

Technical Design Review of Duke Robotics Club's Oogway: An AUV for RoboSub 2024

Will Denton, Michael Bryant, Lilly Chiavetta, Vedarsh Shah, Rico Zhu, Philip Xue, Vincent Chen, Maxwell Lin, Hung Le, Austin Camacho, Raul Galvez, Nathan Yang, Nathanael Ren, Mathew Chu, Amir Ergashev, Saagar Arya, Kaelyn Pieter, Ethan Horowitz, Maanav Allampallam, Patrick Zheng, Mia Kaarls, June Wood

Abstract—The Duke Robotics Club is proud to present our robot for the 2024 RoboSub Competition: Oogway. Now in its second year, Oogway has been dramatically upgraded in both its capabilities and reliability. Oogway was built on the principle of independent, well-integrated, and reliable subsystems. Individual components and subsystems were tested and designed separately. Oogway's most advanced capabilities are a result of the tight integration between these subsystems. Such examples include a re-envisioned controls system, an entirely new electrical stack, advanced sonar integration, additional cameras and system monitoring, a new marker dropper, and a watertight capsule mechanism. These additions enabled Oogway to prequalify for Robosub 2024.

I. COMPETITION GOALS

At previous RoboSub competitions, we prioritized versatility over reliability. While our previous designs allowed us to complete more tasks, our runs were inconsistent and unreliable. With only three scoring runs, we need to minimize the complexity of our systems to ensure they are reliable.

Our strategic vision consists of *core tasks*, which the robot is programmed to perform every run, and one *alternate task* which may be completed at the end of a successful run if more points are necessary.

RoboSub's design goals focus on four fundamental principles: movement, vision, manipulation, and acoustic tracking. To accomplish our strategic vision we chose to focus on 3 of 4 design goals: movement, vision, and manipulation. This subset of goals aligns with our team's strengths in controls, computer vision (CV), and mechanics. Our focus

on these methods enables us to perform the gate, buoy, and bins (our *core tasks*) reliably, while still maintaining the flexibility to attempt the octagon task (our *alternate task*) if needed.

A. Strategic Vision

Our strategy this year focuses on completing our *core tasks* while having fail-safes in place to ensure a successful run. This year, we upgraded Oogway with new peripherals to ensure that we can always complete our *core tasks* using multiple methods.

Reflecting on last year's competition, we noticed two large issues which prevented us from achieving our strategic vision: the unreliability of our movement and acoustic subsystems. Since movement is essential to all Robosub tasks we spent the first six months of this season completely remaking our control system and testing it in accordance with our testing plan. When discussing acoustics, we realized that we would not have enough time this year to make the system reliable. Instead of redesigning our complex acoustic system, we instead decided to create a simple and reliable marker dropper system for use in the bins task.

With our new controls system and marker dropper, we decided to change our *core tasks* from last year by removing the octagon and adding the marker dropper. In addition to these new systems, we also added multiple new cameras and a second sonar which allow us to fulfill our strategic vision by adding reliability to the gate, buoy, and bins tasks.

B. Enter the Pacific - Gate

Since the gate task is required to start a run, we use multiple subsystems to ensure that we can

reliably move through the gate. With our CV sub-system, Oogway can identify the gate from any angle and get an estimated position of symbols beneath the gate. If the robot cannot find the gate using CV, we can use our sonar to get the gate's approximate position. Finally, if neither of these methods work, our new controls system is accurate enough to dead reckon through the gate. Given these fallback methods, we will elect to do the coin flip for bonus points. We will always pass through the Hot/CW side of the gate to ensure we can drop our markers in the Hot side of the bins task. Our revamped controls also allows us to gain bonus style points by completing a barrel roll while moving through the gate.

C. Hydrothermal Vent - Buoy

The buoy task is a core component of our competition strategy. After we complete the barrel roll, we use our newly added downward-facing camera in combination with our CV system to find and track the path. Once in range, CV identifies the position of the buoy. If unsuccessful, we use sonar to search up to 10m away to get a rough estimate of where the buoy is located. Since our new controls is reliable, we will circumnavigate the buoy in the clockwise direction for bonus circumnavigation points.

Since we do not have a torpedo system, we will forego using the torpedoes to hit the buoy.

D. Ocean Temperatures - Bins

The final component of our core tasks is dropping both markers in the Hot side of the bin. After the buoy, we use the path to navigate towards the bins until the downward camera can identify the red side of the bin. We then line up, hold our position a foot away from the bins, and drop both markers. In case we go off-course from the path, we perform a search around the area until Oogway can find the bins.

When weighing the trade off between last year's final core task, the octagon, and the bins, we decided that since our movement and camera tracking is much more reliable than our acoustic tracking. Thus, completing the bins task will reliably get us more points in this year's competition

E. Collect Samples - Octagon

With limited reliability in our acoustic tracking systems, we made the octagon task our *alternate*

task. If we need more points at the end of a run we can use our current acoustics and sonar systems to attempt to find and surface inside of the octagon. After the bins task, we can enable our acoustics and perform large sonar sweeps until we have a "guess" at the octagon location. Since this is not reliable, we decided to leave this task until the end in case we surface outside of the octagon and end our run.

Since our actuators are still in development, we opted to skip the sample collecting task.

F. Mapping - Torpedoes

We are not planning on attempting the torpedoes task since Oogway does not have a torpedo system. However, we have this system in development and plan to use it next year.

G. Trade Offs Between Complexity and Reliability

Given last year's competition experience, we realized that many of our systems were too complex and did not have enough testing. We attempted to perform tasks at competition that were untested, resulting in failed runs. For this year's competition, we decided to remake many of Oogway's systems to decrease overall complexity while increasing reliability. Additionally, we developed an open-source GUI to interface with Oogway while in the water which allowed us to test individual components quickly. By narrowing our focus to 3 out of the 5 tasks during the design process for this year's competition, we allocated more testing time for each task to ensure that our integration can work reliably.

II. DESIGN STRATEGY

A. Mechanical Design

The main mechanical structure of Oogway has not changed from last year. However, we fixed major issues with leaks in the capsule and utilized our modular design to add peripherals such as a marker dropper and an adjustable buoyancy system.

1) *Reworked Capsule Design*: During last year's competition, we had multiple small leaks inside our capsule. The leaks were caused by the four brass inserts in our capsule slowly stripping out from the surrounding plastic due to the stress of holding a seal. To make the capsule watertight, we created a new locking mechanism that secures the capsule using a custom 0.125" Grade V titanium

plate, chosen for its corrosion resistance and high yield strength to weight ratio [5].



Fig. 1: New Capsule and Plate Design

The hexagonal design reduces weight while still maintaining rigidity under pressure [6]. The plate is held by precision-cut hexagonal aluminum standoffs that are connected to Oogway's frame.

In addition to a significant reduction in leaks, the new system is also simpler to secure and remove as all screws are top-facing. Installation time has been reduced from several minutes to under a minute.

2) *Modular Peripheral Components:* We utilized Oogway's modular design to add several new sub-systems. The first is our precise adjustable buoyancy system. We transitioned from a rudimentary zip-tie system to a rod-based one, which consists of 3D-printed pillars that the buoyancy blocks slide onto (Fig. 2a). A shaft collar is placed at the top to secure the blocks in place. This mechanism retains the advantages of our previous design—it is light, easy to fabricate, and can be placed anywhere along Oogway's rails. Moreover, the two points of contact hold the buoyancy blocks rigidly in place, preventing changes in buoyancy while the robot moves.



(a) Rod-based Buoyancy System



(b) Marker Dropper (body translucent)

Fig. 2: New Components Added to Oogway

Another key addition to the robot this year is the marker dropper (Fig. 2b). It uses a crescent-shaped

design that fits two spherical markers, one on each side. The markers naturally roll to the center outlet when the center paddle is rotated by the waterproof servo. There are also several openings at the top of the device: two large holes and three smaller ones. The two larger holes are for quick reloading of the mechanism and are plugged with two rubber stoppers to prevent the balls from falling out during a barrel roll. During testing, we found that air would become trapped in the crescent and slowly leak out, gradually altering the robot's buoyancy throughout a run. The three smaller holes allow air to escape when the robot is submerged.

B. Electrical Design

Oogway's new sensing capabilities and ambitious goals for the 2024 RoboSub competition warranted a complete redesign of the electrical stack.

1) *Electronics Stack Architecture:* Our new electrical architecture is based on two independent power grids: one for sensitive low-power and land-safe operations (ex. CPU, cameras), and the other for high draw, underwater-operation-only components (ex. thrusters, DVL). Both grids are fused respective to their own nominal loads, providing added electrical protection to all component failure modes. See Figure 40 for a full system diagram.

The new stack is accessible and adaptable to future components. Expansive busbars and CPU serial busses can accommodate future upgrades. This year alone, the stack has supported three new cameras, an additional sonar, fiber-optic integration, three new sensor types, and two new microcontrollers.

With Oogway's new control system, we found our ESCs susceptible to μs level inaccuracies in PWM timings. We extensively tested these hardware limitations and devised two means of correction, which we used to advise other RoboSub teams.

We determined that these hardware defects were manufacturing-based and we wrote firmware configurable by what version of the ESC was on the robot. This was utilized when we eventually replaced our ESCs and found that the defects differed. As such, we are adaptable to new and even faulty hardware.

Prior to this year, all of Oogway's sensory components focused on external sensing (e.g., robot acceleration, vision systems, orientation). This year, we have made substantial upgrades to Oogway's

internal monitoring, part of which we use to inform our ESC control efforts. Specifically, a new internal voltage sensor is used in dynamically controlling thruster allocations based on available power, using a manufactured-specified and in-house verified thruster power curve [13]. Additional internal monitoring includes a capsule-side temperature sensor and a humidity sensor for leak detection.

