Cornell University Autonomous Underwater Vehicle: Design and Implementation of the Artemis and Apollo AUVs

Nicholas Akrawi, Mustafa Ansari, Tennyson Bardwell, John Bartlett, Zander Bolgar, Mark Brancale, Timothy Braren, Jonathan Chan, Richard Chen, Eric Chiang, Cuyler Crandall, Julia Currie, Eric Dalpe, Kerri Diamond, Nancy Ding, Joseph Featherston (lead), Felipe Fortuna, Jennifer Fuhrer, Dhruv Gaba, Ganesh Gopal, Kevin Guo, Horace He, Brendon Jackson, Mira Kim, Tae Kyung Kong, Sanika Kulkarni, Joshua Lederman, Alan Lee, Laura Lee, Junlan Lu, John Macdonald, Artina Maloki, Mihir Marathe, Kristina Nemeth, Alex Ozer, Kai Pacheco, Cora Peterson, Noel Picinich, Danny Qiu, Alexander Renda, Nicholas Sarkis, Chesley Tan, Ian Thompson, Samuel Tome, Anthony Viego, Patrick Wang, Zaibo Wang, James Wu, Angela Yang, Kevin Ying, Ray Zeng, Mei Zhang, and Evan Zhao.

***Abstract*—Artemis and Apollo are the 2016-2017 autonomous underwater vehicles (AUVs) created by CUAUV, a team of approximately 50 undergraduate students at Cornell University. Both vehicles were designed, manufactured, and tested over ten months. Artemis and Apollo draw heavily on past experience building AUVs, bringing together advanced technology to build our most robust underwater vehicles to date.**

I. Design Strategy

The Cornell University Autonomous Underwater Vehicle team's primary objective is to design and build an autonomous underwater vehicle (AUV). The foremost goal of this endeavor is the education of the team's members. Additionally, the team seeks to compete in the AUVSI RoboSub competition each year. This competition provides both the engineering constraints and supplementary motivation to the team in order to aid in the learning experience. Our overall design philosophy reflects a balance between the goals of seeking interesting and educational projects for the team's members and maximizing our chances for success in the RoboSub competition.

This year, we decided to continue with our previous strategy of developing two cooperative vehicles. Last year, our mini sub, Loki, marked the first instance of a second vehicle being deployed at RoboSub. While we consider Loki a success for being able to pass through the gate, we have much higher aspirations for the mini sub this year. Many of the shortcomings with Loki were a result of limited testing time due to being made a lower priority than the main sub and the increased workload of debugging two separate electrical systems. To counteract these issues, this year, we sought to get our new mini sub, Apollo, "in-water" before our new main sub, Artemis, and we designed both vehicles to have nearly identical electrical systems. In order to achieve our goal of having more testing time, it was decided to reuse key mechanical systems from Thor and Loki, CUAUV's 2015-2016 vehicles. This decision allowed for significantly decreased time spent in both the design and manufacturing phases. This was the first time the main hulls were reused, which was a large trade off as it strongly limited the designs of other projects, such as the frame and racks. This year also saw the development of more first-time projects than previously attempted. Some of these new projects include a model-based controller, custom thermal heat sinks, and an active ballast system. These projects promote innovation and present unique engineering challenges. Unfortunately, due to the nature of first-time projects, design and manufacturing problems often arise. In order to balance additional capabilities of the vehicles with reliability, many of these projects have been partially implemented and focus shifted to more critical components of the vehicles. Having gone through a first design cycle, these projects are expected to be revisited in the future without compromising reliability.

II. Vehicle Design

A. Mechanical

Artemis and Apollo were designed with robustness, interchangeability, and ease of manufacturing as priorities. Together they comprise a highly developed and reliable mechanical system. The largest trade-offs that are made to achieve our three goals are increased size and weight. Learning from previous years, the ability for the vehicles to be put under high stresses is a necessity. In order to ensure the vehicles are capable of handling the rigorous tasks expected of them, rigid structural design is implemented on both subs, which adds weight to both vehicles. Interchangeability of components between the vehicles, such as the battery pods and hydrophones enclosures, allows for a highly efficient mechanical system at the expense of the size of both vehicles increasing in order to accommodate these systems. Lastly, the majority of the mechanical systems are manufactured in-house, thus making manual machining of components a must in order to significantly reduce the amount of time and expenses required to manufacture these parts. However, manually machining parts limits their designs to simple geometries, often leading to larger enclosures. Each mechanical component on Artemis and Apollo is designed and optimized with these trade-offs in mind to create our most robust and capable vehicles to date.

The mechanical design cycle begins in the Fall with design reviews held every two weeks for each project. The first design review is the Pre-Preliminary Design Review (PPDR), which is held as a formal collaboration on initial project thoughts. At this design checkpoint, multiple designs are proposed and discussed for each project. The Preliminary Design Review (PDR) is then held with an initial Computer Aided Design (CAD) model of the design discussed during the PPDR. This design review goal is to establish key features in the project and work out any initial problems encountered in the design process. The Mid Design Review (MDR) is intended to be a checkpoint to ensure progress is made towards a final design and to allow for additional collaboration to happen before the design is finalized. The Final Design Review (FDR) is for presenting the completed design to the team and to discuss any last details regarding the project before entering the manufacturing phase. This design review process provides multiple checkpoints to ensure progress is made and to facilitate collaboration between team members.

