Design and Implementation of an Integrated, Autonomous Underwater Vehicle: Dory

E. Fouad, T. Okamoto, F. Zhou, J. Orenstein, K. Agarwal, T. Pailevanian, J. Larson
California Institute of Technology, 1200 E California Blvd., CA 91126
June 19, 2016

***Abstract* – Caltech Robotics Team's new autonomous underwater vehicle (AUV), Dory, is designed for precise movement and high-level autonomy for the 2016 intercollegiate RoboSub competition. Designed and manufactured by over thirty undergraduates at the California Institute of Technology, Dory is the team's most sophisticated and capable vehicle yet, iterating on last year's vehicle Crush. Improvements include a Kalman filter, a more robust LQR controller, a sleek pressure hull with a curved acrylic dome, a gimbaled camera, seven thrusters to provide extra power and redundancy, two gripper mechanisms, and a hydrophone system. 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 magnetic and non-magnetic objects. The vehicle includes a custom inter-process communication library that is the foundation for its advanced autonomous capabilities. Dory represents the state-of-the-art for low cost, actively stabilized autonomous underwater vehicles.**

I. INTRODUCTION

The Caltech Robotics Team's goal is to design and implement an autonomous underwater vehicle for the purposes of research, education, and outreach. The vehicle, Dory, allows the team's student members to develop engineering, communication, and collaboration skills in an exciting and hands-on manner. Dory's design goals are to successfully complete the challenging course in the Association for Unmanned Vehicle Systems International (AUVSI) Foundation and Office of Naval Research's RoboSub Competition. The competition takes place annually in July at the Transducer Evaluation Center (TRANSDEC) pool, at Space and Naval Warfare Systems Center Pacific (SSC Pacific) in San Diego, California. The competition realistically simulates real-world underwater vehicle tasks such as underwater manipulation, navigation, and visual inspection. The RoboSub competition elements include firing a torpedo at a small target, locating a sonar pinger, manipulating small objects, and dropping a marker into a bin. All of these tasks must be completed autonomously with no human in the loop. To complete these tasks, Dory must have accurate and sensitive visual and acoustic sensors, precise guidance, control and navigation, object recognition, and object manipulation.

II. DESIGN OVERVIEW

Dory is a small autonomous underwater vehicle designed for precise and accurate movement in all six degrees of freedom. The team has learned many lessons from last year's vehicle, Crush, and integrated that knowledge into Dory. The pressure hull consists of a custom anodized aluminum shell which was machined on a five-axis CNC. An acrylic dome provides clear visibility of the electronics. A high pressure pneumatic line powers the gripper, torpedoes, and markers. A gimbaled camera allows the vehicle to effectively examine its surroundings and track objects. Seven thrusters give the vehicle all six degrees of freedom, while adding redundancy in the vertical direction.

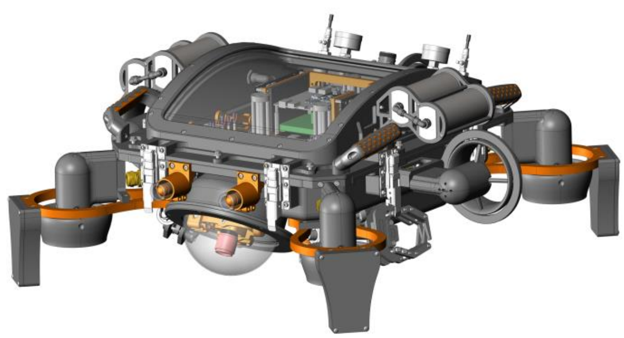

Figure 1: Computer-aided design model of Dory.

The electrical systems have also been improved by integrating all of the custom electronics onto four circuit boards, saving space and reducing clutter while still remaining modular. The electronics are mounted to a two stage platform that raises and lowers, providing access to the electronics on the lower platform. The vehicle is powered by two lithium-polymer batteries giving Dory 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. Dory uses a custom inter-process communication library to implement its vision, actuation, and autonomy systems. This year we have added a hydrophone system to amplify and process the sonar signals.