2) *Acoustics*: A critical component of the acoustics system includes hand-populated 0603 PC boards that offer bandpass filters and amplifiers for the analog hydrophone signals. The data from these boards is pipelined directly into a DAQ where we calculate and report the azimuth of the pinger to a ROS topic. We formalized findings from prior variable test circuits as a custom PCB, which in total is less than half the size of a credit card.




Fig. 3: Final Acoustic PCB

This hardware was iteratively designed and tested in order to detect pings with maximal accuracy. As part of the testing, we added a high-pass backpack to the boards. The result of our testing yielded filter bands and gain values where the pings were visible in the raw data to the naked eye, even from opposite corners of the lane pool (Fig. 30).

Our strategy uses rough guesses collected over time and accurate robot state information to build an aggregate estimate of the absolute position of the pinger. Our hydrophones are located within a few inches of each other to have similar acoustic exposure and use absolute timings to determine the guessed pinger location. Pings are found by first measuring the level of background noise at the desired frequency, and then setting a statistical volume threshold of 4-5 standard deviations from the mean peak levels. Pings can be accurately detected from long distances, with very few false positives. This is done for each hydrophone, using

time difference of arrival (TDOA) to approximate the angle at which the source came from [1]. We turned this problem into a gradient descent problem as well, classifying the ping timing offsets into 45-degree octants. After verifying feasibility in Python, we wrote an optimized library in Rust.

C. Software Design

This year, we have rewritten or refactored every part of our software. This includes controls, graphical user interface (GUI), computer vision (CV), sonar, and task planning. All of our software is open-source and publicly available on GitHub [11], supporting the AUV community. See Figure 46 for Oogway’s overall software control flow.

1) *Controls*: We rewrote our controls system from the ground up this year. We wrote it entirely in C++ for maximum efficiency. (See Figure 47 for a system flowchart.)

The system centers on a set of six Proportional-Integral-Derivative (PID) control loops, one for each axis of movement. We wrote the control loops from scratch, customized for controlling Oogway. They include a second-order Butterworth filter [15] to smooth the derivatives for smoother movement. The PID gains and other important parameters can be changed on-the-fly, without re-launching the software, making tuning easy and fast.

The control efforts output by PID are summed with a dynamic offset that neutralizes the robot’s buoyancy, enabling it to remain submerged with ease. This is fed into the thrust allocator, which computes the speed each thruster needs to spin at to achieve the desired control efforts without exceeding the thruster power limit. This constrained optimization problem is solved using quadratic programming.

The system also supports combining multiple control types – position, velocity, and power – for more versatile and precise control over the robot.

2) *GUI*: This year, we established a new GUI subteam tasked with redesigning our visualization and controls platform using a new, modern interface and architecture. For the past several years, we’ve used RQt, which is outdated and incapable of holding a real-time connection with Oogway.

In order to facilitate efficient testing, we decided to switch to Foxglove Studio [9], a new robotics

visualization platform with plug-and-play ROS integration. Foxglove has a set of built-in extensions (e.g., image and plot panels) which we supplement with custom extensions for our specific use cases. Our new GUI architecture addresses the core limitations of RQt by lowering the bandwidth required by the system, and using TypeScript to develop custom extensions. (See Appendix C).

3) *Computer Vision (CV)*: This year, the primary focus of our CV subteam was to develop a synthetic data model for machine learning training. In prior years, the process of manually collecting, annotating, and reviewing real images took hundreds of hours of manual labor and was a significant burden. By using Unity Perception [10], we developed a CV simulation which can generate a large labeled dataset of 50,000 images in only a few minutes.

The primary design consideration during the development of the simulation was realism: synthetic images must be captured in a competition-like environment with appropriate randomizations.




Fig. 4: Simulated Gate in 3D Modeled Pool

These images were then further modified with color, noise, and blur augmentations to create the final model. Like last year, we used the YOLOv7-tiny model [2] which can process camera data at 20Hz with high accuracy.

To detect simple, monochromatic shapes such as the buoy (red circle) and path markers (orange obround), a complex YOLO model would be overkill and incur a lot of computational overhead. Instead, we detect them using HSV filtering and contour matching. For each camera frame, we mask out the Region of Interest (ROI) of all pixels that are close to the target color in HSV space. We then find the contour of the ROI to classify the object. This proved to be very reliable, as we successfully completed the pre-qualification task using HSV filtering to track the lane markers.

4) *Sonar*: This year, we rewrote the sonar code to improve its reliability and to have multiple methods of completing tasks. The sonar now avoids false positives caused by acoustic reflections and can detect when the target object is missing.

We overhauled the system with a few important changes. First, we physically mounted the sonar directly below the camera to better align with CV data. Second, we apply a filter to get rid of all low-intensity values and use the DBSCAN [16] algorithm to find all clusters in the image. We select the biggest cluster, which represents the buoy, and find its center of mass. In our testing, this method proved to be reliable and ensures the sonar only recognizes one object. It also provides the the normal angle to the buoy which helps to align Oogway in the correct orientation.




Fig. 5: Buoy Detected with Normal Line

We also implemented a new system to detect and identify objects within long range sonar scans. This serves as a fallback method for multiple tasks by finding a rough position of all objects in the pool. We find clusters in the image with DBSCAN and filter based on area and circularity. Using these two measurements, we classify objects into *gate*, *buoy*, and *unknown*. We avoided false positives from buoy reflections by filtering out neighboring clusters with similar size and circularity [3]. See Appendix E for the full sonar pipeline.

5) *Task Planning*: The task planning system receives information from all the robot's sensors, decides the best action to complete the task, and commands the controls system to carry out that action. Each task is written using a Python coroutine. Thus, each child task frequently yields control to its parent task, allowing the parent task to frequently reevaluate the robot's environment and switch between different task pathways. The system is also very modular; small actions are combined to accomplish complex tasks. By adding a simple Python decorator, each task is automatically included in the real-time task monitoring system.

III. TESTING STRATEGY

A. Test Plan

The most important lesson from last year was that testing is the key to success. Ensuring every component and subsystem is reliable will ensure a successful competition run. Oogway's testing suite comes from a combination of simulated and real-world environments. Our test plan is as follows:

- 1) Test all compatible components in simulation. This ensures that each component works properly and connects to the robot without negatively impacting other systems.
- 2) Test each component physically on the robot to ensure all components fit properly. We run component specific tests in the pool and in the sink when possible.
- 3) Once all components work separately, test the integration of components in the pool to mimic a real competition environment. Pool testing allows us to test the robot on completion tasks and ensure we could complete each task reliably.

Following this strategy ensured that every component was individually reliable and the robot was able to complete all *core tasks*.

B. Simulated Testing

1) *Acoustics*: While developing the algorithm underlying the acoustics stack, we created a simulated acoustics program. This involved simulating a pinger source and calculating the induced pressure waves. This, combined with encoded hydrophone geometries and artificial noise, provided realistic data to test the algorithm with.

2) *Software*: We utilized our custom sim environment in CoppeliaSim—which includes models of Oogway, the gate, buoys, and the octagon—to test controls and task algorithms. Before empirical testing, we ran each of our *core tasks* in the sim to ensure our code completed the task correctly.

3) *CV*: As described above, we use Unity to generate images and test our CV model.

C. Empirical Testing

1) *Mechanical*: All water-facing parts underwent extensive functionality and water-resistance testing. This often involved long-duration submersion tests in a sink. Once passed, the testing would escalate to

a diving pool. This comprehensive approach guarantees all vital components will remain watertight during operation. Non-water-facing components go through the standard prototyping process, which involves testing and iteration on dry land.

2) *Electrical*: Each electrical component was initially ran for extended periods of time to test long-term performance. All subcon and epoxy connections, as well as waterproof components, underwent 24 hour submersion testing. Once all of the components were tested individually, we performed full system integration tests to ensure accurate data.

3) *Software*: As mentioned above, we tested our new controls and task-planning systems through our Foxglove GUI which allows us to set PID constants and observe setpoints. We visualized the task flow of the robot to verify the robot's state and actions.

4) *CV*: To test the object detection algorithm, we set up the gate and manually moved the robot around to test all vision angles. We also tested CV-sonar integration, by matching the two data streams. We then tested movement autonomously, verifying the reliability of our gate task.

D. Testing Results and Integration Testing

Towards the end of our season, we tested the robot with each of our *core tasks* to verify that we could compete in a competition environment.

In total, we spent over 500 hours in simulated and empirical testing environments, resulting in Oogway prequalifying for the 2024 RoboSub competition.

IV. ACKNOWLEDGEMENT

Duke Robotics Club is primarily sponsored by Duke University's Pratt School of Engineering. We are grateful to Ali Stocks and the Foundry staff for housing and supporting us. We are also indebted to Director of Undergraduate Student Affairs, Tarina Argese, for helping us plan events, our advisors, Professor Michael Zavlanos and Dr. Boyuan Chen, for technical assistance, and the Engineering Alumni Council, for overall club advice. Furthermore, we owe our success to our longtime sponsors: The Lord Foundation, Duke Student Government, General Motors, and SolidWorks. Finally, we thank RoboNation, whose commitment to RoboSub empowers student engineers like ourselves to explore our passion for robotics.

