Design and Implementation of an Integrated, Autonomous Underwater Vehicle: *Flo*

T. Okamoto, J. Orenstein, K. Agarwal, H. Sun, B. Hollaway, F. Zhou, A. Darbinian, C. O'Neill, A. Batzorig, J. Deacon
California Institute of Technology, 1200 E California Blvd., CA 91126
June 20, 2017

Abstract – **Caltech Robotics Team's new autonomous underwater vehicle (AUV), *Flo*, is designed for precise movement and high-level autonomy for the 2017 intercollegiate RoboSub competition. Designed and manufactured by 25 undergraduates at the California Institute of Technology, *Flo* is the team's most sophisticated and capable submarine yet, iterating on last year's vehicle *Dory*. Improvements include a stereoscopic vision system, a refined pressure hull design, seven new thrusters to provide extra power and reliability, a servo-driven gripper, and magnetically-actuated torpedoes and markers. With these improvements, the vehicle can maintain an arbitrary orientation and velocity, navigate the obstacle course under many lighting conditions, locate an ultrasonic pinger, fire torpedoes and markers, and grip various objects. *Flo* represents the state-of-the-art for low cost, actively stabilized autonomous underwater vehicles.**

I. INTRODUCTION

The Caltech Robotics Team's mission is to help its members develop engineering, communication, and collaboration skills in an exciting and hands-on environment. We do this by designing and building autonomous underwater vehicles for research, education, and outreach. Our newly designed vehicle, *Flo*, was built to complete the challenging RoboSub competition course, cosponsored by the Association for Unmanned Vehicle Systems International (AUVSI) and the Office of Naval Research. The competition takes place annually in July at the Space and Naval Warfare Systems Center Pacific's San Diego Transducer Evaluation Center (TRANSDEC). The competition simulates real-world underwater vehicle tasks such as manipulation, navigation, and visual inspection. Competition elements include firing torpedoes at small targets, locating sonar pingers, manipulating small objects, and dropping markers into bins. All of these tasks must be completed autonomously with no human in the loop. To complete these tasks, *Flo* must have accurate visual and acoustic sensors, precise navigation, control and guidance, and robust object recognition and manipulation.

II. DESIGN OVERVIEW

Flo is an 82-pound autonomous underwater vehicle designed for reliable operation and precise control. The team has learned many lessons from last year's vehicle, *Dory*, and integrated that knowledge into *Flo*. The pressure hull consists of two custom-designed anodized aluminum shells that were machined on a five-axis CNC. An acrylic dome allows us to clearly see the electronics. Electromagnets are positioned in strategically located regions of the hull to actuate external torpedo and marker launchers without requiring openings for pneumatics lines. Powerful servos are positioned outside the hull to actuate a 3-DoF gripper. Seven thrusters give the vehicle control over all six degrees of freedom, and redundancy in certain axes. We replaced the gimbaling camera on *Dory* with two wide-angle cameras arranged as a stereo pair. Each have the same field of view as the entire gimbal did last year, which allows the vehicle to see all of its surroundings simultaneously.

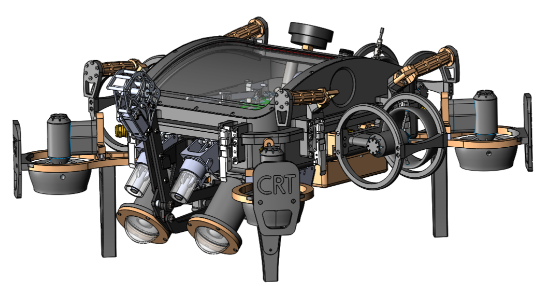

Figure 1: Computer-aided design model of *Flo*.

The electrical systems have also been improved by reducing the overall board count, which saves valuable space whilst retaining modularity. The electronics are mounted to a two stage enclosure that swings open on a 4-bar linkage to ensure easy access to all components. The vehicle is powered by two lithium-polymer batteries giving *Flo* a runtime of over two hours. The electronics have separate power and battery systems for the actuators and for the computer and sensors. The vehicle's software runs on a smaller Intel NUC™ computer system, which saves mass and volume. *Flo* uses a custom inter-process communication library to implement its vision, actuation, and autonomy systems. This year we have refined our sensor systems and simplified inter-board communication protocols.

