Vision-Based Autonomous Surface Ship

6th Generation

Authors: Moscicki, Travis; Koenig, Justin; Wiard, William; Simmons, Joshua; Wampler, Jared; Williams, Cullen; Long, Andrew; Miranda, Mario; Davis, Harold; Mesa, Bianca;

Abstract: The Florida Atlantic University (FAU) Marine Robotics Club RoboBoat Team is looking to continue moving forward from last year's successes. The vehicle that FAU is bringing to the 2016 AUVSI RoboBoat Competition is the 6th generation of the Vision-Based Autonomous Surface Vessel (V-BASS). The 6th generation of V-BASS aims to improve upon the difficulties of previous generation while leaving in place the tried and true onboard systems. In pursuit of the team's goal of overall victory, the team has enhanced the onboard vision, overhauled the acoustics system, implemented a secondary vehicle, and tidied the communications system.

I.   Introduction

The field of Autonomous Surface Vehicles (ASVs) is continuously perused by researchers and, as a result, ASV technology is in a state of perpetual evolution. ASVs are becoming more reliant on sensors which obtain information from their surrounding environment as opposed to vehicles that rely only on waypoint navigation. The state of the art for environmental sensors is computer vision and, as such, ASVs which are guided by vision sensors, can struggle to overcome certain atmospheric lighting conditions. A vision based system also has to be able to detect desired objects such as buoys and disregard objects of no interest. Taking this into consideration, the vision system onboard the vessel must be robust and reliable. This iteration introduces a LiDAR first system, detecting objects which are normal to the

LiDAR and then cross-references these with a color based system to ensure accuracy and collect further data including shape and color when desired. The vehicle will then process these inputs and determine the necessary actions needed to complete the objective for that specific challenge.

The challenges that must be completed for this competition will illustrate the vehicles speed and power, ability to navigate a buoy field, ability to detect the correct dock-slip and perform auto-docking, ability to launch and recover a submersible-imager, ability to locate and circle a buoy based on an acoustic pinger, and ability to return to the starting dock as quickly as possible without disturbing the course.

Given the challenges that the vessel must complete, a vision based system is the only feasible choice. Waypoint navigation would not compensate for the extremely dynamic environment in which the boat is required to operate in.

The sixth generation of FAU's V-BASS is an adaptive system that can update its mission particulars given various input parameters. The system is the result of a collaboration of a multi-disciplinary group that is comprised of focused and dedicated engineering students from the Ocean and Mechanical, Electrical, Computer Engineering, and Computer Science departments here at FAU as well as the vast amount of experience at FAU when it comes to developing maritime systems.

II.   Design Description

The vessel's navigation is based upon the inputs of the LiDAR, color detection system, acoustic system, compass, and GPS. The primary source of propulsion for the vehicle is two brushless DC motors, one affixed to each hull. Maneuvering for the vehicle is achieved using differential thrust. The power source for the vessel is two banks of batteries, one bank located in each hull. The autonomous commands onboard the vessel are executed by a single board computer. The vessel is also outfitted with a secondary remotely operated vehicle. Other features include an RF Modem, Wi-Fi router, and RC Receiver for data transmission and manual operation.

A.   Design Strategy

The design strategy of the Vision-Based Autonomous Surface Ship is to innovate where we can, renovate where we should, and always push for on water testing. The V-BASS platform is a catamaran type vessel with a length overall of 58'' and a beam length of 36''. V-BASS is propelled by two (2) brushless DC Motors that are affixed to the underside of both hulls just forward of the transom. The platform is maneuvered using differential thrust which keeps the system simple and reliable. The electrical power onboard the vessel is provided by two banks of Nickel Cadmium batteries, one bank in each hull. The system has a variety of well tested sensors onboard including a 270° LIDAR, RGB cameras and a hydrophone array.

B.   Vehicle Design
I.   Hulls and Superstructure

Even though the V-BASS platform is returning for another year of competition, many innovations and redesigns have been implemented. The primary motivations for these redesigns is to reduce the overall weight of the vehicle as much as possible and accommodate a system which will launch and recover the secondary vehicle.