Following the design phase, each project begins manufacturing, which lasts approximately four months. To reduce the dependence on CNC machines, which are shared between many users at the Cornell machine shop, many components were created with manual machining in mind. This was a large improvement over previous years since it allowed for a much faster production of parts. Following the machining of parts, sanding of each part is done in order to achieve proper surface finishes for sealing surfaces and surface aesthetics on external parts. Finally, the custom parts were sent out to be anodized to ensure they do not corrode underwater and extend their lifetime.

In order to more accurately plan our deadlines and to assist in project planning for future years, the time spent both machining and sanding the parts was monitored. The team logged 721 machining hours (of an original 1,100 hour estimate) and 686 sanding hours for the completion of Artemis and Apollo.

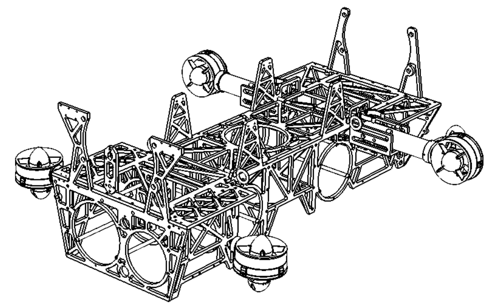

Fig. 1: The Artemis Frame is waterjet from $\frac{1}{4}$" 6061-T6 aluminum. The overall frame measures approximately 34" long, 12" tall, and 12" wide.

1) Frame: The frame constitutes all the structural pieces that hold the vehicles and all their enclosures together. This year, with the addition of many new capabilities, such as two downward manipulators, active ballast, and vector thrust, problems arose with keeping the frames light and compact on both Artemis and Apollo. Both the main and mini sub saw an increase in overall size this year due to the added components. A change in mounting of various enclosures occurred this year to facilitate fast swapping and easy debugging. Rather than all the enclosures being directly screwed into the frame, the battery pods and valve enclosure employ a slide-in mounting, secured with shock-cord, which eliminates the need for external brackets and fasteners. Each enclosure on the frame is strategically placed to achieve a neutral roll and pitch moment, as well as to allow for easy access for debugging. Artemis' eight thrusters and Apollo's six, are located around the frame to optimize for controllability as well as water flow around the vehicles.

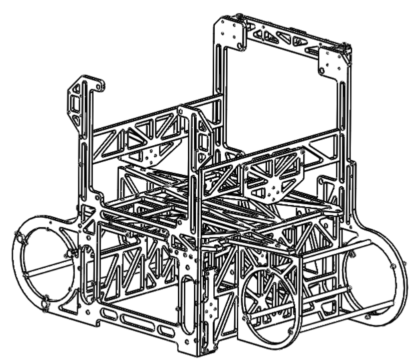

Fig. 2: The Apollo Frame is waterjet from $\frac{1}{4}$" 6061-T6 aluminum. It measures approximately 16" long, 16" tall, and 20" wide.

Both the Artemis and Apollo frame were rigorously tested using Finite Element Analysis (FEA) to ensure they retained structural integrity whilst optimizing for low weight. Utilizing water-jet cutting of the frame pieces allowed for rapid manufacturing, significantly decreasing the time spent on in house CNC machining.

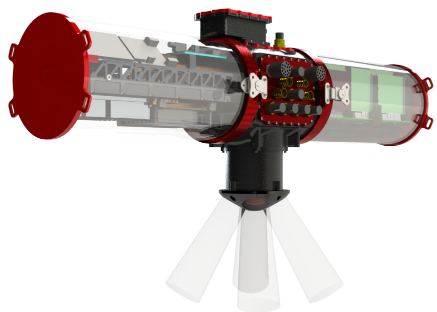

Fig. 3: The Artemis main hull has a dual-hull design. The midcap (central piece) provides mounting for the DVL, SEACON panels, sensor boom, and acrylic hulls. The overall length of the hull is approximately 34” long.

2) Main Hull: The main hulls are the largest enclosures on both Artemis and Apollo which house the custom electronics boards, computer, and many on board sensors. The main hull project is traditionally the most CNC intensive due to the numerous sealing surfaces, interfaces with other projects, and SEACON connectors. In our efforts to decrease the manufacturing time of the vehicles, the central part of the main hull (the midcap) from our previous vehicle, Thor, was chosen to be reused for Artemis. This decision was made possible by having designed for interchangeable SEACON connector panels on the midcap last year, allowing us to only have to remanufacture the connector panels for this year’s requirements. Conversely, the choice to reuse the midcap drove a lot of the hull design this year by constraining it to remain as a dual hull and limiting the overall size. The form factor of the vehicle relies heavily on the design of the main hull due to its size and central location, which resulted in the overall vehicle design not changing significantly this year. A new feature in the Artemis hull is the addition of conductive thermal paths from the electronics boards to the midcap and two endcaps.

The main hull was reused from the previous mini sub in order to complete the vehicle sooner and allow for more testing. The most significant change was moving the battery system from being inside the main hull, which required completely unsealing the hull to change a battery, to external battery pods interchangable with Artemis’ pods. The partial hydrophones system previously housed in the Loki hull was also transferred to an external enclosure in order to give the mini sub a full hydrophones array.