III. MECHANICAL SYSTEM

Flo's mechanical systems consist of a main pressure hull, external electrical equipment, and task-interaction hardware. Some external equipment (the camera modules and battery pods) are integrated into the main pressure hull; their enclosures are sealed to prevent water ingress and mount to the main pressure hull. All parts were designed by team members using SolidWorks computer-aided design software. Several pressure and structurally critical parts were simulated using ANSYS finite element analysis software. Parts were machined by hand, and on 3-and-5-axis CNC machines. All in-house machining used HSMWorks CAM software, and all 5-axis machining used MasterCAM software.

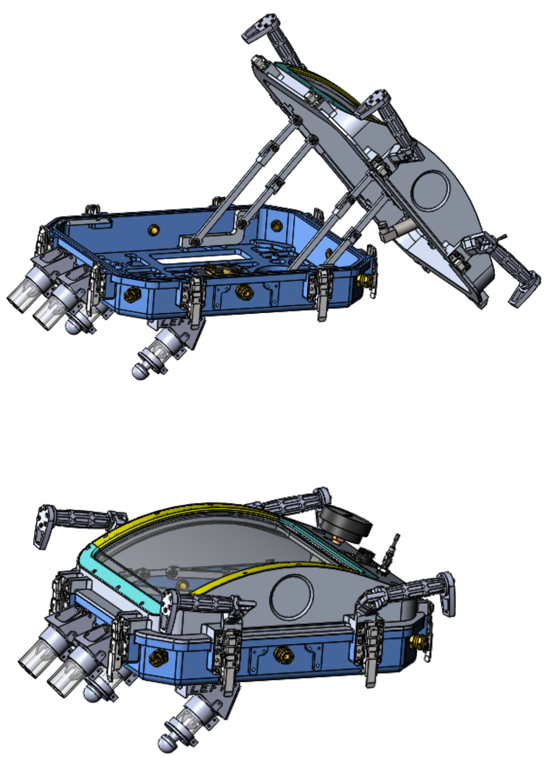

Figure 2: **The main pressure hull.** The five-bar linkage (shown open at top) keeps the top connected to the bottom when open. Some parts not shown to display hull structure.

Main pressure hull

The mechanical design of Flo is centered around a main pressure hull that all pressurized enclosures connect to. See Figure 2 for a view of the main pressure hull. The main hull assembly consists of two large machined aluminum parts and an acrylic dome. Each hull section was machined from single blocks of aluminum using 3-and-5-axis CNC milling. Both feature extensive ribbing to reduce the weight of the vehicle and tune the center of mass, allowing for optimum control and maneuverability. The dome consists of a custom-molded sheet of acrylic that seals to the upper section of the pressure hull using a 3D O-ring seal and a series of curved clamps that ensure a good seal without sacrificing aesthetics. This system ensures that the acrylic can be replaced quickly and easily in the event of damage. The clear dome enables the user to easily view the state of the entire vehicle electrical system.

The two halves of the pressure hull seal to each other using a single O-ring. To ensure that the internal hardware is easily accessible, the hull is sealed using eight strong latches. The latches were integral to the overall design of the pressure hull and the layout of external accessories, since the user needs to be able to access every latch with ease. In order to select an appropriate choice of latch, an extensive testing program was conducted where the clamping force of a variety of latches was measured. When the pressure hull is not in use, the upper half can be lifted up and then rotated back using a five-bar linkage that permanently connects it to the lower hull.

