# RoboSub 2026: Pushing The Frontier

Mecatron
Nanyang Technological University, Singapore

***Abstract*—Team Mecatron enters Robosub 2026 with a new motto "Pushing The Frontier", featuring two new Autonomous Underwater Vehicles (AUVs), *Hydra* and *Kraken*, each with different operational requirements to decentralize risks and optimize competition performance. The technical architecture integrates upgraded mechanical and electrical platforms for enhanced reliability, a revised autonomy stack prioritizing spatial perception and precise navigation, and a newly implemented, fully custom acoustic communication system. To validate these generational upgrades within the competition timeline, we followed a structured, parallelized testing and systems engineering strategy to ensure platform stability and mission compliance.**

Fig. 1: *Hydra* (left) and *Kraken* (right)

## I. Competition Strategy

For RoboSub 2026, Mecatron continues to build on the foundational principles of reliability and simplicity that shaped its 2025 debut, while officially adopting a new motto for this campaign: "Pushing the Frontier". This year, our team deploys two specialized AUVs: *Hydra* and *Kraken* (Figure 1), with the goal of achieving full task completion with higher accuracy and mission success rate.

### *A. Multi-Vehicle Approach*

The deployment of a second AUV marks a strategic shift from a single-platform approach to a task-oriented framework after taking stock of the previous year's iteration. During RoboSub 2025, the team relied exclusively on Kevin AUV, a single all-purpose vehicle; while that strategy gave the team a robust and practical foundation, it placed all mission requirements on one platform, thus limiting time efficiency and overall task coverage.

This year, tasks are systematically allocated based on the functional capabilities of each vehicle. This division allows each vehicle to be tuned to a narrower mission profile, having more deployment time per task. This strategic pivot ensures higher overall completion and significantly increases task capacity and tactical flexibility during competition.

Integrating a dual-AUV architecture reshapes our complexity and reliability trade-offs and parallel development pipelines through a strategic division of tasks. This approach gives the team a refined back-up plan: because each vehicle operates independently within its designated scope, a critical failure on one platform will have a mitigated impact on the scoring potential of the overall mission.

Furthermore, this dual-platform framework accelerates prototype testing and mitigates risks by attempting modifications systematically across two platforms. Hydra, an evolution of the 2025 Kevin AUV architecture, utilizes a modular hull for easier assembly and serves as the frame for electrical and software iterations. Concurrently, Kraken is used to evaluate mechanical redesigns; engineered with a configurable acrylic skeleton, it allows for quick iterations to adapt to regular changes in design requirements.

The complete team roster, subteam breakdowns, and specific engineering contributions are detailed in Appendix B.

### *B. Tasks Analysis and Approach*

This year, we divide tasks into *Essential* and *High-Stakes* tasks. Essential tasks are Gate, Slalom, Bin, whereas *High-Stakes* tasks are Torpedo and Picking, which are more challenging and risky due to the need for well-tuned control and robust perception for alignment. To decentralize risks and increase testing time for essential tasks, we allocate one *High-Stakes* task for each vehicle; the remaining tasks are split based on typical task duration, requirement and AUV capability. More specifically, we identify that both Bin and Picking tasks require clear down-facing vision; hence, we designed Hydra to have a down-facing camera, together with down-facing actuators such as dropper and gripper. Kraken will focus on Slalom and attempt Torpedo tasks, both of which require horizon-only perception. The allocation of the tasks aligns with key requirement for each tasks, shown in Table I.

*1) Hydra—Agility and Precision:* Hydra is assigned to Gate, Bin, and Picking, which require precise navigation and vertical positioning. Optimized for agile, multi-axis operations, Hydra leverages precise vertical mobility and downward perception to execute its tasks: utilizing short-range vision for marker deployment at the Bin and granular control for picking.

TABLE I: Task Allocation and Strategic Emphasis

| Task | Vehicle | Key Requirement | Priority |
|---|---|---|---|
| Gate | Both | Detection + alignment | High |
| Bin | Hydra | Precision positioning | High |
| Picking | Hydra | Manipulation + Stability | High |
| Slalom | Kraken | Tracking + yaw control | High |
| Torpedo | Kraken | Stable aiming | High |
| Acoustics | Conditional | Bearing estimation | Medium |

*2) Kraken—Stability and Manipulation:* Kraken is assigned to Gate, Slalom, and Torpedo, where hydrodynamic stability and fine-tuned velocity control are critical. Kraken utilizes a low-drag frame and precise position alignment to tackle its tasks: maintaining stable horizontal control through the Slalom debris channel and ensuring steady alignment to accurately fire through the Torpedo openings.

*3) Inter-Vehicle Communication:* Both Hydra and Kraken will be equipped with identical acoustic communication systems. To guarantee data integrity and eliminate transmission errors, we follow a robust 4-Way Handshake protocol for inter-vehicle communication. Consider a synchronized mission abort initiated by Kraken: Kraken first triggers a red LED visual cue and transmits an abort packet. Hydra receives the signal, acknowledges it, and echoes the exact message back to Kraken for verification. Kraken then evaluates the echoed data and transmits a final confirmation that the message matches the original intent. Only upon successful validation of this complete 4-way handshake will Hydra execute a parallel mission abort.

### C. Technical Masterplans

*1) Mechanical:* To successfully execute the dual-AUV deployment strategy, the mechanical architecture requires a complete overhaul, prioritizing modularity, rapid serviceability and control stability.

To achieve this, Hydra's skeletal frame deliberately exposes the internal volume to eliminate entrapped water to drastically improve the vehicle's acceleration and transient control response while simultaneously providing unobstructed access for wire routing and maintenance. Kraken's acrylic frame aligns the main enclosure across all three axes, which centers its Center Of Gravity (COG) and Center Of Buoyancy (COB) while minimizing rotational inertia. This directly translates into highly stable control and precise movement manipulation. Further details can be found in Appendix C and D for interested readers.

Both frames reduce the complexity of the mechanical design and incorporate additional penetrator holes, allowing both vehicles to easily adapt to tactical changes in the plan and experiment new features while minimizing risks. Redesign of thruster configuration for both vehicles also improves control stability, predictability, and less troubleshooting when tuning controllers. This refinement aids with every task, especially ones that require precise control and navigation such as picking and aiming tasks.

*2) Electrical:* The introduction of two vehicles necessitated a robust inter-vehicle link, motivating the electrical team to engineer a custom, in-house acoustic system. This acoustic development will lay the foundation for our 3-year plan for AUV design with the focus on acoustics localization and communication—an essential, mission-critical and differentiating element for underwater robotics.

The electrical system also undergoes significant upgrades with a better Controlled Power Distribution Unit (PDUC), allowing better monitoring, power management, and power budget. Using the same electrical architecture for both robots allows less testing time, less debugging time, and hence mitigates risks in operation.

*3) Software:* To manage the growing complexities associated with dual-AUV deployment, we embrace Docker containerization, prioritizing Docker images compatible with both `x86_64` and ARM-based platforms to ensure a seamless transition from development to hardware deployment and between the two AUVs. Furthermore, we design two variants for Docker image, one for perception and one for non-perception (Figure 2); this leads to greater separation of concerns and freedom of development for members without a GPU.

Apart from containerization, this cycle marks a critical paradigm shift toward spatial perception and navigation. In contrast to the purely vision-based, reactive approaches utilized in RoboSub 2025, our current strategy focuses on 3D pose estimation of objects within the competition arena, enabling prior path planning and thus finer control authority for more challenging tasks such as the Torpedo and Dropper missions. We also employ ROS2 Composition [1] in tandem with Docker containers to minimize the number of processes and leverage intra-process communication, thereby reducing latency and CPU usage (illustrated in Figure 2).

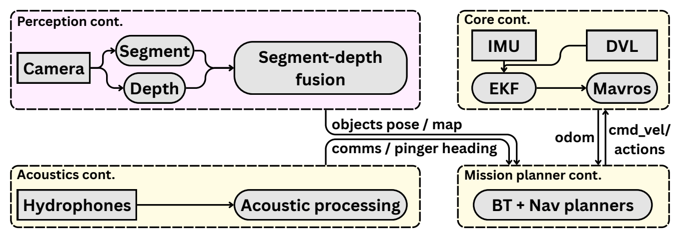


Fig. 2: Software diagram showing node composition and relation between processes. Purple and yellow colors are used to indicate perception and non-perception containers.

### D. Risk Management

Transitioning to a dual-vehicle strategy inherently increases system complexity, as does the reliability risks associated with aiming to execute all mission tasks with both AUVs. This explains why we enforce strict task boundaries for each vehicle to enable independent validation and parallel testing, significantly mitigating risks associated with either AUVs.

This risk decentralization extends to our software architecture. To manage complexity, we adopt the open-source Nav2 framework [2] for navigation and control, rather than engineering a fully proprietary solution. Furthermore, strictly standardizing communication interfaces (e.g., 3D poses and costmaps) between the perception and navigation stacks ensures high fault isolation and enables parallel subteam development.

Ultimately, our objective is to pursue targeted innovations while maintaining a strict focus on resolving the clear limitations observed during RoboSub 2025. By aggressively isolating hardware and software faults, we accept calculated short-term risks while establishing a highly reliable, scalable platform for future innovations.

## II. Design Strategy

### A. Mechanical Overhaul

To lay a strong foundation for software integration and specific task requirements, we design Hydra and Kraken with emphasis on enhanced stability, hydrodynamics, and modularity. Hydra is built on a tried and tested design, whereas Kraken introduces a novel design with a new electrical enclosure.