3) Racks: The Artemis and Apollo racks were designed with two main considerations: interchangeability and thermal conductivity. The racks provide structural support for all of the larger electrical components and custom PCBs in the vehicles, while still maintaining the ability to remove each board and modify it. In addition, this year’s racks effectively route the heat away from the components and into the cooler endcaps.

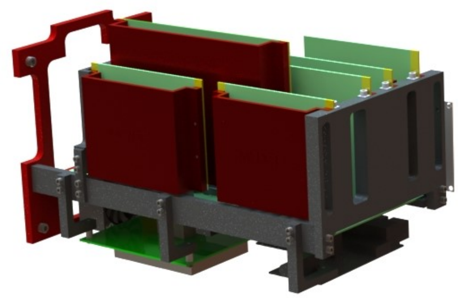

Fig. 4: The Apollo racks mount to the mini sub hull with four contact points (left red piece). The thermal management system consists of custom heat sinks (red blocks mounted to the boards) and a thermal tower (black block on the right) to disspate heat to the endcap. The custom electronics boards are mounted on the top layer of the racks and the computer is mounted underneath.

The custom electrical boards this year were designed to fit and operate in either vehicle. This decision saves time and allows for easier debugging, but also led to great constraints in the rack designs. Since Artemis and Apollo have dramatically different hull dimensions, it was difficult to design structures that would allow for the same boards in both vessels. In the future, the hull dimensions of both vehicles may need to be more similar in order for ease of design and assembly between vehicles.

The Artemis fore racks did not significantly change from Thor’s design. This decision was made since none of the electrical components in the fore rack were changed. However, some optimization was done to remove some bulkier items, thus leaving space for other boards to fit. The aft racks were completely redesigned for a lighter and more thermally effective system. However, the re-designed sliding mechanism proved more difficult to secure in place.

The Apollo racks were redesigned to both provide effective thermal paths to the endcap while also allowing for an easier user experience. Boards were easier to access, and additional space was provided for fans to circulate air throughout the system.

A completely new project was implemented this year that focused on analyzing the thermal properties of

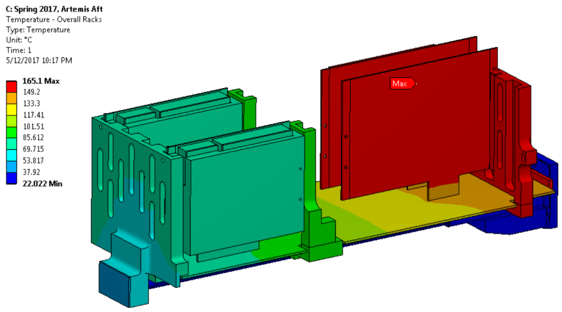

Fig. 5: Thermal analysis was conducted on the Artemis aft racks and Apollo racks using ANSYS (shown is the Artemis analysis). The analysis shows that the addition of custom heat sinks with a thermal path to the endcaps and midcap significantly decrease the temperature of the boards. This can be seen by the light blue color of the boards with heatsinks and the red color of the boards without.

boards on the vehicles and providing custom heatsinks and proper thermal paths to cool the components via conduction. Previously, the racks had used convection as the primary cooling tactic for boards, but convection proved to be limited due to the lack of fresh cool air within the vehicles. Thus, a conduction system was implemented to better route heat to the endcaps and midcap, which has direct contact with the water outside and provide more consistent cooling. One issue that this project encountered was that the board layouts were not finalized until very late into the design cycle, so preliminary heatsink designs would be rendered obsolete or not ideal as a result. In the future, schedules may need to be realigned to provide ideal heatsinks before the start of the summer.

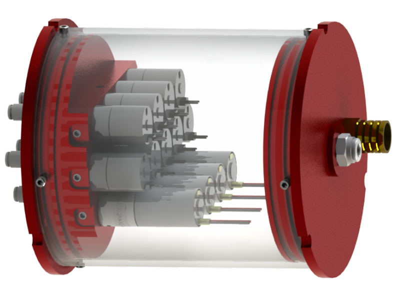

Fig. 6: The valve enclosure (hull rendered to be transparent) distributes air to actuators on board Artemis. The front of the enclosure (left) mounts the thirteen valves and push-to-connects. The back of the enclosure (right) houses the SEACON connector and pressure relief valve.

4) Actuators: The actuators system on Artemis utilizes thirteen valves to actuate two torpedoes, two droppers, two active downward manipulators, and a four point active ballast system. On board Apollo is a similar system scaled down to actuating two torpedoes and two droppers. Both vehicles' actuators system is powered by stored air preloaded in paintball tanks. The largest driving factor in the design of the Artemis valve enclosure was making it manually machinable. Historically, the valve enclosures have been CNC machined to accommodate their complexity and low tolerances. This year, the valve enclosure was simplified and made manually machinable to decrease its fabrication time at the expense of having a larger internal volume due to simpler geometry. Additionally, the concept of an integrated manifold was reused from last year since it was very successful. Though the manifold increases the intricacy of the design, it significantly decreases the volume of the enclosure by eliminating the need for an external manifold. For uniformity and ease of debugging, the valve enclosure was designed with the same slide-in mounting system as the battery pods.

The Loki valve enclosure from last year was repurposed for the use of Apollo. This decision was largely based on the Loki enclosure having not been used much last year. Changes made to the system on Apollo include a new set of valves and switching from a CO2 cartridge to a paintball tank powered system in order to better control the air pressure. These changes make the actuators system on Apollo very similar to that of the successful systems seen on the main vehicles.