The interior of the pressure hull consists of electronics mounted to acrylic plates. The package consists of all non-sensor electronic systems, such as processing, power distribution, and electromagnetic solenoids. Boards are mounted onto one of two levels based on their function and connections to other systems, which allows for efficient use of limited internal volume. To ensure access to hardware mounted on and below the lower level, the upper level is designed to rotate out on straight linkages when the dome is opened. In addition, the package was designed to make cable routing as simple as possible, which reduces clutter and allows for easier maintenance on Flo's electrical systems.

External Accessories

The external accessories consist of all parts whose functionality necessitates being outside of the main hull. The principal systems are the cameras, the thrusters, the DVL, the hydrophones, and task-interaction hardware (markers, torpedoes and gripper). Modules are not permanently attached to the vehicle to ensure that faulty hardware can be replaced. The camera unit features two stereoscopic cameras mounted inside of hemispherical acrylic domes. The wide field of views for the cameras allow them to be used for tasks requiring horizontal and vertical computer vision. This wide field of view also allows us to eliminate the cumbersome roll-pitch search maneuver. The DVL uses the doppler effect to monitor the velocity of the vehicle, which is important to navigation. The hydrophone elements are mounted beneath the vehicle to avoid obstructions and interference from the chassis. Contrasting with last year's vehicle, Flo has no high pressure pneumatic systems, instead using a combination of electromagnetic hardware in order to actuate the markers, torpedoes and gripper. Both the torpedoes and markers are spring-loaded and fire using an electromagnet inside the pressure hull, eliminating the need for through-holes for pneumatic cabling. The single gripper on the vehicle uses 3 servo motors in order to maneuver itself into positions to access all the necessary manipulation tasks during the competition.

Thrusters

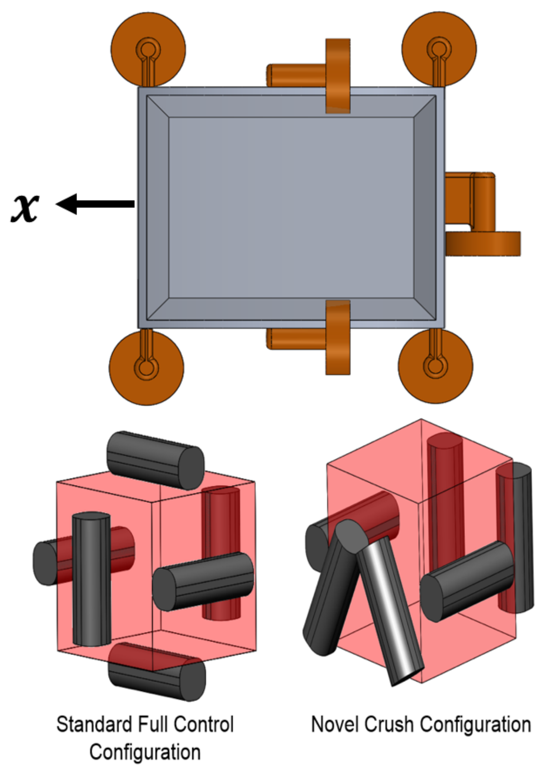

Figure 3: Thruster Configurations. The basic configuration for Flo (top) is shown along with two alternative 6-thruster layouts. Notice that there is redundancy in the vertical thrusters, which enables the vehicle to have a "thruster-out" operating mode.

Flo has inherited *Dory's* iconic thruster configuration, utilizing seven brushless VideoRay M6 brushless thrusters. Because of the linearity of newtonian forces and torques, this configuration can apply any force and torque in any direction. This allows *Flo* to actively control all six degrees of freedom. Although six axis control theoretically only requires 6 thrusters, the seven thruster configuration has excellent handling characteristics without restricting the physical layout of the vehicle (see Figure 3, bottom). Notice that the more traditional 6-thruster layout requires thrusters to be mounted on every surface of the vehicle. By contrast, our system keeps the top, bottom, and front of the vehicle available for other hardware and also adds a redundant vertical thruster.

IV. ELECTRICAL SYSTEM

Flo's electrical system powers and provides a communication platform for all sensors, actuators, and computational elements onboard in a robust and efficient manner. Electronics are implemented modularly; each subsystem is managed by a specific PCB. This minimizes the chance of an error in one part of the system causing failure within the rest of the system. The electronics package is shown in Figure 4.

