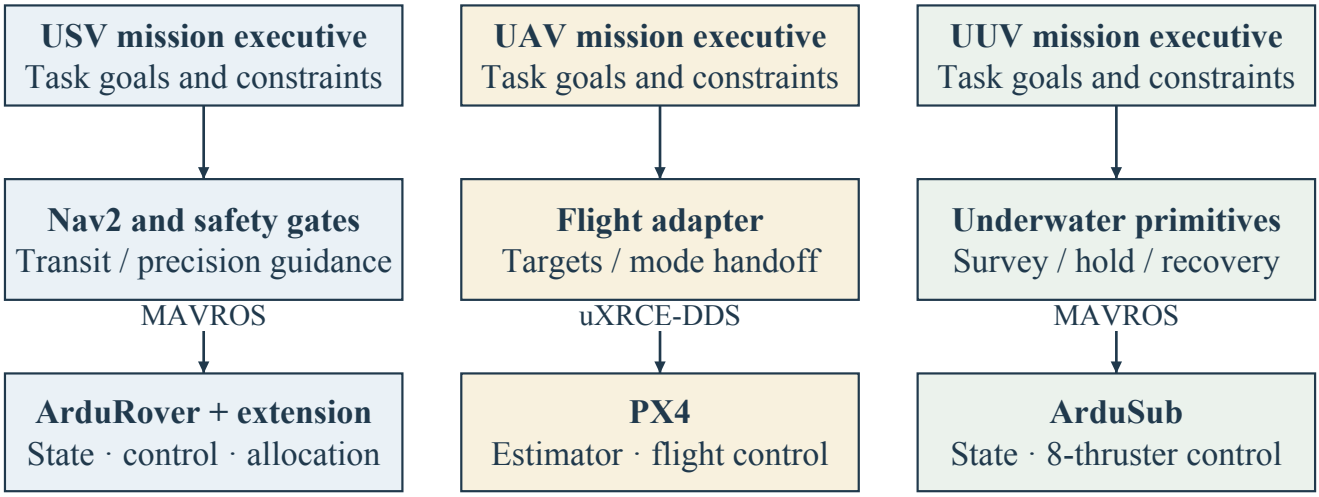


Fig. 5. Selected navigation and control architecture. Mission software supplies guarded objectives; each autopilot retains local feedback and actuator ownership. The USV precision-control extension is under development.

The fleet shares mission information but uses vehicle-specific control interfaces (Fig. 5). This allows navigation algorithms to change without duplicating the autopilot's inner loops. ROS 2 provides the integration layer for observations, objectives, and status [19]; a common framework does not imply that every vehicle uses the same estimator or controller.

1) Surface Transit and Precision Work: On the USV, GNSS position and velocity, moving-baseline GNSS yaw, and inertial measurements feed ArduRover EKF3. The two GNSS receivers form a coupled heading subsystem rather than two independent poses to average [20]. Shore corrections improve global positioning, while the local moving baseline supports heading at low speed. Docking additionally uses fresh relative geometry because a globally correct pose alone does not establish alignment with the bay.

SmacPlanner2D generates transit routes, and Regulated Pure Pursuit follows the accepted path using surge and yaw-rate references [21], [22]. Ordered gate crossings and entry/exit orbits are expressed as navigation objectives. Precision guidance adds sway for docking, station keeping, and firefighting, avoiding repeated turns merely to remove lateral error.

Both profiles pass through one command owner, motion limits, collision monitoring, and an authority gate before MAVROS forwards references to the autopilot [23]. The planned ArduRover extension provides native planar velocity control and bounded allocation for the two stern and two diagonal thrusters. This is an engineering development item; stock differential-drive control is not assumed to provide the required three-axis behavior. Thrust saturation and the resulting unmet force will be considered when admitting a maneuver [24].

2) Aerial and Underwater Control: The UAV's navigation adapter manages position and velocity targets and PX4 mode transitions through uXRCE-DDS [25]. It provides one autonomous command owner and explicit handoff between native flight modes and offboard operation. Pilot takeover and autopilot failsafes remain available.

The UUV design routes guarded survey, hold, and recovery primitives through MAVROS to ArduSub. ArduSub retains its estimator, inner loops, and eight-thruster allocation. A peer-link failure does not transfer feedback control to shore. Continuation, holding, or recovery depends on local navigation, energy, clearance, and the admitted task state.

Across all domains, a controller handover revokes the previous owner's command. A restored link cannot cancel a manual takeover or clear a hardware stop. Course boundaries, current collision geometry, and Task 4 restrictions are checked as distinct constraints, so an old map cannot certify that the present path is clear.

E. UAV 3.0: Overhead Observation and Delivery



Fig. 6. UAV 3.0, Becca, on the Holybro X650 airframe. The custom electronics enclosure and downward-looking sensor arrangement support observation, delivery, and recovery over water.

The UAV reads top-facing cues and delivers resources requested after underwater repair or surface firefighting. These duties determine the design: observations must be geolocated for another vehicle, the correct payload must reach the correct platform, and sufficient energy must remain for recovery.

Becca (Fig. 6) combines a Holybro X650 frame, Pixhawk 6X, Jetson Orin NX, and downward-looking ZED 2i. A custom enclosure integrates electronics placement, cable routing, and mounting. Selected water-resistant propulsion components reduce exposure risk. The design includes a recovery handle and a water-activated locating light; recovery is performed with propulsion disarmed. This maritime packaging extends the team's previous aerial-development experience [26].

1) Perception for Cooperative Tasks: Learned detection identifies candidate buoys, platforms, and resource tins. Calibrated geometry and temporal color observations then determine the information needed by the mission. Pixel coordinates are converted using camera calibration, vehicle attitude, and the relevant observed surface; proximity in an image alone cannot establish correspondence with a surface vehicle's map [8].

The platform's known planar structure supports pose estimation [1]. The estimator must therefore use measured geometry and its uncertainty. Relative circle sizes alone do not uniquely determine platform tilt. A smaller central downward-looking camera is planned to preserve target visibility during the final descent when the main camera's view may be obstructed.

Stereo depth is used only within its validated range and confidence; reflective water or an unseen surface does not provide a reliable depth constraint. Each aerial observation includes capture time, calibrated pose, uncertainty, and coverage. The USV performs the cross-vehicle identity association. It does not interpret a bare color or an aircraft-local track number as a globally unique buoy.

2) Flight and Payload Handoff: The navigation adapter separates mission objectives from PX4 messages and mode details. Native takeoff, landing, and return behavior is coordinated with offboard position or velocity control. Offboard health signaling is maintained separately from the target stream, following PX4's mode requirements [27]. Mode changes clear obsolete targets before a new command owner is admitted.

Geographic targets also require a qualified conversion into the current PX4 local frame [25]. The adapter tracks the origin and reset state and converts coordinate conventions once. A target remaining numerically unchanged after an estimator reset may refer to a different physical location, so affected objectives are revalidated.

Delivery requires both the correct request and evidence of the physical payload outcome. The UAV must also arbitrate between continued cue observation and a delivery commitment; accepting one can remove coverage needed by another vehicle. Appendix F illustrates the perception, control, and avionics concepts. Health checks, descent visibility, and recovery remain part of physical qualification.

The resource tins are aluminum, and the gripping mechanism remains under development. The design requires observed attachment, retention, and release, rather than assuming magnetic pickup. The aircraft may queue or reject a valid request when reserve or an existing observation commitment prevents execution. For continuing aerial cue support, a diversion requires a handoff or suspension of the dependent surface maneuver.

F. UUV 1.0: Underwater Survey and Repair

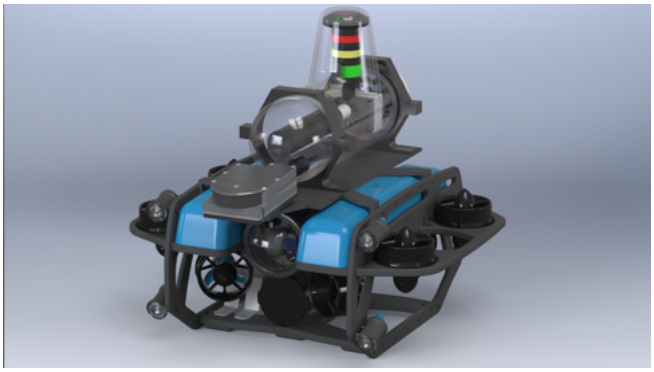


Fig. 7. UUV 1.0 CAD model, derived from the BlueROV2 Heavy. The sensor and payload arrangement supports pipeline inspection while preserving the eight-thruster control architecture.

Task 2 requires an ordered inspection and a distinction between initial pipeline condition and repair outcome. AAV selected the BlueROV2 Heavy (Fig. 7) for low-speed translation and attitude control near infrastructure. Companion computing and untethered communication extend the platform while ArduSub retains local control.

The selected sensing package combines a Water Linked A50 Doppler velocity log (DVL) [28], pressure-based depth sensing, a Sonoptix ECHO imaging sonar, a camera, and four Lumen lights. Bottom-track velocity constrains motion, pressure constrains depth, sonar supplies structural range and azimuth, and images identify targets and light sequences.

1) One Navigation Measurement Path: The selected BlueOS-side DVL interface validates velocity, quality, and timing before injecting the measurement once into ArduSub's estimator. The Jetson Orin NX consumes the accepted state through MAVROS rather than injecting a second copy of the same observation. This avoids double-counting one sensor.

The measured velocity must refer to the vehicle frame and the same physical reference point used by navigation. Mounting rotation and the rotational velocity at the DVL's offset are accounted for once [24]. A valid bottom-track solution is required; an old velocity packet or water-relative return cannot be relabeled as current bottom-relative motion.

Surface GNSS provides an initial anchor and a recovery reference; it does not provide continuous absolute position underwater. Bottom-track loss, poor geometry, or a degraded depth source therefore changes which motion primitives can be admitted. The DVL requires submerged acoustic faces and unobstructed beams; the simulation checks these mounting constraints.

2) Task Evidence and Recovery: The selected sonar pipeline suppresses range-dependent background, applies thresholding and morphology, and groups returns into structural candidates. Robust local line hypotheses and multi-frame association form a connected pipeline representation. Following that structure supports reacquisition at bends and prevents repeated views from being counted as new segments.

The imaging fan supplies range and azimuth, with unresolved elevation [29]. Camera observations add identity and temporal light evidence after in-water calibration and time alignment. A missing echo does not establish the physical terminus, and a forward-looking fan cannot certify overhead ascent clearance.

The task design records the ordered pre-repair survey before performing an admitted repair. At higher tiers, it captures the post-repair code during the finite signaling window and forms a delivery request with the original capture time [4]. A delayed acoustic acknowledgment cannot change the observed color order or make an expired observation current. Repair contact and successful repair feedback are separate from reaching a commanded pose.

The M16 link carries selected events through the USV gateway, while full IP access is available when surfaced. Loss of the acoustic link alone does not command an automatic ascent: the safe action depends on local clearance, navigation health, energy, and task state. Recovery is treated as a controlled vehicle behavior, including the possibility that a direct ascent is obstructed.

III. Testing Strategy

A. Gazebo as a Mission Development Environment

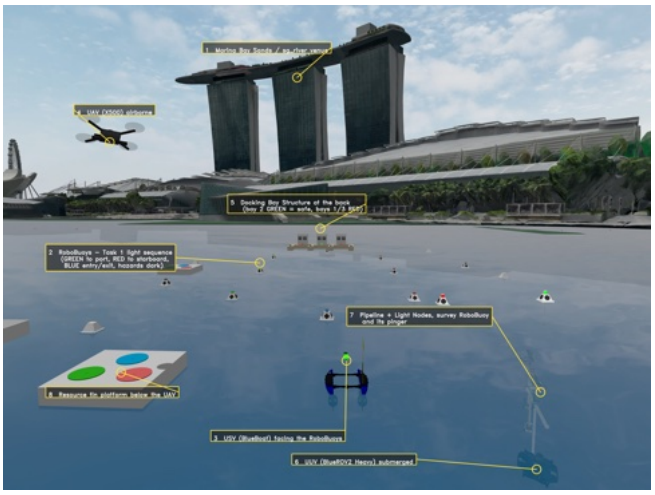


Fig. 8. RobotX 2026 Gazebo environment under development, showing the Marina Bay scene, course objects, and three vehicle domains. The X500-class UAV is a simulation proxy for Becca.

We are building the RobotX 2026 environment in Gazebo Harmonic (Fig. 8) so that navigation, sensing, cooperation, and failure recovery can be exercised before scarce water sessions. Five generated worlds cover Proof of Readiness, individual Tasks 1–3, and their combined preliminary course. Selectable Core, Advanced, and Disruptive configurations change cue behavior and clutter. The aim is to expose the same constraints that shape a real mission, not merely to display three moving vehicles.

The simulator runs separately from vehicle autonomy. The target production profile connects the autonomy stacks over ROS 2 to ArduRover and ArduSub software-in-the-loop (SITL), with PX4 providing the aerial firmware path [30], [31]. Surface commands traverse the mission executive, navigation, safety gates, MAVROS, MAVLink, and firmware before reaching simulated thrusters. End-to-end validation checks command ownership and excludes simulator ground truth from production autonomy; evaluators alone may use it for scoring.

Each scenario begins with a nominal run, then changes a defined disturbance or fault while retaining the same task objective. Camera tasks must use rendered imagery through the detector, not simulator-provided object identities. LiDAR obstacles enter the geometric safety path. Biased or unavailable DVL measurements must propagate through the underwater estimator and its health checks. This tests whether the complete system responds to degraded evidence, rather than whether an isolated planner follows an ideal route.

The world is anchored at The Promontory, with Marina Bay represented as freshwater [32]. Environmental configuration draws on local climate information and the World Magnetic Model [33], [34]. Generated course dimensions are tied to published build guides; unspecified values remain labeled approximations. Hydrodynamics, wave forcing, camera visibility, LiDAR returns, underwater sensing, and energy loads are varied to challenge the design. Gazebo's physics interfaces and the VRX wave implementation provide the simulation foundation [35], [36].

Sensor models include camera calibration, LiDAR sampling and range noise, scene-dependent DVL availability, depth, sonar, lighting, and the pinger's propagation geometry. Weather cases vary waves, rain, and visibility; acoustic cases add noise and multipath. These cases expose mounting occlusion, intermittent observations, and conflicting task cues before hardware trials. Their parameters are model assumptions to be checked against measurements, not claims of exact Marina Bay conditions.

Independent evaluators check mission behavior: a gate must be crossed through its opening in the correct order and direction; a survey needs suitable range, viewing direction, and speed; a landing is evaluated relative to deck motion. Evaluators observe outcomes without commanding the vehicles. Recorded software versions, scenario settings, seeds, and simulation time support repeatable comparisons, with replay consistency checked rather than assumed across all hardware.

Simulation informs physical design through propulsion-clearance checks, sensor visibility, DVL immersion, buoy draft, and energy budgeting. Numerical limitations, including an unstable added-mass formulation, prompt model review. Appendix H describes the model boundaries. Bench and field measurements establish their correspondence with the deployed fleet.

Model comparisons guide the mount and AFE designs. Physical correlation and regression checks follow changes to geometry, circuitry, or firmware. Test time is reserved in dependency order: components, single vehicles, then cooperative rehearsals. Failed acceptance cases take priority over a new task tier. Appendix I defines the retained evidence and release gates.