Fig. 7: The downward manipulator parallel linkage arm in its extended state. The links gradually increase in length towards the bottom. The mounting and links are made from 6061-T6 and the gripper (grey at bottom) is 3D printed with ABS plastic.

The active manipulators on Artemis are for actuation of mission elements. An off the shelf piston was chosen this year instead of a custom pistonm which is consistent with our goals of decreasing manufacturing time making complex parts. The design of the manipulators are primarily completed prior to the mission tasks being released, thus necessitating the need for flexibility. Last year saw success with the usage of a parallel linkage extension system for the manipulator which provides a highly compact design while allowing for a long reach when extended. Building upon what was learned last year, experimentation was conducted on changing the sizes of the linkages in order to achieve even more reach. The trade-off of this design is that the links do not all lay flush with each other when retracted. A second manipulator was also added to Artemis in order to complete each task faster than before.

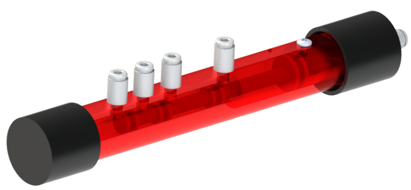

Fig. 8: Custom piston alternator for active ballast system. Based on the position of the piston, a different combination of the push-to-connects are allowed to flow, which allows for the fill, bleed, and hold configurations of the ballast bags.

5) Active Ballast: A new active ballast system was developed this year to provide dynamic buoyancy control of Artemis. The long term goal of this project was to eliminate the need to manually trim the vehicles with weights and foam, as well as to aid the thrusters in control. The biggest challenge with this system was creating a working mechanical system that was compact and easily controllable. In order to achieve this, a custom four state pneumatic piston was designed to allow for filling, bleeding, and holding of air into and out of four ballast bags. Multi-state pneumatic systems are often large, which led to the development of a completely new and unique piston system based on a ball-point pen mechanism. Custom ballast bags were also created in order to allow the active ballast system to change the internal volume of the vehicle. These bags needed to be able to bleed air using only water pressure, change volume rapidly, and retain a high failure point. Ballast bags were manufactured in house from a special plastic material that was melted to form bags with tube fittings and mounting to the frame.

6) Battery Pods: The new battery system returned to a removable battery pods in order to facilitate efficient battery changes and charging. The previous vehicles required unsealing of the battery enclosures each time the batteries were to be changed, which introduced the

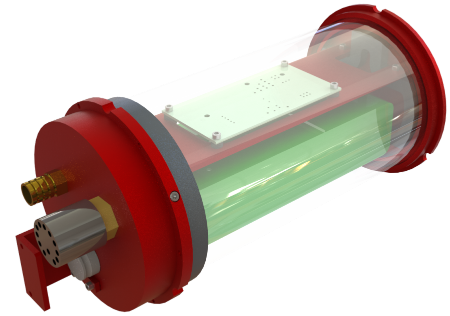

Fig. 9: The battery pod design for Artemis and Apollo provides mounting of the battery (green block inside hull) and Battery Management Board (grey plate inside hull). For manufacturability and weight reductions, the main hull is made from acryllic.

risks of improper resealing and water damage. The pod system allows for each battery to remain sealed inside its own compact enclosure with in-enclosure charging and battery monitoring. An easy to use slide-in mounting system with no external removable parts allow for quick and easy mounting of the pods. Another priority with the battery pods design was ensuring they could be used on both Artemis and Apollo interchangeably. This cross-compatibility between the vehicles reduces the time spent designing and manufacturing separate battery systems.

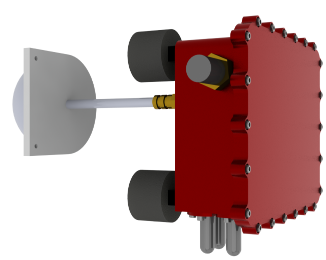

Fig. 10: The hydrophones enclosure on Artemis and Apollo feature a full transducer array (bottom) and transmit transducer (left). The mounting of the enclosure also allows for rubber dampers to be used (black cylinders) for mechanical isolation from the vehicle frame.

7) Hydrophones: Similar to the battery pods, the dual hydrophones system between Artemis and Apollo are interchangeable between the vehicles. The hydrophones system allows for acoustic tracking as well as communication between the two subs. The prior iteration of the hydrophones enclosure was designed specifically for the vehicle, making it highly space inefficient. This year's enclosure was made with future compatibility at the forefront. The transducer element are sealed to a separate collar to allow for removal from the main enclosure without needing to cut the wire. This allows for easier debugging of the system as well as the ability to reuse the transducer elements without losing length through the years. The size of the enclosure is significantly decreased and the mounting to the frame is simpler.

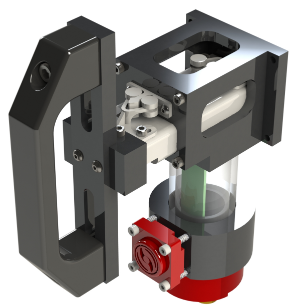

Fig. 11: The kill switch assembly features a full handle for easy kiling of the sub (left) and mission start button (red rectangle in bottom engraved with CUAUV logo). A fully sealed hull allows for removal of the board.