The Hulls used for the 2016 AUVSI RoboBoat Competition are of a catamaran configuration. The parent hull that the current vessel was designed after was the bailey hull. The vessel is designed with a watertight transverse bulkhead at the midsection of each hull. These transverse bulkheads' primarily serve as structural reinforcement for the hulls, but they also serve as a watertight barrier to the battery compartments. The decking for the hulls, which also serves as the access point to the forward sections and the battery compartments, is made of aluminum sheets and polycarbonate. Polycarbonate is used to provide visual access to a battery voltage readout on a seven-segment LCD inside each battery compartment. The motor controllers for the propulsion thrusters are located within each hull of the V-BASS platform. In order to maintain the watertight seal needed for the motor controllers and the batteries, silicon is used to seal the deck. IP-67 rated charging ports and data connectors are used to connect the battery chargers and top-side control box to the batteries and motor controllers inside each hull. This allows the hulls to be permanently sealed.
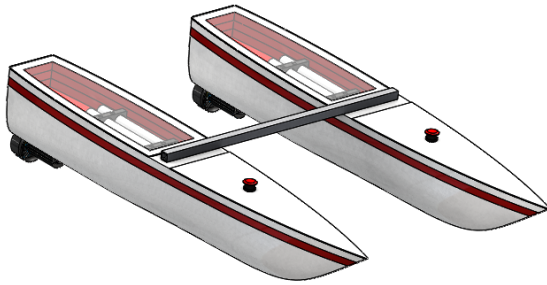
Figure 1: Hulls with Bridge

A removable superstructure was designed to allow easy mounting of V-BASS' many systems and direct access to the systems during bench testing. The superstructure consists of 3/4-inch square aluminum tubing welded together in a minimalist footprint. Components are bolted to the superstructure which is then secured to the hull-bridge with coder pins. To ensure that stress on the bridge is reduced, the removable portion of the structure rests on the top side of the hulls, making contact with four feet. This novel design affords the security of a permanently fixed superstructure with the access speed of a removable one.

II.   Propulsion System

V-BASS is driven by two Seabotix outboard thrusters, each connected to its respective hull with FAU designed and fabricated mounts. The compact yet powerful thrusters spin at 6000 rpm, supplying V-BASS with reliability and power in a lightweight package. The motors receive their signal from the main electronics box through Roboteq SDC1130 motor controllers.

III.   Electronic System
i.   Main System

The electronics system in the 5th Generation V-BASS is far more advanced than previous generations. Introducing a NVIDIA Jetson TK1 embedded computer with 192 onboard CUDA cores allows all high level and vision programming to be handled by a dedicated internal platform, as opposed to a laptop.

The issue of Wi-Fi network saturation at competition has been addressed by moving nearly all of V-BASS' ship-to-shore communication to the 900 MHz X-Tend RF Module in preference to the more common Wi-Fi. However, in order to handle passing of larger packets, such as video and images, the vehicle maintains a long range wireless router.

The Starboard and Port motor controllers are designed to take in a RC Pulse Position Modulated Signal from either an RC Receiver or from the control of the low level TS7800 computer; at the same time reporting important information back to the system, such as battery voltage, temperature inside the hulls, and current draw by the motors. This data is then logged internally on the high level controller and sent digitally to the shore base station to be displayed and monitored.
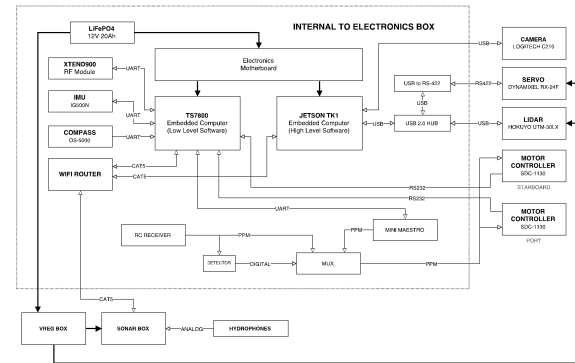

Figure 2 – Electronics Block Diagram

ii.   Power

The advances in computational power also require a higher electrical power capacity. Therefore, we decided to integrate a number of power improvements to the system. These improvements begin with a new battery system and a number of high efficiency switching regulators to boost the overall power usage efficiency of the system.