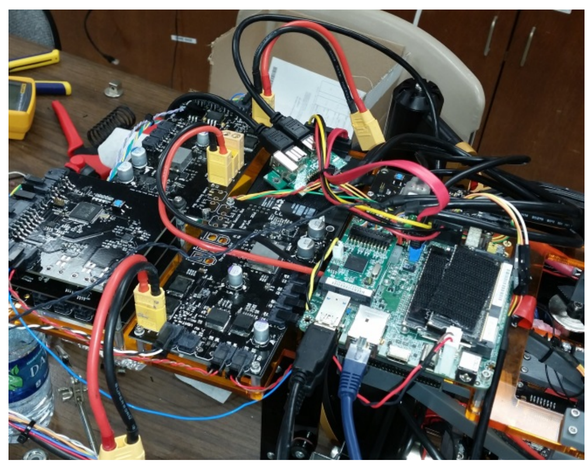

Figure 4. Electronics package outside the pressure hull.

A. Mechatronics System

The thrusters and pneumatic solenoids use a separate power system to avoid noise and interference with the rest of the electronics.

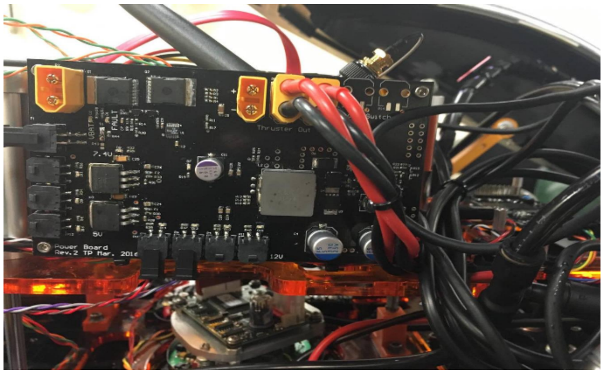

Figure 5: Motor power board. The upper left contains the protection circuitry while the rest of the board contains voltage regulators. The board is modular in that they can power the motors or the computer depending on which components are soldered on.

B. Power System

Flo uses the same circuit topology for both the motor and computer side electronics. The circuitry consists of a hot swap controller at the input which provides over/undervoltage, reverse voltage and overcurrent protection. The hot swap uses two external power MOSFETs and sense resistors to respond to fault conditions. The controller features a capacitive timer to adjust the timing that triggers a fault condition, and series sense resistors to set the overcurrent fault. The motor power and computer power boards have two different overcurrent set points. The motor power has a current limit of 40A, largely dominated by the 36A thruster current at full throttle, and the computer power board has a current limit of 20A.

During an over/undervoltage fault condition, the hot swap controller will cut power to the system by driving the gates of the external MOSFETs low. An overcurrent fault is triggered if the system draws more current than the current limit for a period of time that is greater than the timer timeout. This type of fault requires a power system reset in order to clear. The hot swap controller has an active low shutdown pin that is used to cut power to the electronics without removing the batteries. The board has a solder jumper that configures the behavior of the kill switch. The computer power board remains on until its kill switch is closed, while the motor power board will remain in shutdown until its kill switch is closed.

The power boards have 12V, 7.4V and 5V outputs. The 12V and 5V power rails are produced using efficient switching buck regulators to minimize power loss. The 12V power rail can handle a 15A output current, while the 5V rail can handle a 3A output current. The 7.4V regulator is only populated on the computer power board. Both power systems are powered by a 10 amp-hour, 36V lithium-polymer battery.

C. Microcontroller system

The microcontroller system provides low-level communication to sensors and actuators that cannot be controlled directly by the main computer. The microcontroller system consists of five boards: the core board, the sensor board, a pneumatics board, a servo board, and the hydrophone board. Each board is equipped with a STM32F407VG microcontroller as the STM32F407 provides many integrated peripherals. These boards communicate over an RS485 bus network in order to efficiently pass messages between subsystems.