REFERENCES

- [1] M. Johnson, "Underwater Acoustic Pinger Locator," EEWeb, Feb. 03, 2011. [Online]. Available: <https://www.eeweb.com/underwater-acoustic-pinger-locator/>. Accessed: Jun. 23, 2024.
- [2] "YOLO - Object Detection," OpenCV tutorial. [Online]. Available: <https://opencv-tutorial.readthedocs.io/en/latest/yolo/yolo.html>. Accessed: Jun. 25, 2024.
- [3] "Understanding and Using Scanning Sonars," Blue Robotics. [Online]. Available: <https://bluerobotics.com/learn/understanding-and-using-scanning-sonars/>. Accessed: Jun. 26, 2024.
- [4] Y. Lee, S. Oh, and S. Park, "Enhanced Control with a General Cascade Control Structure," *Ind. Eng. Chem. Res.*, vol. 41, no. 11, pp. 2679–2688, May 2002, doi: 10.1021/ie010157f.
- [5] Engineering ToolBox, "Young's Modulus, Tensile Strength and Yield Strength Values for Some Materials," Apr. 4, 2024. [Online]. Available: https://www.engineeringtoolbox.com/young-modulus-d_417.html. Accessed: Jun. 20, 2024.
- [6] Christmc, "Hexagonal vs. Square Holes in Mild Steel Perforated Metal Mesh: A Comprehensive Comparison," The Mesh Company, Jan. 31, 2024. [Online]. Available: <https://theshcompany.com/hexagonal-vs-square-holes-in-mild-steel-perforated-metal-mesh-a-comprehensive-comparison/>. Accessed: Jun. 25, 2024.
- [7] "How to Anodize Aluminum," The Federal Group USA, Feb. 23, 2024. [Online]. Available: <https://www.tfgusa.com/how-to-anodize-aluminum/>. Accessed: Jun. 25, 2024.
- [8] J. V. Ecker, A. Haider, I. Burzic, A. Huber, G. Eder, and S. Hild, "Mechanical properties and water absorption behaviour of PLA and PLA/wood composites prepared by 3D printing and injection moulding," *Rapid Prototyping Journal*, vol. 25, no. 4, pp. 672–678, 2019, doi: <https://doi.org/10.1108/RPJ-06-2018-0149>.
- [9] "Foxglove - Visualizing and Debugging Your Robotics Data," Foxglove. [Online]. Available: <https://foxglove.dev/>. Accessed: Jun. 28, 2024.
- [10] "GitHub-Unity-Technologies/com.unity.perception: Perception Toolkit for Sim2Real Training and Validation in Unity," [Online]. Available: <https://github.com/Unity-Technologies/com.unity.perception>. Accessed: Jun. 24, 2024.
- [11] "Duke Robotics Club · GitHub," [Online]. Available: <https://github.com/DukeRobotics>. Accessed: Jun. 25, 2024.
- [12] "Greenroom-Robotics/ros-typescript-generator," Greenroom Robotics, Jun. 01, 2024. [Online]. Available: <https://github.com/Greenroom-Robotics/ros-typescript-generator>. Accessed: Jun. 25, 2024.
- [13] "T200 Thruster," Blue Robotics. [Online]. Available: <https://bluerobotics.com/store/thrusters/t100-t200-thrusters/t200-thruster-r2-rp/>. Accessed: Jun. 21, 2024.
- [14] Mechanical Integration Considerations; Alignment; Figure 20. "Transducer Alignment Reference Points," Teledyne PATHFINDER Manual, Apr. 2018.
- [15] "Example: Second-Order Butterworth Lowpass." Accessed: Jul. 01, 2024. [Online]. Available: https://ccrma.stanford.edu/ios/filters/Example_Second_Order_Butterworth_Lowpass.html
- [16] M. Ester, H.-P. Kriegel, and X. Xu, "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise".

APPENDIX A: COMPONENT LIST

Component	Vendor	Model/Type	Specs	Custom/Purchased	Cost	Year of Purchase
Buoyancy Control	Blue Robotics	Subsea Buoyancy Foam	R-3312	Purchased	200.00	2023
Frame		Aluminum	821	Custom	2,000.00	2023
Waterproof Housing		Polycarbonate	-	Custom	345.00	2023
Waterproof Casing	Blue Robotics	Watertight Enclosure	4" Cast Acrylic Plastic	Purchased	124.00	2023
Waterproof Connectors	MacArtney	Subconn/Seaconn Connectors	-	Purchased	1600.00	2019
Waterproof Connectors	Blue Robotics	WetLink Penetrators	4.5MM-HC	Purchased	400.00	2023
Propulsion	Blue Robotics	T200 Thruster	-	Purchased	1,790.00	2023
Power System	Turnigy	16000mAh	4S 12C	Purchased	210.00	2023
Motor Controls	Blue Robotics	Basic ESC	-	Purchased	300.00	2023
High Level Control	Arduino	Arduino Nano Every (x2), Arduino Nano (x1)	-	Purchased	40.50	2023
CPU	ASUS	MiniPC PN80	-	Purchased	1,300.00	2022
Internal Fiber Comm Network	Mini Industrial	4 Rj45 + 1 SFP Switch	-	Purchased	56.00	2024
External Comm Interface	NetGear	8-Port 1G/10G Gigabit Ethernet Unmanaged SFP Switch	-	Purchased	100.00	2024
Inertial Measurement Unit (IMU)	VectorNav	VN-100 Rugged IMU	-	Purchased	1,335.00	2020
Doppler Velocity Logger (DVL)	Teledyne	Pathfinder 600	-	Purchased	10,000.00	2023
Stereo Cameras (x2)	Luxonis	OpenCV AI Kit: OAK-D W PoE	-	Purchased	1,000.00	2023
Hydrophones	Aquarian	H1a Hydrophones	-	Purchased	600.00	2021
Imaging Sonar	Blue Robotics	Ping360 Scanning Imaging Sonar	-	Purchased	2,650.00	2023
Waterproof Servo	Blue Trail Engineering	SER-2020	-	Purchased	495.00	2024
Mono Camera (x2)	DWE	exploreHD 2.0	-	Purchased	630.00	2023
Single-Beam Echosounder	Blue Robotics	Ping2 Sonar	-	Purchased	435.00	2024
Vision	OpenCV, DepthAI					
Localization Mapping		Extended Kalman Filter, SLAM		Custom		
Autonomy		Custom task planner		Custom		
Open-Source Software Programming Language(s)		Docker, ROS, Ubuntu C++, Python, Rust				

APPENDIX B: TEST PLAN RESULTS

A. Mechanical

Mechanical testing is arguably the most critical aspect of verifying Oogway's performance. A failure here would disable all other subsystems. As described previously, water-facing components were put through multiple rounds of testing.

1) *Waterproofing the Stack:* The titanium plate discussed in the body, for instance, underwent several prototypes, each subject to rigorous testing. The goal of the first prototype was to prove the feasibility of the concept and used leftover 0.125" 6061 aluminum stock (Fig. 8). Testing on dry land showed that the thin aluminum was far too flexible and bent severely while under pressure. Part of this was also due to the open window design, which future designs improved on.



Fig. 6: Full Assembly Render of Oogway



Fig. 7: New Capsule Plate Design

The second prototype used 0.4" 6061 aluminum stock, also leftover from a separate project (Fig. 9). To compensate for the weight added by the increased thickness, the plate only stretched across the center. This prototype performed much better than the previous: there was no bend in the aluminum



Fig. 8: First Plate Prototype

when fully tightened, but it did weigh significantly more. Moreover, because it only covered the center, pressure was not applied evenly across the entire oring.

Even so, we were able to trial this design in the pool with some added precautions. In addition to screwing down the four screws that keep the plate in place, we also screwed in the four corner screws of the capsule as extra support. We also put the robot in slowly and monitored closely for bubbles.




Fig. 9: Second Plate Prototype

The final prototype combined the benefits from both of our previous prototypes: the design stretches across the entire capsule to provide even pressure, and the chosen material, Grade V titanium, is both light and rigid (Fig. 10) [5].

This new titanium plate design initially showed a substantial increase in the watertightness of the capsule's seal. While we could not trial this design in the sink, we took steps to ensure that the electrical stack would not be damaged during the first few tests. Like with the previous prototype, this involved the use of the four corner screws in addition to the four screws used to secure the top plate. Later, a humidity sensor was added to the bottom of the stack such that leaks could be recognized instantly.



Fig. 10: Final Plate Design

While this solution has performed much more consistently than the previous one, some minor leaks still occasionally occur, likely due to wear-and-tear on the capsule, wear-and-tear on the o-ring, or loose ports in the stack plate. However, the humidity sensor allows these leaks to be identified before they become catastrophic.

There are still areas for improvement. As discussed in the body, moving to more standard parts or utilizing in-house fabrication would lessen our reliance on distributors and allow us to replace worn parts easily.



Fig. 11: Aluminium base plate built for adaptability

The aluminum plate at the bottom of the stack likewise went through rigorous waterproof testing. It was built to incorporate many sensors and electrical attachments and contains extra holes to be adaptable to additional sensors that need to be attached (Fig. 11). We tested the watertightness of the plate using blanks for all the plugs and screwing in a capsule on the four corners. Dye packets were placed inside the capsule and it was then submerged in the sink for 24 hours to stress test the system. No leaks were detected. Moreover, we anodized the plate following a professional guide to prevent corrosion [7].

2) *Waterproofing 3D prints:* Due to the large number of 3D printed parts on the robot, one problem we noticed during pool tests was that our

PLA 3D prints would fill with water [8]. Although usually not a problem with PLA plastic, the low infill density (15%) and gaps between the layer lines meant water could easily fill the large cavities within the print. This resulted in the robot losing its finely tuned buoyancy gradually, making robot movement impossible to correctly tune.

To fix this, we tried several different solutions that either provided a seal around the outer layer of the 3D print or fused the layer lines. For each of the solutions, we took data at regular intervals over an hour. The results are given in Table I.