a.   Batteries

VBASS has 2 new LiFePO4 Battery packs this year. LiFePO4 Cells are best known for their superior power density to almost all mainstream battery chemistries, while at the same time remaining safer than LiPo Cells. The move from Lead Acid 12V to LiFePO4 cells gives VBASS not only a longer run time but also a 35% decrease in weight. The new batteries weigh 1.4Kg per pack, while the original batteries weighed 2.18Kg each. This gives us a weight saving of 1.56Kg or 3.44lbs. The LiFePO4 packs have a nominal voltage of 12.8V and a current capacity of 10Ah, which gives them a power capacity of 128Wh. These values give the LiFePO4 packs a power density of 91.42Wh/Kg, in comparison to the Lead Acid cells at 38.53Wh/Kg the LiFePO4 cells offer 2.43 times the power density; allowing us to run

longer with a higher power requirement than the new system demands.

b.   Power Regulation

To deliver power to V-BASS' dramatically upgraded sensor suite, a voltage regulation box was implemented to effectively accommodate the new LiFePO4 cells and its pack voltage at the end of a discharge cycle. Towards the end of the discharge cycle the pack can drop below 12V, which could potentially under supply the LiDAR and servo. Inside the regulator box is room for up to 3 SEPIC buck/boost voltage regulators. These are high efficiency switching DC to DC voltage regulators which can maintain an output of any voltage between 0.5V to 30V DC, regardless of if the input voltage is above or below the output setting. The input voltage can be anything between 3V to 15V, which means any state of charge in the discharge curve can meet the input voltage requirements. The box in which these regulators are housed and the electrical connection into and out of the box are engineered to at least an IP-67 rating to ensure the optimal water requirements of the project are met.
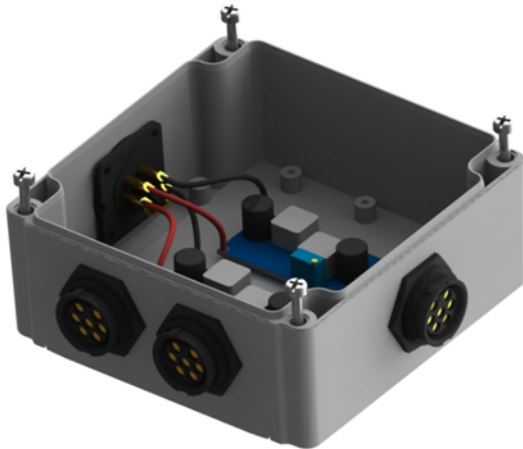

Figure 3 - Voltage Regulation Box

IV.   Hydrophone Array
i.   Ultra-Short Baseline SONAR System

The FAU USBL (Ultra-Short Baseline) SONAR system is a split-beam array that contains a horizontal and vertical baseline which provides an azimuthal and vertical bearing to the source, respectively.

The system uses an inferometric process known as a split-aperture correlator (Figure 4) to determine the analytical phase difference between two hydrophone signals on any given baseline and thus bearing to the

source. This is accomplished by exploiting the real form of the complex wave signal to produce analytical imaginary form of the signal with the use of the Hilbert Transform. This type of filtering is a discrete convolution of the digitized data with the designed filter. The Hilbert transform for this application is a linear time invariant (LTI) system with a transfer function defined as Equation (1)[2]:

$$\hat{x} = H(f) = -j \, sgn(f) \tag{1}$$

The Hilbert Transform shifts the phase of all signals with negative frequencies forward by 90° and all signals with positive frequencies back by 90°, such that the phase shift can be represented as Equation (2) [2],

$$\sigma_H(\omega) = \begin{cases} i = e^{j\pi/2}, & for \; \omega < 0 \\ 0, & for \; \omega = 0 \\ -i = e^{-\frac{j\pi}{2}}, & for \; \omega > 0 \end{cases} \tag{2}$$

The superposition of the real form and imaginary form allows one to extract the signal phase by using a phasor representation of the signal
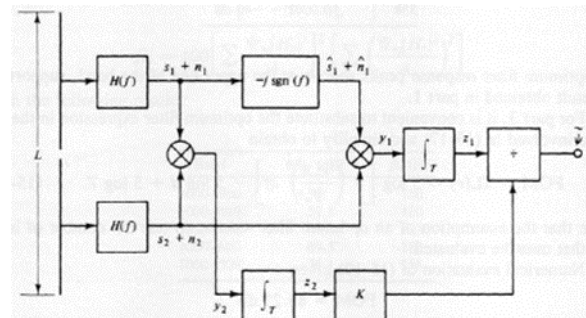

Figure 4 - Split-aperture correlator

The phase shift at the center frequency of the narrowband signal is inserted so that the correlator output is responsive to the electrical phase difference.