III. MECHANICAL SYSTEM

Dory's mechanical systems consist of a main pressure hull, external electrical equipment, and pneumatic hardware. Some external equipment (the camera dome 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.

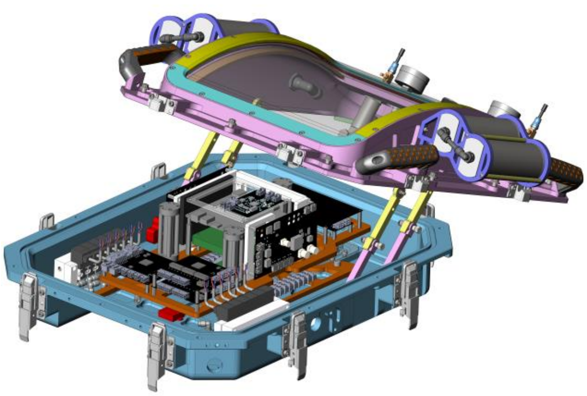

Figure 2: Main pressure hull. The four-bar linkage keeps the top connected to the bottom when open. The top electronics platform rises to provide convenient access to the boards below.

A. Design Strategy

Dory is centered around a single main pressure hull that all pressurized enclosures connect to. See figure 2 for a view of the main pressure hull. The main hull consists of 2 pieces of machined aluminum and an acrylic dome. The two halves of the pressure hull were each machined from single blocks of aluminum using 3-and-5-axis CNC milling. Both feature extensive ribbing to lower the weight of the vehicle and tune the center of mass. 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. External enclosures and a clear dome are used to reduce weight and complexity while also enabling the user to easily view the state of the internal electronics.

The two halves of the pressure hull seal to each other using a single o-ring. To ensure accessing the internal hardware is as easy as possible, the hull can be sealed using 8 strong latches. The latches were integral to the overall design of the pressure hull and 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. The same rig was also used to identify the correct spacing of the mounting holes on the pressure hull. When the pressure hull is not in use, the upper half can be lifted up using a 4-bar linkage that permanently connects it to the lower hull. Two of the bars feature linear motion, which allows the upper hull to pivot back to improve access to the interior of the 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 pneumatic solenoids. Boards are mounted onto one of two levels based on their function and connections to other systems. To ensure access to hardware mounted on and below the lower level, the upper level is designed to raise about two inches on linear slides. In addition, the package was designed to make cable routing as simple as possible, which reduces clutter.

B. External Accessories

The external accessories consist of all parts whose functionality necessitates being outside of the main hull. The principal systems are the camera, the thrusters, the DVL, the hydrophones, and pneumatics hardware. Modules are not permanently attached to the vehicle to ensure that faulty hardware can be replaced. The camera unit features a gimbaled camera mounted inside of a hemispherical acrylic dome. The gimbal has two degrees of freedom, which enables the single camera to be used for tasks requiring horizontal and vertical computer vision. The camera is mounted at the optical center of the dome to minimize distortion. Using a single gimballed camera simplifies mechanical design and enables the cumbersome roll-pitch search maneuver to be eliminated. 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. The pneumatics hardware consists of air tanks, grippers, torpedoes, and markers. There are two gripper designs: a claw and a magnetic gripper. Both mount onto a standard arm design, which can stow the unit against the hull when not in use.