B. Physical Testing, Regression, and Competition Readiness

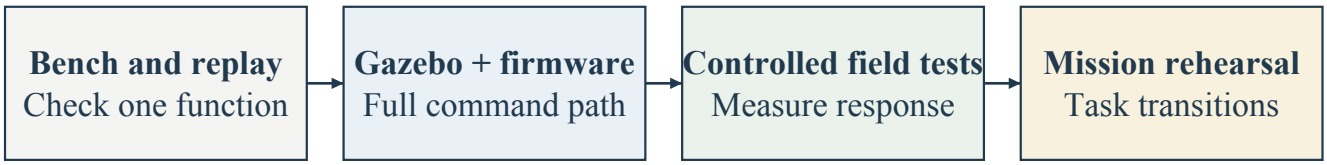


Fig. 9. Verification progression. Component checks and simulation precede controlled field tests; observed discrepancies feed back into the models and the next mission rehearsal.

TABLE II. Tests that connect design decisions to mission readiness.

Area	Verification focus	Decision informed
Electrical and safety	Rail behavior, load transitions, stop timing, reset, interference, and retained service power.	Whether the assembled vehicle can be energized and recovered safely.
Perception and navigation	Recorded-data replay, frame/time checks, obstructed sensing, route tracking, and relative approach.	Which observations and maneuvers can be admitted.
Cooperation	Loss, delay, duplication, stale observations, and cancellation during physical actions.	Whether a task transition remains correct under link degradation.
Integrated missions	Ordered gates, survey and repair evidence, docking, firefighting, and delivery handoffs.	Which task tier is ready for a full rehearsal.

The team has completed **10 days of on-water testing**. Further testing is organized around the dependencies in Table II, following the progression in Fig. 9. Each session addresses a defined question and retains its configuration, conditions, outcome, and logs.

Navigation trials progress from manual operation and basic motion to ordered routes and close approaches. Docking and firefighting add relative geometry and disturbance rejection; underwater trials add bottom-track loss and clearance-aware recovery. Aerial tests establish mode handoff, pilot takeover, final descent visibility, and payload handling before delivery is coupled to another vehicle's request.

Our contribution is the integration of time-aware task evidence with bounded local control across three domains. The design builds on established sensor fusion, navigation, and marine-craft dynamics; its research questions concern the reliability gained from separating object identity, changing task cues, and physical completion evidence [9], [22], [24].

Regression scenarios retain nominal and degraded cases. A repaired fault becomes a repeatable test before new capability is admitted. Perception data are partitioned by acquisition session, keeping paired views together to reduce train/test leakage. Field time addresses sealing, thermal rise, interference, optical contrast, and acoustic multipath [3].

ACKNOWLEDGMENTS

Team Archimedes' completion of USV 2.0, UAV 3.0 and UUV 1.0 would not be possible without the help from various organizations and people: Associate Professor Xie Ming – for his supervision of the team and technical advice on autonomous systems and machine vision. The team would also like to thank Dawn Tan for producing, directing, and writing our video; Megan Tan and Dix Ong for videography, lighting, and sound; Megan Tan for photography; Teoh Sim Kee for design assets; and Jaime Aw Yong for editing. Team Archimedes expresses our appreciation to our sponsors listed in Appendix A for their continuous support.

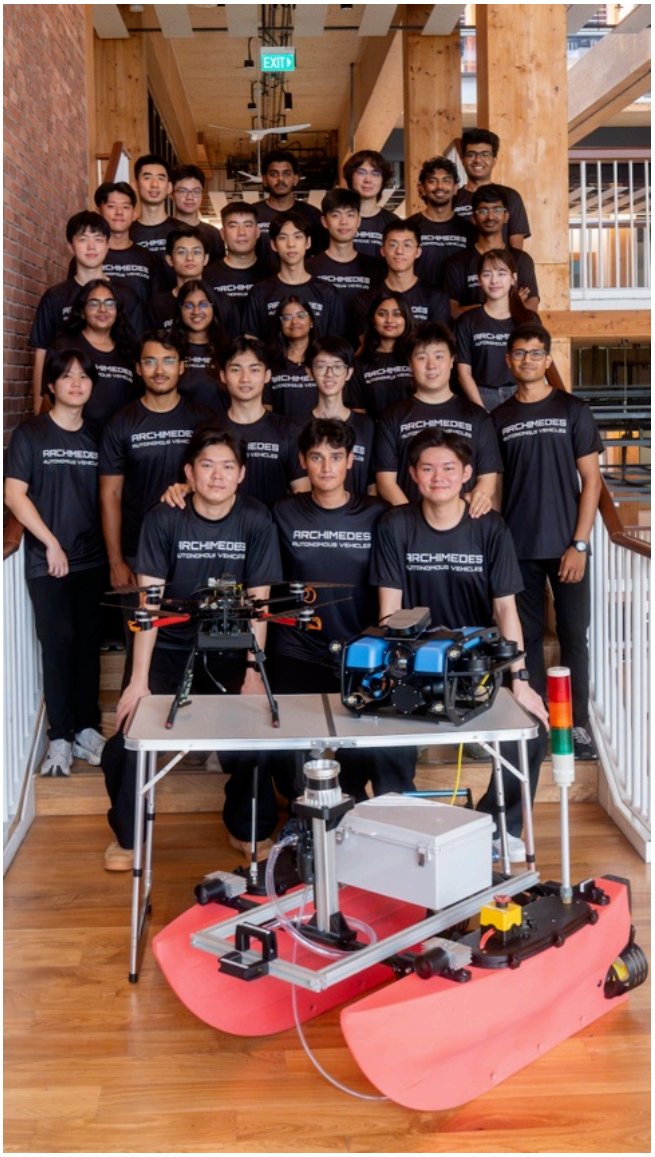


Fig. 10. Team Archimedes.

REFERENCES

[1] RoboNation, "RobotX 2026 Team Handbook," https://robonation.gitbook.io/robotx-2026-team-handbook, 2026, accessed: Sep. 28, 2026.

[2] H. W. Aw *et al.*, "RobotX 2024 Technical Design Paper," Nanyang Technological University, Archimedes Autonomous Vehicles, Tech. Rep., 2024, available: https://www.aav.sg/docs/robotx-2024-tdr.pdf.

[3] RoboNation, "RobotX 2026 Team Handbook, Sec. 3.1: Proof of Readiness," https://robonation.gitbook.io/robotx-2026-team-handbook/section-3-autonomy-challenge/3.1-proof-of-readiness, accessed: Sep. 28, 2026.

[4] ——, "RobotX 2026 Team Handbook, Sec. 3.3.3: Mission Task 2—Infrastructure Survey and Repair," https://robonation.gitbook.io/robotx-2026-team-handbook/section-3-autonomy-challenge/3.3-mission-task-descriptions/3.3.3-mission-task-2-infrastructure-survey-and-repair, 2026, accessed: Sep. 28, 2026.

[5] ——, "RobotX 2026 Team Handbook, Sec. 5.2: Safety Requirements," https://robonation.gitbook.io/robotx-2026-team-handbook/section-5-rules-and-requirements/5.2-safety-requirements, 2026, accessed: Sep. 28, 2026.

[6] Ultralytics, "Object detection with Ultralytics YOLO," https://docs.ultralytics.com/tasks/detect/, accessed: Sep. 29, 2026.

[7] R. B. Rusu and S. Cousins, "3D is here: Point Cloud Library (PCL)," in *2011 IEEE International Conference on Robotics and Automation*. Shanghai, China: IEEE, May 2011, pp. 1–4, available: https://doi.org/10.1109/ICRA.2011.5980567.

[8] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge University Press, 2004, available: https://www.cambridge.org/9780521540513.

[9] L. Stanislas and M. Dunbabin, "Multimodal Sensor Fusion for Robust Obstacle Detection and Classification in the Maritime RobotX Challenge," *IEEE Journal of Oceanic Engineering*, vol. 44, no. 2, pp. 343–351, 2019, available: https://doi.org/10.1109/JOE.2018.2868488.

[10] T. Clunie, M. DeFilippo, M. Sacarny, and P. Robinette, "Development of a Perception System for an Autonomous Surface Vehicle using Monocular Camera, LIDAR, and Marine RADAR," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. Xi'an, China: IEEE, May 2021, pp. 14 112–14 119, available: https://doi.org/10.1109/ICRA48506.2021.9561275.

[11] D. V. Lu, D. Hershberger, and W. D. Smart, "Layered Costmaps for Context-Sensitive Navigation," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Chicago, IL, USA: IEEE, Sep. 2014, pp. 709–715, available: https://doi.org/10.1109/IROS.2014.6942636.

[12] Aquarian Audio & Scientific, *AS-1 Scientific Broadband Measurement Hydrophone: User Manual*, available: https://www.aquarianaudio.com/AqAudDocs/AS-1_manual.pdf, accessed: Sep. 28, 2026.

[13] Texas Instruments, *OPA165x Ultra-Low-Noise, Low-Distortion, FET-Input, Burr-Brown Audio Operational Amplifiers*, Texas Instruments, Sep. 2022, rev. C. Available: https://www.ti.com/lit/ds/symlink/opa1656.pdf, accessed: Sep. 28, 2026.

[14] Measurement Computing, *USB-1808 Series: High-Speed, High-Precision, Simultaneous USB Devices*, Measurement Computing, available: https://files.digilent.com/datasheets/DS-USB-1808-Series.pdf, accessed: Sep. 28, 2026.

[15] RoboNation, "RobotX 2026 Team Handbook, Sec. 3.5.4: Acoustic Pingers," https://robonation.gitbook.io/robotx-2026-team-handbook/section-3-autonomy-challenge/3.5-task-build-specifications, 2026, accessed: Sep. 28, 2026.

[16] C. H. Knapp and G. C. Carter, "The Generalized Correlation Method for Estimation of Time Delay," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 24, no. 4, pp. 320–327, Aug. 1976, available: https://doi.org/10.1109/TASSP.1976.1162830.

[17] ArduPilot Development Team, "SiK Radio: Advanced Configuration," https://ardupilot.org/copter/docs/common-3dr-radio-advanced-configuration-and-technical-information.html, accessed: Sep. 28, 2026.

[18] RoboNation, "RobotX 2026 Team Handbook, Sec. 3.4: Communications Protocol," https://robonation.gitbook.io/robotx-2026-team-handbook/section-3-autonomy-challenge/3.4-communications-protocol, accessed: Sep. 28, 2026.

[19] S. Macenski, T. Foote, B. Gerkey, C. Lalancette, and W. Woodall, "Robot Operating System 2: Design, Architecture, and Uses in the Wild," *Science Robotics*, vol. 7, no. 66, p. eabm6074, May 2022, available: https://doi.org/10.1126/scirobotics.abm6074.

[20] ArduPilot Development Team, "GPS for Yaw (Moving Baseline), Rover Documentation," https://ardupilot.org/rover/docs/common-gps-for-yaw.html, accessed: Sep. 28, 2026.

[21] Nav2 Development Team, "Smac 2D Planner, Jazzy Documentation," https://docs.nav2.org/jazzy/configuration_and_development/configuration_guide/planners_plugins/smac/smac_2d/configuring_smac_2d/, accessed: Sep. 28, 2026.

[22] S. Macenski, S. Singh, F. Martín, and J. Ginés, "Regulated Pure Pursuit for Robot Path Tracking," *Autonomous Robots*, vol. 47, no. 6, pp. 685–694, Aug. 2023, available: https://doi.org/10.1007/s10514-023-10097-6.

[23] Nav2 Development Team, "Collision Monitor Node, Jazzy Documentation," https://docs.nav2.org/jazzy/configuration_and_development/configuration_guide/core_servers/collision_monitor/configuring_collision_monitor_node/, accessed: Sep. 28, 2026.

[24] T. I. Fossen, *Handbook of Marine Craft Hydrodynamics and Motion Control*, 1st ed. John Wiley & Sons, Ltd., 2011, available: https://doi.org/10.1002/9781119994138.

[25] PX4 Development Team, "ROS 2 User Guide, PX4 v1.15," https://docs.px4.io/v1.15/en/ros2/user_guide, accessed: Sep. 28, 2026.

[26] J. Zhao, D. Suen, T. Y. Li, T. L. Li, A. Chia, and T. Hu, "SUAS 2025 Technical Design Report," Nanyang Technological University, Archimedes Autonomous Vehicles, Tech. Rep., 2025, available: https://www.aav.sg/docs/suas-2025-tdr.pdf.

[27] PX4 Development Team, "Offboard Mode, PX4 User Guide v1.15," https://docs.px4.io/v1.15/en/flight_modes/offboard, accessed: Sep. 28, 2026.

[28] Water Linked, "Water Linked DVL API," GitHub repository, available: https://github.com/waterlinked/waterlinked_dvl, accessed: Sep. 28, 2026.

[29] Blue Robotics, "Sonoptix ECHO Multibeam Imaging Sonar," https://bluerobotics.com/store/sonars/imaging-sonars/sonoptix-echo/, accessed: Sep. 29, 2026.

[30] ArduPilot Development Team, "Using SITL with Gazebo," https://ardupilot.org/dev/docs/sitl-with-gazebo.html, accessed: Sep. 28, 2026.

[31] PX4 Development Team, "Gazebo Simulation, PX4 User Guide v1.15," https://docs.px4.io/v1.15/en/sim_gazebo_gz/, accessed: Sep. 28, 2026.

[32] PUB, Singapore's National Water Agency, "The Urban Reservoir Experience," https://www.pub.gov.sg/-/media/Images/Feature/Content-Pages/Resources/Publications/Pdfs/MB-Marketing-DL_2210.pdf, accessed: Sep. 28, 2026.

[33] Meteorological Service Singapore, "Climate of Singapore," https://www.weather.gov.sg/climate-climate-of-singapore/, accessed: Sep. 28, 2026.

[34] NOAA National Centers for Environmental Information, "World Magnetic Model (WMM)," https://www.ncei.noaa.gov/products/world-magnetic-model, accessed: Sep. 28, 2026.

[35] Gazebo Development Team, "Gazebo Sim 8: Physics Engines," https://gazebosim.org/api/sim/8/physics.html, accessed: Sep. 28, 2026.

[36] Virtual RobotX Development Team, "VRX Wavefield Implementation," https://github.com/osrf/vrx/blob/main/vrx_gz/src/Wavefield.cc, accessed: Sep. 28, 2026.

[37] Measurement Computing, *USB-1808X User's Guide*, available: https://files.digilent.com/manuals/USB-1808X.pdf, accessed: Sep. 28, 2026.

[38] Analog Devices, *LT3042: 20V, 200mA, Ultralow Noise, Ultrahigh PSRR RF Linear Regulator*, Analog Devices, Nov. 2017, rev. B. Available: https://www.analog.com/media/en/technical-documentation/data-sheets/3042fb.pdf, accessed: Sep. 28, 2026.

[39] ——, *LT3094: -20V, 500mA, Ultralow Noise, Ultrahigh PSRR Negative Linear Regulator*, Analog Devices, Jan. 2020, rev. B. Available: https://www.analog.com/media/en/technical-documentation/data-sheets/LT3094.pdf, accessed: Sep. 28, 2026.

[40] Blue Robotics, "T200 Thruster: Performance Charts and Public Performance Data," https://bluerobotics.com/store/thrusters/t100-t200-thrusters/t200-thruster-r2-rp/, accessed: Sep. 28, 2026.

[41] ——, "Technical Reference: Products and Software," https://bluerobotics.com/learn/technical-reference/, accessed: Sep. 29, 2026.

[42] Compulab, "EdgeAI-ORN Hardware Reference Guide," https://cdn.edge.compulab.com/EdgeAI-ORN/documents/EdgeAI-ORN_Hardware_Reference_Guide.pdf, accessed: Sep. 29, 2026.