The electrical phase difference is proportional to the target bearing. The phase difference, $d\varphi$, between the two given signals is a component, along with the signals carrier frequency, f_c, baseline length, b, and ambient acoustic sound speed, c, to estimate the direction of arrival of the incident plane wave relative to the baseline, as shown in Equation (3).[3]

$$\theta_i = \arccos\left(\frac{\frac{d\Phi}{2\pi f_c b}}{c}\right)$$

(3)

One of the most common and sometimes most difficult issues to overcome in bearing estimation is reducing the presence of multipath or reverberated acoustic waves caused by the environmental wave guide; the sea surface, nearby side walls or structures, and the sea bottom. These unanticipated acoustic waves cause destructive interference to the received signals of the hydrophones in the form of unwanted signal noise when trying to acquire a direct path ray from a received data set.

In this particular design, the array has been configured for a forward-down looking cone of sensitivity. This allows several features: 1.) The array has a narrow cone of sensitivity to force signals in the array from a desired direction, 2.) The array depression angle allows for optimal sensitivity of the system in the forward and downward directions relative to the vehicle for when homing far away and near/on top of the source. To reduce the multipath issue, many features have been incorporated including auto-variable gain tuner, increased detection gains to force only direct rays to be heard such that multipath rays are attenuated enough to be just below the threshold to be dismissed, enhanced coarse and finely tuned detection windows to extract first wave arrival information.

The current version of the system has a data acquisition component, digital filter, signal detector and bearing estimator. A dedicated data acquisition unit (National Instruments DAQ), is used to collect voltage samples from four separate channels simultaneously at 100,000 samples/sec. Each channel is the input for a hydrophone from the USBL array. The data are band pass filtered to remove ambient noise outside of the desired received bandwidth so as to maximize the accuracy of signal detections and bearing estimations. The digital filter is an FIR band pass filter with user selectable pass-band frequencies and 96 dB of stop band attenuation. The signal detector algorithm uses peak detection component operated with blocks of 10 ms at a time. If a signal maintains a peak above a provided detection threshold, then the signal is deemed to be "detected", while signals below this threshold are discarded. The

signal detection threshold can be automatically scaled based on the ambient noise power and is recalculated every 8 seconds, thus providing a time-variable gain tuner for when the vehicle moves into new geographic areas with different ambient noise characteristics.
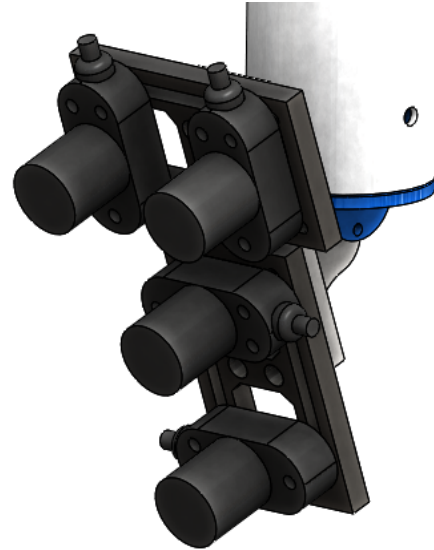

Figure 5 - USBL positioning system

Figure 5 shows the current system assemblies. The mechanical packaging included restraining the NI DAQ system and TCS (Intel i5) processing board, as well as development of an in-house power regulation and signal interface board. This board is particularly designed to provide power to the necessary electrical components and harness the input signals from the hydrophone bus.

V.   Layered Software Architecture
i.   High Level Planner

A major focus of this year's project is to implement a high level system capable of multiple communication protocols and processes, robust data logging, obtaining feedback from auxiliary devices and the sensor suite, as well as modular mission planning. The extensive open source framework of ROS provides the perfect tool for developing such a

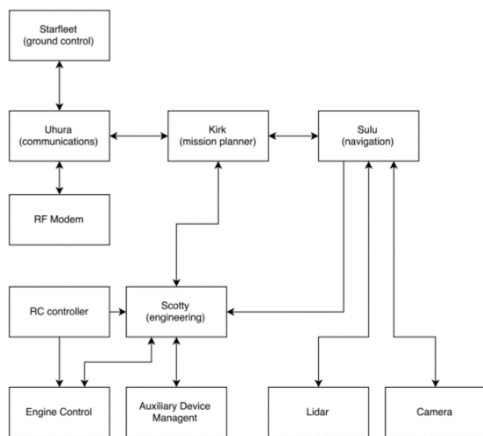system. The logical flow of information is as follows:

Figure 6 – High Level Planner

a.   Uhura

The communications officer is a ROS node that has both input and output capabilities. This output process wraps the heartbeat and task specific information in JSON format and publishes over RF or WiFi to the ground station. The input process takes in all run-time information from the technical director's network and publishes the appropriate data on ROS topics for Kirk to adjust to accordingly.

b.   Sulu

Our navigation node contains all video and LiDAR programming. The information that Sulu passes to Kirk is mission dependent.

c.   Scotty

The vehicles engineering node provides Kirk with feedback containing temperature and power information from the motor controllers and handles whether the boat should respond to RC or autonomy commands.

d.   Kirk

Our fearless mission planner takes a completely modular approach to implementing missions. On previous vehicles, the mission was contained in a single script. While developing, and even at competition, we encountered problems with troubleshooting due to difficulties pinpointing where problems lie. By separating each objective into its own discrete mission, we are able to isolate and independently refine each mission.

Kirk is designed to communicate with the officers to ensure that all systems are active and we are successfully logging and transmitting data prior to the start of each mission. To navigate from one obstacle to the next, Kirk plots a series of GPS waypoints with a desired heading upon arrival. Once the resting GPS point and heading are achieved, Kirk begins the objective specific mission loop.
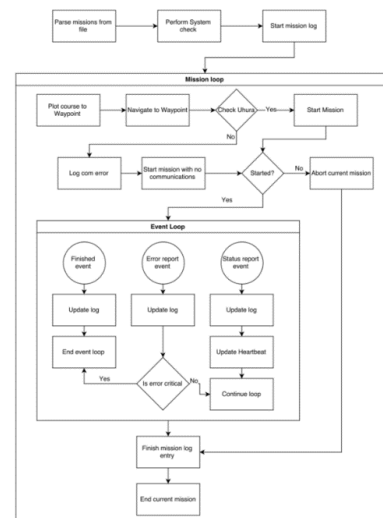

Figure 7 – Kirk's Structure

VI.   Vision System
i.   Owl Eye

One of the most meaningful integrations to this year's vehicle is a fully redesigned optics system dubbed "Owl Eye." Owl Eye is a 3 component platform containing a Hokuyo UTM-30LX LiDAR, Logitech c210 webcam, and RX-24 Dynamixel servo. The housing for these components is built with a mixture of 3-D printed PLA and water-jetted aluminum components. This hybrid structure helps to keep weight as low as possible while maintaining its structural integrity.
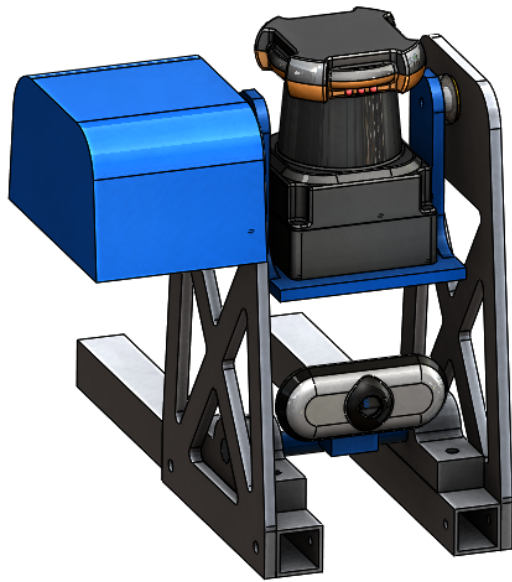
Figure 8 – Owl Eye

ii. Computer Vision

This year's RoboBoat challenge places a heavy emphasize on detecting a wide range of colored objects, multicolored shapes, and hexadecimal figures in a sometimes saturated, dynamically lit environment. To address these myriad challenges, Team Owltonomous has implemented a multifaced computer vision suite designed using OpenCV, C++, ROS, and the Jetson TK1 embedded computer. The first stage of detection for all challenges is a pre-processing routine that uses Grey-Edge normalization for color constancy. The filtered image is then ran through a secondary operation specific to the current challenge at hand.

iii. Preprocessing Routine

The preprocessor for V-BASS' vision system performs a Grey-Edge normalization followed by a smoothing routine.