*1) Stability:* Both Hydra and Kraken utilize an over-actuated, 8-thruster configuration to improve control stability. Compared with the 6-thruster design from RoboSub 2025, this increased redundancy provides greater control separation between the vertical and horizontal axes. Center of mass and center of buoyancy are also carefully analyzed and adjusted to make the control more stable against external perturbations.

*2) Mass optimization:* During the design of Hydra, we employ topological optimization [3] to maximize the frame's strength-to-weight ratio while minimizing both material mass and hydrodynamic drag. Specifically, this computational method iteratively removes low-stress material from the structural ribs, generating an organic, low-volume lattice capable of rigidly supporting the central electrical housing and dual battery cylinders without compromising overall strength (supporting details in Appendix C).

*3) Hydrodynamics:* On the other hand, Kraken uses an airfoil-inspired outer shell that streamlines the AUV and minimizes drag. As shown in Appendix E, these design decisions help improve the overall hydrodynamics and optimize energy consumption during longer testing sessions.

*4) Modularity:* Both AUVs feature a dual-battery system with a large middle bay area for mounting various sensors and actuators. Hydra is composed of a large 3D-printed body featuring a compact cap and leg, with specialized mounts for thrusters, sensors, and actuators attached directly to the frame (Figure 3a). Kraken consists of an acrylic chassis paired with a 3D-printed shell and custom mounts for the battery, camera, and torpedo. This modular approach allows the positions and orientations of parts to be adjusted at runtime for fine-tuning, or to be replaced quickly in the event of a failure.

*5) New Electrical Enclosure:* Kraken is fitted with an experimental rectangular aluminum enclosure designed to maximize spatial efficiency by packing electronic components inside a compact space. As shown in Figure 3b, the electronics are mounted on modular trays for easy access from the top, allowing individual components to be adjusted or replaced rapidly when needed.

Conversely, Hydra retains the cylindrical pressure vessel from our previous iteration. While this leverages a proven design to mitigate unknown risk, the cylindrical geometry acts as an optimal pressure vessel, evenly distributing hydrostatic loads to ensure structural integrity.

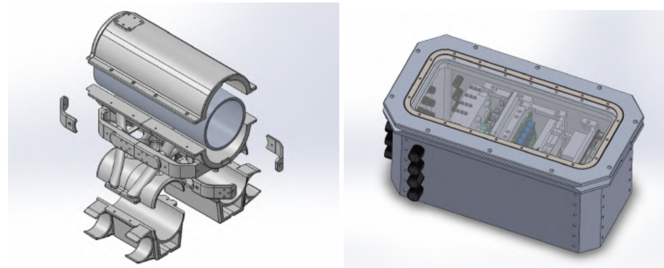

(a) Exploded view of Hydra external frame (b) Isometric view of Kraken electrical enclosure

Fig. 3: Mechanical overhaul

### B. Improved Actuation Systems

*1) Electromagnetic Dropper:* To guarantee the reliability of our marker deployment, we critically evaluate and resolve the corrosion risks inherent in our previous design. Our dropper uses an electromagnetic repulsion system to repel and drop the marker. However, prior iterations suffered from water ingress that exposed the bare components to rapid corrosion and subsequent electrical degradation. To eliminate this failure mode, the new electromagnet is fully encapsulated within a custom, epoxy-potted housing, providing a robust, watertight seal against the corrosive underwater environment. A polyurethane tube extends from this sealed housing to securely route and protect the actuator wiring.

To guarantee consistent deployment, the electromagnet is upgraded from a $8\,\mathrm{kg}$ to a $15\,\mathrm{kg}$ rating for stronger repulsive force. In addition, the torpedo-shaped marker (Figure 4a) is hydrodynamically optimized for a straight, stable descent. We significantly lower its center of gravity using a hollow upper cavity and a concentrated tungsten weight to mitigate pitch and roll, while $2\,\mathrm{mm}$ helical grooves and angled wings help stabilize axial spin.

*2) Torpedo Launcher:* This year, our goal is to improve the reliability of the torpedo launcher by mitigating premature deployment. As seen in Figure 4b, we upgrade the friction-based retention system in RoboSub 2025 to a physical mechanical lock. The new design utilizes a threaded bolt and nut assembly, where the back of the torpedo is threaded into a fixed T8 linear screws and nut. This provides a physical mechanical lock that secures the torpedo in place.

Upon activation, the rotation of the propeller generates the necessary torque to un-thread the torpedo, allowing a clean and controlled release. To achieve this, we also switch from a front propeller design to a rear propeller design.

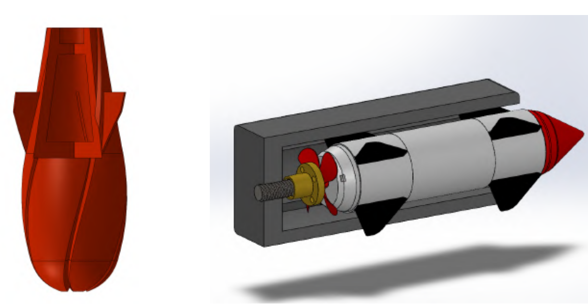

(a) Marker Design (b) Torpedo Assembly Sectional View

Fig. 4: Improved Actuation Systems

### C. Electrical System

Both AUVs employ safety-focused power distribution architectures to ensure reliable operation of all onboard subsystems. Each vehicle comprises an Actuation Electrical Subsystem (AESS) and a Compute Electrical Subsystem (CESS). The AESS distributes and monitors power for all actuators, while the CESS supplies non-actuating, compute-centric subsystems such as sensors and onboard computers (see Appendix I).

Both subsystems are built around a shared, custom-designed PDUC platform. In the AESS, the PDUC delivers high-current power to the Electronic Speed Controllers (ESCs) and interfaces with the actuator daughterboard. In the CESS, it supplies power to multiple voltage regulators that generate stable rails for sensitive electronics. A detailed render of the respective assemblies can be found in Appendix K.

Upstream, a dedicated board called ReRoute, housed within the battery enclosure, provides system-level power switching and protection, including overvoltage, undervoltage, and overcurrent safeguards.

This modular architecture replaces the previous generation of specialized power boards, offering a higher power budget, reduced footprint, integrated monitoring, and a consistent distribution strategy across both subsystems.

### D. Acoustics Systems

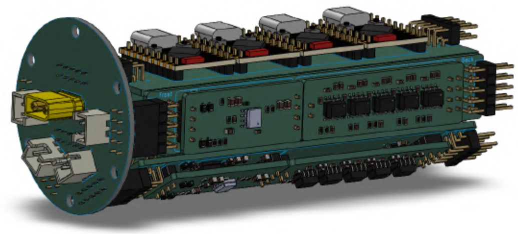

Fig. 5: Acoustics Module

To satisfy stringent real-time constraints and achieve low end-to-end latency, the acoustics system is implemented as a fully edge-based architecture. In the previous system iteration, acoustic data acquisition was performed by a Teensy Microcontroller (MCU), while downstream signal processing was offloaded to the Single Board Computer (SBC). In the current iteration, both data acquisition and signal processing are consolidated onto ESP32 MCUs. Each ESP32 performs real-time acoustic sampling, local signal processing, communication, and localization entirely on-device. This architectural unification eliminates inter-processor data streaming, thereby reducing end-to-end latency to the order of a single processing frame. A detailed system model can be found in Appendix J. To save space, the acoustics module is designed to fit into a standalone cylinder as seen in Figure 5.

*1) Audio Acquisition:* In pursuit of a unified acoustic system, the audio acquisition pipeline begins by continuously sampling the input signal and storing it in a triple circular buffer. This buffering strategy enables simultaneous acquisition, processing, and transmission without data loss. For signal processing, overlapping Hann windows [4], [5] are applied to the buffered audio frames prior to transformation into the frequency domain using the Fast Fourier Transform (FFT).

*2) Communications:* We employ Frequency Shift Keying (FSK) [6], specifically 4-FSK, with carrier frequencies at 10, 12.5, 13.75, and 15 kHz. Following the FFT, single-instruction multiple-data (SIMD) acceleration [7] is used to efficiently identify the frequency bin with the highest magnitude. This dominant frequency is then used for symbol detection and cross-correlation-based decoding.

*3) Localization:* We compute two sets of pairwise generalized cross-correlation with phase transform (GCC-PHAT) [8]. These are applied to the front–back and left–right hydrophone pairs, respectively. By analyzing the relative Time-Difference-Of-Arrival (TDOA) estimates from these orthogonal pairs, we infer a relative bearing vector. To improve reliability and filter multipath acoustic noise, the final localization output is quantized into discrete 45-degree angular intervals, ensuring a robust, actionable directional estimate for mission navigation.

### E. Spatial Perception

To overhaul the perception stack and enable precise 3D pose estimation, we focus on achieving three important criteria.

*1) High-Fidelity Depth-Segmentation Fusion:* The stack must extract accurate pixel-level segment masks rather than loose bounding boxes, which introduce prohibitive depth estimation noise for non-rectangular objects. To achieve this, we deploy YOLO Segmentation [9] for robust feature extraction and fuse its output with metric depth maps generated by Depth Anything 3 (DA3) [10]. A custom fusion node then merges the segment and depth masks to compute the exact 3D coordinates of target objects relative to the camera frame (Figure 6).