C. Thrusters

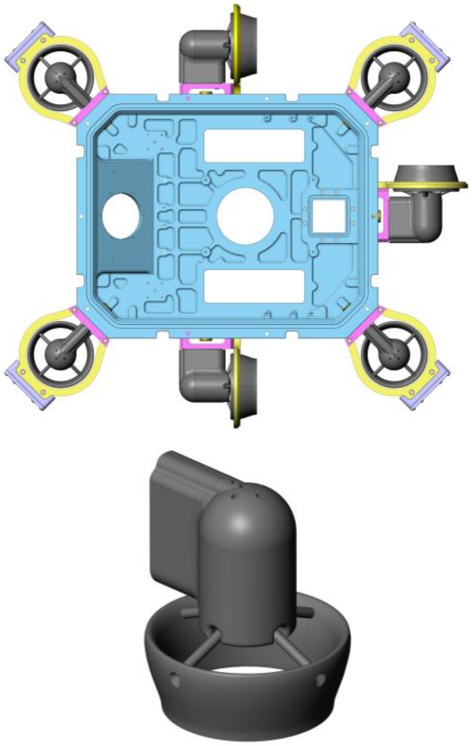

Figure 3: Thrusters and thruster configuration. The thruster configuration (top) uses seven VideoRay M5 brushless thrusters (bottom). The thruss for Dory (above)). The configuration consists of four vertical thrusters, two forward-facing thrusters, and one lateral thruster.

Dory uses a unique thruster configuration of seven brushless VideoRay M5 brushless thrusters. Due to the linearity of Newtonian forces and torques, this configuration can apply any force and torque in any direction. This allows Dory to actively control all six degrees of freedom. Although six axis control theoretically only requires six thrusters, the seven thruster configuration has good handling characteristics without restricting the physical layout of the vehicle. Locating the four vertical thrusters on the corners of the pressure hull opens up more space for accessories. In addition, it ensures that they interfere as little as possible with the flow from the planar thrusters. Note that the traditional thruster configuration (Figure 3, bottom) requires thrusters to be placed on every surface of the vehicle, which severely restricts mechanical design. In addition, using four vertical thrusters provides redundancy; the vehicle has been run with one thruster disabled.

IV. ELECTRICAL SYSTEM

Dory's electrical system provides power and communication for all sensors, computation, and actuation onboard in a robust and power efficient manner. The electrical system uses a combination of custom and off-the-shelf components to achieve the low cost and performance requirements of the system. The electronics package is depicted in Figure 4.

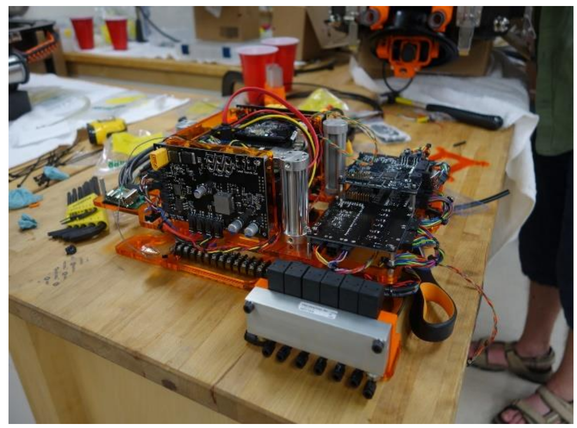

Figure 4: Electronics package outside the pressure hull.

A. Mechatronics System

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

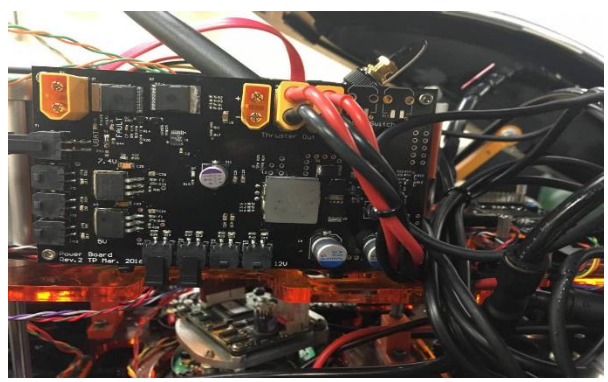

Figure 5: Motor power board. The upper left contains the protection circuitry while the rest of the board contains voltage regulators. The board is modular in that they can power the motors or the computer depending on which components are soldered on.

B. Power System

The sub 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 since it is required for the servos that control the camera gimbal. Both power systems are powered by a 10 amp-hour, 22.2V 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, two pneumatics boards, and the hydrophone board. Each board is equipped with a STM32F407VG microcontroller as the STM32F407 provides many integrated peripherals. These boards communicate over a CAN bus network in order to efficiently pass messages between from board to board.