8) Kill Switch: The kill switch is the main human interface with the vehicles in water. This year, emphasis was put on the ergonomics and user interface of the switch. The kill switch on Artemis and Apollo features a full sized handle for ease of killing the sub underwater. There is also an easily accessible mission start button. Previous kill switch designs relied on epoxy to waterproof the electronics board to eliminate the need for an O-Ring seal. This method reduced the size of the kill switch but required a new board and kill switch hull to be remade if the board stopped working. This year's kill switch design features a hermetically sealed main enclosure to allow access to the board and sensors housed inside.

9) Sensor Boom:

In accordance with our goal to create wide cross compatibility between Artemis and Apollo, the sensor boom created for Apollo mimics the one used on Artemis to house the inertia measurement unit (IMU), accelerometers, gyroscope, and compass. These sensors were moved to an external enclosure a few years back to decrease noise interference with success.

B. Electrical

The electrical systems of Artemis and Apollo are designed to provide power, drive actuators and thrusters, and interface with various sensors. The electrical team follows a similar design cycle as the mechanical team, with the fall semester spent designing and the spring spent manufacturing and testing.

The vast majority of printed circuit boards (PCBs) are shared between Artemis and Apollo. The boards within the main hull are connected via a custom PCB backplane using slot-edge connectors. Each vehicle has a unique backplane to accomadate the different hull designs. In previous years, the backplane connector pinout was identical for every PCB. This made the routing of the backplane easier and also allowed for boards to be swapped between positions. However, since the pinout of this connector needed to be changed each year for various changes and new features, it was impossible for boards to be reused from year to year. This year each board connects to its own specialized connector, which gives us the ability to reuse some of this years PCBs in next years vehicles.

1) Power: The power system of our vehicles consists of four unique boards. There is the Battery Management Board (BMB), Merge Board, Power Distribution Board and DCDC board. The vehicles are each powered by two 4-cell 14.8V Lithium Polymer batteries. Our vehicles have previously been powered using 6-cell 22.2V batteries for a number of years while last year Loki was powered by a 3-cell 11.1V battery. We chose to switch to 4-cell batteries to enable both the Blue Robotics T-100 and T-200 thrusters to be run near their optimal voltage.

The first board in the power path of the vehicles is the BMB which accompanies each battery inside the battery pod. The project was created because, on multiple occasions, our batteries have been over-discharged to a point that damaged both the batteries themselves and electrical components on the submarine.

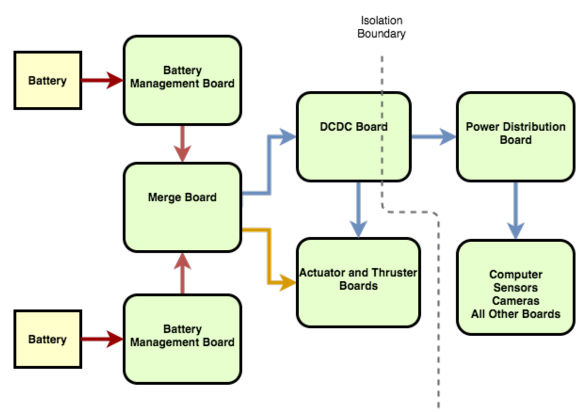

Fig. 12: System level overview of Artemis and Apollo's power system. The blue lines only are active if the vehicle on switch is activated. The yellow line is only active if the on switch is activated and the vehicle is not "hard killed."

The Merge board handles swapping between the two batteries as well as responding to the On and kill switches. The merge functionality is accomplished by connecting each battery across a transistor controlled by an Ideal Diode Controller. This enables the vehicles to be powered by either of the two batteries, which allows continuous vehicle up-time during battery swaps. Due to the large bulk capacitance on the thruster controllers last year we experienced issues with high inrush current. This year we attempted to swap to new control chips which provide in-rush current protection. However, we could not properly test these in the required time so we reverted to the prior design for the final iteration of the board. The board contains a set of Field Effect Transistors (FETs) which are controlled by the vehicles On switch. When these FETs are off, only the BMB and Merge boards receive power. This year we have added the ability for the microcontroller on the merge board to also turn off these FETs, providing a way for the vehicle to be remotely restarted via software. The kill switch functionality is implemented in the same manner as the On switch except only the power to the thrusters and actuators board is controlled.

Due to the large number of different sensors and systems on the vehicle, an array of different voltage rails are needed. Additionally, due to the electrical noise caused by the thrusters and actuators systems, we implemented an isolation boundary between the mechatronics boards and the rest of the vehicle. Artemis requires 6 different DC-DC converters to create the unisolated 5V, as well as the isolated 5V, 12V, 12V CPU, 24V, and 48V rails. These components have traditionally been spread between the Merge and Power Distribution boards. Since the converters are very large and expensive relative to other components on our boards, we decided to place them all together on one simple board. This allows for the other boards to shrink in size and for multiple revisions of those boards to be made without needing to desolder and resolder the DC-DC converters, potentially damaging them. Additionally, since this board has very little custom circuitry, it is an ideal candidate for reuse in future years. In order to efficiently fit inside the hull of the vehicle, the DC-DC board was split into two boards with 3 DC-DC converters on each board. Because of the fewer number of components on the mini sub, only one of the unisolated and isolated 5V and the 12V rails are necessary, which allows it to only use one of the DC-DC boards. In order to save on board printing coasts, the two boards were designed to be identical with only different components populated to differentiate.