*2) Task-Specific Modularity:* The pose-estimation pipeline should be able to accommodate the different objects in the scene, as each object requires a different method to interpret orientation from segment or depth mask (for instance, estimating orientation for a Torpedo canva and a flat-lying object are different). We address this requirement by implementing a C++ plugin architecture; dynamically loadable at runtime, these plugins execute task-specific localization algorithms such as dedicated packages for the Torpedo, Dropper, and manipulation tasks, providing extensive modular flexibility while leveraging the execution speed of native C++.

*3) Globally-Consistent Pose:* To fuse objects' position (in pixels), depth and orientation (relative to camera), we design a tracker node using an Unscented Kalman Filter (UKF) [11] to estimate their pose in the global frame. The UKF is chosen over the Extended Kalman Filter (EKF) for its superior performance in this highly nonlinear problem [12]. The tracked pose—corresponding to the common interface—will be sent to the navigation and planning stack for mission execution.

### F. Navigation and Mission Planning

To establish full three-dimensional (3D) spatial navigation, our 2026 architecture shifts from reactive bounding-box control to a unified global planning and execution paradigm.

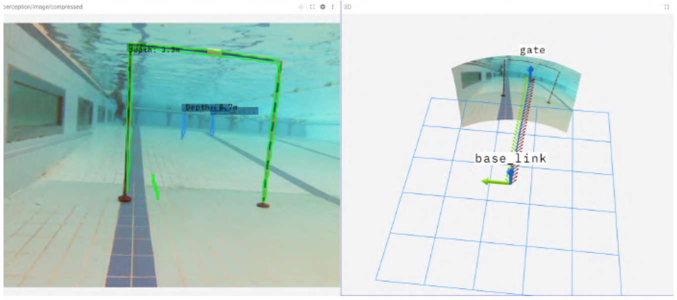


Fig. 6: Spatial Perception and Navigation

Following a comprehensive review of alternative architectures, such as Aerostack2 [13] and experimenting with custom in-house navigation implementations, we adopted the open-source Navigation2 (Nav2) framework [2]. This selection was driven by the framework's operational maturity, extensive library of baseline planners and controllers, and native compatibility with Behavior Tree design methodologies.

To extend this framework into underwater environments, we developed Nav3—a specialized suite of custom 3D planners, controllers, and recovery behaviors engineered as modular Nav2 plugins. This approach ensures stable, six-degree-of-freedom navigation for the AUVs while establishing a highly scalable software abstraction layer that can be ported directly to other autonomous mobile platforms.

The Behavior Tree (BT) framework remains the core of the mission planning system. Analysis of the previous year's performance revealed that overly granular nodes resulted in excessively complex tree logic that was difficult to design and monitor. This year, we encapsulate low-level control decisions into higher-level actions (Figure 7), improving the debugging process and adhering more closely to the original principles of designing behavior trees [14].

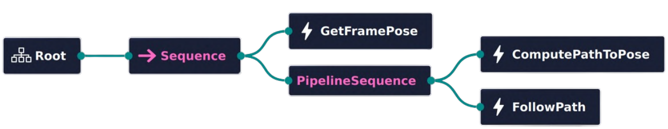


Fig. 7: A typical Behavior Tree for a task

## III. Testing Strategy

Our testing paradigm similarly adheres to a structured Systems Engineering approach [15], progressing sequentially from component-level verification to full-system validation. Initial testing focuses on characterizing individual sensors and components, followed by compartmentalized subsystem validation to evaluate functional performance in isolation (Section III-A, Section III-B and Appendix F). Final validation occurs during full-systems integration within pseudo-deployment environments (Section III-C), both in a SITL simulation setup called *UnityMDS* [16] and in semiweekly pool tests. This bottom-up verification pipeline ensures that new hardware and software features can be introduced safely without compromising backward compatibility or system stability.

To execute this testing strategy within a compressed sub-two-month competition timeline, the development cycle relies on parallel testing enabled by strict architectural decoupling. Designing subsystems with well-defined interfaces prevents critical-path bottlenecks and eliminates cross-team dependencies. Subteams can validate their respective modules concurrently, significantly accelerating development speed and ensuring that field deployments focus on improving inter-systems interactions rather than debugging individual subsystems.

### A. Acoustic Communication Testing

The validation of the acoustic subsystem starts with the characterization of hydrophones and transmitters. For the choice of the waterproof AS-01 hydrophone, we assessed its sensitivity by analyzing the signal produced by a Teledyne Benthos pinger with $4\,\mathrm{ms}$ pulses at frequencies ranging from $25\,\mathrm{kHz}$ to $45\,\mathrm{kHz}$; preliminary results on the oscilloscope demonstrated that this hydrophone model was capable of picking up $25\,\mathrm{kHz}$ to $100\,\mathrm{kHz}$ with high reliability. Higher frequencies were detected due to harmonics generated in our small lab pool. The initial voltage readings also gave us insight into how much gain the pre-amplifiers needed to provide so that the signal can be readable to a Teensy MCU.

On the transmitter side, we used a low-cost, epoxied piezo-transducer to establish an operational baseline. Using a Teensy MCU to generate a square wave and modulating the output frequencies between certain channels, we were able to create an initial prototype capable of sending simple bit-size messages to the hydrophone.

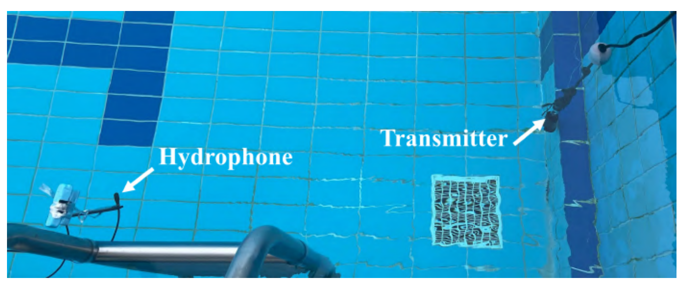


Fig. 8: Validation of Acoustics Comms in the swimming pool

To move beyond the initial prototyping stage to a reliable acoustic bridge, we investigated the acoustic profile of our swimming pool to select appropriate channels for communication (Figure 8). Initial results showed that the swimming pool generated substantial noise below $7\,\mathrm{kHz}$ and our vehicle thrusters generate noise between $20\,\mathrm{kHz}$ and $25\,\mathrm{kHz}$; these frequencies should be avoided to prevent false triggers.

Using these inisights, we selected four frequencies as our transmission channels for further tests in multiple environments. First, in a small water tank located inside our laboratory where echo and signal noise dominated, we focused on iterating error correction and noise rejection algorithms. Once consistent transmissions inside this water tank is achieved, we moved on to a competition-size pool. With signal attenuation becoming a greater concern, we investigated if larger transmission power or receiver gain was needed for signals to be transmitted reliably, while balancing with greater multipath reflection effects as a result of higher energy signals. Analyzing data-logs produced by the Teensy led us to conclude that signal integrity remained the more important reliability factor compared to signal amplitude in the $10 - 25\,\mathrm{m}$ range.

### B. *Software Autonomy Testing*

Individual ROS2 packages are systematically validated within the simulation environment to assess operational suitability and isolate software bugs prior to integration into the autonomy stack. To evaluate the localization pipeline, we generated synthetic acceleration and velocity datasets with predefined noise parameters within UnityMDS; discrepancies between simulated and empirical performance were analyzed to refine the characterization of the physical IMU and DVL sensors. To determine the viability of the Nav2 framework [2] for underwater applications, the stack was executed concurrently with a high-fidelity digital copy of Hydra in UnityMDS to replicate full 3D kinematics. The evaluation confirmed the framework's suitability for long-term deployment, provided custom 3D planners and controllers are integrated. Spatial perception packages underwent similar rigorous validation in UnityMDS (Figure 9), ensuring that these custom perception packages were reliable and bug-free before they reach deployment maturity; consequently, the main focus of pool testing for perception was fine-tuning machine learning models to mitigate the sim-to-real domain gap. Examples of software testing in simulation are archived in Appendix H.

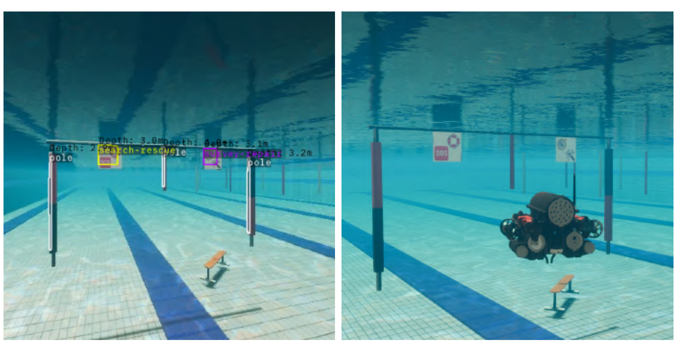

Fig. 9: Navigation and perception testing in UnityMDS (left and right are robot and 3rd-person point-of-view respectively)

Subsequently, individual software subsystems (localization, navigation, perception, etc.) are validated and integrated into the overall autonomy stack in a gradual manner. While localization and navigation were paired to tune the odometry-guided control in the real environment, we experimentally deployed the perception stack and collected additional image data to improve its reliability in subsequent sessions. Once the core control and perception pipelines were independently validated, they were combined to perform the Gate and Slalom tasks, giving software team initial opportunities to assess operational viability with RoboSub mission requirements.