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. Serial computer to core messages are parsed by the core board and handled by sending the data to another microcontroller board as a CAN packet. CAN packets received by the core board are sent to the computer through serial.

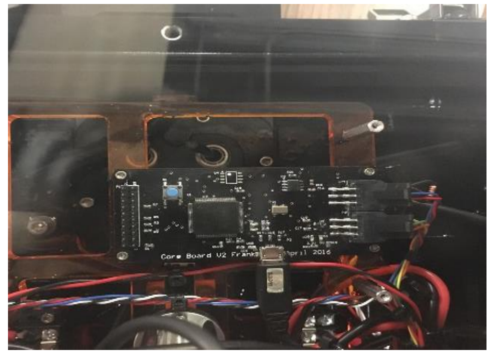

Figure 6: Core board. This standalone board handles communication between the Intel NUC™ computer and the rest of the microcontroller system.

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.

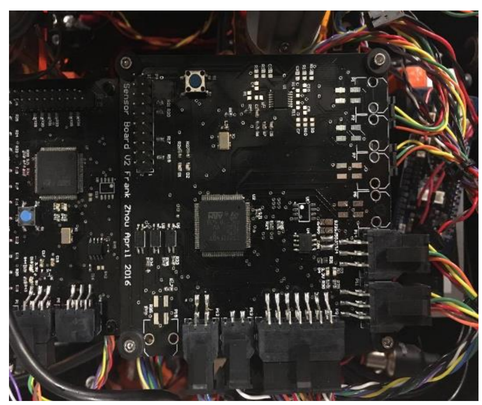

Figure 7: Sensor board. The sensor board polls the depth sensor, motor kill switch, and controls the feedback LED strips.

The pneumatics boards (Figure 8) use optoisolated GPIO signals to control pneumatic solenoids. Each pneumatics board is responsible for two normal pneumatic solenoids and four latching pneumatic solenoids.

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

D. Sensors

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

1. Camera

The vehicle has a single AVT GuppyPro camera on a gimbal to perform visual recognition of the competition elements. The GuppyPro provides a detailed view of the vehicle's surroundings and the gimbal allows the vehicle to quickly locate and identity objects.

2. Attitude and Heading Reference System (AHRS)

A VectorNav VN-100 Rugged uses 3-axis accelerometers, 3-axis gyros, and 3-axis magnetic sensors to provide accurate heading and attitude information. Dory uses the AHRS to determine and actively maintain its orientation in the water.

3. Doppler Velocity Log (DVL)

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

4. Depth Sensor

An Omega PX-319 pressure sensor filtered using a second order Bessel filter and processed by an ADC on a microcontroller is used to provide depth information to the computer.

5. Hydrophones

Four Teledyne TC4013 piezoelectric transducers are 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

Dory runs Ubuntu, the most popular Linux distribution, on an Intel NUC™ with a Core i3 processor and a 120 GB solid-state hard drive. This setup provides high performance in a small, rugged, low-power form-factor. For convenience, external connections can be made either via WiFi or Ethernet.

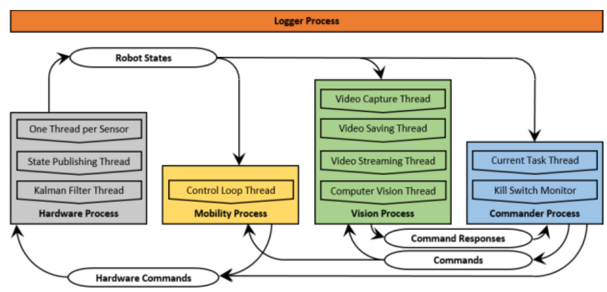

Figure 9: Software architecture. The five main processes that run on the vehicle and the interprocess communication channels are shown. Each process is illustrated by a large box, with threads inside. Each circular box represents a communication channel.