[43] Holybro, "H-RTK ZED-F9P Rover," https://holybro.com/products/h-rtk-zed-f9p-rover, accessed: Sep. 29, 2026.

[44] Basler, "a2A2448-23gcIP67: Product Documentation," https://docs.baslerweb.com/a2a2448-23gcip67, accessed: Sep. 29, 2026.

[45] Ubiquiti, "airMAX Rocket AC Lite Technical Specifications," https://techspecs.ui.com/uisp/wireless/r5ac-lite, accessed: Sep. 29, 2026.

[46] ——, "AMO-5G10 Antenna Technical Specifications," https://techspecs.ui.com/uisp/wireless/amo-5g10, accessed: Sep. 29, 2026.

[47] CubePilot, "Herelink Overview," https://docs.cubepilot.org/user-guides/herelink/herelink-overview, accessed: Sep. 29, 2026.

[48] Holybro, "SiK Telemetry Radio V3," https://docs.holybro.com/radio/sik-telemetry-radio-v3, accessed: Sep. 29, 2026.

[49] Water Linked, "Modem M16: Technical Documentation," https://docs.waterlinked.com/modem-m16/modem-m16/, accessed: Sep. 29, 2026.

[50] Holybro, "X650 Kits: Frame Specifications," https://holybro.com/products/x650-kits, accessed: Sep. 29, 2026.

[51] T-Motor, "MN4110 Motor: Specifications and Matching Guide," https://store.tmotor.com/product/mn4110-motor-navigator-type.html, accessed: Sep. 29, 2026.

[52] ——, "P17*5.8 Carbon-Fiber Propellers," https://store.tmotor.com/product/polish-carbon-fiber-17x5_8-prop.html, accessed: Sep. 29, 2026.

[53] ——, "AIR 40A 2–6S No-BEC ESC," https://store.tmotor.com/product/air-40a-6s-esc.html, accessed: Sep. 29, 2026.

[54] Holybro, "PM02D Power Module," https://holybro.com/products/pm02d-power-module, accessed: Sep. 29, 2026.

[55] ——, "Pixhawk 6X Technical Specification," https://docs.holybro.com/autopilot/pixhawk-6x/technical-specification, accessed: Sep. 29, 2026.

[56] Seeed Studio, "reComputer J4012: Orin NX 16 GB," https://www.seeedstudio.com/reComputer-J4012-w-o-power-adapter-p-5628.html, accessed: Sep. 29, 2026.

[57] Holybro, "H-RTK ZED-F9P Ultralight," https://holybro.com/products/h-rtk-f9p-ultralight, accessed: Sep. 29, 2026.

[58] Stereolabs, "ZED Camera Depth Ranges and Lens Configurations," https://support.stereolabs.com/articles/9175329825-for-the-depth-sensing-what-is-the-range-for-the-zed-cameras-in-terms-of-measuring-distances-and-the-accuracy-of-the-detection, accessed: Sep. 29, 2026.

[59] Intel, "Wi-Fi 6E AX210 Specifications," https://www.intel.com/content/www/us/en/products/sku/204836/intel-wifi-6e-ax210-gig/specifications.html, accessed: Sep. 29, 2026.

[60] Molex, "204281 Series Cabled Antennas," https://www.molex.com/en-us/products/series-chart/204281, accessed: Sep. 29, 2026.

[61] RadioMaster, "RP4TD ExpressLRS 2.4 GHz True Diversity Receiver," https://www.radiomasterrc.com/products/rp4td-expresslrs-2-4ghz-diversity-receiver, accessed: Sep. 29, 2026.

[62] Blue Robotics, "BlueROV2: Technical Details," https://bluerobotics.com/store/rov/bluerov2/, accessed: Sep. 29, 2026.

[63] Seeed Studio, "Getting Started with reComputer Mini," https://wiki.seeedstudio.com/recomputer_jetson_mini_getting_started/, accessed: Sep. 29, 2026.

[64] Water Linked, "DVL A50: Technical Documentation and Datasheet," https://docs.waterlinked.com/dvl/dvl-a50/, accessed: Sep. 29, 2026.

[65] Ubiquiti, "airMAX Bullet AC IP67 Technical Specifications," https://techspecs.ui.com/uisp/wireless/bulletac-ip67, accessed: Sep. 29, 2026.

[66] Phoenix Contact, "ANT-OMNI-2459-04 Antenna, 1284780," https://www.phoenixcontact.com/en-pc/products/antenna-ant-omni-2459-04-1284780, accessed: Sep. 29, 2026.

[67] F. R. Menter, "Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications," *AIAA Journal*, vol. 32, no. 8, pp. 1598–1605, 1994, available: https://doi.org/10.2514/3.12149.

[68] J. P. Van Doormaal and G. D. Raithby, "Enhancements of the SIMPLE Method for Predicting Incompressible Fluid Flows," *Numerical Heat Transfer*, vol. 7, no. 2, pp. 147–163, 1984, available: https://doi.org/10.1080/01495728408961817.

[69] MathWorks, "MATLAB Function: Include MATLAB Code in Simulink Models," https://www.mathworks.com/help/simulink/slref/matlabfunction.html, accessed: Sep. 29, 2026.

[70] T. I. Fossen, "Fossen's Marine Craft Model," https://www.fossen.biz/html/marineCraftModel.html, accessed: Sep. 29, 2026.

APPENDIX A
SPONSORS

A. Platinum Sponsors

Nanyang Technological University, School of Mechanical and Aerospace Engineering (MAE) – For the lab facilities to our project.

Future Systems & Technology Directorate (FSTD) – For providing funding that allows Archimedes to reimagine what's possible.

Ouster – For providing the latest LiDAR sensor.

SAF Yacht Club – For enabling Archimedes to conduct water test to validate our systems.

B. Gold Sponsors

Nanyang Technological University, College of Engineering (COE).

Singapore Maritime Foundation (SMF) – For media features and exposures.

C. Silver Sponsors

WaterLinked – For their hardware support.

D. Bronze Sponsors

Basler – For their hardware and technical advice.

E. Supporters

Blue Robotics – For their hardware discounts.

APPENDIX B
ARCHITECTURE BLOCK DIAGRAMS

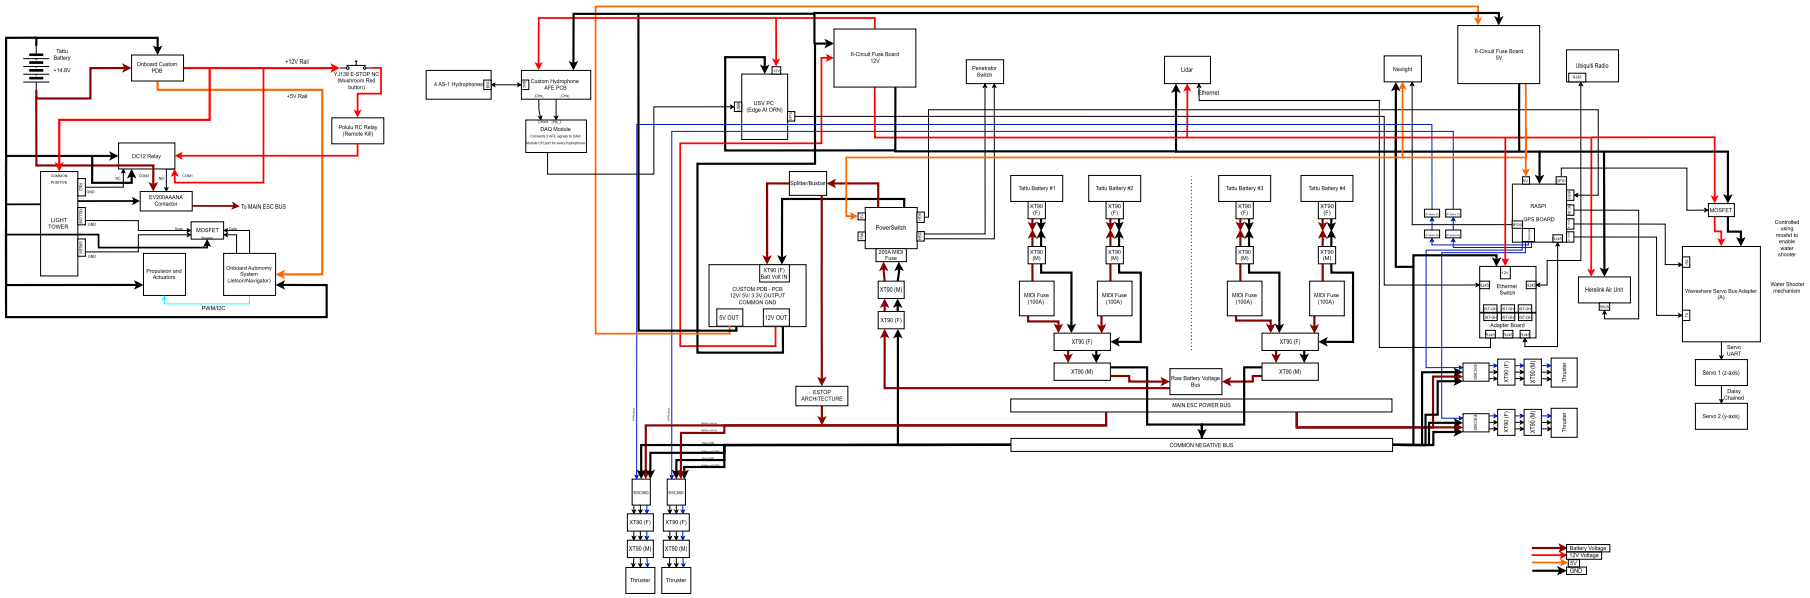


Fig. 11. USV electrical architecture. Protected distribution, retained electronics power, and actuator isolation form distinct integration paths.

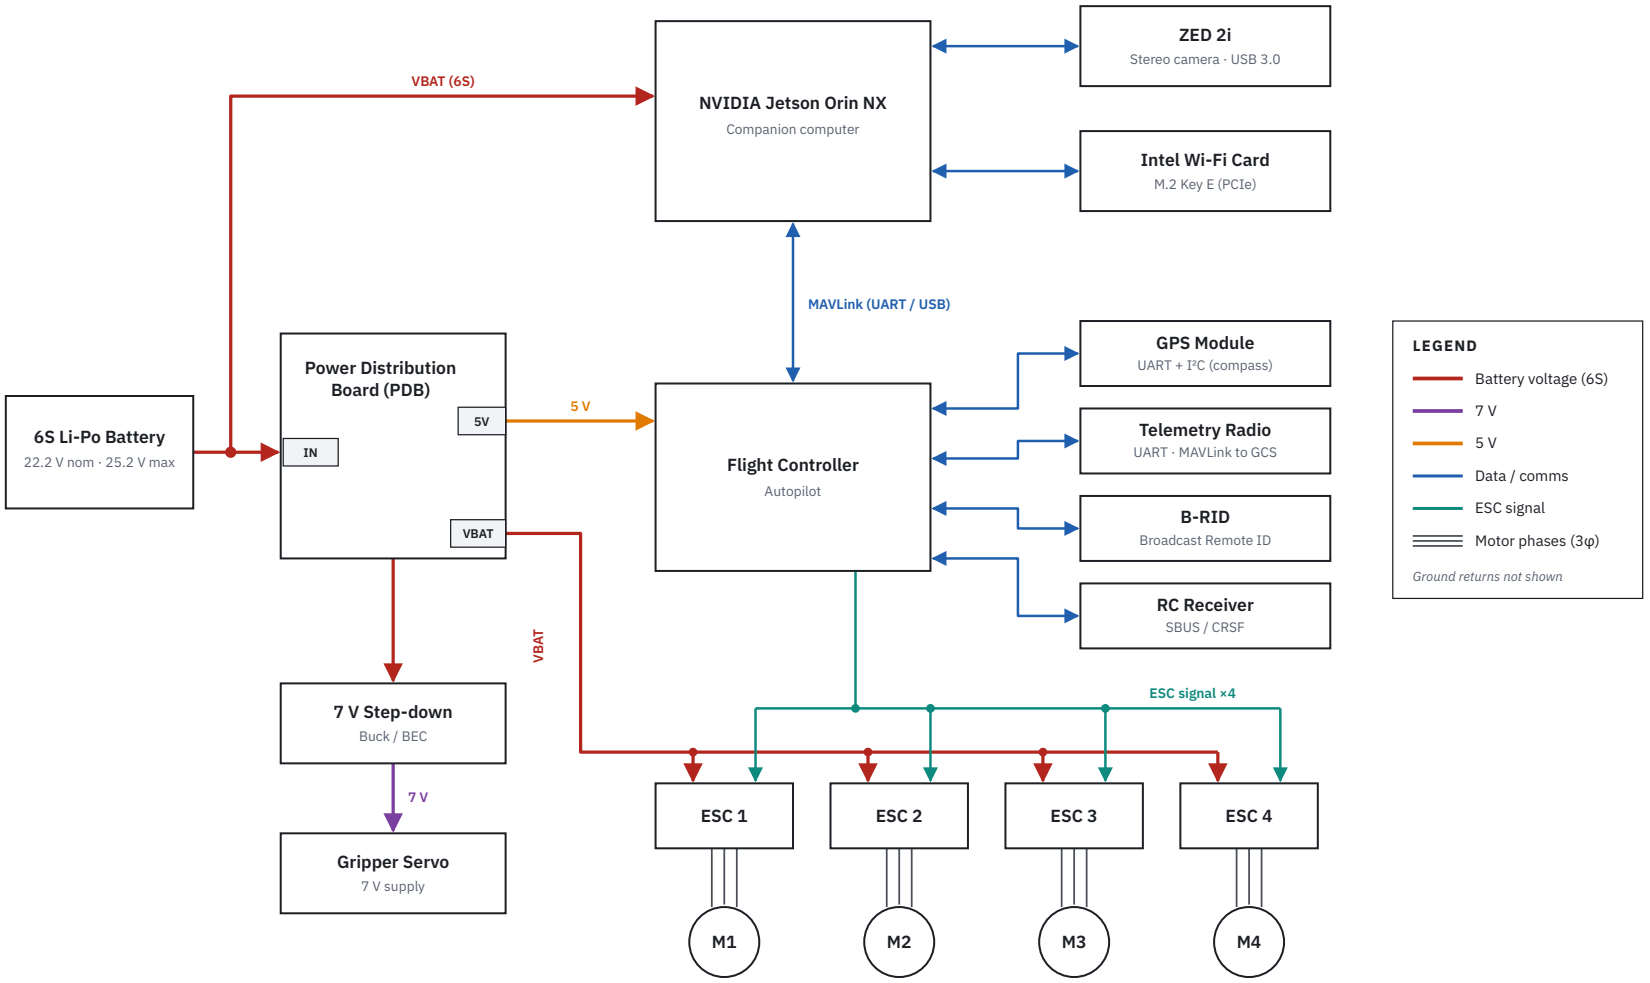
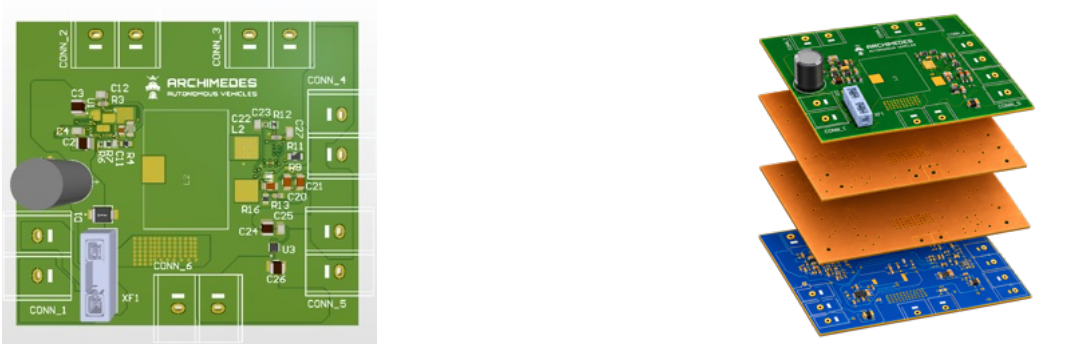


Fig. 12. UAV electrical architecture showing power distribution and avionics interfaces.