The Power Distribution board takes in the isolated voltage rails from the DC-DC boards and provides rails to the rest of the components on the vehicle. It allows us to enable the power to monitor the current draw of each component on the vehicle individually. This enables us to selectively power cycle individual components, which is very useful as any IT professional can tell you "turning it off and on again" fixes many strange issues.

2) Mechatronics: The actuators board serves as a driver for each of the pneumatic valves on the submarine. The board uses the battery voltage to power each of the valves. Since the batteries reach 16.8V at full charge and the valves are only designed for 12V, when a valve is enabled the board uses Pulse-Width Modulation (PWM) to operate the valve at a rate simulating 12V operation to prevent damage to the valves.

The thruster board houses the electronic speed controllers (ESCs) for each of our thrusters. The board receives direction and speed commands for each thruster. The thrusters board also handles ramping between different thruster speeds to prevent motor acceleration which can cause excessive current draw that damages the electrical system.

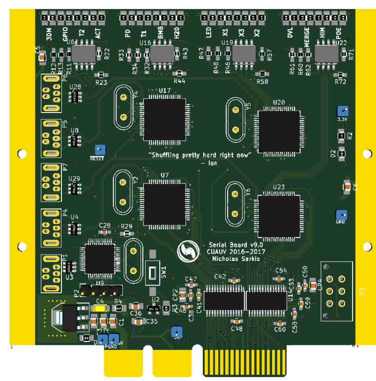

Fig. 13: Rendering of Serial Board. The USB ports are along the left edge with the four FTDI USB to UART interface ICs in the middel of the board and the RS-232 level shifters along the bottom edge directly above the solt-edge connector.

3) Serial: The Serial board is responsible for interfacing between the vehicles main computer and its custom boards and on board sensors. The serial board connects to the computer over the Universal Serial Bus (USB) and communicates with the custom electronics via RS-232 Serial. The board was originally designed to include a USB hub to allow the board to connect via only one USB port. However, the custom hub circuitry could not be stabilized in a timely manner so the board instead has several USB ports that are connected to an off-the-shelf USB hub. In addition to RS-232, the serial board also provides a Controller Area Network (CAN) interface. This is implemented using a STM32 ARM microcontroller with custom firmware.

4) Sensors: Each vehicle contains an array of different sensors. For orientation, the vehicles use a Lord Microstrain inertial sensor. Artemis uses the 3DM-GX1, while Apollo uses the 3DM-GX4. The decision to equip mini sub with the newer of the sensors was made because the 3DM-GX4 is more immune to electromagnetic interference from the thrusters and due to the compact nature of the mini sub the sensors are located nearer the thrusters. In addition, the vehicles use a custom Heading and Inertial Measurement (HIM) board to provide redundancy. Artemis uses a Teledyne Explorer Doppler Velocity Log (DVL) to provide accurate velocity. A pressure depth sensor is the primary method of determining the vehicles depth. However, the DVL altimeter is sometimes used when precise distances from the floor of the pool are needed. A notable example of this is removing objects from the tower during the recovery mission.

In addition to these navigational sensors, the vehicles also feature sensors for self monitoring. The main hull of each vehicle is depressurized to a slight vacuum to allow us to notice leaks before the vehicle enters the water. An on board pressure sensor gives us feedback on whether the hull is properly holding the vacuum. This has traditionally been our only method of remotely detecting leaks and is only applicable to our main hull. Leaks are more common, and harder to notice, in the smaller external enclosures. This year we implemented simple leak sensors to detect leaks throughout out the vehicle.

5) Hydrophones: In order to test the pinger tracking features of our hydrophones system, we have traditionally used off-the-shelf pingers. Due to the high cost of these devices, it has been a focus of ours to develop a custom pinger solution. Additionally, the Pinger board features the same transmit circuitry as the Acoustic Transceiver board which allows it to act as a prototype for the transmit functionality of the Acoustic Transceiver board.

The Acoustic Transceiver board is the newest iteration of our custom hydrophones system. It is designed to retain the passive pinger tracking functionality while also adding transmit capabilities to enable acoustic communication between vehicles. The signal processing for both receive and transmit is computed on an on board Field Programmable Gate Array (FPGA). The receive transducers are Teledyne TC-4013 hydrophones while the transmit transducers are made by potting a piezo element inside epoxy.

C. Software

Our vehicle software stack is structured as a series of daemons which communicate via shared memory, running on the Debian 8 distribution of Linux. The software stack runs on a ADLINK cExpress-HL-i7-4650U on Artemis and Intel NUC5i7RYH on Apollo. To streamline development, the AUV software stack is sufficiently general such that Artemis and Apollo can share almost identical software stacks. The differences between the vehicles are almost entirely specified by a configuration file for each vehicle. These mostly describe hardware-dependent parameters, such as the vehicles centers of mass, which sensors should be used to measure which quantities, and the positions of the thrusters.

1) Communications: Fast and thread-safe inter-process communication is handled by our shared memory system (SHM). SHM consists of many variables organized into groups and automatically protected from synchronization issues by mutexes, which can be written to and read from via simple API. SHM enables modularity and extensibility by splitting the functionality in our software stack into multiple processes.