Dory's autonomy relies on its robust all-C++ software architecture, consisting of five processes diagrammed in Figure 9. This architecture is modular in order to aid development, bug isolation, testing, and logging. Each of these five processes run independently while communicating with each other using a custom, interprocess communication (IPC) library based on a publisher-subscriber model. Processes can publish data to topics (illustrated as rounded rectangles) and subscribe to receive all data published to relevant topics. 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 reducing the learning curve for new members.

A. Hardware Process

The hardware process interfaces with the low-level serial devices including the core board (depth sensor, kill switch, and solenoids), Doppler velocity logger (DVL), attitude heading reference system

(AHRS), VideoRay M5 thrusters, and Dynamixel servos.

The output is orientation and position. Orientation is provided by our AHRS, which internally fuses accelerometer, gyro, and magnetometer measurements. The magnetometer is calibrated to account for the distortion of the fields caused by the vehicle using the onboard live HIS calibration. Position is a little more problematic because the DVL only pings once per second, and if the water column is not tall enough, no data is returned. With these limitations we can neither run tight translational control loops nor do tasks below us. To deal with both of these issues, we implemented a custom 12-dimensional Extended Kalman Filter. While many teams use an Unscented Kalman Filter, because this is not analytic method, it is plagued by the curse of dimensionality and sampling issues. Instead, an Extended Kalman Filter linearizes the process around the current state, and these linearizations of quaternion and matrix expressions are all computed from analytic formulae. Using quaternions instead of Euler angles minimizes the non-linearities. The key insight of the Kalman Filter is that it incorporates the thrust outputs in order to estimate position when the DVL is unable to take a measurement or in between pings. The output of the Kalman filter is a unified state estimate including 3D position, quaternion-based orientation, and translational velocities, and angular rates, which is periodically published to the robot state topic.

B. Mobility Process

The mobility process is responsible for setting the speeds of all seven thrusters based on the current state and target state, received from the hardware and commander processes respectively. Last year we used a custom 6 DOF dynamic-inversion-based PID control loop. However, because the center of mass and center of buoyancy were not lined up perfectly our vehicle was not stable, which was a major performance bottleneck.

An engineer at NASA JPL suggested using a linear quadratic regulator (LQR), which is an optimal control scheme for systems that can be modeled as

$$x = Ax+Bu,$$

where x is the state of the system, u is a vector of thrust forces, and A and Bare matrices. An LQR controller, unlike a PID controller, computes the optimal path to minimize a quadratic cost function depending on x(t). We chose the state, x, as an 18-dimensional vector including translational and angular errors, integrals of errors, and velocities. This way our cost function could penalize absolute error, steady state error, and high speeds.

However, rotations in 3D make the vehicle dynamics nonlinear. So we had to first linearize the system around the target state. This was done using an analytic formula which returned the linearized versions of A and B for the current state using complicated matrix and quaternion expressions. The result was given to MIT's Drake Library in order to solve for the gain matrix and optimal thrusts on the fly. All the constants used in the dynamic modeling are generated from the CAD including the center of mass, center of buoyancy, inertia tensor, thruster positions, and thruster directions. To get accurate thrust control, we profiled our thrusters using a strain-gauge with high-speed software-controlled measurement.

Like our last controller, LQR can function at all pitch and roll angles. But because LQR accounts for vehicle dynamics and plans an optimal path, it can uniquely deal with non-diagonal inertia tensors and misaligned centers of mass and buoyancy. For this reason, its stability is leagues ahead of any PID-style control system. In addition, it requires no tuning: we transitioned vehicles, including changing thruster type and configuration without touching a single tuning parameter, whereas re-tuning a PID controller would have taken a month.

C. Vision Process

The vision process converts images into actionable information. It asynchronously reads new frames from an Allied Vision Guppy Pro camera using libdc1394. Useful features of this camera are the high quality optics and CCD, exposure control, white balance control, and on-demand shutter, which help us prevent reflective objects from saturating and now exactly where the vehicle and gimbals were when each frame was taken.