Throughout this integration process, we continually analyzed performance discrepancies such as localization drift, perception noise and hydrodynamic drag between UnityMDS and physical reality. These insights allowed us to modify the digital copy of the AUVs, driving the simulator toward a true digital twin framework for long-term development.

### C. *Systems-Integration Validation*

During full-system, hardware-software validation within real deployments, we prioritize system exposure and task coverage over granular task fine-tuning. Given the compressed two-month runway leading up to RoboSub, this strategic focus allows us to simultaneously evaluate multiple subsystems, isolating architectural design flaws early and maximizing our timeline for timely corrective adjustments. This testing strategy has proven highly effective across several scenarios: (1) periodically evaluating acoustic communication while the AUVs execute unrelated mission profiles helps identify specific zones in the pool characterized by high acoustic noise and multipath reflections; (2) combining Slalom and Torpedo tasks or Gate and Dropper tasks strategically couples an actuation-heavy task with a non-actuation task, making our parallel development and debugging timeline more manageable. More examples are given in Appendix G.

Towards the final three weeks of the testing cycle, we shift the focus towards granular task fine-tuning to maximize mission success rate and ensure that *Essential* tasks (Section I-B) can be accomplished with high confidence. To achieve this, we increase pool testing frequency to three-to-four sessions per week, aiming to cover every *Essential* task alongside at least one *High-Stakes* task per vehicle. This elevated testing frequency and heavy deployment load provide the critical runway necessary to isolate and rectify any remaining or previously undiscovered system reliability issues before competition.

## IV. Conclusion and Future Direction

RoboSub 2026 marks a generational leap in Team Mecatron's development cycle, featuring a dual-AUV setup, overhauled mechanical designs, custom-built acoustic capabilities, and novel spatial perception and navigation pipeline. For RoboSub 2027 and beyond, we continue to commit to the multi-AUV strategy with greater emphasis on inter-vehicle communication and collaboration, explaining why short-term risks induced by the current year's upgrades are acceptable trade-offs for our long-term vision and competitiveness.

## Acknowledgment

Mecatron extends its deepest gratitude to our distinguished sponsors and supporters, detailed in Appendix A. They provided invaluable technical insights and resources which significantly accelerated our technological advancements this year. From Nanyang Technological University (NTU), Mecatron would especially like to acknowledge Ms. Era Carmen Chan from the College of Engineering and Associate Professor T. H. New from the School of Mechanical and Aerospace Engineering for their sustained administrative support and mentorship throughout the competition cycle.

## References


[1] S. Macenski, A. Soragna, M. Carroll, and Z. Ge, "Impact of ros 2 node composition in robotic systems," *IEEE Robotics and Automation Letters*, vol. 8, no. 7, pp. 3996–4003, 2023.

[2] S. Macenski, F. Martin, R. White, and J. Ginés Clavero, "The marathon 2: A navigation system," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.

[3] L. Siva Rama Krishna, N. Mahesh, and N.Sateesh, "Topology optimization using solid isotropic material with penalization technique for additive manufacturing," *Materials Today: Proceedings*, vol. 4, no. 2, Part A, pp. 1414–1422, 2017, 5th International Conference of Materials Processing and Characterization (ICMPC 2016). [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2214785317301633

[4] F. Harris, "On the use of windows for harmonic analysis with the discrete fourier transform," *Proceedings of the IEEE*, vol. 66, no. 1, pp. 51–83, 1978.

[5] A. Nuttall, "Some windows with very good sidelobe behavior," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 29, no. 1, pp. 84–91, 1981.

[6] M. Simon and M. Alouini, "A unified approach to the performance analysis of digital communication over generalized fading channels," *Proceedings of the IEEE*, vol. 86, no. 9, pp. 1860–1877, 1998.

[7] M. J. Flynn, "Some computer organizations and their effectiveness," *IEEE Transactions on Computers*, vol. C-21, no. 9, pp. 948–960, 1972.

[8] C. Knapp and G. Carter, "The generalized correlation method for estimation of time delay," pp. 320–327, 1976.

[9] G. Jocher, J. Qiu, and A. Chaurasia, "Ultralytics YOLO," Jan. 2023. [Online]. Available: https://github.com/ultralytics/ultralytics

[10] H. Lin, S. Chen, J. H. Liew, D. Y. Chen, Z. Li, G. Shi, J. Feng, and B. Kang, "Depth anything 3: Recovering the visual space from any views," *arXiv preprint arXiv:2511.10647*, 2025.

[11] S. J. Julier, J. K. Uhlmann, and H. F. Durrant-Whyte, "A new approach for filtering nonlinear systems," in *Proceedings of 1995 American Control Conference-ACC'95*, vol. 3. IEEE, 1995, pp. 1628–1632.

[12] J. J. LaViola, "A comparison of unscented and extended kalman filtering for estimating quaternion motion," in *Proceedings of the 2003 American Control Conference, 2003.*, vol. 3. IEEE, 2003, pp. 2435–2440.

[13] M. Fernandez-Cortizas, M. Molina, P. Arias-Perez, R. Perez-Segui, D. Perez-Saura, and P. Campoy, "Aerostack2: A software framework for developing multi-robot aerial systems," 2023.

[14] M. Colledanchise and P. Ögren, *Behavior Trees in Robotics and AI: An Introduction*. CRC Press, 2018.

[15] I. S. E. Handbook, "A guide for system life cycle processes and activities," 2015.

[16] T. L. Vu, M. C. Togamulia, and S. D. Upadhyay, "UnityMDS: A High-Fidelity Multi-Drone, Multi-Domain Simulator for Maritime Robotics," *Proceedings of the IEEE/MTS OCEANS Conference*, in press.


## Acronyms

**AESS** Actuation Electrical Subsystem.
**AUV** Autonomous Underwater Vehicle.

**BT** Behavior Tree.

**CESS** Compute Electrical Subsystem.
**COB** Center Of Buoyancy.
**COG** Center Of Gravity.

**DA3** Depth Anything 3.
**DVL** Doppler Velocity Log.

**EKF** Extended Kalman Filter.
**ESC** Electronic Speed Controller.

**FFT** Fast Fourier Transform.
**FSK** Frequency Shift Keying.

**IMU** Inertial Measurement Unit.

**MCU** Microcontroller.

**Nav2** Navigation2.
**Nav3** Navigation3.

**PDUC** Controlled Power Distribution Unit.

**ROS2** Robotic Operating System.

**SBC** Single Board Computer.
**SITL** Software In The Loop.

**TDOA** Time-Difference-Of-Arrival.

**UKF** Unscented Kalman Filter.

## Appendix A
## Sponsors


### A. Gold Sponsors

Future Systems & Technology Directorate (FSTD) — For providing critical funding that empowers Mecatron to push technological boundaries.

Singapore Maritime Foundation (SMF) — For their generous monetary support and invaluable networking opportunities.

### B. Silver Sponsors

DSO National Laboratories — For their generous monetary support.

VectorNav — For their generous equipment sponsorship.

WaterLinked — For their generous equipment support.

### C. Bronze Sponsors

Zen4Blue — For their constant support and trust in Mecatron's mission.

Espressif Systems — For their generous equipment support.

Aquarian Audio — For their generous equipment support.

DeepWater Exploration Inc. — For their generous equipment support.

Solidworks — For their generous software support.

### D. Supporters

NTU College of Engineering — For providing essential institutional backing within the University. We would especially like to thank Ms. Era Carmen Chan, Senior Executive, Student Experience, College of Engineering, and Associate Professor New Tze How, Daniel, School of Mechanical & Aerospace Engineering, for their continuous administrative support, technical guidance, and tireless student advocacy.

The James Dyson Foundation — For supporting Mecatron's engineering development through the Dyson-NTU Studio.

National Institute of Education Physical Education and Sports Science (PESS) Sports Facilities Office — For enabling Mecatron to conduct crucial pool tests to validate our systems.

## Appendix B
## Team Roster and Contributions

### A. Team Leads

Han Zhou Arthur Lim — CEO
Tien Luc Vu — CTO, Technical Documentation IC
Jennifer Zhen Ni Goh — COO, CFO, Website IC
Scott Jia Xu Cheng — CRO, Electrical Lead
Zi Xuan Tan — CWO, Brand Design Manager, Website IC
Maximus Zheming Yeo — Mechanical Lead
Gia Bao Bui — Software Lead

### B. Members

Akil Rafid — Mechanical
Do Phuoc Truong Nguyen — Mechanical, Video IC
Gerardo Nandavardhana Abdipranoto — Mechanical
Yi Liang Soon — Mechanical
Aung Ko Oo — Electrical
Cheng Jing Hao — Electrical
Hai Duong Trinh — Electrical
Chuanhua Liu — Software
Gia Minh Nhan — Software
Poh Hian Chee — Software

## Appendix C
## Hydra Frame Geometry

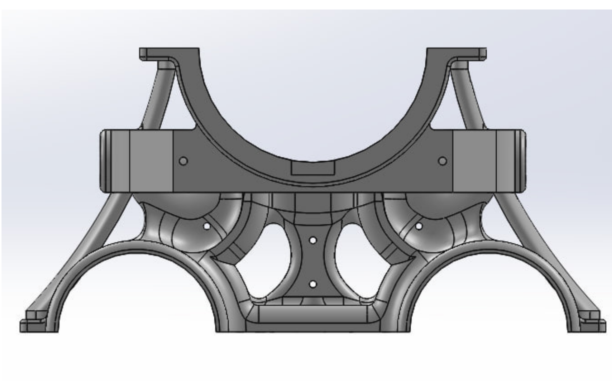

(a) Front view

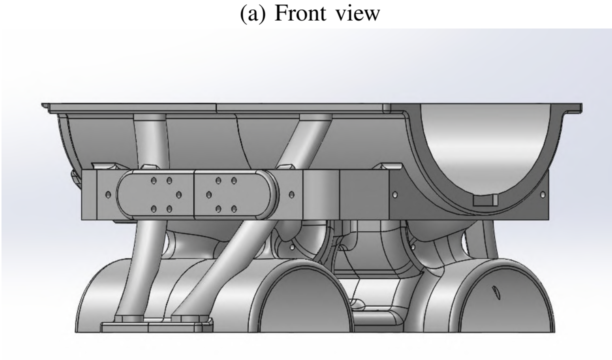

(b) Perspective view

Fig. 10: Hydra main frame

1. Hydra's chassis departs from traditional bulky frames through the deliberate application of topological optimization. This computationally driven process generated an organic, low-volume lattice framework. The resulting structure provides highly rigid, dedicated load-bearing pathways that perfectly contour to support the central electrical housing and dual lower battery cylinders, maximizing the frame's overall strength-to-weight ratio without relying on excess material.