TABLE I: 3D Print Water Absorption Data

Time	Control White	Control Black	Wood Glue	Clear Nail	Blue Nail	Clear Sealer Gloss
0	66	56	72	62	28	66
5	68	58	74	62	28	68
10	68	58	76	62	28	70
15	68	58	74	62	28	68
20	68	60	76	62	30	70
25	68	60	76	62	30	68
30	68	58	74	62	28	68
45	68	58	74	62	28	68
60	68	58	76	62	30	68
% gain	3.03	3.57	5.56	0	7.14	3.03

The clear nail polish performed best: the acetone in the nail polish fuses the layer lines and can be thoroughly applied. It was unclear why the blue nail polish performed much worse, but may have to do with the dye used or variation in the quality of the print that it coated.

During pool tests, the nail-polished prints have been successful in minimizing water intake. In previous pool tests, bubbles were noticeable coming out of 3D prints as the robot was dropped in water, indicating the print filling with water, but this no longer occurs. Moreover, buoyancy has been found to remain stable across the entire period.

3) *Buoyancy:* One of the main priorities for the mechanical team is to ensure that the robot's buoyancy is level and can be easily adjusted as new parts are added. Thus, we tested different buoyancy mount designs before settling on the final design. One such design was the capsule-based system pictured below.



Fig. 12: Buoyancy Piston Prototype

We decided to retrofit an old, damaged acrylic capsule into a piston-like mechanism that allows its volume to be adjusted when needed, providing higher granularity than the discrete buoyancy blocks we previously relied on (Fig. 12). Moreover, a future design could leverage actuators to automatically compress and depress the piston based on a feedback loop.

However, this system posed a few disadvantages over the old design. It was much larger and heavier than the simple 3D-printed mounts, meaning there were only a few spots where it could be mounted on the robot. While having high granularity was important, testing showed that positional flexibility was much more important in ensuring a neutrally buoyant robot. Therefore, we moved to the rod-based system discussed in the body.

4) *Mounting Brackets*: This year also saw the update of all of our old mounting solutions. During the competition last year, several 3D-printed mounts failed under stress. However, we quickly identified the issue for the camera capsule and battery capsule mounts: the two-part mechanism was weak at the interface, causing the plastic to strip out under pressure. Therefore, we adopted a unibody mount for the camera capsule, eliminating the weak point (Fig. 13, Fig. 14). This system was stress-tested during pool tests and we noted no failures this year.

However, we could not do the same for the battery capsule, as the ring needs to rotate separately from the body to allow the battery to be quickly and easily replaced. Therefore, we opted to instead increase the width, the infill density, and the wall thickness on the 3D print (Fig. 15, Fig. 16). This kept the design



Fig. 13: New Unibody Camera Capsule Mount




Fig. 14: Camera Capsule Subassembly

light and easy-to-fabricate, while also increasing the overall robustness of the mount. While much less prone to break than the old design, we still have occasional fractures of the mount at the same interface - the 3D-printed mount comes under stress when loading and unloading the heavy battery. An ideal design would eliminate the two-piece bracket (similar to the new camera mount) and is a priority next year.

5) *Flexible Mounting*: Per our goal of maximizing the modularity of the robot, one design element we trialed this year was the use of standardized mounting blocks (Fig. 17). To ensure minimal



Fig. 15: New Battery Capsule Mount




Fig. 16: Battery Capsule Subassembly

interference with already working parts, we only tested this system on new components added this year. These 3D-printed blocks clamp around the aluminum frame and are fitted with brass inserts to allow components to be easily face-mounted.




Fig. 17: Modular Mounting Block

We trialed this system on new components that were in flux, such as the hydrophone mounts (Fig. 18, Fig. 19). As the acoustics subteam was testing different mounting locations and styles, the modular

mounting block system allowed the hydrophone orientation to be modified easily.



Fig. 18: Old Hydrohone Mount. Required four different mounting locations.




Fig. 19: New Hydrohone Mount. All hydrophones grouped together in one spot.

We utilized the same system for the 1D sonar, another new addition to the robot (Fig. 20). By having an easily adjustable location, we were able

to find and mount it in the most effective spot on the robot.

The trial was a clear success for the new standardized mounting solution: it is compact, and parts are easy to attach and detach. Old mounts will be adapted to this new design next year.



Fig. 20: 1D Sonar Mount. This traded spots with the hydrophone mount a few times when finding the optimum location.

6) *Torpedo Mechanism:* One of the components that underwent the most iterations this year was the torpedo mechanism. The testing process was three-fold: first, a prototype would be tested on dry land to ensure the torpedoes could be safely discharged without causing harm. We would then use the sink to test underwater performance. While the length of the sink was short (only around two feet), we could identify any obvious problems with the prototype quickly without needing to go to the pool. Finally, the prototype would be tested in the pool to gauge its accuracy and distance in a realistic environment.

This testing plan has allowed us to fine-tune the torpedo tube design: we were able to pinpoint the minimum safe compression length for the spring using the dry land test, and the sink test showed that we needed holes to allow the torpedo to push

water out of the way as it exits the chamber (Fig. 21).



(a) First Torpedo Tube Prototype

(b) Current Torpedo Tube Prototype

Fig. 21: Torpedo Tube Design Evolution

The dry land test also allowed us to optimize the design of the firing mechanism. Original prototypes were initially too weak due to their PLA construction and the shallow angle meant that misfires were common (Fig. 22a). We corrected this by giving the trigger a larger surface area and steeper angle, ensuring it could keep the spring compressed and release it correctly (Fig. 22b). While this mechanism works well, there is still work to be done integrating it with a servo motor. We are currently designing a system with rods offset 180° from each other on the servo horn. These will individually push down on each trigger, allowing the torpedoes to be fired independently.

The sink test also allowed us to improve on the design of the actual torpedo. We initially experimented with fins to maintain a straight line, but this failed to show results in the sink (Fig. 23a). We then introduced a spiral pattern to help maintain their path, which once again failed to show results (Fig. 23b). Eventually, we designed the torpedoes as cylinders with spherical tops tapering towards the bottom (Fig. 23c). While the current prototype occasionally travels in a straight line, it is not consistent. Moreover, the 3D-printed torpedoes tend to float, giving them an undesirable upward trajectory. Future work consists of running fluid simulations to optimize torpedo shape and determine why the torpedo sometimes varies from its path.

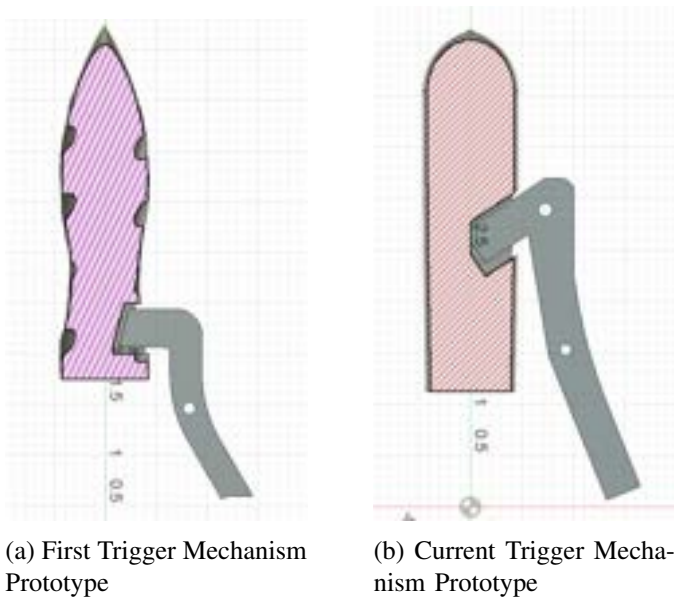


Fig. 22: Trigger Mechanism Design Evolution




Fig. 23: Torpedo Design Evolution

7) *Robot Assembly:* With all the new components this year, much more focus was put on maintaining an up-to-date full robot assembly. This has become an integral part of our prototyping process and has allowed us to better understand how components fit together and identify conflicts before fabrication. Moreover, as we further develop the task planning simulation next year, the accurate assembly will provide a starting point to roughly

tune movement constants before putting the robot in the pool.

This assembly has also allowed our team to construct higher-quality media, such as all the mechanical diagrams found in this paper and the render in Figure 6. For the first time this year, we also leveraged Blender to convert the assembly into the compact glb file format, allowing us to embed a 3D view of the robot directly on our website (Fig. 24).



Fig. 24: Full Robot Assembly in Blender.

B. Electrical

Rigorous testing is the foundation of ensuring electrical reliability. Our testing surrounds the two principal responsibilities of the electrical subsystem: power distribution and data acquisition, in addition to research and development surrounding our acoustics system.

1) *Power:* The electrical team must assume that the entire electrical stack will undergo extensive use under very demanding circumstances. This includes significant power, high heat, and surviving mechanical imperfections such as vibrations and controlled water ingress.

Due to the limited battery life of the batteries, the team constructed a dedicated testing rig to stress test the entire system for durations far exceeding those on a single battery charge. This rig involved a modified power supply serving as a battery emulator, which we configured to have the same power characteristics as our 4S 12C lithium battery. During our 6-hour testing, we logged many system vitals, such as capsule temperature and sensor drift.

The results of our testing showed a dramatic accumulation of heat, with internal temperatures nearing 70°C. After inspection, we deduced that the main cause was the POE network switch, which quite inefficiently drew excess power that it did not deliver

over unconnected ports. We took the opportunity to replace it with a lower-voltage (12V versus the original 53V) switch and saw notable improvements as temperatures stabilized at around 55°C. We also took the opportunity to upgrade to a 1G SFP switch, which supported the team goal of refactoring to the improved water-resistance of a fiber optic tether with added network speeds required by an advanced GUI framework with additional camera feeds.