The core board (Figure 6) is responsible for handling communication between the Intel NUC™ computer and the rest of the microcontroller system. Messages between the computer and the core board consist of a message ID and five bytes of data. Computer to core messages are parsed by the core board and handled by sending the data to another microcontroller board as an RS485 data packet. Data packets received by the core board are sent to the computer through a USART.

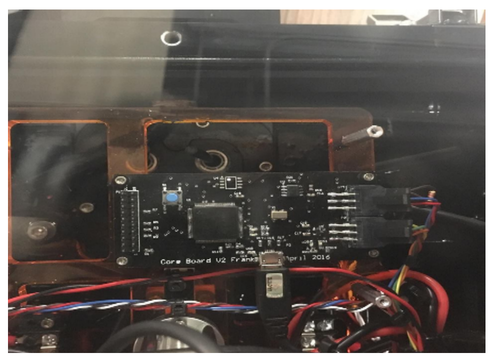

Figure 6: Core board. This standalone board handles communication between the Intel NUC™ computer and the rest of the microcontroller system.

The sensor board (Figure 7) uses an ADC to poll the Omega PX-319 depth sensor and a GPIO to poll the motor kill switch in order to provide continuous updates to the computer about the state of the vehicle. The sensor board is also used to control the WS2812B LED strip that is used to provide visual feedback about vehicle.

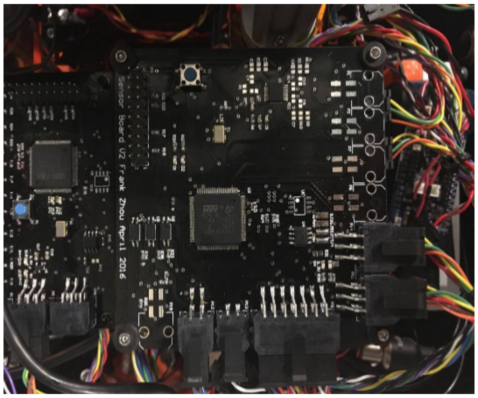

Figure 7: Sensor board. The sensor board polls the depth sensor, motor kill switch, and controls the feedback LED strips.

The servo board drives the three servo motors which control the gripper. It converts angular position information to pulse widths and asserts the PWM signals.

The pneumatics board (Figure 8) uses optoisolated GPIO signals to control four electromagnets.

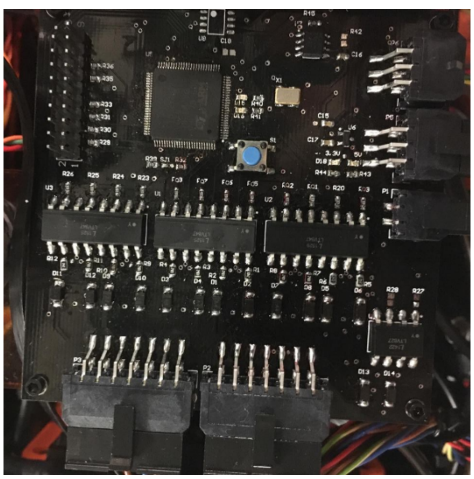

Figure 8: Pneumatics board. Optoisolated GPIO signals control pneumatic solenoids which actuate the torpedoes and markers.

D. Sensors

Flo uses a variety of sensors to gather the crucial information it needs to navigate and understand the underwater environment.

1. Depth Sensor

An Omega PX-319 pressure sensor filtered using a second order Bessel filter and processed by an ADC on a microcontroller measures the vehicle's depth to within an accuracy of 1-2 inches.

2. Attitude and Heading Reference System (AHRS)

A VectorNav VN-100 Rugged measures the vehicle's heading and attitude using a 3-axis accelerometer, gyro, and magnetometer. The magnetometer is calibrated with a custom hard-soft-iron calibration routine. Orientation measurements help *Flo* actively maintain its underwater orientation.