Communication between the electrical system and the software stack is accomplished by seriald, a daemon which either sends the value of a specified SHM variable to a specified electrical board periodically, or reads the value of a variable being sent to the software system over serial periodically, and writes it to SHM. For example, the GPIO board will periodically send the value of the depth sensor to the main computer, which seriald will receive and write to the depth SHM variable. Likewise, when a desired motor speed is written to the motor desires SHM group, that value is sent to the thrusters board by seriald.

2) Control: On both Artemis and Apollo, sensor filtering is done by kalmand, a daemon which implements a Kalman filter. We use a Unscented Kalman Filter implementation (UKF), which is written in Python using the NumPy library for efficient matrix multiplication. Kalmand removes noise from the orientation estimate from our IMUs. On Artemis, it also estimates vehicle position from DVL velocity, and fuses DVL depth velocity with pressure sensor data to provide a smooth depth velocity estimate.

The controller is responsible for taking an input of desired velocity and orientation and computes the set of thruster speeds to achieve that state. The controller has full 6-degrees of freedom control. Based upon the desired state, the controller first estimates the different forces and torques that need to be exerted to reach the state. These forces are then mapped into the optimal set of thruster speeds using the known thruster positions and torque curves. In the three orientational degrees (pitch, roll, yaw) as well as depth the controller directly tries to achieve a specific position. However, for the surge and sway directions, the control is based off of velocity since that is the data available from the DVL.

In many cases it is more useful for the position of the sub to be controlled directly rather than through velocities. To accomplish this there is a mode to the controller which targets specific coordinates rather than velocities. An example of the differences between the behavior of the vehicle in velocity vs positional control mode is if you shoved the vehicle while it is trying to stay still. In velocity control, the sub would immediately try to stop moving after being shoved whereas in positional control it will return to its original location before stopping.

The implementation of positional control allows more advanced navigation to be written in a layer above the controller. The Navigator takes an input of points, which contain both a location and an orientation, and dynamically calculate an optimal route between the points. The Navigator can also use information about known obstacles to avoid collisions. This feature makes writing missions to go between areas of the pool as well as precise movement much simpler.

3) Mission: The top level code running on the vehicles is the “Mission.” This is the program that dictates what the vehicle is attempting to accomplish and how to do it. Missions are organized into different tasks which can contain subtasks. For example a simple buoy ramming task could include a search task to find the buoy, a task to align the vehicle with the buoy and a task to surge into the buoy. Tasks can be organized together using different “combinators” with the most common being sequential and concurrent. Using the previous buoy ram example, the search task would need to occur before the aligning and surging so the sequential combinator would be used. However, when surging towards the buoy, the vehicle must continue to be aligned with the buoy so the surge and align task should be run concurrently.

In order to facilitate code reuse, a large internal repository of shared tasks called the mission framework is maintained. This includes many small tasks that are common to many missions. An example of this is the align target task. Most mission tasks require aligning to a target during some portion. By making this a shared task between missions, the amount of code that needs to be written and tested is greatly reduced. Each year this library is expanded upon as new tasks are created.

4) Vision: Most of the tasks in the RoboSub competition are highly reliant on computer vision. The vision system allows for the efficient processing of camera images. The abstract system allows input from not only the vehicles cameras, but for prerecorded video files or input from the visualizer. The system provides the data to both the vision processing modules as well as log recording and debugging outputs. The vision system itself manages the concurrent and parallel problem of transferring large images from many sources to many modules, where each part can produce and consume at varying rates.

The CUAUV Automated Vision Evaluator (CAVE) is a tool that allows for the playback of annotated videos. Videos can have mission elements tagged in each frame, and CAVE can compare the existing tags to live vision recognition. A new component on CAVE this year tracks objects between frames allowing a single click to tag an object across hundreds of frames.

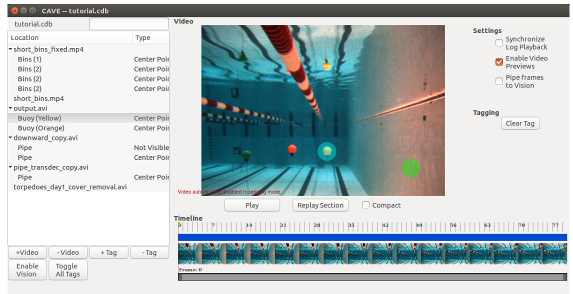

Fig. 14: Screenshot of the CUAUV Automated Vision Evaluator with the yellow buoy annotated.

5) Infrastructure: The current iteration of the WebGUI is the most capable version yet. It was designed for both software, and general team members. A main goal is to allow any team member to test basic functions of the sub without the need for a terminal. This includes seeing concise or full sensor status, and the ability to trigger tests of thrusters or actuators.

A common problem for our software team is a lack of in pool testing time. This has motivated the software team to place a higher focus on simulated testing. The largest component of of our simulation systems is Fishbowl, a network enabled physics based simulator. A simulator allows each member to quickly test changes to missions in a live environment, all contained on a personal computer.

A related but separate part of the offline, local infrastructure is the visualizer. The visualizer is primarily used to display the current state of the simulator, but its uses go far beyond just that. The visualizer can render what cameras attached to the submarine would see, and feed the video back into the vision system. This allows for the basic performance of vision modules to be tested on simulated footage. While modules cannot be as finely tuned on the simulated footage as with real video, it has the powerful advantage that missions and vision modules can be run together in a fully simulated environment whereas with prerecorded footage you can only observe the performance of the vision module following the path that the vehicle took when capturing the footage. Lastly, the visualizer can display the believed state of the submarine. This is invaluable for testing sensor calibration - any discrepancy between the sub movement, and the displayed movement is immediately clear in real-time.