The migration to this new switch did prompt an engineering trade-off decision. Specifically, since we have two POE cameras on Oogway, we then had to create our own POE injectors, per the IEEE 802.3a Class A standard. In making this decision, we consulted with a former Duke Robotics member who currently works professionally as an electrical engineer on POE-based networking, and concluded that this approach would simplify the stack and improve thermal performance.

Further, this long-duration power test was repeated with the robot in various orientations, notably completely inverted. The main usage of this test was to determine if mechanical attachments and adhesives would fail with heat accumulation. We also tested the robot while undergoing intense vibrations through use of a dolly. This allowed the team to pinpoint mechanical attachments that would eventually fail. This included nuts that were replaced with lock nuts, components that would slide out of 3D printed mounts, and connectors that would expose their conductors below their shielding.

The battery testing rig that we developed now serves an additional purpose of running the robot during typical lab testing and development without diminishing the capacity of our existing batteries.

2) *Data:* As described earlier, the team noticed hardware defects within our ESCs. Table II outlines our first-difference based calculations for various PWM signals, showing both our input efforts, their corresponding PWM durations, and the measured offsets calculated with an oscilloscope's trigger. The result of these calculations for this specific ESC yielded a very consistent offset of 31 μ s, which we then utilized in our firmware configuration. This process was applied for each ESC.

An additional result of the extended power testing described earlier also unveiled what ended up being a rare, but catastrophic memory leak in our




Fig. 25: Plotted differentials between sent and measured PWM shows linear offset, which we then applied the inverse of in our corrective software.

microcontroller firmware bridging ROS and the ESC interface. As indicated by our extensive logging within ROS, we were able to detect these outages before the microcontroller was able to reboot. Upon correction, we found autonomous performance of the robot to improve to a noticeable degree. This included general smoother movement and faster response times. We concluded that, while the microcontroller would only occasionally crash, the memory leak buy would gradually degrade performance on the memory-constrained device, which is why correcting this bug improved performance overall.

Another data reliability issue we encountered during our testing pertains to our POE cameras. We noticed that the streams would drop after many hours of use, and they would only return functionality once rebooted. Despite the outages, the issue only occurred after at least an hour, which is not an expected runtime during official runs. That, combined with double-coverage on most of cameras and sensors, and the determination that the issue is likely a memory leak within the camera's internal firmware, led the team to consider the issue to not impact Oogway's performance.

It is inevitable that trace amounts of water will enter the capsule during an extended pool test. This is due to the various reasons described earlier in the mechanical sections. While not catastrophic, we have been able to validate our humidity sensor data by removing the robot from the pool after a spike in the data and subsequently verifying that trace amounts of water has entered. Specifically, we use humidity sensing because some of the water evaporates upon entering the warm capsule. Like-

int8	PWM Sent	PWM Measured	Sent - Measured	Measured Diff.	Test Diff.	PWM Test	Test Diff.
-128	1131	1100	31			1100	31
-120	1156	1120	36	20	25	1125	31
-110	1187.25	1160	27.25	40	31.25	1156.25	31
-100	1218.5	1180	38.5	20	31.25	1187.5	31
-90	1249.75	1220	29.75	40	31.25	1218.75	31
-80	1281	1240	41	20	31.25	1250	31
-70	1312.25	1280	32.25	40	31.25	1281.25	31
-60	1343.5	1300	43.5	20	31.25	1312.5	31
-50	1374.75	1341	33.75	41	31.25	1343.75	31
-40	1406	1360	46	19	31.25	1375	31
-30	1437.25	1400	37.25	40	31.25	1406.25	31
-20	1468.5	1420	48.5	20	31.25	1437.5	31
-10	1499.75	1460	39.75	40	31.25	1468.75	31
0	1531	1480	51	20	31.25	1500	31
10	1562.25	1520	42.25	40	31.25	1531.25	31
20	1593.5	1541	52.5	21	31.25	1562.5	31
30	1624.75	1580	44.75	39	31.25	1593.75	31
40	1656	1608	48	28	31.25	1625	31
50	1687.25	1640	47.25	32	31.25	1656.25	31
60	1718.5	1680	38.5	40	31.25	1687.5	31
70	1749.75	1700	49.75	20	31.25	1718.75	31
80	1781	1740	41	40	31.25	1750	31
90	1812.25	1760	52.25	20	31.25	1781.25	31
100	1843.5	1800	43.5	40	31.25	1812.5	31
110	1874.75	1820	54.75	20	31.25	1843.75	31
120	1906	1860	46	40	31.25	1875	31
127	1927.875	1880	47.875	20	21.875	1896.875	31

TABLE II: Measured PWM offsets for a single ESC

wise, our other internal monitoring sensors, voltage and temperature, have also been empirically verified during dedicated pool test runs. All such sensors have proved very useful and practical.

3) *Acoustics Hardware*: Our acoustics PCBs were designed the thought that we would only need low-pass filters. However, when mounted to Oogway, low frequency noise from the thrusters and other inductive loads deteriorated the signal-to-noise ratio (SNR). To surmount this, we devised a bandpass filter (BPF). Remarkably, this did not require any modification to the existing circuitry as a backpack was designed to attach directly to the already-made boards.

With the improved filters, we constructed a testing setup wherein the gain of the amplifiers could be changed in real time. Through a series of tests in a pool, a gain on the order of magnitude of $10^1 V/V$ provided visual indication of the ping in the raw data, even from tens of meters away.

4) *Acoustics Software*: The acoustics software is composed of Python and Rust scripts working together, and the choices made are a reflection of




Fig. 26: Acoustics Filter and Amplifier Boards with BPF Backpack

our development process.

Our first attempts included tests with three hydrophones, using a target with a known x and y location in the plane. This required exact measurements within the pool relative to the position and rotation of the robot. Gathering this data was slow so there were only a few data points to work with.




Fig. 27: Magnitude of frequency response of filters in pinger-less pool. Bandpass 'hump' visible at $\sim 40\text{kHz}$ is about the same as the frequency range of our pinger.

Among those few data points received, the data was too noisy to be reliably predicted through ideal mathematical means. A tensorflow neural network trained to predict distance and angle on this data produced wildly varied results, as there weren't enough data points to accurately train the model. This led to the idea of collecting data in discrete angular quantities.

Splitting the angular space into octants allowed the team to gather much more data in the pool, as we could drag the pinger around the pool without worrying about its exact location. Predictions within octants would also be fine enough to steer the robot in the desired direction.




Fig. 28: Finding pings amidst background noise

The first task in the pipeline is to detect pings from a single hydrophone. Timing is important, so the goal is to find the ping in the audio source

when it started. Finding the ping using noise level alone didn't work with different noise levels, and the derivative of the sound wave curves produced too many false positives. Using the envelope of the sound wave increased accuracy, but it was too computationally expensive. In the end, we used a dynamic threshold based on the mean peak levels of the background noise. This made it adaptive to background noise levels, with few if any false positives.




Fig. 29: Classifying the acoustics subspace. Each dot is a ping measurement across three hydrophones. (a) rotation 1 (b) same graph different rotation

Once we could successfully detect pings as they started, we started trying to train a classifier into octants. Soon after we started, we noticed the dependence shown in Fig. 29: due to geometric constraints across three hydrophones, the measured relative ping timings are dependent. The three sets of relative measurements across the three hydrophones (1-2, 1-3, and 2-3) is rank-2, meaning that only two of those attributes are needed to contain all the information. Thus, all the data points collapse onto a plane, shown in gray on the graphs. This is useful, as when training we were able to use reduce the size

of each input to only the (1-2, 1-3) measurements.




Fig. 30: Training a tensorflow model on the data

Finally, we trained a neural network on the hundreds of examples we were able to collect from all around the pool. After preprocessing the data to include only what was needed, the network achieved 86% accuracy on the test dataset. With guesses this accurate, the robot can quickly and effectively find the location of the pinger and the probability of determining the global optimization in pinger location is greatly increased.

Lastly, once this pipeline was finished, it was too slow. Python itself is an interpreted language with large overhead, so the code itself could only be optimized so much. We set out to compile the code closer to the hardware with Rust. Rust is a compiled language that has about the same speed as C. One-by-one, the lines of code were translated from Python to Rust, making use of rust-specific optimizations as well. With the whole setup, the time to process and classify data from the hydrophones was dramatically decreased - from 25 seconds to 1.5 seconds. This code was then packaged into a Python library using pyo, allowing it to be integrated with ROS through python.

The acoustics system is currently still in development. This is mostly due to difficulties mounting the hydrophones properly where they are not blocked by Oogway's other components and do not suffer from thruster RF noise. Additionally, the test data was collected off of the robot on a separate custom rig. So to account for this, we would need to re-record the data on the robot in a possibly different layout. Further, the

We also considered using the MUSIC algorithm, which uses phase differences to determine the

source location of a sound. However this requires extremely precise phase information, which we have not yet been able to produce. Additionally, the hydrophones would need to be moved to within a wavelength of each other. We look to developing our own DAQ with custom hardware processing to combat this issue as a point of future research.

C. Software

The majority of the individual software components passed component testing without many problems.

1) *State*: Before the controls rewrite, we ran into verification issues with our robot's state. Like most AUVs, we represent our state in 6 variables: x , y , z , *roll*, *pitch*, *yaw*. When manually moving Oogway forward the x and y components of the state would increase drastically once the robot stopped moving. When debugging, we mainly looked at the output from our three positional sensors, the IMU, DVL, and pressure sensor. We quickly realized that the data from the IMU and DVL were in different coordinate frames; the IMU is right-handed and the DVL is left-handed. After changing the configurations, our state was more stable. However, we decided to create a testing plan in case of future issues with the state.