3. Doppler Velocity Log (DVL)

A SonTek Argonaut-M DVL measures an accurate 3D velocity vector, which is used to estimate and control the vehicle's position in the water.

4. Camera

The vehicle has two AVT GuppyPro cameras with 120° wide-angle lenses. The cameras are positioned 45° below the horizontal axis to provide full coverage of obstacles in front of and underneath the vehicle. All task accessories are positioned parallel to the cameras to maximize efficiency. The cameras are parallel to within 1°, which allows the use of advanced stereoscopic ranging of field obstacles.

5. Hydrophones

Four Teledyne TC4013 piezoelectric transducers mounted in a square pattern listen for the sonar signal. The current from the transducers is converted to a voltage across a 100 MΩ resistor. The signals then pass through a JFET input stage preamplifier with a DC servo to increase the signal amplitude. A 4th stage differential active filter is used to attenuate signals above 200 kHz to prevent aliasing. A quad channel simultaneously sampling 16-bit ADC is used to digitize the differential analog signals at 400 kHz / channel. The signals are then processed on a STMF407VG and the heading information is sent to the computer via UART. A Fourier transform is used around the pinger frequency to extract the phase information. The heading is computed using the phase differences between the four sensors.

V. SOFTWARE

Flo's autonomy relies on its robust C++ software architecture, which consists of the five processes diagrammed in Figure 9. These processes communicate with one another using a custom inter-process communication (IPC) library based on the publisher-subscriber data model, where processes can publish data to topics (illustrated as rounded rectangles) and subscribe to receive all data on certain topics. This makes the system modular and scalable throughout the lifetime of the code. The team chose to write its own IPC library instead of using an existing framework such as ROS in order to minimize the use of external black-boxes and reduce the learning curve for new members.

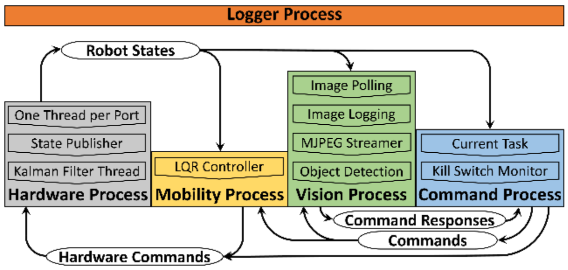

Figure 9: Software architecture. The five main processes that run on the vehicle and the interprocess communication channels are shown. Each process is illustrated by a large box, with threads inside. Each circular box represents a communication channel.

These five processes run on an Ubuntu Linux environment. While testing, we can communicate with them over Ethernet using Secure Shell (ssh).

A. State Estimation

State estimation involves determining the vehicle's orientation and position within a set of world-fixed coordinates, and it is done in the Hardware Process. Orientation is provided by the AHRS. Position can be estimated by integrating the DVL velocity measurements over time. However, the DVL only operates at 1Hz, and even then, only when there is at least one meter of water underneath it. So if we just integrate raw measurements, the update rate would be slow and the vehicle would lose control near the pool floor.

To deal with both of these issues, we implemented a custom 11-dimensional extended Kalman filter (EKF), whose state includes the vehicle's translational velocity, position, translational drag coefficients, inverse virtual mass, and net buoyant force. The prediction step uses an algebraically-linearized Newtonian dynamics model with linear drag.

A Kalman filter is well suited to this task because it can first estimate the drag coefficients of the vehicle by comparing the applied thrust to the measured velocity. Then it can use these drag coefficients to estimate the vehicle's position based on the known thrust being applied.

The output of the Kalman filter is a unified state estimate that includes the vehicle's 3-D position, orientation, translational velocity, and angular rates, all of which are periodically published to the Robot State IPC channel.