2. Hydrodynamic Optimization and Added Mass Reduction Beyond structural efficiency, this skeletal architecture serves a critical hydrodynamic function. Traditional flooded hulls and heavily contoured chassis often suffer from the "bucket effect," trapping significant volumes of dead water within internal cavities. When the AUV attempts to accelerate or turn, the thrusters must overcome the inertia of both the vehicle and this heavy entrained water (added mass). By deliberately utilizing a flow-through, open-frame architecture, Hydra prevents water entrapment. Eliminating this localized added mass drastically reduces the vehicle's inertial payload, directly resulting in the superior acceleration, rapid deceleration, and sharp transient control response required for Hydra's agile task execution.

3. Accessibility and Systems Integration Strategically, this open-architecture geometry aligns with our focus on systems engineering and rapid iteration. The deliberate absence of an enclosing outer shell or deep internal pockets provides the electrical team with immediate, 360-degree physical access to the core components. This unobstructed spatial layout significantly streamlines complex wire routing, facilitates rapid poolside hardware debugging, and minimizes the downtime required for modular sub-assembly replacements during the compressed testing phase.

## Appendix D
## Kraken Mass Analysis

1. Geometric Optimization and Mass Reduction The transition from a traditional cylindrical pressure vessel to a custom CNC-machined rectangular aluminum enclosure was a primary driver in optimizing Kraken's overall spatial efficiency. Cylindrical hulls inherently enclose unutilized "dead" air volume when housing flat, planar components like printed circuit boards (PCBs). This excess internal volume generates significant positive buoyancy, which historically forced the mechanical team to add heavy physical ballast to achieve neutral buoyancy. By engineering a custom, form-fitting rectangular enclosure, we drastically minimized the vehicle's total volumetric displacement. This proportional reduction in buoyancy eliminated the need for excessive ballast, resulting in a significantly lighter overall AUV with a vastly improved thrust-to-weight ratio and transient agility.

2. CG and CB Co-Location for Passive Stability As illustrated in the vehicle's orthographic projections (Front, Side, and Top views), the rectangular enclosure architecture enables perfectly symmetrical mounting along both the longitudinal and transverse axes. Because the heavy electrical components are distributed evenly across the flat bottom of the enclosure, we were able to perfectly centralize the vehicle's Center of Gravity (CG). Furthermore, the symmetrical external geometry of the enclosure ensures the Center of Buoyancy

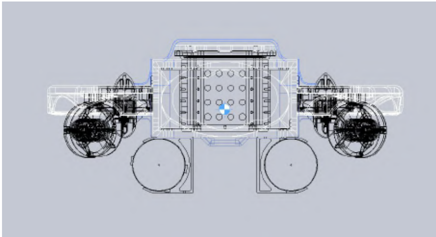

(a) Front view

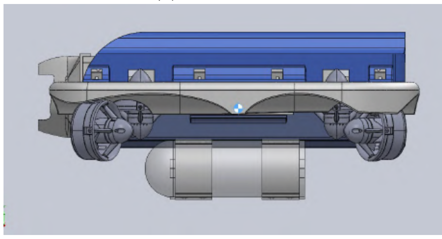

(b) Side view

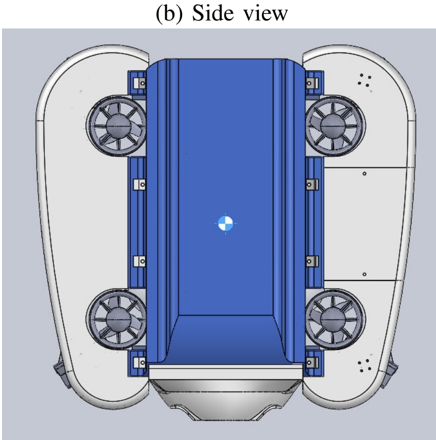

(c) Top view

Fig. 11: Kraken CG analysis: (a) front view, (b) side view, and (c) top view.

(CB) perfectly aligns with the CG at the absolute geometric center of the AUV. By eliminating the spatial offset between the CG and CB, the design passively eradicates parasitic pitching and rolling moments, ensuring inherently stable flight dynamics and simplifying the active control requirements for the software team.

3. Modular Electrical Architecture and Serviceability Beyond hydrodynamic and mass optimizations, the planar interior of the rectangular enclosure establishes the foundation for a highly accessible, modular electrical layout. We engineered swappable, standardized internal mounting plates that allow the entire electronics stack to be extracted as a single, cohesive unit. This completely eliminates the restrictive space constraints and cumbersome wiring harnesses typical of cylindrical slide-out trays. Consequently, the electrical subteam can perform rapid poolside debugging, seamless hardware modifications, and highly parallelized maintenance with unobstructed access to all internal components.

## Appendix E
## Kraken CFD Analysis

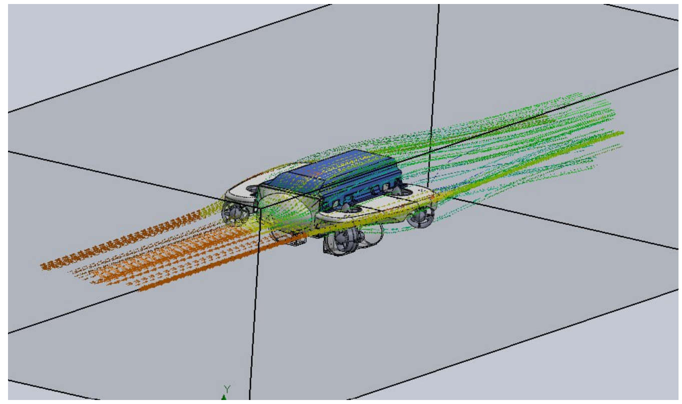

Fig. 12: Kraken's CFD Analysis

| Parameter | Without Shell | With Shell |
|---|---|---|
| Drag during Sideways motion | 36.1N | 33.2N |
| Drag During Forward Motion | 25N | 17N |

Hydrodynamic Optimization: The vehicle's lateral wings are explicitly modeled after aircraft airfoils to delay flow separation, reduce overall form drag, and minimize turbulent wake. Computational fluid dynamics (CFD) analysis validates this design, demonstrating a 32% reduction in forward-motion drag (from 25N to 17N) and an 8% reduction in lateral drag (from 36.1N to 33.2N) when the outer skin and airfoil profiles are fully integrated. This severe reduction in hydrodynamic resistance optimizes thruster power consumption and extends overall mission duration. Additionally, the streamlined profile is achieved using a detachable front nose cone.