APPENDIX C
ELECTRICAL BOARDS AND ACOUSTIC ANALOG FRONT-END DESIGN

A. Protected Power and Actuator Isolation

The PDB centralizes protected distribution and regulated service power. The safety function governs permission to energize the external actuator contactor; secured modules on a soldered carrier form the initial prototype, with the illustrated custom PCB retained as a later packaging option. The separation supports fault diagnosis and preserves sensing and logging after a stop. Branch protection is selected against the connected wiring and credible load conditions; connector labels or nominal converter ratings alone do not establish an acceptable continuous load.



(a) PDB top view. (b) PDB exploded view.

Fig. 13. Power distribution board. The top view identifies the physical connection layout; the exploded view illustrates its layered construction.



(a) Safety-board top view. (b) Safety-board exploded view.

Fig. 14. Safety-board packaging concept and layer arrangement. Independent stop inputs and deliberate reset behavior govern actuator-power permission.

The USV distribution design distinguishes the nominal 14.8-V battery bus from the regulated 12-V service rail. High-current actuator returns are routed to avoid carrying motor or contactor transients through sensor references. Fuse selection, copper and connector capacity, regulator derating, and enclosure heat dissipation are checked together at maximum charge, minimum service voltage, and credible simultaneous load. Retained compute, communication, sensing, and logging rails allow a stopped vehicle to report its condition.

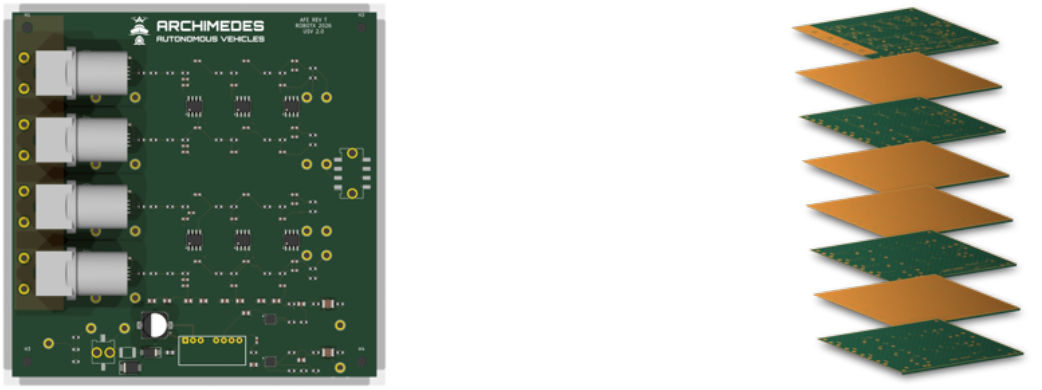
The safety design separates hard-fault permission from remote-run permission. Both initialize OFF after power-up or brownout. A physical stop or hard fault requires deliberate local reset; a remote-only stop may permit a deliberate handheld recovery only into accepted manual, disarmed, neutral operation. Restoring permission and arming are separate actions. RF return, a released button, or a

software restart must not re-energize actuators automatically. The external contactor interrupts every actuator supply, including alternate BEC or auxiliary feeds, while the carrier handles only logic and coil control. Contactor feedback and bus-voltage decay are checked because a command to open is not proof of isolation.

Acceptance is performed at the assembled vehicle: stop latency, transmitter power loss, contactor behavior, auxiliary actuator isolation, retained service rails, and reset/arming sequences are checked together against the safety requirements [5]. The board renders show design intent rather than completed safety qualification.

B. Acoustic Front End and Bearing Estimation

Four passive hydrophones require low-loading inputs and closely matched channel responses. The selected FET-input amplifier family limits input-current noise at the capacitive source, while preamplification and low-pass filtering condition the signal for simultaneous acquisition [12], [13], [37]. Layout separates the input region from conversion and high-current paths. Low-noise positive and negative post-regulation supports the analog rails [38], [39].



(a) AFE top view. (b) AFE exploded view.

Fig. 15. Four-channel acoustic AFE design. The retained CAD illustrations show the channel and layer arrangement; they predate a corrected protection-diode polarity and are not assembly instructions.

For source position $\mathbf{p}$, sensor positions $\mathbf{r}_i$ and $\mathbf{r}_j$, and sound speed c, the arrival-time difference is

$$\tau_{ij} = t_i - t_j = \frac{\|\mathbf{p}-\mathbf{r}_i\| - \|\mathbf{p}-\mathbf{r}_j\|}{c}, \qquad |\tau_{ij}| \le \frac{\|\mathbf{r}_i - \mathbf{r}_j\|}{c}. \tag{1}$$

The bound follows from the triangle inequality and constrains generalized-correlation delay candidates [16]. Four sensors provide six pairs but only three independent time differences. Narrowband cycle ambiguity, reflected arrivals, array geometry, and sound-speed uncertainty must be resolved or rejected before reporting a bearing. The 5-μs sample interval at 200 kS/s is not a bearing-accuracy specification.

TABLE III. Selected AFE model results and their scope.

Quantity	Retained result	Interpretation
Total gain at 30 kHz	19.53 dB	Nominal electrical signal-chain model.
Filter response at 30 / 40 kHz	-0.338 / -0.448 dB	Passband attenuation relative to the filter input.
Filter response at 100 kHz	-21.282 dB	Finite Nyquist-frequency rejection; not a complete aliasing guarantee.
Output noise, 25–40 kHz	7.04 μV rms	Component-model noise, excluding DAQ, acoustic, switching, and layout contributions.

The signal chain uses an approximately 20-dB OPA1656 preamplifier followed by three unity-gain Sallen–Key sections. The realized sixth-order response is a retuned low-pass network, not an exact 60-kHz Butterworth filter. Lower-Q sections precede the more resonant section to reduce its exposure to out-of-band energy. The quoted gain and noise values belong to the loaded component-model chain, including finite loading; they are not inferred solely from ideal filter equations.

At 200 kS/s per channel, the 100-kHz Nyquist frequency is only one anti-aliasing checkpoint. An input between 160 and 175 kHz folds into the 25–40-kHz receive band, so rejection must be compared with actual interference amplitude over the relevant alias bands. Large low-frequency disturbances can also overload the preamplifier before later filtering. Simultaneous acquisition preserves interchannel timing, but cable delay, channel phase mismatch, and array geometry still require calibration. Positive and negative post-regulators reduce supply contamination; regulator headroom, startup sequencing, and thermal dissipation must be evaluated with the complete converter and amplifier load [13], [37]–[39].

Table III summarizes the nominal response. Additional studies address channel matching, loading, supply response, and coupling sensitivity. An assumed 0.10-pF coupling case fails the provisional crosstalk screen, so final isolation remains unresolved. Hardware release requires the approved stackup, full supply startup and thermal checks, interference testing, channel calibration, and wet localization trials. Selected model results guide these tests; they do not establish assembled-board or array performance.

APPENDIX D
COMPONENT SPECIFICATIONS

Quantities are per vehicle, excluding spares, replacements, and shore equipment. A set is one onboard assembly. Vendor identifies the manufacturer or brand, not the supplier. Specifications describe the selected configuration and manufacturer ratings; assemblies under development are marked. A component rating does not qualify the complete vehicle.

A. USV 2.0

TABLE IV. USV 2.0 principal component specifications.

Component	Vendor	Model / Type	Key specifications	Qty.
Platform	Blue Robotics	BlueBoat	Catamaran hull; differential stern propulsion with two added bow thrusters.	1
Battery	Grepow	4S pack using GRPA088190 cells	Selected pack: nominal 14.8 V, 33 Ah (488.4 Wh); four packs. Nominal energy is voltage times capacity, not usable mission energy.	4
Stern propulsion	Blue Robotics	M200 / weedless propeller	Two stern motor–propeller assemblies for surge and differential yaw.	2
Bow propulsion	Blue Robotics	T200	7–20 V operating range; diagonal port/starboard mounts for low-speed sway and yaw [40].	2
Motor control	Blue Robotics	Basic ESC 500	One bidirectional controller per motor; selected Basic ESC 500 configuration [41].	4
Power and safety	Archimedes	Protected distribution and actuator isolation	Retained electronics supply; independent actuator contactor; regulated service rails (prototype assembly).	1 set
Navigation I/O	Blue Robotics	Navigator	Inertial sensing and PWM/serial interfaces for the Raspberry Pi autopilot [41].	1
Low-level computer	Raspberry Pi	Raspberry Pi 4 Model B	Selected 2 GB RAM variant; BlueOS and ArduRover.	1
Autonomy computer	Compulab	EdgeAI-ORN	Selected Orin NX 16 GB / 1 TB NVMe configuration with EB-EDGEPOE expansion: four 2.5 GbE PoE ports [42].	1

Continued on next page

USV 2.0 component specifications (continued)

Component	Vendor	Model / Type	Key specifications	Qty.
GNSS / heading	Holybro	H-RTK ZED-F9P Rover	Two UART receivers [43]; moving-baseline heading for low-speed navigation.	2
LiDAR	Ouster	OS0-128	128 channels; 360° horizontal coverage; one sensor in use (Rev. 7 loan unit).	1
Vision	Basler	ace 2 a2A2448-23gc IP67	2448 × 2048 color global shutter; 20.5 fps at default settings; GigE; selected 12 mm lenses [44].	2
Hydrophones	Aquarian Audio	AS-1	Passive four-element acoustic array for 25–40 kHz pinger bearing estimation.	4
Acoustic front end	Archimedes	Four-channel AFE	20 dB preamplification and sixth-order low-pass filtering; matched acquisition paths (design).	1
Acoustic acquisition	Measurement Computing	USB-1808X	Eight simultaneous analog inputs; four differential channels selected at 200 kS/s per channel [37].	1
Shore radio	Ubiquiti	Rocket 5AC Lite + AMO-5G10	5 GHz airMAX radio with a 10 dBi omnidirectional antenna; 20–27 V passive PoE radio supply [45], [46].	1 set
Manual-control link	CubePilot	HereLink 1.1 air unit	2.4 GHz operator command and manual-control link [47].	1
UAV telemetry	Holybro	SiK Telemetry Radio V3	Selected 433 MHz, 100 mW variant for serial telemetry to the UAV [48].	1
Underwater link	Water Linked	M16 Extended	10 bit/s acoustic communication; RS-422 interface to the USV [49].	1
Firefighting payload	Archimedes	Pump and steerable nozzle assembly	Water-delivery payload with actuated nozzle aiming; integration in progress.	1 set

B. UAV 3.0

TABLE V. UAV 3.0 principal component specifications.

Component	Vendor	Model / Type	Key specifications	Qty.
Airframe	Holybro	X650 V2	650 mm wheelbase carbon-fiber frame with a custom electronics enclosure [50].	1
Battery	Tattu	22 Ah 6S LiPo	Selected pack label: 22.2 V, 22 Ah, 25C; one flight pack. Usable capacity and installed current limits require measurement.	1
Motors	T-Motor	MN4110 KV400	400 rpm V^{-1} brushless motors; selected protection variant; propeller/load qualification pending [51].	4
Propellers	T-Motor	P17 × 5.8	17-inch diameter, 5.8-inch pitch; carbon fiber; two CW and two CCW [52].	4
Motor control	T-Motor	AIR 40A	40 A brushless ESC; 2–6S input; no BEC; one per motor [53].	4
Power sensing	Holybro	PM02D	Digital voltage/current monitoring and flight-controller supply; cable/connector current limits govern the installation [54].	1
Flight controller	Holybro	Pixhawk 6X v2A	STM32H753 processor [55]; selected PX4 v1.15 integration baseline.	1
Autonomy computer	Seeed Studio	reComputer J4012	Jetson Orin NX, 16 GB RAM; 128 GB SSD [56]; selected ROS 2 Humble stack.	1
GNSS	Holybro	H-RTK ZED-F9P Ultralight	RTK-capable GNSS receiver for position estimation [57].	1
Downward vision	Stereolabs	ZED 2i	Stereo camera with 4 mm lenses and a 120 mm baseline; range depends on calibration and correspondence [58].	1
Shore Wi-Fi	Intel	Wi-Fi 6E AX210	2×2 radio; selected 5 GHz IP link for mission data [59].	1
Wi-Fi antennas	Molex	204281-1150	150 mm cabled flex antennas; two ports of the AX210 radio [60].	2

Continued on next page

UAV 3.0 component specifications (continued)

Component	Vendor	Model / Type	Key specifications	Qty.
Manual-control link	RadioMaster	RP4TD	2.4 GHz ExpressLRS receiver; true diversity, CRSF interface [61].	1
USV telemetry	Holybro	SiK Telemetry Radio V3	Selected 433 MHz, 100 mW variant for serial telemetry to the USV [48].	1
Delivery payload	Archimedes	Servo-actuated gripper	Custom printed mounts and release mechanism; integration in progress.	1 set

C. UUV 1.0

TABLE VI. UUV 1.0 principal component specifications.

Component	Vendor	Model / Type	Key specifications	Qty.
Platform	Blue Robotics	BlueROV2 Heavy, acrylic enclosure	Eight-thruster configuration; manufacturer platform depth rating 100 m; modified assembly requires qualification [62].	1
Battery	Blue Robotics	4S lithium-ion battery	Nominal 14.8 V, 18 Ah; one onboard pack [41].	1
Propulsion	Blue Robotics	T200	Four horizontal and four vertical thrusters for six-degree-of-freedom control.	8
Motor control	Blue Robotics	Basic ESC	One bidirectional brushless controller per thruster [41].	8
Navigation I/O	Blue Robotics	Navigator	Inertial sensing and PWM/serial interfaces for the Raspberry Pi autopilot [41].	1
Low-level computer	Raspberry Pi	Raspberry Pi 4 Model B	Selected 2 GB RAM variant; BlueOS and ArduSub.	1
Autonomy computer	Seeed Studio	reComputer Mini J4012	Orin NX 16 GB; selected 128 GB NVMe; Mini J401 carrier without the extension board [63].	1
Velocity sensing	Water Linked	DVL A50 Standard	0.05–50 m altitude under ideal conditions; maximum velocity 3.75 m/s [64].	1
Depth sensing	Blue Robotics	Bar30	30 bar pressure sensor; depth derived using pressure offset and water density [41].	1
Imaging sonar	Sonoptix	ECHO	400/700 kHz multibeam sonar; 90/120° horizontal field of view, respectively [29].	1
Inspection camera	Blue Robotics	Low-Light HD USB Camera	1920 × 1080 at 30 fps; forward inspection imagery [41].	1
Illumination	Blue Robotics	Lumen Subsea Light	Dimmable underwater lighting; up to 1500 lm per light [41].	4
Surface GNSS	Holybro	H-RTK ZED-F9P Rover	UART GNSS receiver for surfaced position updates [43].	1
Surface radio	Ubiquiti	BulletAC-IP67	Dual-band airMAX AC radio; selected 5 GHz surfaced link; 20–27 V passive PoE [65].	1

Continued on next page

UUV 1.0 component specifications (continued)

Component	Vendor	Model / Type	Key specifications	Qty.
Surface antenna	Phoenix Contact	ANT-OMNI-2459-04	Omnidirectional; 2.5 dBi in the 4.9–5.9 GHz band; N-male connection [66].	1
Underwater link	Water Linked	M16 Standard	10 bit/s acoustic communication; 3.3 V TTL UART interface to the UUV [49].	1

The original platform rating is 100 m; higher ratings of individual sensors do not increase it. New penetrators, enclosures, seals, and payload interfaces require separate pressure qualification before assigning the modified vehicle a depth rating.