B. 6-DoF Control

The Mobility Process receives the current state and the target state from the Hardware and Commander Processes respectively. It then runs a 6-DoF controller, which decides what power level to send to each thruster in order to move the vehicle from the current state to the target state. Most RoboSub teams (and even commercial vehicles) do this using PID loops. However, on *Flo* we use a Linear Quadratic Regulator (LQR). LQR is an optimal control scheme for multidimensional linear systems. Unlike PID control, a LQR is particularly well suited to situations like ours where the systems ha modelable, but coupled dynamics. This is because LQR requires a model of the system, but once given that model, it takes full advantage of the system dynamics.

Our 18-dimensional LQR controls the six translational and angular errors, the six integrals of those errors, and the six rates of change of those errors. Since the vehicle's dynamics are non-linear, we must locally linearize them around the target state before applying LQR. The resulting linearized system, computed anew every time the target state changes, is given to the LQR solver from MIT's Drake Library

Experimentally, this controller was far superior to even our best-tuned PID control systems. In addition, tuning the cost matrices, Q and R, is borderline trivial; it took only twenty minutes to tune our LQR controller compared to PID controllers, which took many months on previous vehicles.

The final step is gracefully addressing thruster saturation. We address this by imposing a prioritization scheme. The output of the LQR controller is decomposed into four components: (1) forces required to keep the sub static, (2) other vertical forces, (3) all other torques, and (4) all other forces. Given this breakdown, the software sums up, in this order, as much as possible of each component that can be added while staying below the thrust caps.

Finally, we profiled our thrusters using a strain gauge to determine the mapping from motor power level and battery voltage to physical thrust. Inverting this relation allowed us to generate the exact thrusts requested by the LQR controller. The entire controller is diagramed in Figure 10.

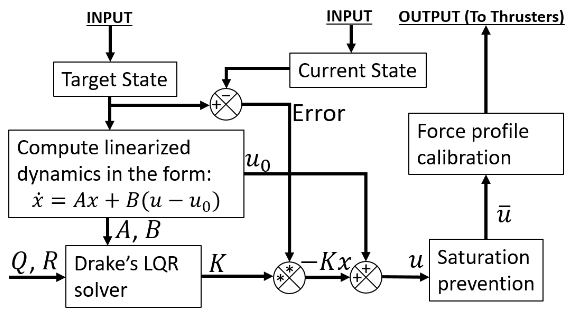

Figure 10: 6-DoF controller diagram. The dynamics are linearized about the target state to compute the LQR gain K, which is applied to the error term to compute the control output.

C. Visual Object Detection

The vision process converts images into actionable information. It synchronously reads new frames from both Allied Vision Guppy Pro cameras using libdc1394. Each frame is corrected for distortions and then fed through a custom detection pipeline. One of these pipelines is a rectangle detector (used for the torpedoes, path, marker bins, and the covers). Our novel algorithm is based on Gauss-Newton minimization and it determines the 3D orientation and position of any rectangle given the pixel positions of its corners. The remaining detectors use an ad hoc mix of edge detection, color thresholding, contour analysis, and Haar features. An example of this pipeline in shown in Figure 11:

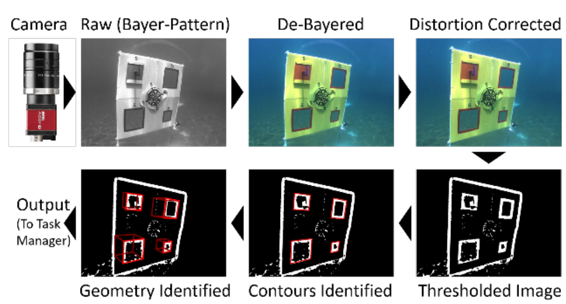

Figure 11: Typical vision processing pipeline. Applied to last year's torpedo task.

D. Navigation and Mission Planning

The Commander Process is responsible for knowing which task we are on, how to accomplish it, and how to get to the next task. It directs the other processes in order to accomplish this. Each task is encoded as a state machine, which makes development modular and the behavior more intelligent.

We address the problem of navigation by manually driving the vehicle through the course once to measure and save the coordinates of each task. Still, the system is only accurate enough to arrive within visual range of each task, at which point the vision system takes over.