## APPENDIX F
## ELECTRICAL TEST PLAN

| Components | Checklist |
|---|---|
| PDUC | 1) Able to operate safely across the full input voltage range from 14.4 V to 16.8 V<br>2) Able to handle a maximum continuous current load of 50 A<br>3) Multi-stage ON/OFF sequence verified functionally:<br>• ON (Physical): Master Switch powers on-board voltage regulators<br>• EN (External): An external signal from another PDUC to enable the load switch<br>• EN (Physical): A physical Emergency Stop switch to enable the load switch<br>• EN (MCU): An internal signal from the on-board MCU to enable the load switch |
| ReRoute | 1) Normal Operation verified under nominal input voltages and standard operational loads<br>2) Under-Voltage Protection (UVP) functional sequence verified:<br>• Cutoff: Load is successfully disconnected when the battery voltage drops down to 14.0 V.<br>• Latch: System remains in an OFF/disconnected state as voltage hovers below recovery levels.<br>• Recovery: Load connection recovers once the battery is charged back above the 14.4 V threshold.<br>3) Over-Voltage Protection (OVP) functional sequence verified:<br>• Cutoff: Voltage cutoff triggers successfully at 19.2 V.<br>• Inductive Transient Recovery: If triggered by inductive flyback voltage, ReRoute resets and restores power within a few seconds once the voltage stabilizes back to nominal levels.<br>• Overcharge Latch: If triggered by an overcharged battery supply, the system latches OFF and requires physical battery removal and discharge below the recovery threshold before reactivation is permitted.<br>4) Over-Current Protection (OCP) safety margins verified:<br>• Continuous Rating: XT60 connectors maintain safe thermal stability under continuous loads up to 60 A.<br>• Fuse Coordination: In-line 50 A fuse reliably clears short circuits and overcurrent events before reaching the thermal structural limits of the board connectors. |
| Actuator | 1) Able to power entire board with a single 4S battery pack (range: 14.4 V to 16.8 V)<br>2) Turning off logic power does not cause unexpected actuator firing or state changes<br>3) Able to handle voltage fluctuations and sag during high-current actuator spikes<br>4) Able to regulate power for on-board logic circuits<br>5) Able to output stable, current-limited power to all connected actuator channels<br>6) Able to drive and control status indicator LEDs<br>7) Able to run all actuators simultaneously at maximum peak current load<br>8) Able to safely power cycle and independently enable/disable all actuator channels |
| Acoustics | 1) Power Stage: Able to provide stable, low-noise 12.0 V, 10.0 V, 5.0 V and 3.3 V power rails from an input range of 14.4 V to 16.8 V, suitable for robust and low-noise audio acquisition.<br>2) Preamp: Able to correctly bias the hydrophone.<br>3) Preamp: Able to clearly observe the hydrophone output on an oscilloscope and reliably detect signal discontinuities when the hydrophone is known to be operating normally.<br>4) Preamp: Provides sufficient gain and bandwidth to preserve signal integrity across the expected frequency range of the hydrophone.<br>5) Tone Decoder: Able to reliably detect and identify a tone at a specified target frequency when an acoustic ping is applied to the hydrophone, with adequate noise immunity. |

TABLE II: Finite Element Analysis For Components.

| **Category** | **Testing Steps** |
|---|---|
| Communication | 1) Pinging PDUCs, Actuator, Acoustics<br>2) Check if there is connection |
| Power Passive Test | • CESS-EN = $MCU \wedge EmergencyStop$<br>• AESS-EN = $CESS - EN \wedge MCU \wedge EmergencyStop$<br>$\Longrightarrow$ If all conditions are met, AESS is ON. CESS must be ON for AESS to be ON. |
| Power Active Test | 1) SBC $\Longrightarrow$ Actuator board activates dropper launch $\Longrightarrow$ Dropper pin is triggered.<br>2) SBC $\Longrightarrow$ Actuator board activates the gripper $\Longrightarrow$ Gripper pin is triggered.<br>3) SBC $\Longrightarrow$ Actuator board activates torpedo $\Longrightarrow$ Torpedo pin is triggered. |
| Acoustics | 1) Connect TX- to SBC and Pinger. Connect RX- to hydrophones and oscillator.<br>2) Put Pinger and hydrophones on two sides of the pool.<br>3) SBC $\Longrightarrow$ pinger pings.<br>4) Acoustic model detects the signal.<br>5) Observe peaks in $V_{\text{rms}}$. |

TABLE III: Integration Testing Features.

## APPENDIX G
## POOL TESTS AND ADJUSTMENT PLAN

Our in-water tests are conducted in the swimming pool of National Institute of Education (NIE) inside Nanyang Technological University (NTU). Logistics for pool testing include replicas of competition tools built by our mechanical team (Figure 13), tools box, charged batteries, latest software, and quality-checked thrusters and other sensors such as Inertial Measurement Unit (IMU), Doppler Velocity Log (DVL) and cameras.

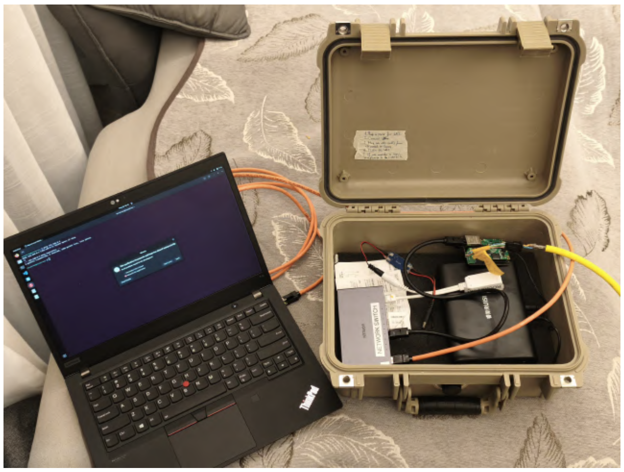

(a) On-shore station

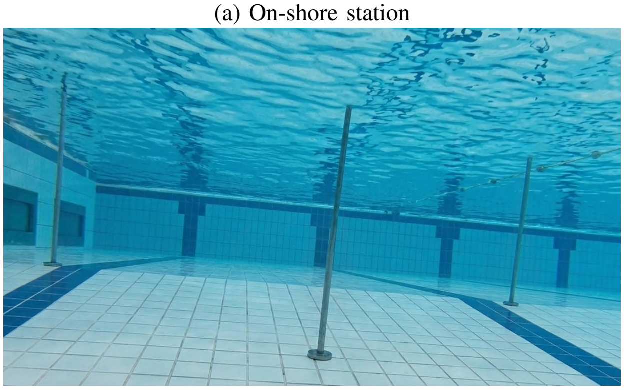

(b) Props for slalom

Fig. 13: Props and On-shore station used in pool tests.

On the day before, the team members attending the pool test have a meeting to set testing goals and prepare any preliminary checks. Before arrival at NIE's swimming pool, spare electrical components and mechanical tools are prepared following the checklist for any upcoming small scale problems. An individual testing session of AUV spans most of the day, with swimmers to watch the robot's in-water performance and act on any underwater emergency problems and mechanical and electrical members on standby to deliver any component replacements, so that any minor issues can be fixed on the spot. Meanwhile, software members always work as pairs at the on-shore station, with both laptop connected to the vehicle. One software operator monitors and executes the main tasks, while the other one prepares rosbag for immediate debugging, and acts as a second brain to support the main operator. Other variables such as lightning and weather conditions are taken into account.

Pool tests vitally reveal blind spots that can only be detected when the AUV is fully deployed underwater. Some major findings that heavily modify our strategy are as follows:

- Unstable Depth Control by Erroneous EKF3 Fusion: Early pool tests exhibited severe depth control instability, characterized by significant vertical oscillations. A rigorous investigation into the localization and control pipelines revealed a flaw in the ArduPilot EKF3 fusion logic: the system automatically integrated unreliable DVL Z-velocity data into its internal loop despite this fusion being explicitly disabled in the parameter configuration. Leveraging our ArduPilot-backend simulation, we successfully replicated the problem and implemented a software patch, resolving the issue without consuming valuable future in-water testing time.
- Instability in the IMU data: During underwater operations, the IMU data exhibited severe data jumping, which heavily degraded real-time localization and navigation accuracy. To mitigate this environmental interference, we reorganized the physical layout of the electronic components within the enclosure to isolate the sensor from electromagnetic noise. Additionally, we enabled relative heading mode and recalibrated the IMU in a magnetic-free environment, successfully restoring signal stability.

To minimize risks and ensure safety, we have lifeguards at every test that need swimmers, schedule the right amount of sessions per week to prevent unnecessary exposure of our system. Moreover, identified dangers and risks are notified and communicated broadly across all sub-departments. Every stage of the in-water test is reported instantly through Telegram group chat for other members who are not at the pool to get updated and for them to give testing members any precautions or suggestions before moving on to next stage, providing another layer of safety.

## APPENDIX H
## SIMULATION TESTING

Simulation serves as our primary validation environment prior to pool deployment, helping us reduce pool time consumption and hardware risk. Standard Robotic Operating System (ROS2)-compatible simulators such as Gazebo provide adequate rigid-body physics for basic motion testing, but fall short in three areas critical to our development. First, they do not replicate flight controller behaviour at hardware fidelity, limiting the extent to which firmware-level issues can be diagnosed in simulation. Second, Gazebo's rendering pipeline does not sufficiently replicate real-world lighting conditions, shadows or material textures, introducing a sim-to-real gap that undermines the reliability of perception models trained or validated within it. Third, hydrodynamic forces acting on the vehicle hull are not accurately modelled, reducing confidence in motion response predictions when evaluating design changes.