APPENDIX E
SUPPLEMENTARY INTEGRATION VIEWS

A. Perception and Information Ownership

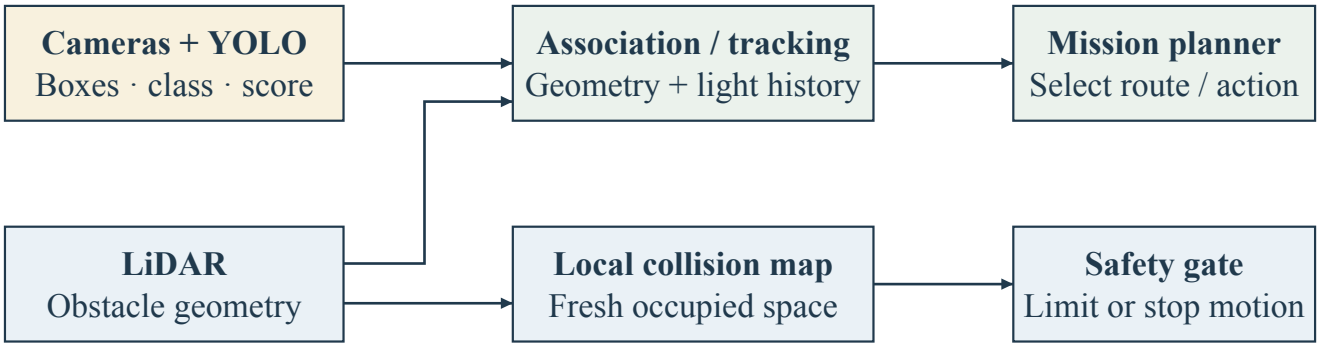


Fig. 16. Selected USV perception responsibilities. YOLO detections and LiDAR clusters feed object association; geometric collision sensing remains independent of semantic classification.

An observation carries its source identity, capture time, coordinate frame, uncertainty, and validity period. Association merges compatible evidence about a physical object without confusing its identity with a changing light state. Task decisions use accepted observations; immediate collision checks require current geometric coverage.

B. Communication and Control Checks

IP, compact radio, and acoustic paths are tested with delay, loss, duplicate delivery, and reconnection. Receipt, task acceptance, physical completion, and official publication are separately observable. Mission OCS is the sole official protocol adapter, and vehicle-specific OCSs supervise their corresponding systems [18]. A stale request must remain stale after retransmission.

USV GNSS corrections and moving-baseline heading have different responsibilities: shore correction data supports global positioning, while the onboard baseline supplies yaw when valid [20]. The UAV uses uXRCE-DDS for PX4 integration [25]. The UUV has a single DVL injection path into ArduSub. These ownership rules prevent duplicate measurements and conflicting actuator commands.

For planar USV allocation, the requested body force and yaw moment satisfy the standard relationship

$$\boldsymbol{\tau} = \mathbf{B}\mathbf{f}, \qquad \mathbf{B}_i = \begin{bmatrix} \cos\alpha_i \\ \sin\alpha_i \\ x_i \sin\alpha_i - y_i \cos\alpha_i \end{bmatrix}, \tag{2}$$

where $\mathbf{f}$ contains signed thruster forces, (x_i, y_i) is the ith thruster location in the body frame, and α_i is its force direction [24]. Force bounds, actuator dynamics, and loss of a thruster limit the achievable command. The proposed precision-control path must be checked against these limits before docking or station-keeping capability is claimed.

APPENDIX F
UAV DESIGN STRATEGY

A. UAV Perception

The UAV's perception stack supports the RobotX 2026 observation, delivery, and recovery tasks [1]. Its outputs provide evidence for mission decisions; a detection is not a completed action. The required capabilities are:

- detecting and locating the buoys, and extracting their status (Tasks 1 and 2);
- detecting and locating the UAV delivery platform (Tasks 2 and 3);
- detecting and locating the metal cans, and identifying their color (Tasks 2 and 3); and
- detecting and locating the provided landing platform.

Currently, the UAV is equipped with a single downward-facing stereo depth camera, which is responsible for all detection, localization, and identification tasks. However, this camera is not mounted centrally, which makes it difficult to keep the landing platform and the metal cans in frame as the UAV descends close to ground level. Hence, we plan to mount a smaller downward-facing monocular camera at the center of the UAV to maintain line of sight with the landing platform and the metal cans. The UAV is also equipped with a Jetson Orin NX, which handles all onboard computation, including perception. An overview of the perception pipeline is shown in Fig. 17.

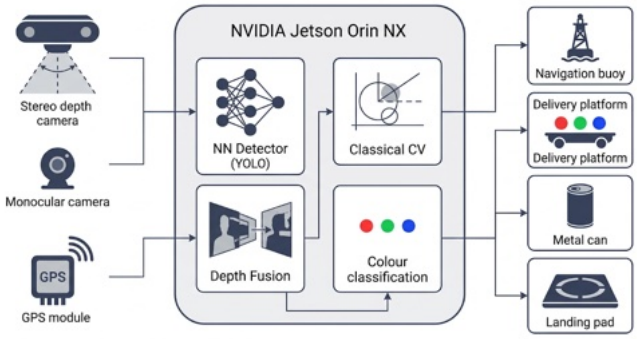


Fig. 17. Overview of the UAV perception pipeline, from sensors to mission outputs.

1) Detecting, locating, and extracting the status of the buoys: We built competition-dimension buoy mock-ups and collected camera observations in field and pool settings to support the YOLO training pipeline. Data collection uses the downward-facing competition camera and varied viewpoints so that object scale, water reflections, glare, and high ambient brightness are represented. The design separates buoy detection from beacon-state interpretation: a detector supplies a bounding region and confidence, whereas subsequent geometry and temporal classification establish position and light state. YOLO's detection interface supplies class and bounding-box predictions; it does not by itself recover three-dimensional position or a flashing sequence [6].

Training and perception integration remain in development. Evaluation reserves independent physical capture sequences, rather than splitting neighboring video frames between training and test sets. This limits optimistic results caused by nearly identical images appearing in both sets. Range, view angle, sun direction, occlusion, and exposure are varied separately to identify failure conditions. Saturated or poorly resolved beacons remain unknown. No precision, recall, or localization-error value is claimed until the corresponding held-out evaluation is available.

Long-range localization is limited by stereo disparity resolution, vehicle-position uncertainty, and capture-time attitude. Small disparity errors produce larger depth errors at greater range; reflective water can also invalidate stereo correspondences. The design therefore uses stereo only within its qualified envelope and retains an explicit geometric alternative for distant buoy observations.

The selected alternative to long-range stereo depth intersects a calibrated camera ray with the local water plane. For an undistorted detection pixel (u, v), let $\mathbf{q} = K^{-1}[u\ v\ 1]^{\mathsf{T}}$, $\mathbf{d} = R_{WC}\mathbf{q}$, where R_{WC} rotates camera-frame vectors into the world frame, and let $\mathbf{c}$ be the camera center in that local metric frame. For the assumed target plane $z = z_w$,

$$\lambda = \frac{z_w - c_z}{d_z}, \qquad \mathbf{p}_{\text{buoy}} = \mathbf{c} + \lambda\mathbf{d}. \tag{3}$$

The camera intrinsics, distortion, extrinsic translation, and capture-time attitude are required [8]. Rays near parallel to the plane or with $\lambda \leq 0$ are rejected. Height must be relative to the modeled target plane, not uncorrected GNSS altitude; wave motion and beacon-face height contribute uncertainty. The familiar $h\tan\theta$ relation is only the aligned, level-camera special case. Figure 18 illustrates the concept, while the full ray-plane expression defines the general calculation.

Aerial observations are associated with the USV's physical object tracks using common frame, source time, uncertainty, and course/run context. Nearest position is only a candidate match; ambiguous associations remain unresolved until additional geometry or appearance supports

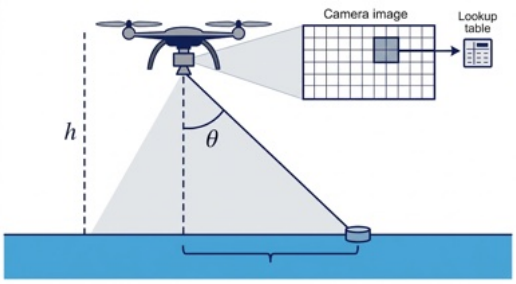


Fig. 18. Conceptual height-based buoy localization. The general calculation uses the calibrated camera pose and an explicit target-plane height; plane and attitude errors are propagated into the observation uncertainty.

identity. This avoids attaching a valid top-beacon state to the wrong buoy. The conceptual correspondence is shown in Fig. 19.

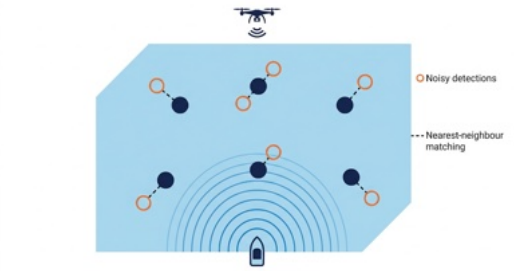


Fig. 19. UAV–USV cooperation: aerial detections are associated with the USV's physical tracks using spatial, temporal, and uncertainty checks; nearest-neighbor proximity alone is insufficient.

Top-beacon classification uses calibrated image regions and temporal observations to distinguish the permitted color and flashing states (Fig. 20). Evaluation varies sun angle, glare, exposure, distance, and background; uncertain or saturated observations remain unknown and cannot authorize a route change. Color classification does not require exact RGB recovery, but still requires measured confusion rates and freshness bounds. Thermal contrast is not assumed to establish visible LED color.

2) Detecting and locating the UAV delivery platform: Delivery-platform detection uses the three colored circles and their known layout as geometric features. The handbook specifies a circle-diameter range of 60–70 cm; the actual target dimensions must therefore be measured or represented as an uncertainty interval rather than assumed exact [1]. The selected classical vision approach segments candidate color regions and fits circle or ellipse boundaries before checking their

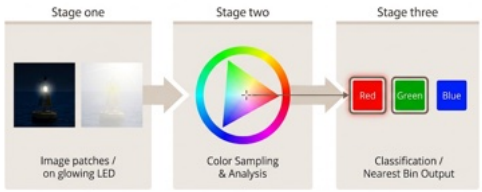


Fig. 20. Beacon-color extraction concept. Temporal evidence additionally distinguishes off, steady, and flashing states; color alone does not identify a mission state.

joint layout. Oblique views, partial occlusion, glare, and similarly colored background objects are rejection cases. This is a pose-estimation design under qualification, not evidence of achieved autonomous delivery accuracy.

To identify its location, we compare the apparent sizes of the three circles against their actual size to estimate the distance to the UAV delivery platform. For a circle of known diameter D that appears with a diameter of d pixels, the optical-axis depth for an approximately fronto-parallel view is estimated as

$$Z \approx \frac{f\,D}{d}, \tag{4}$$

where f is the camera's focal length in pixels. The known circle layout provides correspondences for calibrated planar pose estimation (Fig. 21). The size-based estimate is a fronto-parallel approximation. General platform pose requires calibrated projection of the known planar layout with distortion correction and a pose-ambiguity check; relative apparent diameters alone do not determine tilt. Near-range stereo may corroborate the estimate only within its measured validity envelope [8].

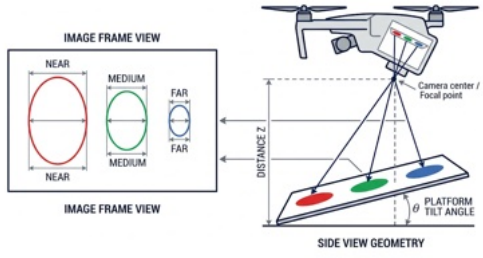


Fig. 21. Conceptual delivery-platform geometry. Apparent diameter provides an approximate range; general orientation requires calibrated pose estimation from the known layout.

3) Detecting, locating, and identifying the color of the metal cans: Resource-tin detection is restricted to the qualified near-platform viewing envelope. Known dimensions, calibrated image geometry, and valid stereo depth support relative position, while illumination-aware classification supplies the requested color. A detection alone does not establish payload acquisition or release; the mission requires separate mechanism feedback and observed payload outcome.

4) Detecting and locating the provided landing platform: GNSS supports the approach to the landing region; final descent uses current platform-relative observations and a qualified landing controller. Platform motion, relative velocity, attitude, clearance, and settled contact are evaluated explicitly. A detected pattern or proximity to its global coordinates alone cannot establish a safe landing. Missed detections, excessive deck motion, or inadequate reserve invoke the separately tested hold, abort, or recovery policy.

B. UAV Navigation and Control

The UAV's navigation and control stack is required for the following:

- executing PX4-native flight phases (arming, takeoff, landing, and return-to-launch);
- flying to position targets requested by the mission nodes, in either GPS or local coordinates;
- providing fine velocity control for precision tasks, such as picking up and dropping the metal cans and precision landing; and
- ensuring the UAV never moves unexpectedly when control is handed between PX4, the safety pilot, and our onboard software.

Low-level flight control runs on a Pixhawk 6X flight controller with PX4 v1.15, while all high-level autonomy runs on the Jetson Orin NX with ROS 2 Humble on Ubuntu 22.04. The two are bridged by the Micro XRCE-DDS Agent over a serial link on the Pixhawk's TELEM3 port, which exposes PX4's internal topics as ROS 2 topics. However, using these PX4 topics directly in every mission node would tightly couple our mission logic to PX4's message layouts and flight-mode behavior, and a single mistake in any one node could leave the vehicle in an unsafe state. Hence, we developed a dedicated navigation abstraction layer, uav_nav, that sits between our mission nodes and PX4 (Fig. 22).

1) Navigation abstraction layer: uav_nav_node is the sole autonomous setpoint and flight-mode command owner in the selected stack; read-only telemetry consumers remain permitted. Mission nodes never import px4_msgs; instead, they only use the small set of ROS 2 topics and services listed in Table VII. This way, mission nodes can be written, tested, and replaced without needing to understand PX4's internals, and autonomous flight-mode transitions can be reviewed at one interface, while PX4 failsafes and pilot authority remain independent.

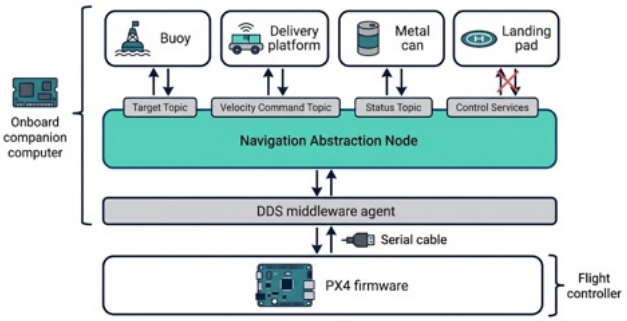


Fig. 22. UAV navigation software architecture. Mission commands pass through uav_nav_node; PX4 state feedback and pilot authority remain separate. The adapter is the sole autonomous command owner, while read-only telemetry consumers are permitted.