After the image is polled and corrected for distortions, they are fed through a custom detection algorithm. One of these is a rectangle detector (including the torpedo, path, navigation channel, marker bins, and the orange covers). Our novel algorithm is based on Gauss-Newton minimization and determines the 3D orientation and position of any rectangle given the pixel positions of its corners, even in the depth dimension and even while looking from very skew angles. The validation gate use Haar-like features to detect the three segments of the gate. Finally, the remaining detectors (buoys, handle, and doubloon) use a mixture of edge detection, color thresholding, and contour analysis.

Lastly, the vision process controls the gimbals. If nothing was detected, the gimbals begin a search phase. Once an object is detected, the gimbals lock onto it so it remains centered in the field of view as the vehicle moves. This eliminated many problem last year related to finding field elements.

D. Commander Process

The Commander process is responsible for knowing which task we are on, how to accomplish it, and how to move to the next task and directs the other three processes. It has two threads. The first runs the task while the second resets the first if the kill switch is pulled. The tasks are programed in a state-machine-like architecture.

E. Logger Process

In the rare case things go wrong, the logging process helps us debug this by recording every message on every topic, in both a binary Google ProtoBuf and JSON formats (used for playback and manual inspection respectively). One advantage of this type of logging is that it allows us to visualize any logged run via a Mathematica GUI (figure 10).

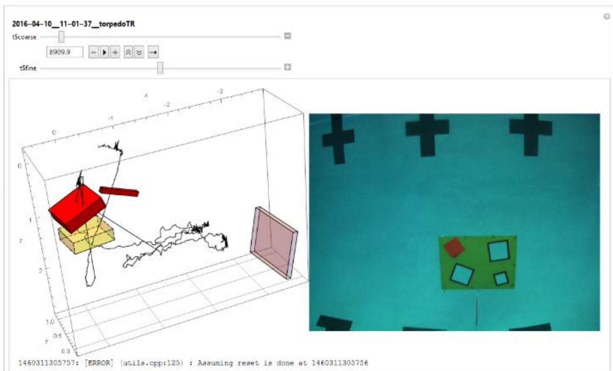

Figure 10: Log file analysis GUI. The GUI displays the vehicle's position and gimbal orientation (red), the target state (yellow), the complete vehicle path, the detected object, and the camera image.

VI. VEHICLE TESTING

From October 2015 through April 2016, pool testing was conducted on the development vehicle, a modular vehicle constructed to test the sensors and motor controllers as Dory was being designed and built. This allowed us to tune the Kalman filter and capture underwater images, and test the pneumatic system. Dory first entered the water under power on May 15th, 2016 for controls testing and tuning. Testing and tuning is still being completed in preparation for the RoboSub competition. Dory's marker deployment test is illustrated in Figure 11.

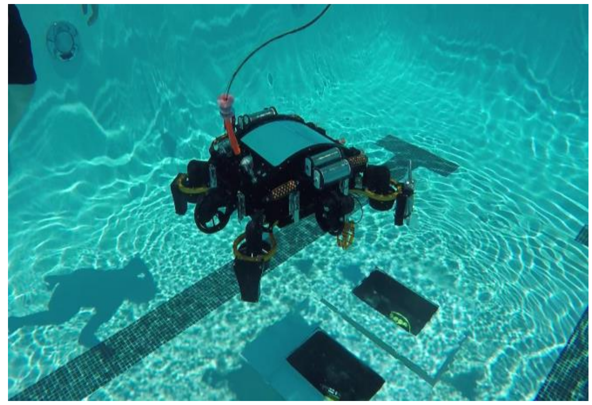

Figure 11: Pool Testing. Dory fires markers into the bins. Pool time graciously donated by the Caltech Athletics Department.

VII. CONCLUSION

Dory 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 2015-2016 academic year.

Ocean Level: Thales, Caltech Moore-Hufstedler Fund and Caltech Housner Fund, SBB Research Group, Rhonda MacDonald, Glendale Community College, Honor, VideoRay
Sea Level: Impresa Aerospace
Lake Level: Danco Anodizing, Teledyne Impulse, Marvin Engineering, Azeus Convene, Brad Foreman