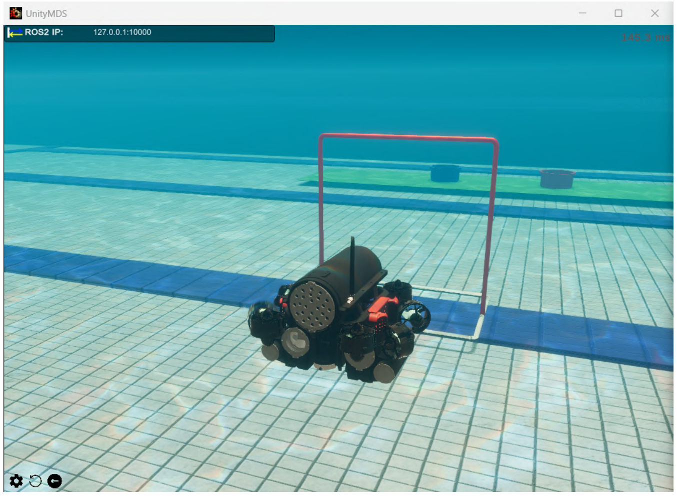


Fig. 14: UnityMDS Simulation

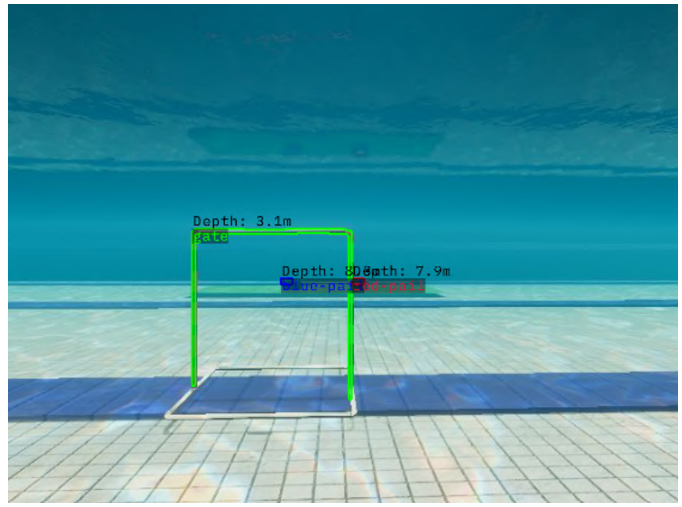


Fig. 15: Perception View in the Unity Simulation

To address these limitations, a custom Unity-based simulation environment was developed on an ArduPilot Software In The Loop (SITL) backend. ArduPilot SITL executes native autopilot code on a desktop computer. This replicates the EKF sensor fusion, MAVLink parameter handling and thruster PWM-to-force correlation of the physical vehicle, enabling the identification of firmware issues prior to pool testing. Unity reduces the simulation-to-reality gap by rendering realistic visual environments and generating simulated sensor data for the DVL and IMU. Furthermore, hydrodynamic forces are modelled through a mesh-based computation of drag coefficients and added mass. This ensures realistic vehicle motion when hull and shroud configurations are altered.

The simulation also identified system-level failures and guided critical development decisions, reducing the reliance on physical pool trials. The major findings are as follows:

- Unstable depth control from ArduPilot EKF and DVL fusion: Simulation testing traced this instability to a known ArduPilot bug where the EKF incorrectly fuses vertical DVL measurements, causing depth estimate divergence. The firmware was updated and validated in the simulation prior to physical pool testing. This allows simulation testing to serve as a reliable gate before hardware deployment, ensuring that software faults are identified and resolved in a controlled environment.
- Refinement of navigation and mission planning logic under real-world perception characteristics: The simulator provided a realistic environment that included sensor constraints such as low-frequency perception updates and transform (TF) jumping. Full mission sequences were executed with concurrent navigation, localization and perception. This enabled iterative tuning against positional uncertainty, resulting in consistent performance and fewer failures during pool deployment.
- Validation of eight-thruster configuration over six-thruster predecessor: During pool testing of the previous six-thruster configuration showed instability during position hold. This was caused by coupled roll and pitch during lateral translation, which deviated the vehicle from straight-line motion. Since the simulator maps individual PWM outputs to force vectors, the six- and eight-thruster configurations were directly compared without physical modifications. The eight-thruster setup eliminated the coupled roll and pitch, improving positional accuracy and stability. This validated the hardware upgrade prior to physical implementation.

# APPENDIX I
# ELECTRICAL ARCHITECTURE BLOCK DIAGRAM

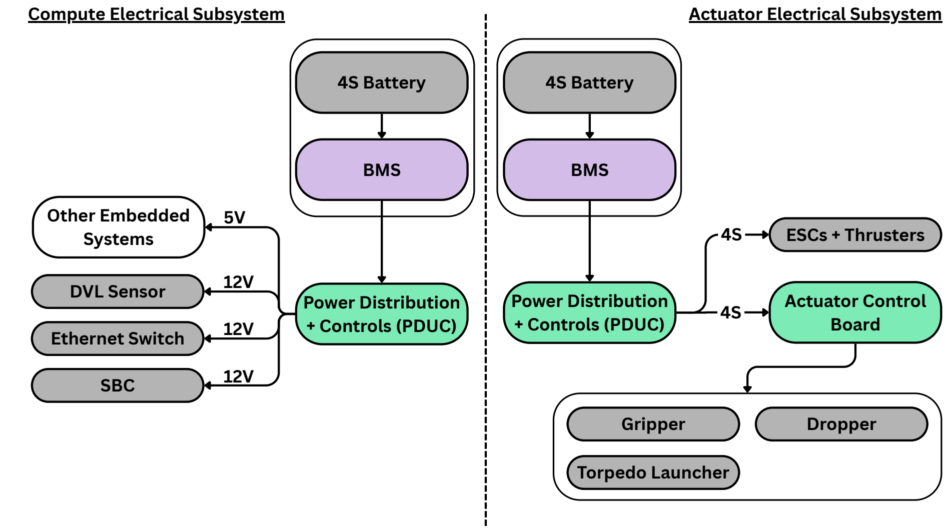


Fig. 16: Hydra and Kraken Power Architecture.

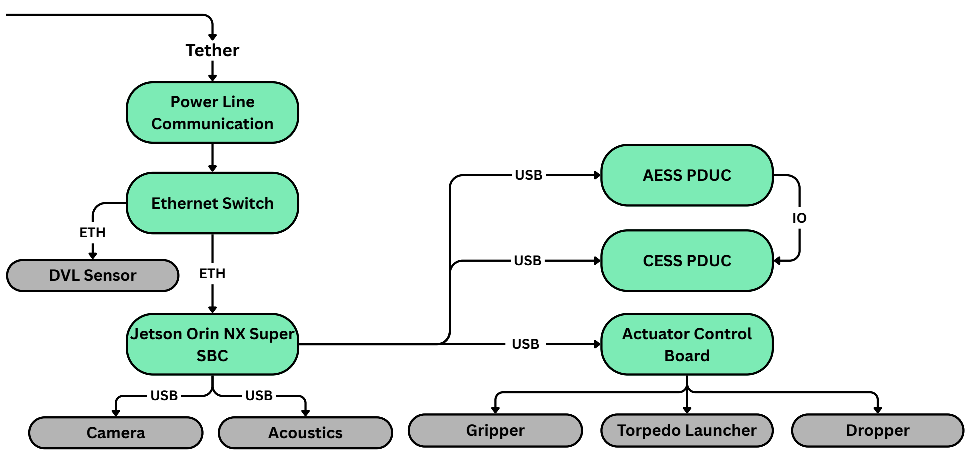


Fig. 17: Hydra and Kraken Communication Architecture.