Our testing plan for issues with the state is as follows:

- 1) First, restart the robot, all Docker containers, and ensure that the Arduino is connected. Then turn the power DVL switch on.
- 2) Verify that all of the DVL, IMU, and pressure sensor are outputting to the correct ROS topics.
- 3) Verify that the robot-localization ROS package is running.
- 4) Manually move the robot forward, to the left, and up. Verify that the DVL's velocities in each of these directions is positive.
- 5) Manually rotate the robot counter-clockwise about the z -axis. Verify that the IMU's yaw value is increasing.
- 6) Leave the robot stationary and observe the frequencies of the DVL, IMU, and pressure sensor. Verify that they are all above 20Hz.
- 7) Leave the robot stationary and observe the `/state` topic's covariance matrices. Verify they are not increasing over time.

Later in the season, after the controls rewrite, we ran into similar issues with the state. Following our test plan, we realized that the pressure sensor would only publish in bursts of data and not at a constant rate. This causes the state to over-rely on the other two sensors for inaccurate depth information. After updating the Arduino code to resolve this issue, we observed that while the robot was stationary, the positional covariance would sometimes increase rapidly. This prompted a deep dive into our sensors configurations.

After many hours of testing we realized that there were two large issues with our IMU's configuration. The first was that ROS's robot-localization expected all quaternion data to be in ENU (East North Up) format while the IMU was outputting data in NED (North East Down) format. This caused Oogway's state to sometimes have a mismatch between the DVL's coordinate frame and the IMU's coordinate frame leading to drift over time. Another issue we found was with magnetic interference from Oogway's thrusters. The IMU does not have hard/soft iron adaptive filtering enabled by default, and once we enabled it, we observed a large increase in state stability.

Further, our testing also involved ensuring compliance with the manual for each of the three positional sensors. This included referencing diagrams such as Figure 31, which imposed mechanical constraints. In the case of the DVL, where we were not able to mount it in the direct center of the robot, we applied a corrective static transform.




Fig. 31: Correct alignment of the IMU with the DVL [14]

2) *Controls*: The rewrite of the controls system demanded careful yet through testing of the new code. We took a step-by-step approach, correcting bugs along the way, and covered a wide variety of scenarios. To prevent the robot from making any

sudden, unintended moves, we disabled the thrusters during these tests.

- 1) We first checked that when controls received a setpoint, it computed the correct error values, which determine where the robot will attempt to move. We fixed a bug that performed an incorrect transform from global to local coordinates.
- 2) Once the errors were accurate, they were fed into the PID loops, which computed their integrals and derivatives. We made sure that these were the correct values, and tuned the parameters of the Butterworth filter for optimal smoothing of the derivative.
- 3) We then checked that the control efforts output by the PID loops would lead to the robot achieving its desired state.
- 4) We checked the dynamic offset added to the PID outputs to counteract the robot's positive buoyancy. We rolled, pitched, and yawed the robot many different directions and ensured that the offset vector always pointed straight up, regardless of the robot's orientation.
- 5) We then verified that the thrust allocations obtained from our quadratic programming solver matched our expectations. We found that a few values were unexpectedly negated, which we corrected.
- 6) We modified the PID gains and other values on-the-fly, and ensured that they updated the system's behavior as expected and were saved to disk so they could be reused in future runs.
- 7) With all parts of the system independently tested and verified to be working as expected, we finally enabled the thrusters and allowed the robot to move using the new controls system. The robot behaved as we expected, and didn't perform any movements that would damage itself.

3) *Foxglove GUI*: Of particular note this year is our emphasis on tight integration of ROS with our Foxglove GUI developer tools. By automatically converting ROS message schemas to TypeScript types [12], we are able to enforce static type checking to validate the connections between our landside GUI and onboard ROS system. This means that *at development-time*, TypeScript will alert us if any concurrent work done on our ROS stack (e.g., a

change to our PID service API) will interfere with our Foxglove GUI stack. This allows us to catch numerous bugs early, before any actual testing has occurred.

4) *CV & Sonar*: The CV object detection worked in the pool on the first attempt. After some initial data collection, the sonar correctly identified the buoys and the gate with no additional testing. However, the integration between CV and sonar required many hours of testing in the pool. We found that the angles of the CV camera and sonar did not always match so the sonar would scan empty pool space. Additionally, we found that the sonar has a minimum range of about a meter so the robot would run through the buoy without stopping and backing up. Both of these are limitations in the hardware of the sonar so we may consider buying a different model for future competitions.

APPENDIX C: CUSTOM FOXGLOVE EXTENSIONS & LAYOUTS



Fig. 32: Standard Foxglove GUI Layout. Our GUI allows us to easily visualize and control Oogway during pool tests. For example, operators can view live bounding boxes, monitor the status of subsystems, and fine-tune PID constants on the fly.

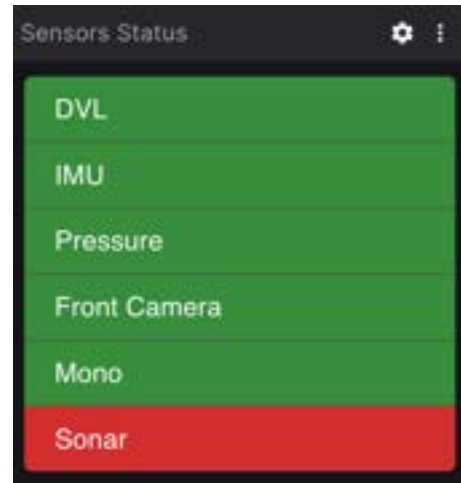


Fig. 35: The Sensors Status Panel continuously checks when the last message from each sensor was published. Sensors currently active are highlighted in green, while those not publishing appear in red.




Fig. 33: The System Status Panel provides real-time information about system components, such as the overall temperature and humidity of the electrical stack. Subsystems operating outside their expected parameters are highlighted, enabling operators to address potential issues proactively.




Fig. 36: The Toggle Joystick Panel maps inputs from a Thrustmaster T.Flight HOTAS X joystick to a desired power message, allowing operators to move Oogway in six degrees of freedom.




Fig. 34: The Toggle Controls Panel serves as a software kill-switch for our robots' control systems, ensuring immediate response to emergency situations or unexpected behavior.



Fig. 37: The PID Panel allows an operator to read and set PID gains. Operators can thus view real-time data using the PID Position Layout (see Fig. 43) and adjust PID constants using the PID Panel to instantly see how they affect Oogway's movement.

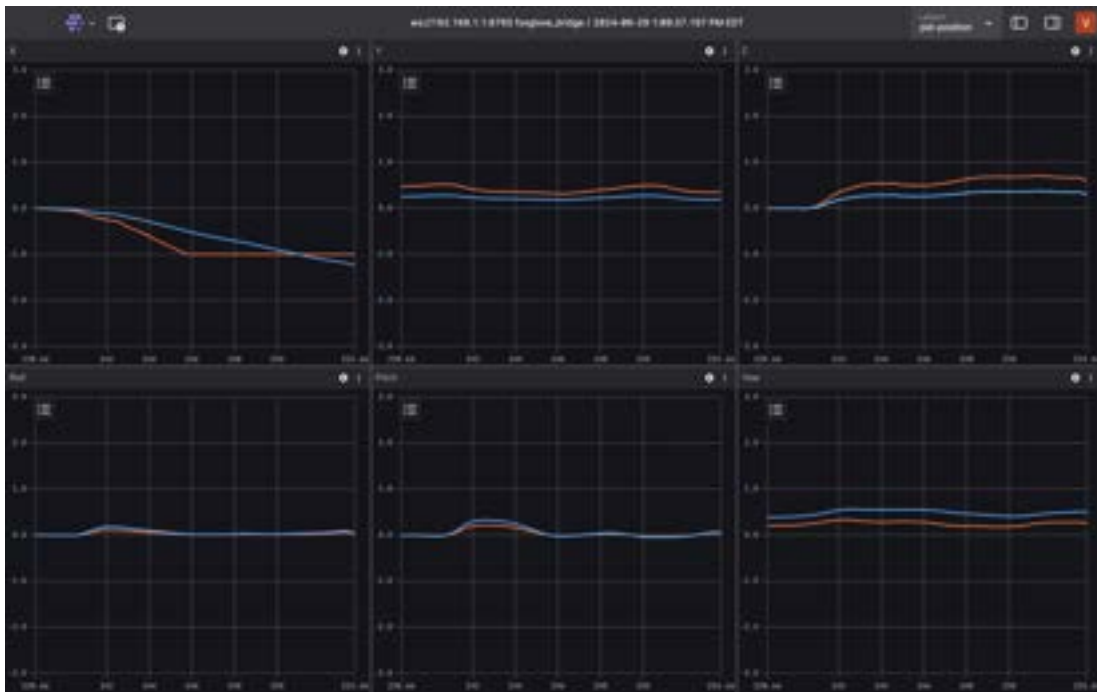


Fig. 38: PID Position Layout. Error is plotted in blue; control effort is plotted in orange. Each subplot corresponds to one of the six degrees of freedom: x, y, z, roll, pitch, and yaw.