We have made an additional push towards standardizing our development and live environments. We have tested using NixOS to create a standard container that each team member can use to ensure they are working in an equivalent environment that the sub will use. NixOS has also been tested directly on the sub. In the future, if we choose to switch to NixOS completely, we can have the same exact environment (whose state is described in a repository tracked configuration) on every installation of our software. Lastly, this would allow us to trivialize the process of building our open source software.

III. Experimental Results

During the academic year, CUAUV holds pool tests every Sunday night. In the Fall semester through the beginning of the Spring semester, the previous year's vehicles are used to continue to test software improvements. In March, the old vehicles are retired to reuse parts for the incoming vehicles. Because of the staggered launch of our two vehicles this year, number of weeks without any vehicles was limited to only one. During the summer break, the number of pool tests increases dramatically. In previous years we have sought to test for 10 hours per day every day in the weeks leading to competition. Due to numerous constraints this year, as well as the recognition that the team has gone way past the point of diminishing returns in previous year, the testing plan plan this summer is much more modest. The goal is for 6 hour tests, 4 days per week. At the time of this writing, Artemis and Apollo have approximately 80 and 100 hours of testing time in the pool respectively.

At the time of writing, the vast majority of mechanical and electrical systems are online and functional. There are still a few projects that are still being worked on and are close to completion. We expect most of these to be working in the coming weeks. These projects include the: kill switch, the second downward-manipulator, leak sensors, PoE, and the acoustic transceiver. Unfortunately some projects were not able to be completed in a timely manner and due to reduced manpower over the summer will likely not be able to be completed before RoboSub. This includes active ballast, vectored thrusters, and the battery management board. The software team is working on testing and tuning the vision and mission code. Currently, Apollo is able to follow paths and ram buoys with some reliability. Artemis is currently tuning its manipulators to grab the tubes for Collect and Classify.

Acknowledgment

CUAUV would like to thank all the individuals who have supported us over the past year, including Cornell's MAE, ECE, and CS departments, and Cornell alumni who have supplied our team with needed lab space, materials, tools, and funding. CUAUV would especially like to thank our advisers: Professors Ross Knepper, Bruce Land, and Robert Shepherd; Director of Cornell Aquatics Brigitta Putnam, the MAE Director of Instructional Laboratories Matt Ulinski, the Swanson Director of Engineering Student Project Teams Rebecca Macdonald, Cornell Engineering's administrative staff, especially Patricia Gonyea, Emily Ivory, Carol Moss, and Maureen Letteer, and last but not least the Cornell Rapid Prototyping Lab (RPL).

We would also like to thank all of our corporate sponsors, without whom we would not be able to compete:

Platinum Sponsors: Cornell University; Monster Tool Company; VideoRay; SolidWorks; Nexlogic; General Plastics; Mathworks; SEACON; Shell; and American Portwell Technology.

Gold Sponsors: Teledyne RD Instruments, LORD Microstrain, SeaBotix, Seaperch, General Motors Foundation, Adlink Technology Inc., IEEE Areospace & Electronics Systems Society and SGS Tool Company, Shaw-Almex.

Silver Sponsors: CadSoft USA; Intel Corporation; LeddarTech; Molex; Advance Energy Inc.; iDS; ConocoPhillips; Lockheed Martin; PNI Sensor Corporation; Sparton; Allied Vision Technologies; Surface Finish Technologies; Exxon Mobil; General Electric; Bittele; SKB Industrial; Teledyne Reson; XIMEA Corporation; Advanced Circuits; and congatec.

Bronze Sponsors: Georgia Case Company; Adata; 3D Connexion; Sensata Technologies; Digi-Key; Advanced Circuits; Polymer Plastics Corporation; The Mines Press Inc.; Cayuga Xpress; MaxAmps; Pneumadyne; Domino's Pizza; DIAB; Drexel University; RJE International; Asian Circuits; and Fast Circuits.

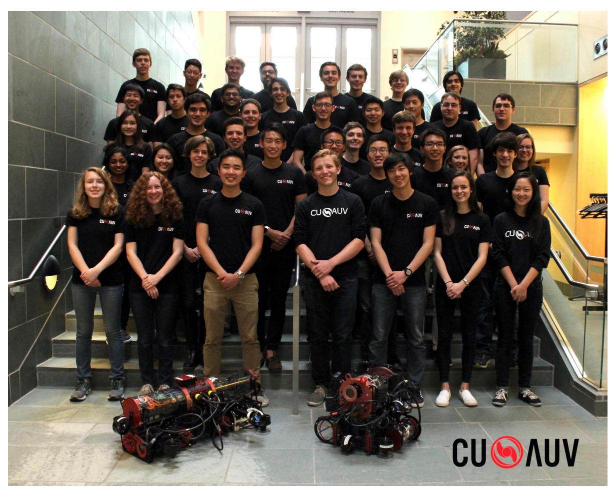

Fig. 15: The 2016-2017 CUAUV team with this year's vehicles Artemis and Apollo.