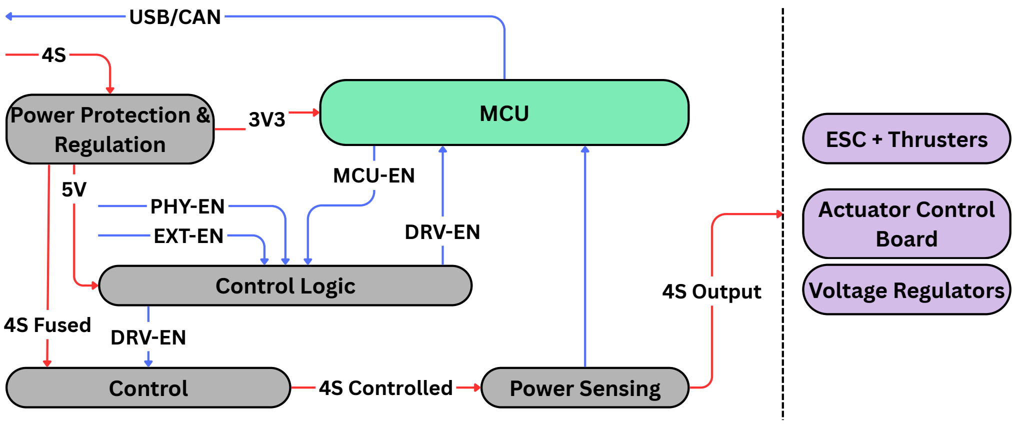


Fig. 18: PDUC Overall Architecture

## Appendix J
## Acoustics Systems Model

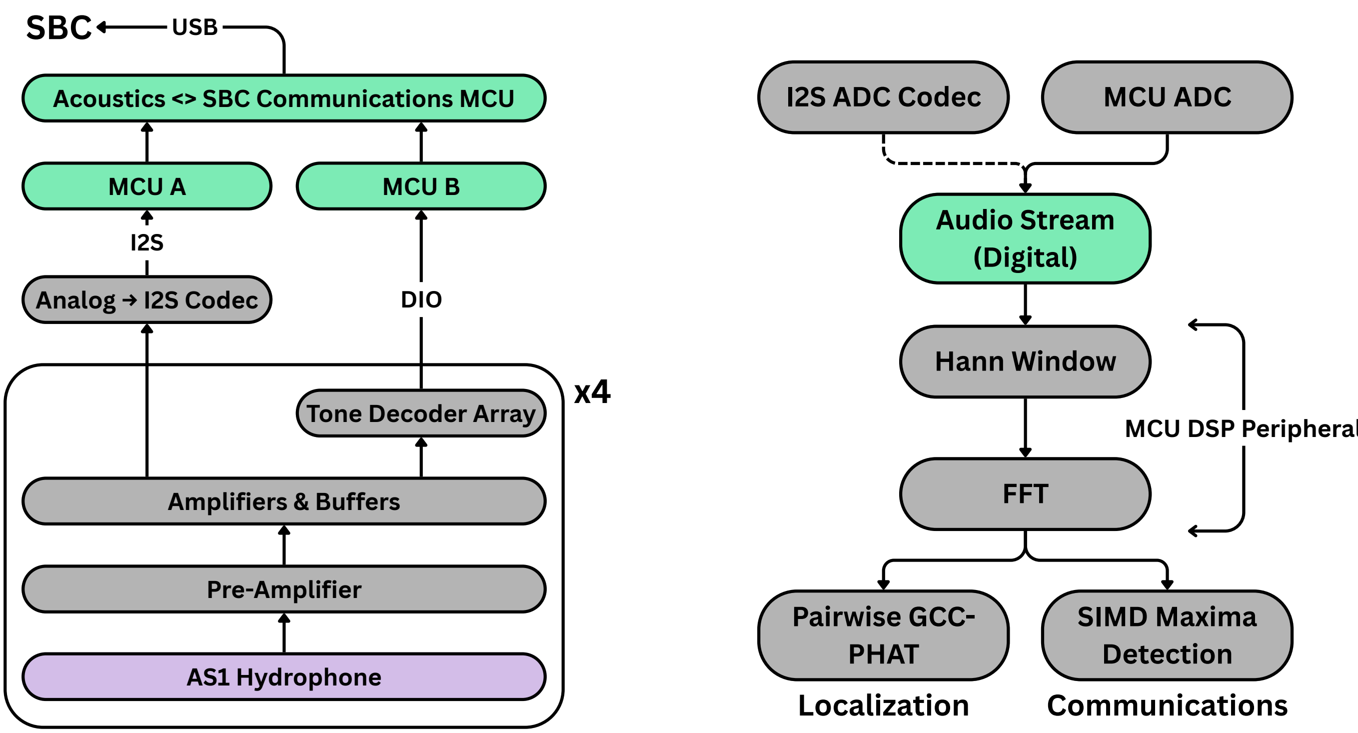


(a) Acoustics System Block Diagram (b) Acoustics Software Architecture

Fig. 19: Acoustics System Models

# Appendix K
# 3D Models of Electronic Boards

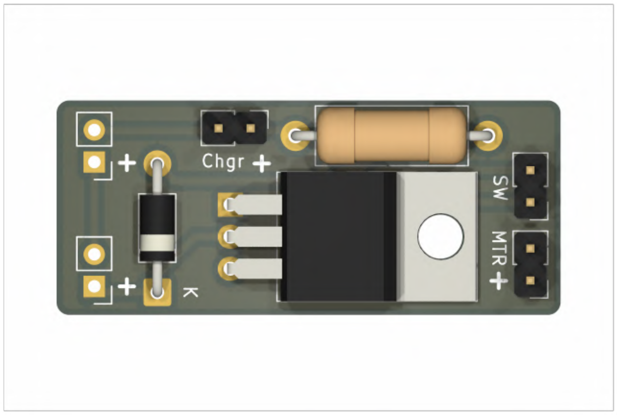


(a) Torpedo Board

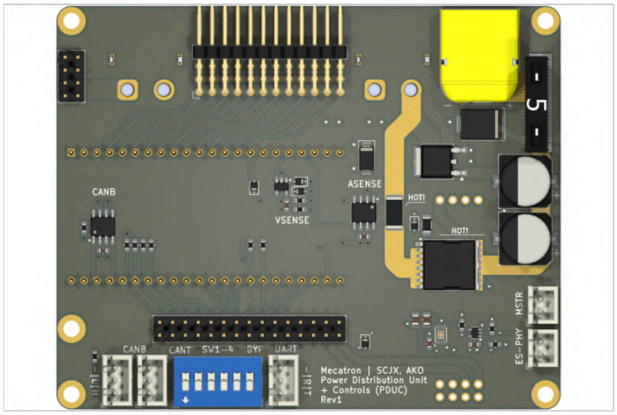


(b) Controlled Power Distribution Unit (PDUC)

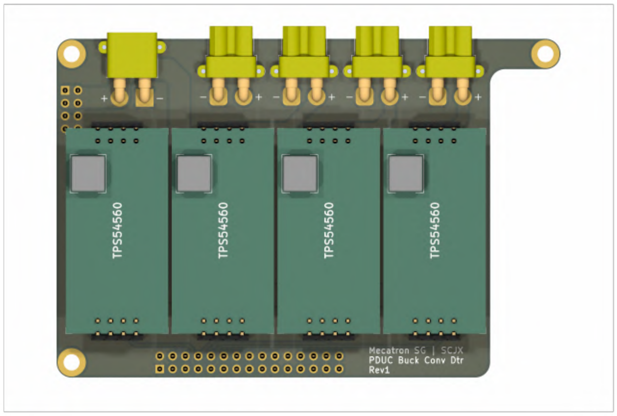


(c) PDUC Buck Converter Board

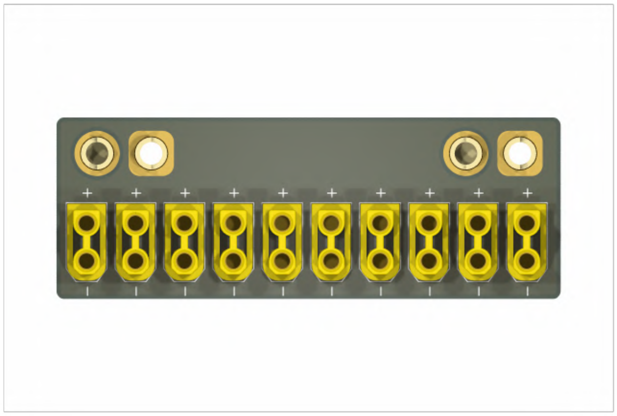

(d) PDUC Power Breakout Board

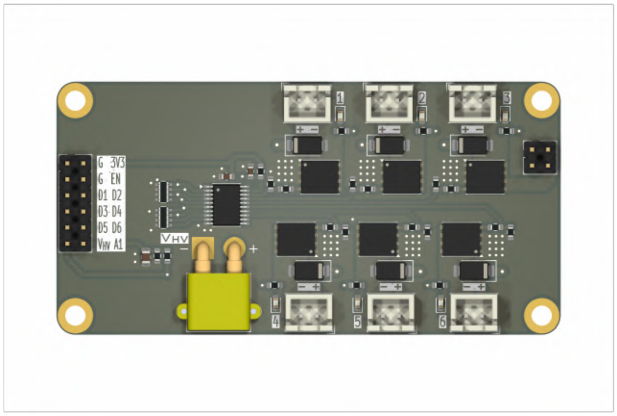


(e) Actuator Board

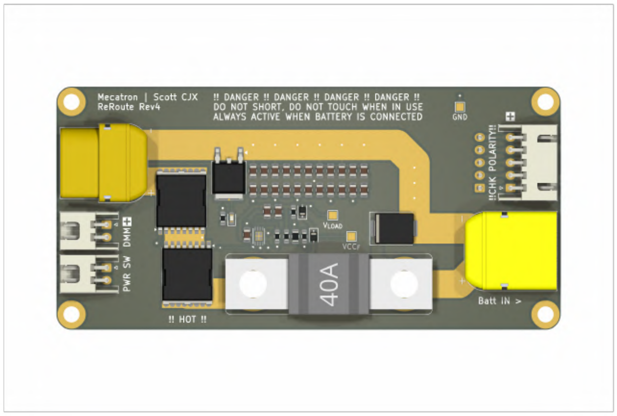


(f) ReRoute BMS Board

Fig. 20: 3D renders of custom PCBs

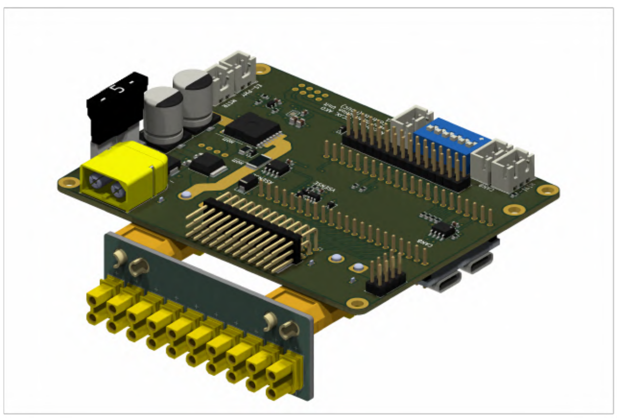

(a) AESS Controlled Power Distribution Unit (PDUC)

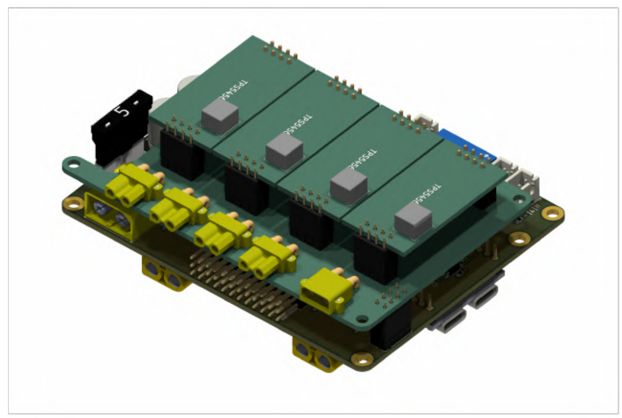

(b) CESS Controlled Power Distribution Unit (PDUC)

Fig. 21: 3D renders of assemblies