E. Logging and Playback

In developing an autonomous system, we cannot prevent the vehicle from ever making a mistake; however, we make sure that we are always able to piece together what went wrong after the fact. This is made possible by the Logging Process, which records every message sent send over every IPC channel, in both a binary and a human-readable JSON format. The binary allows us to replay logged messages into the original C++ code, while JSON allows us to parse and analyze the logs in a scripting language (Figure 12).

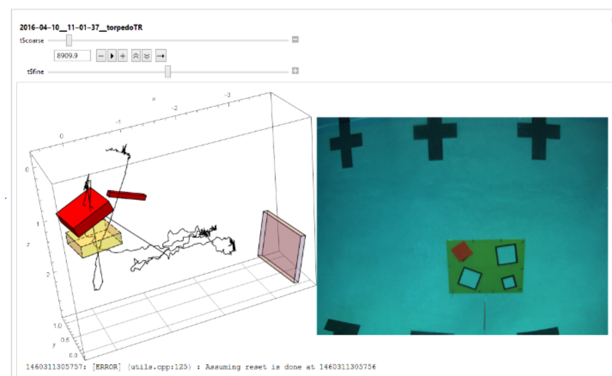

Figure 12: Logging and playback GUI. This Mathematica utility shows what the vehicle is thinking at any user-specified time, which includes the latest camera frame, the current vehicle state (red), the vehicle's target state (yellow), the vehicle's position history (black), and any object detections (pink-gray).

VI. VEHICLE TESTING

Flo first entered the water on April 6th, 2017 for controls testing and tuning. Testing and tuning is still being completed in preparation for the RoboSub competition.

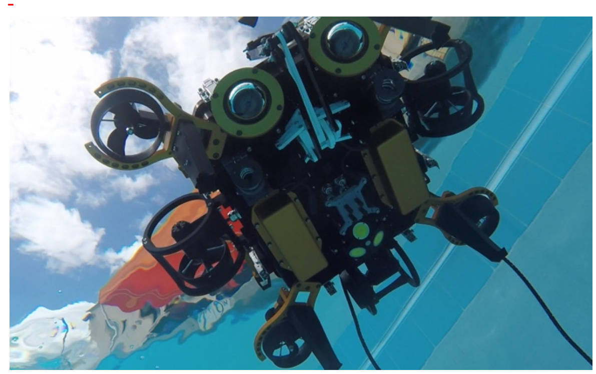

Figure 13: Pool Testing. Flo takes a close look at our camera on her way to the red buoy.

VII. CONCLUSION

Flo is an integrated, performance autonomous underwater vehicle designed to be a cost-effective tool for research and education. The vehicle can perform a number of real-world tasks efficiently, reliably, and completely autonomously such as visual inspection and underwater manipulation. Further work includes adding additional autonomy, logging, and debugging capabilities to further enhance and tune the vehicle's autonomous capabilities.

VIII. ACKNOWLEDGEMENTS

The CRT would like to thank Caltech for all its support during the past academic year, particularly within the Engineering and Applied Science division. We would like to specially thank our faculty advisor and mentor, Professor Joel Burdick, who has given us direction through our many hours of development and testing. Additional thanks goes out to Perry Radford, Jannah Maresh, Matt Burkhardt, John Van Deusen, Professor John Dabiri, Heather Morgan, Betsy Mitchell, John Leichty, Glen George, Richard Murray, Tom Mannion, and all the gym and pool staff that have helped us through numerous hours of testing. Finally, we express our immense gratitude to all of our gracious sponsors for the 2016-2017 academic year.

Ocean Level: Caltech Moore-Hufstedler Fund and Caltech Housner Fund, L3, Rhonda MacDonald, Glendale Community College, Honor, VideoRay, Rigol, Northrop Grumman
Sea Level: Impresa Aerospace, Caltech Student Investment Fund
Lake Level: Danco Anodizing, Teledyne Impulse, Marvin Engineering, Azeus Convene, Brad Foreman