As PX4 and px4_msgs are matched structurally over uXRCE-DDS, a version mismatch between the firmware and the message definitions can silently break fields. The integration design follows the PX4 ROS 2 interface guidance [25]. To limit this risk, we pin px4_msgs to the release/1.15 branch to match our PX4 v1.15.x firmware. Additionally, any vehicle command that PX4 does not accept is logged as a warning through the command acknowledgment, making it easy to diagnose why a service call did not take effect.

2) Control modes and targets: uav_nav_node runs a loop at a configurable rate that continuously streams offboard control mode and trajectory setpoint messages to PX4. While OFFBOARD control is admitted, the ROS 2 OffboardControlMode proof-of-life stream must exceed 2 Hz and be established for more than one second before entry. This stream is distinct from TrajectorySetpoint; loss triggers the configured OFFBOARD-loss timeout and action [27]. Continuing to publish is therefore not a substitute for health and command-lease validation. Internally, the node tracks one of three control modes (Fig. 23):

- **HOLD**: no target is set, and the UAV hovers

TABLE VII. Interfaces exposed by uav_nav_node to the mission nodes.

Interface	Type	Purpose
/uav/nav/target	Topic (sub)	Position target in wgs84 or local_enu frame
/uav/nav/cmd_vel	Topic (sub)	ENU velocity and yaw-rate override for fine control
/uav/nav/status	Topic (pub)	Arrival, distance, control mode, arming, and OFFBOARD state
/uav/nav/arm, /disarm	Service	Arm or disarm the vehicle
/uav/nav/takeoff	Service	PX4-native takeoff, followed by automatic OFFBOARD handoff
/uav/nav/set_offboard	Service	Manually request OFFBOARD mode
/uav/nav/land, /rtl	Service	PX4-native landing or return-to-launch, ending in disarm

at its current position;
- **POSITION**: the UAV flies to a target published on /uav/nav/target; and
- **VELOCITY**: the UAV follows a velocity override from /uav/nav/cmd_vel, reverting to POSITION (if a target is set) or HOLD after 0.5 s without a new message.

Position requests distinguish geographic coordinates from metric local coordinates. A wgs84 request carries latitude, longitude, and an explicit altitude datum; it is not an ENU vector. A local_enu request requires an accepted metric origin and valid local position. The adapter converts once into PX4's NED convention using the deployed firmware's projection and origin, or a qualified transformation into that frame. PX4's local origin is not automatically the launch/home position. Origin resets invalidate targets tied to an earlier reference, and height above water or a landing deck is not interchangeable with AMSL or ellipsoidal altitude [25]. Local coordinates alone do not provide GPS-denied navigation.

The interface also separates automatic heading selection from an explicit fixed yaw. The prototype encodes automatic yaw with an identity quaternion, which is also a valid zero-yaw rotation. This ambiguity must be removed or explicitly resolved in the released request contract before relying on a fixed zero-heading command. Attitude and velocity conversions require the same reviewed ENU/NED and body-frame sign conventions as position.

The prototype arrival thresholds consider a target reached when the UAV is within 0.3 m horizontally and 0.3 m vertically of it. Once this holds for 0.5 s, the target is automatically cleared and the node reverts to HOLD. These are software thresholds rather than demonstrated position accuracy. Clearing the target is intended to prevent it from surviving a land, re-arm, and takeoff cycle and being flown to again as soon as OFFBOARD is regained. The velocity override bypasses position targeting entirely and is used mainly for fine control tasks, such as picking up and dropping the metal cans and precision landing. Mission nodes must keep publishing to sustain velocity control, so expiry can select a qualified hold while the adapter and local position remain healthy. Adapter or companion failure instead relies on the independently tested PX4 OFFBOARD-loss response. Travel speed toward position targets is not governed by our node, but by PX4's own cruise and maximum velocity parameters.

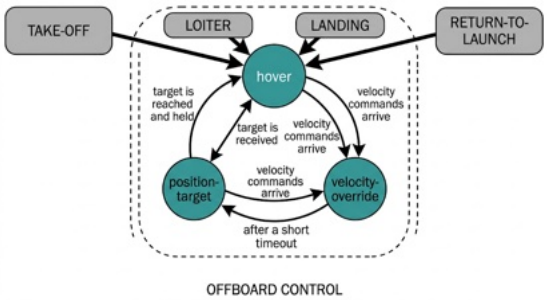


Fig. 23. Control mode state machine of uav_nav_node, alongside the PX4-native modes that take over outside OFFBOARD.

3) OFFBOARD mode switching: There is exactly one automatic path into OFFBOARD, and it

only exists because a takeoff was explicitly commanded (Fig. 24):

1) /uav/nav/takeoff is called, and PX4 performs a native takeoff using a target height with an explicitly verified reference and vertical datum;
2) PX4 completes the climb and switches to AUTO_LOITER;
3) AUTO_LOITER is held for 0.5 s; and
4) uav_nav_node automatically requests OFFBOARD mode.

If the takeoff is interrupted (for example, by an RC override), the pending handoff is canceled and control remains manual from then on. We implemented this cancellation with a 0.5 s debounce window rather than canceling on the first unexpected reading. This protects against the first vehicle status sample after the takeoff command still reflecting PX4's state before the command was received, and also ignores a single stale sample mid-sequence. If /uav/nav/takeoff is never used, there is no automatic switching at all; reaching AUTO_LOITER any other way, such as through a pilot's RC mode switch, never triggers an OFFBOARD handoff. /uav/nav/land and /uav/nav/rtl are required to finish without automatic OFFBOARD re-entry. The recorded software test report still flags this transition and incomplete health/hazard integration; these remain acceptance items rather than established flight guarantees.

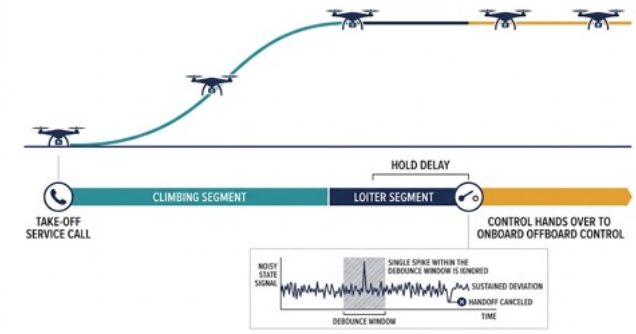


Fig. 24. Takeoff-to-OFFBOARD handoff sequence. Admission additionally requires valid navigation, permission, and a qualified proof-of-life stream; interruption handling requires integration testing.

4) OFFBOARD entry/exit reset: We approached the handover between PX4, the safety pilot, and our software with a focus on safety, and built the node around a single invariant: on every OFFBOARD entry or exit, the adapter invalidates active targets and velocity overrides and initializes any subsequent hold reference from accepted current local state (Fig. 25). This intended contract applies to a post-takeoff handoff, a manual OFFBOARD request, a landing or RTL command, a disarm, a pilot taking over on the RC and handing back control, and even causes with no dedicated service call, such as a PX4 failsafe-triggered RTL. This means:

- regaining OFFBOARD is never itself a trigger for motion, as the UAV initializes a hold reference at the accepted current position, instead of snapping back to a stale hold point or flying off toward a target set before OFFBOARD was lost;
- a target or velocity override can never survive an OFFBOARD dropout and resume the instant OFFBOARD returns; and
- targets and velocity commands are only accepted while OFFBOARD is already active, so there is no way to pre-stage a target ahead of an OFFBOARD entry.

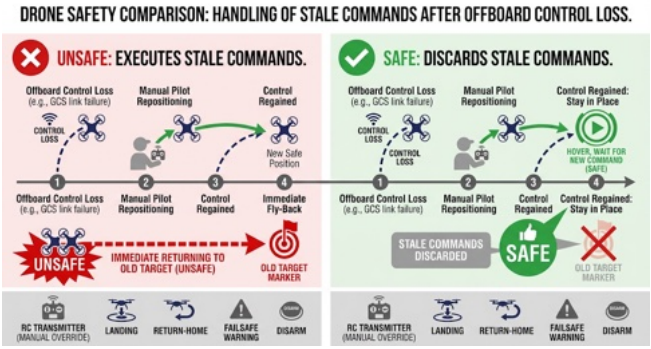


Fig. 25. OFFBOARD entry/exit reset logic. Mode changes discard stale targets and initialize a hold reference only from accepted local state; native flight modes and pilot authority take precedence.

5) Status feedback and failsafes: Mission nodes track progress by polling /uav/nav/status, which reports whether the UAV is at its target, the distance to the target, the arming state, whether OFFBOARD is active, and whether a valid NED reference is available. The control mode field reports POSITION, VELOCITY, or HOLD while our node owns OFFBOARD, and the PX4-native mode in charge (such as TAKING_OFF, LANDING, or RTL) otherwise. This allows the mission nodes to always know who is in control of the vehicle.

PX4 retains native mode, estimator, battery, geofence, and link-loss handling. The released parameter set must define the response to each

failure separately: return-to-launch is not a valid universal response to lost position or exhausted energy. OFFBOARD loss, pilot takeover, and companion-process death are tested with the actual firmware and continued unrelated telemetry [27]. A software motor-stop command is distinguished from independently verified physical actuator isolation. Battery thresholds, takeover latency, and recovery margins require measured evidence for the installed payload and flight envelope; illustrative settings do not establish those outcomes.

The uav_nav software report documents selected simulator checks, while integrated health gating, hazard-zone rejection, and terminal landing/RTL behavior remain incomplete. Release therefore requires command-ownership checks, fresh navigation, validated geofence and target rejection, bounded timeout response, and no automatic mission resumption after manual intervention. The mode diagrams specify intended transitions and are not evidence that all transitions have passed physical flight tests.

C. UAV Platform Design

The UAV platform is designed around the following requirements of the missions:

- operating reliably over water and in exposed maritime weather conditions;
- carrying the sensing and computing hardware required by the perception and navigation stacks (Sections F-A and F-B); and
- maintaining reliable communication with both the ground control station (GCS) and the USV.

The design mass is approximately 7 kg; the release configuration requires a measured mass, center of gravity, and payload-specific energy reserve. The airframe CAD is shown in Fig. 6, and the selected components are summarized in Table VIII.

TABLE VIII. Main components of the UAV platform.

Subsystem	Component
Airframe	Holybro X650
Motors	T-Motor MN4110, selected for wet-environment protection
Propellers	T-Motor P17×5.8 carbon fiber
Electronics enclosure	Custom enclosure, Multi Jet Fusion printed; ingress qualification required
Flight controller	Pixhawk 6X running PX4 v1.15
GNSS	Holybro H-RTK ZED-F9P Ultralight
Onboard computer	NVIDIA Jetson Orin NX, ROS 2 Humble on Ubuntu 22.04
PX4–ROS 2 bridge	Micro XRCE-DDS Agent (serial, TELEM3)
Depth camera	Stereolabs ZED 2i (downward-facing)
GCS link	Intel AX210, configured for the selected standard 5-GHz Wi-Fi cell
USV link	Paired 433-MHz SiK for admitted compact events

1) Airframe and propulsion: The UAV is built on a Holybro X650 frame, which provides sufficient space to mount our payload of computing hardware, sensors, and the electronics enclosure. As the UAV operates over water, it is exposed to rain, sea spray, and splashes during flight, takeoff, and landing. Hence, we selected T-Motor MN4110 motors with environmental protection, paired in the design with T-Motor P17×5.8 carbon fiber propellers. The listed KV400 winding and 17-inch propeller combination requires specific thrust, current, and temperature testing: the current manufacturer guide recommends 15–16-inch propellers for KV400 [51]. Because the motors remain exposed, the motor choice reduces exposure risk. The complete propulsion installation still requires verification of connectors, drainage, corrosion protection, thermal behavior, and operation after the specified wet exposure.

2) Water-resistant electronics enclosure: The flight controller, onboard computer, and supporting electronics are housed in a custom-designed enclosure intended to protect against water ingress. The enclosure was printed using Multi Jet Fusion (MJF), which allowed us to design the enclosure around the exact layout of our electronics and mounting points on the X650 frame, while producing a part whose wall thickness, mass, and mounting geometry can be controlled. Material strength and print consistency remain manufacturing acceptance checks. Designing the enclosure ourselves also meant that the cable routing, connectors, and

mounting of the components could be planned together with the waterproofing, rather than being added on afterward.

In the event that the UAV lands on the water, a water-activated safety LED switches on to make the aircraft easy to locate, and the recovery plan uses a safety boat and marked handle to retrieve the aircraft if flotation is retained, with the motors disarmed before the boat approaches. Buoyancy, visibility, ingress resistance, and safe retrieval must be checked on the loaded aircraft; enclosure material and a recovery handle alone do not qualify water recovery.

3) Avionics and computing: Low-level flight control is handled by a Pixhawk 6X flight controller running PX4 v1.15, with positioning provided by a Holybro H-RTK ZED-F9P Ultralight GNSS receiver. All high-level autonomy, including perception (Section F-A) and navigation (Section F-B), runs on an NVIDIA Jetson Orin NX with ROS 2 Humble on Ubuntu 22.04. The Pixhawk and the Jetson are connected over a serial link on the Pixhawk's TELEM3 port, with the Micro XRCE-DDS Agent bridging PX4 topics into ROS 2. On top of this bridge, our uav_nav package connects PX4 to our mission packages, so that no mission package needs to communicate with PX4 directly (Fig. 26).

For perception, the UAV is equipped with a downward-facing Stereolabs ZED 2i stereo depth camera. As described in Section F-A, we also plan to mount a smaller downward-facing monocular camera at the center of the UAV to keep the landing platform and the metal cans in frame during the final descent.

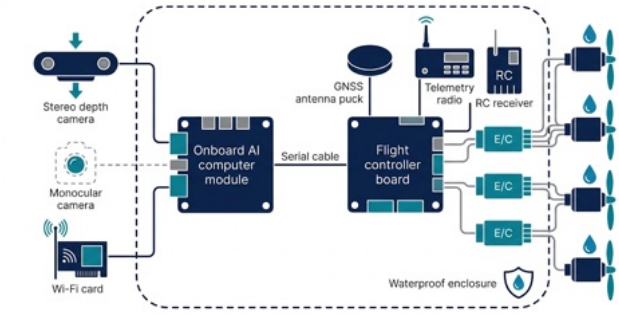
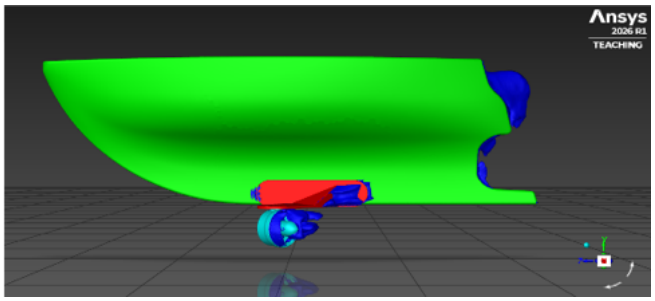


Fig. 26. UAV avionics interface concept. The original propulsion icons are schematic; the competition X650 is a quadrotor, as shown in Fig. 6.