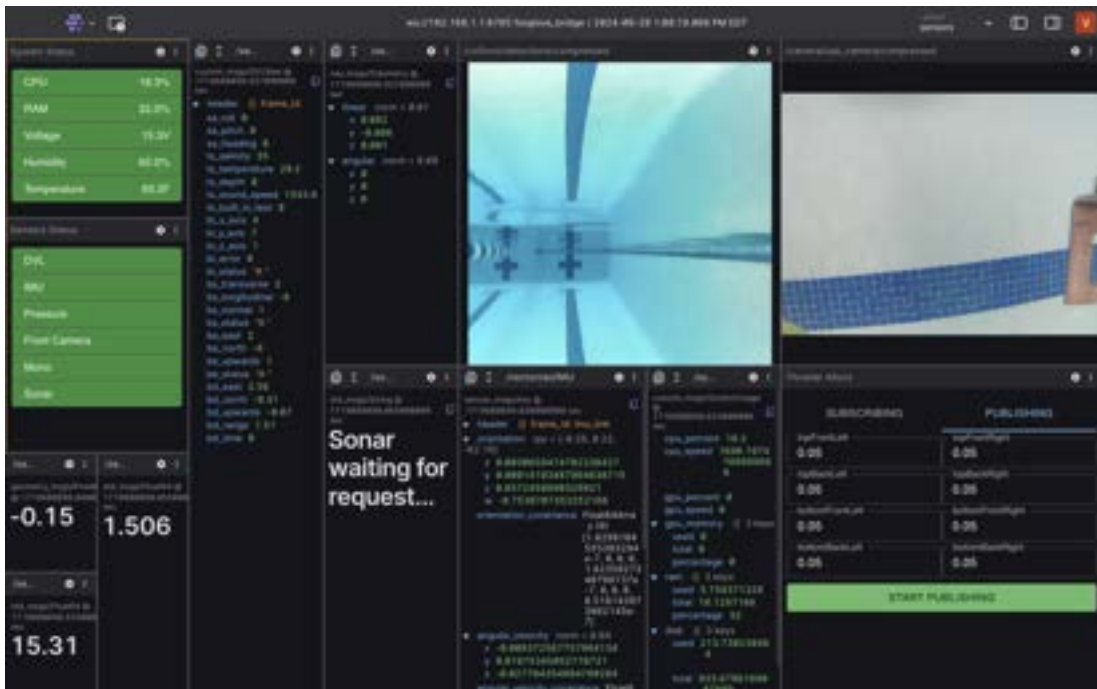


Fig. 39: Sensors Layout. This layout includes the System Status and Sensors Status panels along with panels to observe individual, raw sensor feed.

APPENDIX D: FULL ELECTRICAL SYSTEM DIAGRAMS

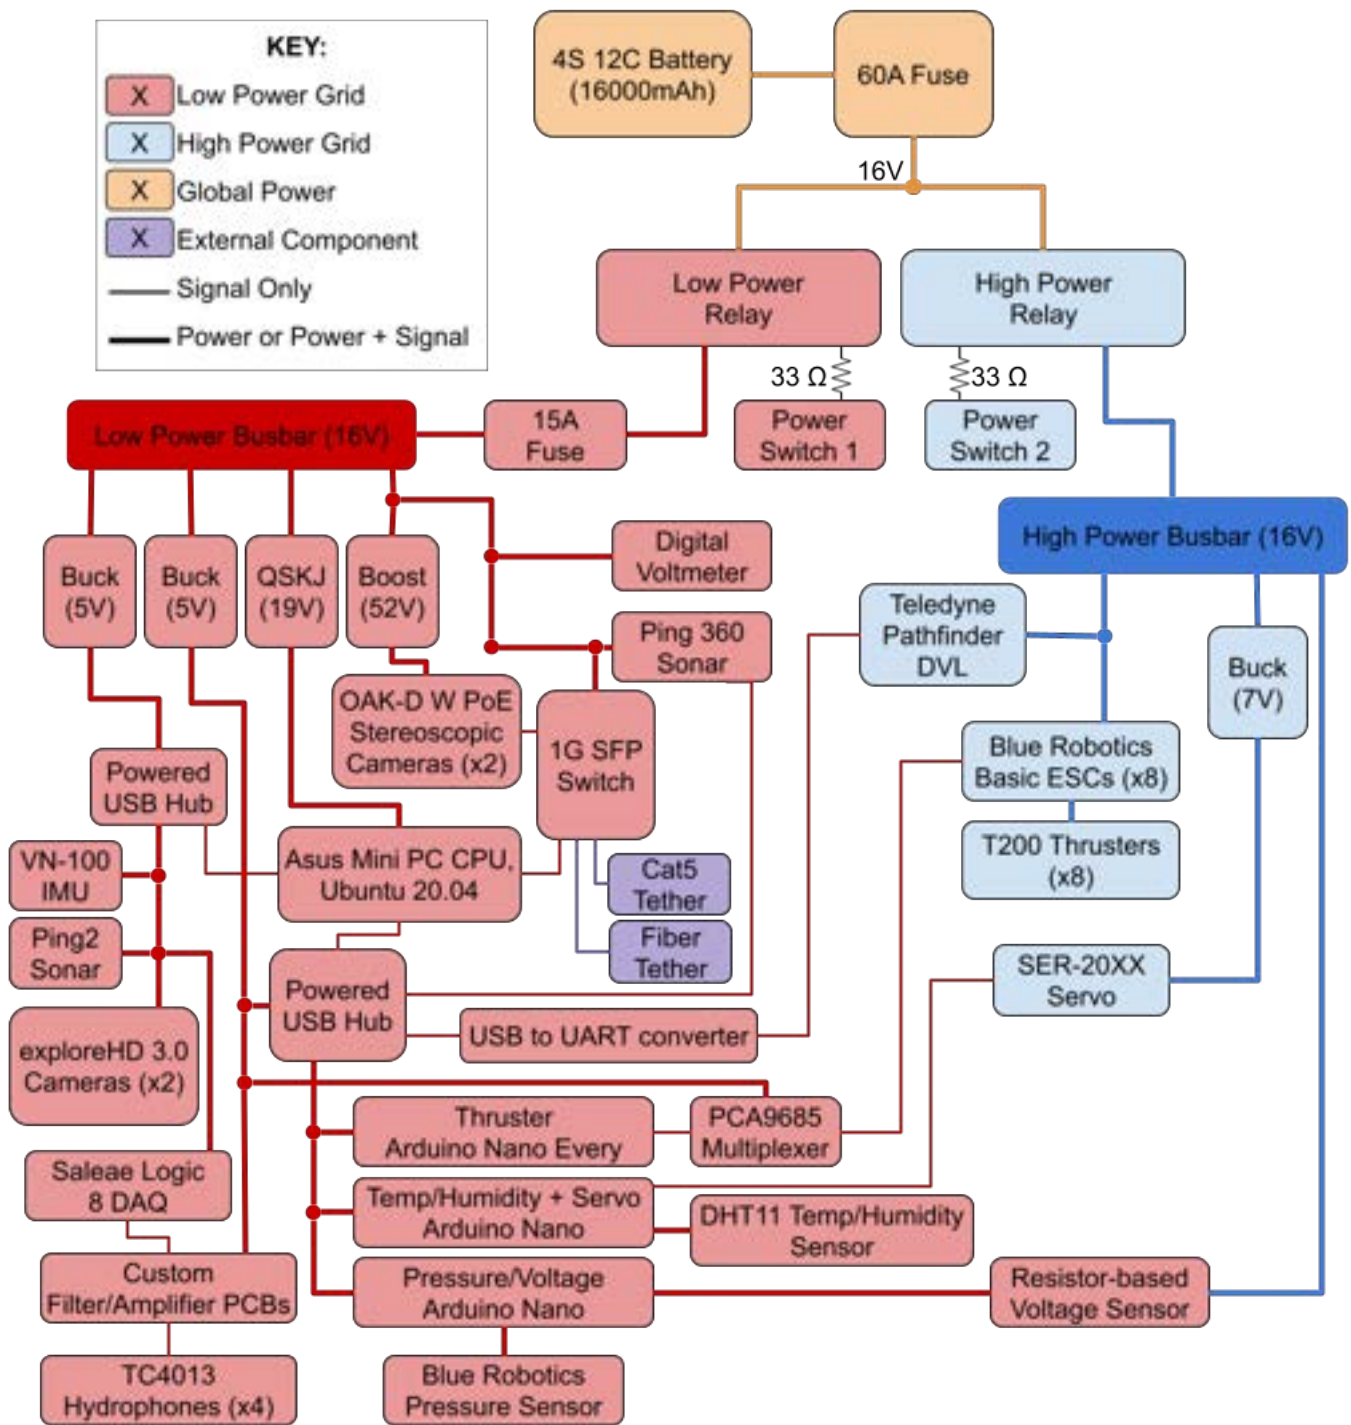


Fig. 40: Full Electrical System Diagram




Fig. 41: New Electrical Stack Mounts – CAD



Fig. 42: New Electrical Stack – Populated, Top View

APPENDIX E: SONAR PIPELINE



Fig. 43: Original sonar sweep data of a flat panel perpendicular to the robot underwater. The panel can be seen as the first light green blob in the image. The following darker green blobs are acoustic reflections.




Fig. 44: Conversion from polar image to Cartesian image. Curved lines in the polar image become straight lines in the Cartesian image which allows for easy classification. The straight line on the right of the image is the wall of the pool.




Fig. 45: Panel identified using cyan circles. The blobs behind the panel are ignored as they are acoustic reflections.