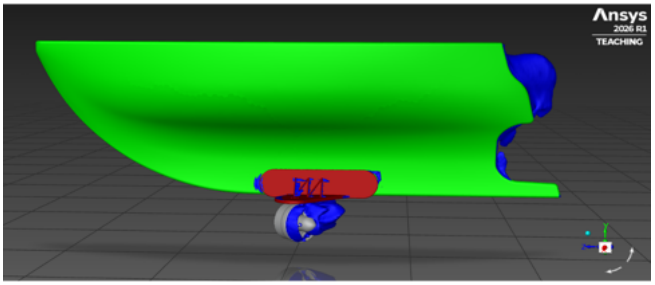
APPENDIX G
THRUSTER MOUNT HYDRODYNAMIC ANALYSIS

A. Introduction

The thruster mount is continuously exposed to the surrounding flow and therefore contributes directly to the hydrodynamic resistance of the unmanned surface vehicle (USV). During the mechanical design process, two mounting configurations were evaluated using ANSYS Fluent to reduce flow obstruction around the thruster while maintaining the required mounting geometry. The original design, denoted as v1, employed a comparatively broad triangular support, whereas the final v2 design replaced this obstruction with multiple slender triangulated ribs. The two configurations were compared using integrated hydrodynamic forces acting on the hull, mount, and thruster, together with qualitative flow visualizations. The objective of the redesign was not only to reduce the drag generated by the mount itself, but also to verify that the improvement remained beneficial when interactions among the mount, thruster, and hull were considered.



(a) Original v1 thruster mount.



(b) Final v2 thruster mount.

Fig. 27. Thruster mount design comparison. The final v2 configuration replaces the broad support of v1 with multiple slender ribs to reduce flow obstruction around the mount while retaining a triangulated support arrangement.

B. CFD Methodology

To provide a consistent basis for comparison, both thruster mount configurations were evaluated using the same CFD setup in ANSYS Fluent. The simulations modeled the submerged BlueBoat operating at an inlet velocity of $3\,\mathrm{m/s}$. The same computational domain, boundary conditions, mesh strategy, and solver settings were applied to both the original v1 and redesigned v2 configurations.

1) Computational Domain and Boundary Conditions

The computational domain was defined around the submerged hull and thruster assembly to provide sufficient distance for the flow to develop before and after the model. The upstream boundary was located at a distance of $5L$ from the model, while the downstream boundary was located at $10L$, where L represents the characteristic length used to define the computational domain.

A velocity inlet boundary condition was applied at the upstream boundary with a velocity of $3\,\mathrm{m/s}$, representing the selected comparison condition rather than an achieved vessel-speed measurement. A pressure outlet was applied at the downstream boundary. The hull, thruster mount, and thruster surfaces were treated as no-slip walls.

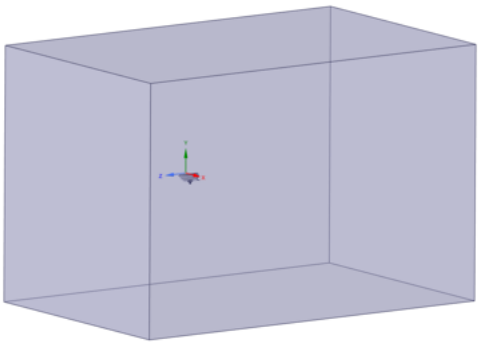


Fig. 28. Computational domain and boundary conditions used for the submerged-flow CFD analysis.

The same domain and boundary conditions were maintained for both configurations to ensure that differences in the calculated forces were primarily associated with the change in thruster-mount geometry.

2) Mesh Generation

A Poly-Hexcore mesh was used for both configurations, with approximately 5.26 million cells. Fifteen inflation layers were applied around the solid surfaces to provide increased mesh resolution near the hull, thruster, and thruster-mount

walls.

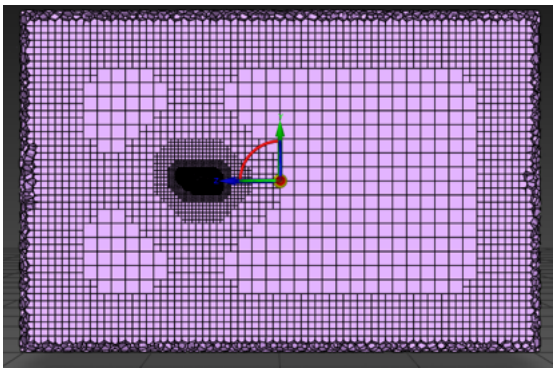


Fig. 29. Poly-Hexcore computational mesh used for the CFD analysis.

The same meshing approach and comparable mesh density were maintained for both v1 and v2 to provide a consistent basis for the comparison.

3) Solver and Numerical Setup

The simulations were performed using a steady-state Reynolds-averaged Navier–Stokes (RANS) approach with the SST k–ω turbulence model [67]. SIMPLEC was used for pressure–velocity coupling [68], with second-order spatial discretization schemes applied in ANSYS Fluent.

The prescribed residual criterion was 10^{-6}, and the study monitored hydrodynamic forces. A residual target is not evidence that a case reached that target or became mesh independent. Release of absolute force predictions requires retained force histories, mass-balance checks, wall-resolution assessment, and mesh/domain sensitivity. The available reports support a matched design comparison, not a validated full operating envelope.

The hydrodynamic force was evaluated in the drag direction used for the simulations. Force reports were obtained separately for the hull, thruster mount, and thruster, allowing the effect of the mount redesign to be evaluated both locally and at the overall system level.

C. Design Iteration

The principal hydrodynamic change between the two iterations is the reduction in the solid area presented by the supporting structure to the incident flow. In v1, the broad triangular support creates a comparatively large blockage between the hull-side attachment and the thruster mounting region. The final v2 geometry divides this support into narrower ribs, preserving a triangulated arrangement while exposing less solid area to the surrounding water.

The CFD visualizations in Fig. 27 show the resulting difference in the local flow region. The v1 configuration exhibits a more extensive disturbed-flow region around and downstream of its broad support, whereas the v2 mount provides a more open flow path through the supporting structure. This qualitative change is consistent with the reduction in pressure-induced force obtained from the integrated force calculations and provided the basis for progressing v2 to the final design.

D. CFD Force Analysis

The hydrodynamic resistance of each configuration was compared using the ANSYS Fluent integrated force reports. Both simulations reported the force in the same drag-direction vector,

$$\mathbf{d} = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T. \tag{5}$$

The archived force reports have negative components along this vector. The positive resistance values below are their magnitudes, with the same convention used for both cases. Pressure and viscous contributions are evaluated over identical hull, mount, and thruster surface groups. Percent changes are computed from the full-precision report values; table entries are rounded independently.

For a hydrodynamic force quantity F, the percentage change between the original and final designs was calculated as

$$\delta_F = \frac{F_{\mathrm{v2}} - F_{\mathrm{v1}}}{F_{\mathrm{v1}}} \times 100\%. \tag{6}$$

A negative value of δ_F therefore represents a reduction in resistance.

As summarized in Table IX, the most significant improvement occurs at the mount itself. The predicted mount resistance decreases from 5.735 N $\rightarrow$ 5.059 N, corresponding to an **11.80% reduction**. The resistance acting on the thruster surfaces also decreases from 25.327 N to 24.060 N, equivalent to a further **5.00% reduction**.

Consequently, the combined resistance associated with the mount and thruster decreases from $F_{\mathrm{mount+thruster,v1}} = 31.062$ N to $F_{\mathrm{mount+thruster,v2}} = 29.118$ N. The redesign therefore produces a reduction of $\Delta F^{\mathrm{reduction}}_{\mathrm{mount+thruster}} \approx 1.944$ N, or approximately 6.26%. The benefit remains when the complete modeled assembly is considered. Although the calculated hull resistance increases slightly from 36.391 N to 36.954 N, corresponding to a 0.563 N or 1.55% increase, the reduction obtained from the mount and thruster is sufficiently large

TABLE IX. CFD hydrodynamic force comparison between v1 and v2.

Component	**v1 (N)**	**v2 (N)**	**Diff. (N)**	**Change (%)**
Mount	5.735	5.059	−0.677	**−11.80**
Thruster	25.327	24.060	−1.267	**−5.00**
Mount + Thruster	31.062	29.118	−1.944	**−6.26**
Hull	36.391	36.954	+0.563	+1.55
Total Modeled Assembly	**67.453**	**66.073**	**−1.381**	**−2.05**

to overcome this penalty.

The resulting total resistance therefore decreases from $F_{\text{total,v1}} = 67.453$ N to $F_{\text{total,v2}} = 66.073$ N, representing a net reduction of $\Delta F_{\text{total}}^{\text{reduction}} = 1.381$ N or approximately $\boxed{2.05\%}$. This comparison shows that the improvement is not limited to an isolated mount-force metric. The hydrodynamic advantage remains after the interaction among the hull, mount, and thruster surfaces is included in the force calculation.

E. Pressure-Drag Reduction

Further examination of the force decomposition indicates that the mount improvement is predominantly associated with the pressure contribution rather than viscous surface friction. For v1, the mount pressure and viscous force contributions were approximately $F_{p,\text{v1}} = 4.833$ N and $F_{v,\text{v1}} = 0.902$ N, respectively.

For the final v2 configuration, these values decrease to $F_{p,\text{v2}} = 4.182$ N and $F_{v,\text{v2}} = 0.877$ N. The percentage reduction in mount pressure force is therefore

$$\frac{4.8332229 - 4.1817049}{4.8332229} \times 100\% \approx \boxed{13.48\%}, \tag{7}$$

whereas the corresponding reduction in viscous resistance is

$$\frac{0.90208393 - 0.87692297}{0.90208393} \times 100\% \approx 2.79\%. \tag{8}$$

The absolute reduction in pressure force is approximately $\Delta F_p = 4.8332229 - 4.1817049 \approx 0.652$ N. Relative to the approximately 0.677 N total reduction in mount resistance, $\frac{0.652}{0.677} \times 100\% \approx \boxed{96.3\%}$. Hence, approximately 96.3% of the predicted mount-drag improvement originates from a reduction in pressure loading rather than viscous friction. This is consistent with the design objective of reducing geometric blockage around the mount: the change in geometry primarily modifies how the incoming flow is redirected around the support structure.

F. Flow Behavior

The velocity-pathline visualizations provide a qualitative assessment of the flow around the two configurations. In the original v1 geometry, the incoming flow passes around the comparatively broad supporting section before interacting with the thruster. This produces a larger obstruction in the immediate mount region.

In contrast, the final v2 geometry introduces several narrower supporting ribs. The spaces between these ribs provide more open passages for the surrounding flow, thereby reducing the solid obstruction presented by the mount.

The pathlines in Fig. 31 are used qualitatively rather than as a direct measurement of drag reduction. The quantitative basis for the design comparison remains the integrated force results in Table IX. Nevertheless, the observed change in local flow behavior is consistent with the predicted reduction in mount pressure loading.

The relationship is particularly evident when the flow visualizations are considered together with the force decomposition. The v2 configuration produces a **13.48% reduction in pressure force**, whereas viscous resistance changes by only **2.79%**. This indicates that the primary hydrodynamic improvement results from reducing the pressure penalty associated with the mount geometry rather than simply reducing wetted-surface friction.

G. System-Level Effect

Optimizing the thruster mount in isolation is insufficient because the mount modifies the local flow encountered by both the hull and thruster. Therefore, the system-level result provides a more representative measure of the effectiveness of the redesign.

The local mount-and-thruster assembly achieves a reduction of 31.0624448 −

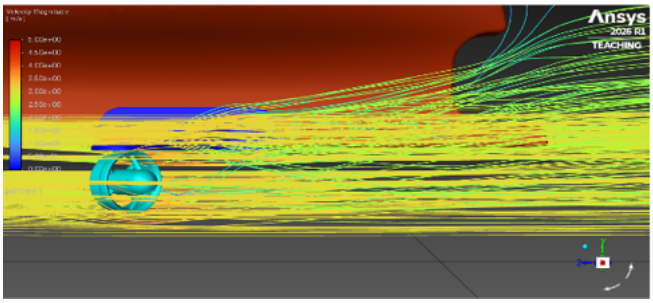
```
Forces - Direction Vector (0 0 1)
                          Forces [N]
Zone                      Pressure        Viscous         Total
hull                      -22.351207      -14.03968       -36.390886
mount                     -4.8332229      -0.90208393     -5.7353068
thruster                  -24.667889      -0.65924916     -25.327138
------------------------- --------------- --------------- ---------------
Net                       -51.852318      -15.601013      -67.453331
```

(a) v1 integrated force report.

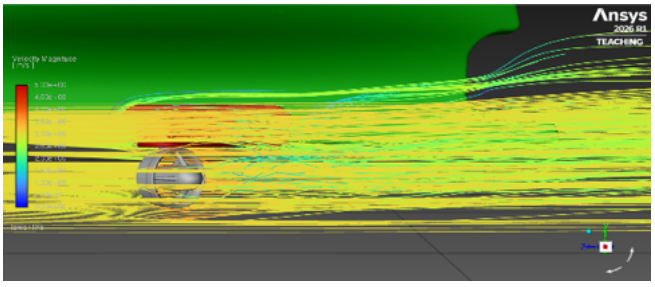
```
Forces - Direction Vector (0 0 1)
                          Forces [N]
Zone                      Pressure        Viscous         Total
hull                      -22.718959      -14.235192      -36.954151
mount                     -4.1817049      -0.87692297     -5.0586279
thruster                  -23.426133      -0.63367104     -24.059804
------------------------- --------------- --------------- ---------------
Net                       -50.326797      -15.745786      -66.072583
```

(b) v2 integrated force report.

Fig. 30. ANSYS Fluent integrated force reports for the original and final thruster-mount configurations. The final v2 design reduces mount resistance from 5.735 N to 5.059 N and total modeled resistance from 67.453 N to 66.073 N.



(a) v1 velocity-magnitude pathlines.



(b) v2 velocity-magnitude pathlines.

Fig. 31. Velocity-magnitude pathlines around the thruster assembly for the original and final mount configurations. The v2 geometry provides a less obstructed flow path through the support region. The pathlines are used as qualitative flow visualizations; the quantitative comparison is based on the integrated ANSYS Fluent force reports.

$29.1184319 \approx 1.944$ N. Over the same comparison, the hull contribution increases by approximately $36.954151 - 36.390886 \approx 0.563$ N. Thus, part of the local improvement is offset by the corresponding change in hull loading. The fraction of the mount-and-thruster improvement offset by the hull increase is approximately $\frac{0.563}{1.944} \times 100\% \approx 29.0\%$. Despite this interaction, approximately 71% of the local resistance reduction remains as a net system-level benefit. This produces the final 1.381 N reduction in total predicted resistance shown in Table IX.

For a vessel operating at a constant forward velocity U, the effective hydrodynamic power required to overcome resistance can be expressed as

$$P_{\text{hyd}} = F_D U, \tag{9}$$

where F_D is the hydrodynamic drag force [24]. At the analyzed 3-m/s condition, the full-precision difference of 1.380748 N corresponds to approximately 4.14 W less effective resistance power. This is a comparison of the modeled hydrodynamic load, not electrical input power. Therefore, for identical operating speed and otherwise identical conditions, a reduction in total resistance directly reduces the effective propulsive power required to overcome the represented hydrodynamic load.

The CFD results should not, however, be in-

terpreted as a direct 2.05% reduction in total electrical power consumption or battery usage because drivetrain efficiency, thruster operating point, auxiliary electrical loads, and environmental disturbances are not represented by the resistance comparison alone. The results instead demonstrate that the final mount geometry reduces the hydrodynamic load that the propulsion system must overcome under the analyzed CFD condition.

H. Final Design Selection

Based on the CFD comparison, v2 was selected as the final thruster-mount geometry. The redesign reduces predicted mount resistance by $\boxed{11.80\%}$, with approximately $\boxed{96.3\%}$ of the absolute mount-drag improvement attributable to reduced pressure loading. When the mount and thruster are considered together, the predicted resistance decreases by $\boxed{6.26\%}$, while the complete modeled hull–mount–thruster assembly retains a net $\boxed{2.05\%}$ reduction in resistance despite the small increase in calculated hull resistance.

The final design therefore demonstrates improvement at both the component and system levels. The narrower support geometry reduces the pressure loading acting on the mount while maintaining a triangulated structural arrangement. More importantly, the reduction remains present after the hydrodynamic interaction among the mount, thruster, and hull is considered.

For the analyzed CFD condition, the final v2 design therefore reduces total predicted resistance from 67.453 N to 66.073 N, while reducing the resistance directly associated with the mount from 5.735 N to 5.059 N. These results provide the quantitative engineering basis for adopting v2 as the final thruster mount for the USV.

I. Interpretation and Follow-Up Verification

The CFD decision is to retain v2 for reduced predicted resistance while verifying mechanical adequacy independently. A triangulated shape does not establish stiffness, fatigue life, or fastener retention. Static and handling loads, vibration, and underwater impact exposure remain mechanical acceptance items. Likewise, the present steady submerged-flow comparison does not establish propeller efficiency, unsteady wake behavior, free-surface resistance, or motion in waves.

The approximately 2.05% system-level difference is small enough that discretization and modeling uncertainty matter. A follow-up study should retain the same surface definitions and drag axis, compare several mesh densities and domain extents, and record whether the ordering of v1 and v2 survives those changes. Wall-normal resolution must be checked against the SST near-wall treatment rather than inferred from the number of inflation layers. Comparison at additional speeds would establish whether the preferred geometry remains beneficial across the intended operating range.

For RobotX, the value of this analysis is a documented design iteration: identify a local obstruction, change its geometry, compare the complete modeled assembly, and retain the associated tradeoff. The hull penalty offsets about 29% of the local mount-and-thruster reduction, showing why the design cannot be judged from mount drag alone. Controlled physical resistance or propulsion measurements are needed to correlate the numerical prediction before using it in a mission-energy budget.

APPENDIX H
SIMULATION ENVIRONMENT AND VALIDATION METHOD

A. Course Coverage and Model Provenance

The simulation is being developed as a repeatable laboratory for the RobotX 2026 system of systems. Five generated worlds cover Proof of Readiness, Safe Passage, Infrastructure Survey, Coordinated Logistics, and a combined preliminary layout. The latter places Tasks 1–3 and shared delivery platforms inside Course B, with adjacent Courses A, C, and D represented. Coverage is tracked against individual handbook requirements [1]; availability of a world does not itself demonstrate successful task execution or complete finals coverage.

The current model set contains 56 course objects, including RoboBuoys, beacons, docking bays, delivery platforms, resource tins, pipeline elements, the pinger, and the launch/recovery deck. Scripts generate Simulation Description Format (SDF) assets from sourced dimensions; generated files are regenerated rather than manually edited. A 73-row source matrix distinguishes published dimensions, CAD-derived geometry, and explicitly labeled RX26 approximations. Core, advanced, and disruptive profiles vary beacon cues and distractors, including top-readable Task 1 cues requiring aerial observation. The vehicles comprise a BlueBoat, a CAD-derived eight-thruster BlueROV2 Heavy, and a PX4 X500-class model. The documented fleet integration checkpoint retains a two-thruster USV model; the measured four-thruster configuration remains a separate release requirement. The X500 model supports integration testing [31]; its dynamics must be matched to the X650 competition aircraft before drawing flight-performance conclusions.

B. Production Interfaces and Evidence Isolation

The simulator owns physics, sensor generation, and evaluation within a dedicated Docker image; separate vehicle images own autonomy. ArduRover and ArduSub SITL use the ArduPilot–Gazebo interface [30], while the UAV uses PX4 SITL [31]. The required production-parity USV transit path runs from the mission executive, Nav2, safety gates, MAVROS, MAVLink, firmware, and simulated thrusters. Other vehicle and precision-control profiles retain their approved interfaces and authority boundaries.

The end_to_end admission contract rejects production subscriptions to exposed Gazebo ground truth, unapproved velocity publishers, and autonomy images without pinned digests. Ground truth is available to evaluators, not field-facing observation products. Controller ownership is checked against the active profile, including explicit transitions between transit and marine precision control. Qualification must demonstrate these guards against deliberate negative cases; a successful open-loop world test does not establish the complete production command chain.

C. Venue, Physics, and Sensor Fidelity

The world origin is configured at $1.281\,467^\circ$ N, $103.855\,616^\circ$ E beside The Promontory, with vertical gravity approximately $-9.78035\,\mathrm{m/s^2}$. Magnetic conditions use WMM2025 evaluated for the run date [34]; atmospheric baselines use the 1991–2020 Changi climate normals [33]. Marina Reservoir is a freshwater environment [32]. The selected density, $995.65\,\mathrm{kg/m^3}$, is a modeling assumption to be checked against water temperature and field measurements, not a measured venue constant.

Gazebo Harmonic uses DART with Bullet collision detection and a projected Gauss–Seidel solver, a supported engine configuration [35]. Interactive and finer-resolution profiles use time steps of 4 and 1 ms, respectively, corresponding to 250 and 1000 physics updates per simulated second. Headless, unthrottled execution changes wall-clock throughput; it does not remove the selected time step. Contact forces, draft, and trajectories must be compared across step sizes before claiming numerical convergence.

VRX's Pierson–Moskowitz wavefield implementation supplies the wave model [36]. Six wind/wave presets span calm conditions to a squall stress case, with associated rain particles affecting camera and GPU LiDAR observations. These are controlled perturbations: their spectra and sensor effects require calibration before they represent measured Marina Bay conditions. Climate normals likewise provide baseline conditions rather than a forecast for an individual run.

The two USV camera models use 2448×2048 images at 20.5 Hz, a 12-mm lens configuration, and a 31.23° horizontal field of view. The pixel array and default frame rate follow the camera documentation [44]; lens and field-of-view values define the selected model. Housing geometry and optical placement are checked for unintended self-occlusion. The OS0-128 configuration uses 128×1024 samples over $360^\circ \times 90^\circ$, a 0.5–75-m range, and a selected 1-cm range-noise parameter.

At 10 Hz, the nominal sampling rate is

$$R = 128 \times 1024 \times 10 = 1.31072 \times 10^6 \text{ s}^{-1}, \quad (10)$$

before invalid returns and filtering. These values define the simulation configuration, not universal sensor accuracy guarantees. IMU, GNSS, and hull-contact observations complete the USV sensor set.

The UUV includes scene-dependent DVL returns using CAD-derived beam directions, pressure/depth sensing, an analytical imaging-sonar model, and four lights driven through the ArduSub output path. The DVL model conservatively requires all four transducers to remain submerged; valid bottom tracking also depends on geometry, range, and return quality. Acoustic pinger observations include slant range, bearing, elevation, propagation delay, and received signal-to-noise ratio. Six noise profiles exercise vessel traffic, rain, thruster vibration, basin reverberation, and multipath. They provide controlled localization stress tests pending calibration with recordings from the intended operating environment.

D. Independent MATLAB and Simulink Model Reference

A separate six-degree-of-freedom load reference addresses a practical integration problem: independently developed force plugins can duplicate buoyancy or added-mass terms while still producing plausible motion. The team implemented MATLAB-compatible reference equations, a numerical comparison harness, and a generator for three stateless Simulink load-reference models. Equivalent NumPy and PyTorch kernels provide an independent implementation comparison. The Simulink generator uses a MATLAB Function block to expose the load calculation for each vehicle type [69]; it is a reference model rather than a second live vehicle simulator.

The governing marine-craft formulation separates rigid-body inertia and Coriolis terms from added mass, damping, restoring loads, and applied forces [24], [70]. Hydrodynamic terms use velocity relative to the water. The reference computes thrust allocation, added-mass Coriolis terms, passive diagonal damping, and supplied immersed buoyancy. Gazebo retains ownership of time integration, rigid-body motion, gravity, and contact. The reference must not add a second weight term or duplicate native added mass. Spatially varying currents and full ocean CFD are outside this reference's demonstrated scope.

The recorded reference checkpoint used Octave to execute the independently coded MATLAB-compatible equations. Three algebra-fixture groups agreed with the comparison implementation within 1.42×10^{-14}; a 58-test fleet CPU suite also passed at that checkpoint. These are software/model checks, not measured vessel accuracy. The MATLAB/Simulink generator and comparison harness are implemented, but licensed MATLAB and Simulink execution is not recorded as completed. Running the generated models in those tools, retaining their outputs, and comparing calibrated vehicle responses remain required before claiming a demonstrated MathWorks simulation result.

The main lesson is that numerical agreement and physical validity answer different questions. The reference can reveal sign, frame, allocation, and force-accounting errors; it cannot identify real drag or thruster parameters without measurements. Parameters are to be fitted using measured mass, inertia, actuator geometry, thrust curves, and response maneuvers, then evaluated on independent maneuvers. The model is frozen before comparing controllers so that tuning improves the controller rather than silently simplifying the simulated environment.

E. Shared-World Integration and Recorded Evidence

The documented fleet checkpoint combines three dynamic vehicle identities in one Gazebo world with an independent DAVE pressure component operating as a read-only reference. It recorded 200 frames over 3.98 simulated seconds and 200 pressure samples. The raised, unpowered UAV moved downward under gravity, demonstrating that this composition contained a dynamic aircraft rather than a static visual fixture. This short open-loop run establishes basic shared-world and sensor integration; it is not autonomous flight, coordinated task completion, or a calibrated fleet trial.

The force-accounting review also identified that two non-overlapping pontoon displacement volumes on one BlueBoat link are legitimate. A guard that merely counted buoyancy plugins rejected that configuration incorrectly; the corrected check distinguishes disjoint volumes from duplicated or overlapping displacement. This finding improved the validation method without changing the physical requirement that each displaced volume be counted once.

Recorded scene geometry and link trajectories support an OpenUSD replay path for visual inspection. Replay keeps physics stopped and must

not feed perfect poses into production estimation. A dependency-isolation correction was required after an incompatible numerical-library environment interrupted conversion. Completion of the earlier physics run therefore does not imply successful textured replay or Isaac rendering. This separation keeps each result attached to the tool, source revision, and test stage that produced it.

F. Scoring, Repeatability, and Acceptance

Nineteen custom Gazebo systems include seven non-commanding course observers. Gate scoring requires passage through the opening in the correct direction and sequence, with failures reported separately. Survey coverage requires suitable range, pointing direction, and traversal speed. Landing evaluation checks position, attitude, descent rate, settling, and dwell time in the moving deck frame. Thresholds are traced to competition rules where specified; team-selected tolerances are labeled and fixed before a comparison.

The test sequence first checks model geometry, buoyancy, sensor frames, and actuator signs; then individual vehicle control; and finally coordinated missions across tiers and disturbance levels. Planned fault cases include stale observations, sensor loss, communication interruption, and conflicting control requests. Each run retains the world/model revision, container digest, firmware and autonomy versions, parameter set, random seeds, physics step, ROS 2 records, and observer results. Simulation-time scheduling reduces dependence on wall-clock load, but does not guarantee identical execution: repeated runs must quantify variation caused by asynchronous middleware and numerical effects. Performance is reported over repeated seeds using completion rate, collisions, tracking error, survey coverage, landing error, recovery time, and energy consumption. No numerical pass rate is claimed until these records exist.

G. Feedback to the Physical Design

Geometry and hydrostatic audits screen propulsion clearance, draft, freeboard, and DVL mounting before fabrication. Static displacement follows $\rho g V = mg$ for the modeled mass m and displaced volume V; trim additionally depends on mass and buoyancy distributions. Propulsion loading is interpolated from published T200 thrust/current data [40], combined with lighting and electronics loads, and integrated as $E = \int V_{\mathrm{b}}(t) I_{\mathrm{b}}(t)\, dt$ to estimate mission energy, where V_{b} and I_{b} are battery-terminal voltage and current. Battery capacity, discharge behavior, and reserve margins require separate measurements.

An observed instability in the force-based added-mass implementation when heave added mass exceeded dry mass is treated as a numerical-model finding. Force formulation, coupling, and time-step sensitivity must be resolved before inferring a physical ballast or trim limitation. Ultimately, simulation narrows hardware choices and prioritizes field tests; measured response and readiness demonstrations determine whether the resulting system is acceptable for deployment.

APPENDIX I
TEST PLAN AND EVIDENCE REGISTER

The following plan connects component checks to the mission capabilities described in the main report. It records the evidence needed for release; a planned test is not a reported pass. The confirmed field-testing total is 10 days on water. Additional trial counts, localization errors, and mission-success percentages are not asserted.

A. Sequence, Resources, and Acceptance

The sequence is dependency-based: complete output-inhibited bench checks before powered subsystem trials, validate each vehicle before cooperative trials, and retain a full mission rehearsal after any change affecting a shared interface. Sessions are scheduled around available laboratory and water access. A subsystem owner prepares the configuration and records; a second reviewer checks the outcome before the next stage. A failed safety or navigation prerequisite returns the system to the earlier stage rather than consuming rehearsal time on an inadmissible configuration.

TABLE X. Planned verification cases and retained acceptance evidence.

Stage	Setup and cases	Evidence and release decision
Electrical bench	Current-limited supply; actuator output inhibited initially. Apply load steps, brownout, local stop, transmitter loss, and reset sequences.	Record rail voltage/current and contactor/bus decay on a common timebase. Wireless disconnection must be below 2 s; no automatic restart is permitted [5].
Acoustic calibration	Electrical injection followed by surveyed wet-source positions; repeat with thrusters and converters active.	Retain gain/phase matching, delay residuals, ambiguous-arrival rejection, and bearing error against source geometry. Choose acceptance bounds before trials.
Perception replay	Separate capture sessions for training and evaluation; preserve paired views, timestamps, calibration, and illumination conditions.	Retain class confusion, false positives, missed targets, geometric association errors, and stale/unknown-state behavior. No detector score substitutes for a complete task result.
Vehicle control	Fixed firmware and parameter set; restrained or controlled-area trials before unrestricted maneuvering. Exercise saturation, lost navigation, handover, and stop.	Retain trajectory, tracking error, command owner, mode changes, and recovery. A fallback must be observable and must respect local clearance.
Fleet simulation	Pinned containers and model revisions; nominal cases plus sensor loss, delay, duplicates, and conflicting requests.	Scorers observe outcomes only. Confirm ground-truth isolation, one command owner, and correct cancellation before admitting a tier.
Water rehearsal	Controlled course, trained recovery crew, manual takeover, and agreed weather/energy limits. Repeat admitted task sequences.	Retain configuration, conditions, logs, outcomes, and discrepancies. Freeze the released configuration after regression and readiness review.

B. Results That Currently Inform the Design

The matched CFD comparison supports selecting the narrower v2 mount: predicted mount resistance is 11.80% lower, while total modeled resistance is 2.05% lower. Appendix G retains the force reports and identifies the remaining numerical and physical correlation work. The AFE model retains the gain, filter, and noise values in Table III; the coupling screen remains unresolved and therefore does not authorize hardware release. The fleet reference checks and short shared-world checkpoint in Appendix H establish their stated software and integration scope, without establishing autonomous mission performance.

These findings change the next test rather than replacing it. Mechanical checks follow the CFD geometry selection; interference and channel calibration follow the electrical model; and closed-loop, production-interface trials follow the open-loop fleet checkpoint. Physical measurements are reserved for effects that the current models cannot establish, including sealing, temperature rise, optical glare, acoustic multipath, and installed propulsion response.