APPENDIX F: CONTROL FLOW DIAGRAM

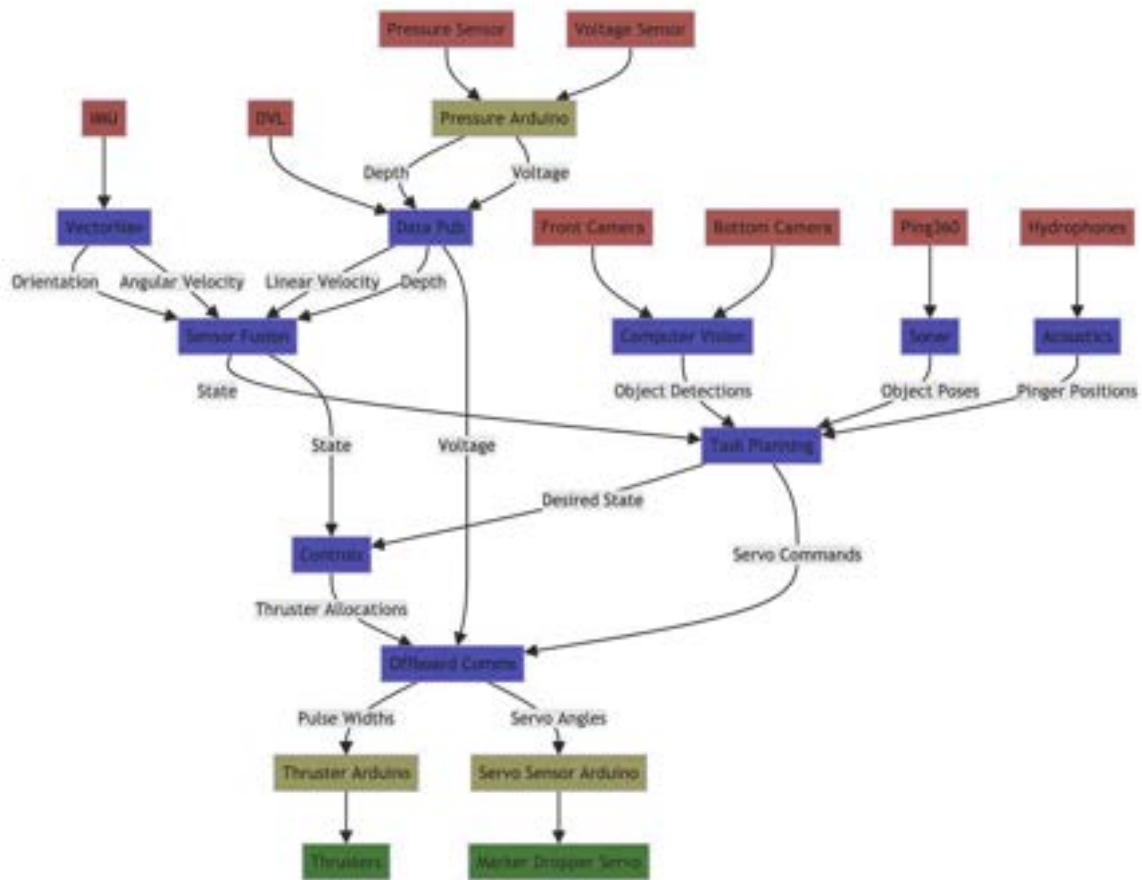


Fig. 46: System Control Flow Diagram

APPENDIX G: CONTROLS OVERVIEW DIAGRAM

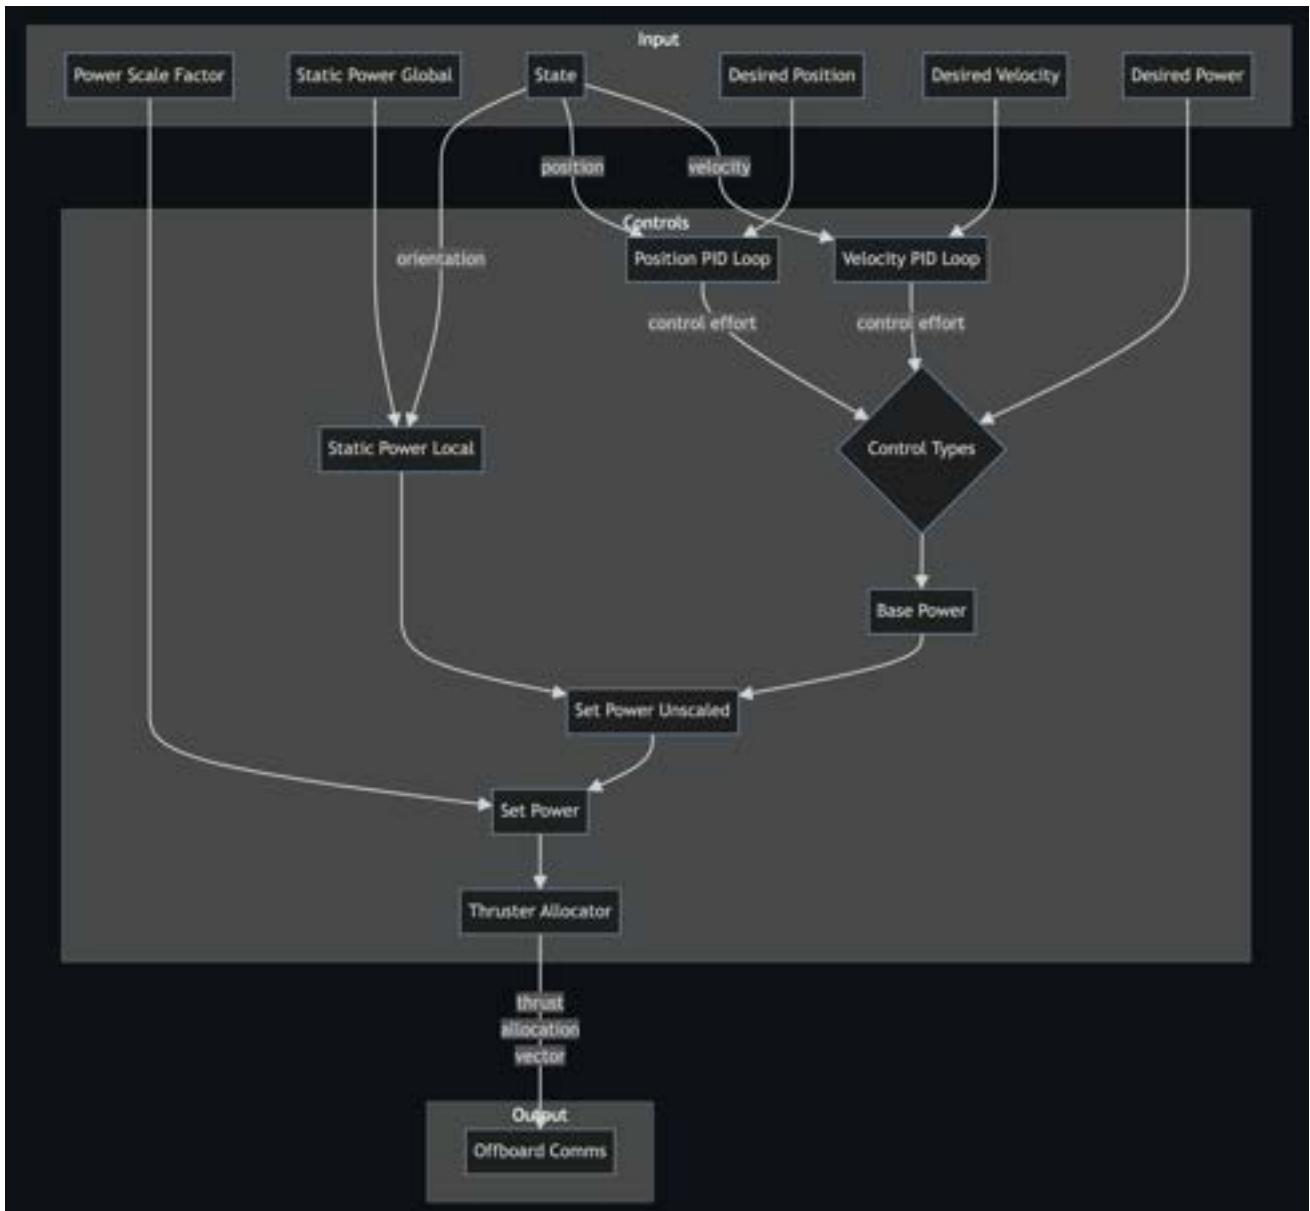


Fig. 47: New Controls System Diagram

APPENDIX H: OOGWAY RENDER AND PHOTO



Fig. 48: Render of Oogway




Fig. 49: Photo of Oogway

APPENDIX I: COMMUNITY OUTREACH

The Duke Robotics Club has a longstanding commitment to improving our community and fostering young students' interest in STEM fields. In the past we have volunteered in middle schools, helped Durham high schoolers learn about coding, and mentored middle and high school robotics teams. This year, multiple Duke Robotics members are proud mentors of Valence Robotics, a FIRST® (For Inspiration and Recognition of Science and Technology) robotics team. The FIRST Robotics Competition (FRC) is an annual competition specifically aimed at high school students with the intent to inspire them to pursue a career in STEM. Duke Robotics members regularly attend Valence Robotics' weekly meetings, helping explain advanced concepts to them, such as machining and software task planning. Many members of our team were part of FIRST themselves, and we value giving back to a community that has impacted our lives so much.

In our second year mentoring the team, it has been inspiring witnessing the students grow as leaders and engineers within the team. Our approach to mentorship is completely hands-off, letting them make mistakes and learn from them—an approach that we have found most valuable to our engineering education in the past. Throughout the season, they have gained valuable insights into the iterative design process, confronting challenges and refining their robot design through practical experience. We always make ourselves available for an in depth conversation on design or just life in general, fostering a deep connection between our teams.

We value our outreach as an important part of our team culture. By actively participating in mentorship and outreach programs, we not only share our passion for engineering but also continuously refine our own skills as engineers and communicators. Moreover, the relationships we build through these initiatives foster a sense of responsibility and commitment to the broader community. As we look to the future, we remain dedicated to expanding our impact and inspiring others to embrace the possibilities of robotics.



Fig. 50: Members of Duke Robotics with Valence Robotics at an FRC competition.