A major challenge in robust outdoor color detection is that the Blue-Green-Red (BGR) values of a color greatly change under different lighting conditions. While OpenCV color spaces such Hue-Saturation-Value (HSV) and Luma-Chroma (YCrCb) attempt to address this issue, experience has shown that these filters leave something to be desired, often requiring frequent calibration. To overcome this issue, the team introduced an instance of Grey-Edge normalization.

Based on the Grey-World Hypothesis, the Grey-Edge hypothesis states that the average of the reflectance differences in a scene is achromatic. [5] This assumption allows for an estimate of the background illumination in a scene by taking the average gradient of each color channel:

$$ke = \frac{\int |f_x(x)| dx}{\int dx} \qquad (4)$$

In order to implement this filter in code, we stepped through each pixel of the 3 channels in our image matrix with a [1, 0, 1] kernel in the x and y directions, taking the magnitude at each point. Each pixel is then multiplied by this value, divided the average color intensity of each channel.

We find that the use of this filter, while not impervious to dramatic changes in lighting conditions, allows for a much wider range of acceptable lighting conditions while using the same RGB thresholding values.

A smoothing filter, in addition to normalization, is advantageous for computer vision tasks in an effort to reduce back-scatter and make binary thresholding easier. After testing a sample filter with different blurring operations, the median filter implemented with the OpenCV function medianBlur() was chosen as the best candidate because of its low computational footprint and success in cancelling noise. The function replaces each pixel in the image with the average of all its neighbor pixels. We performed this function after our normalization to ensure the best possible estimation of background color.

iv. Blob Detection

The speed gate, buoy field and acoustic portion of the competition all require V-BASS to detect different colored buoys in the water. To achieve this, we introduced a class aptly named "BUOY" that houses member functions to perform the necessary operations.

The ROI method was introduced to increase the speed of thresholding and allows the user to select a region from inside of the image stream. The values for each channel in the region of interest are averaged and then used to either directly set the minimum and maximum values for thresholding, or as a basis for trackbar thresholding. trackbars are used to refine the

values gathered from the region of interest, and is used as a final stage of thresholding.

Once the appropriate values for the channels have been obtained the binary thresholder reads them in and converts our HSV image into black and white. This image is then eroded and dialed to further reduce false positive readings.

Once the captured frame is converted to a binary image, the frame is sent into our detection function, trackObject(). This process, calls on OpenCV data types and functions to find all contours in the target binary image and return the on screen x and y coordinates by determining the respective moments.

Figure 9 - Camera Frame after Grey-Edge normalization and HSV transfer

Figure 10 - Green Buoy Thresholding

Figure 11 - Object detection (red thresholding also performed)

     v.   LiDAR

For this project, we make use of the UTM-30LX scanning laser rangefinder by Hokuyo Automatic Co. The UTM-30 is a planar scanning laser range finder, thus its range and bearing measurements are restricted to the plane of rotation of the device's laser emitter. In order to generate 3D point clouds, it is necessary to rotate the plane of the laser by mounting it to a tilting platform. To provide platform actuation, we will be using a Dynamixel RX-24F robotic servo. The RX-24f was chosen for its compact size, as well as its ability to provide accurate and timely feedback of the servo horn's angular displacement. This information will be necessary to transform the two dimensional measurements into three dimensional spatial coordinates.

In developing this application, extensive use has been made of Robot Operating System (ROS). ROS is a Linux based open source collaborative framework for developing robot software. ROS greatly simplifies implementation of inter-process communication using predefined message types and a publish/subscribe model. ROS also provides open source packages that handle low level I/O for both the Dynamixel servo and the Hokuyo sensor.

Figure 12 depicts a diagram of the flow of data in the system. The one directional arrows correspond to named ROS topics. Predefined message types are shared between nodes on topics using a publish/subscribe model. For bi-directional communication, ROS provides a service / client implementation. A service node receives input parameters from the client node, performs some service or computation, and returns the output to the client. In this implementation, the URG Node performs low level I/O for the UTM-30LX sensor and publishes the sensor output to the /scan topic.

The sensor data is published in the form of a predefined message type. Each individual message contains the range, bearing and intensity data from one complete evolution of the sensor as well as a time stamp and coordinate frame ID.

The Dynamixel Manager Node handles low level I/O for the servo, publishes servo state messages on the/tilt_controller/state topic, and listens for position commands on the /tilt_controller/command topic.

The Coordinate Frame Transform broadcaster reads the current position of the servo shaft via the /tilt_controller/state topic and uses this information to determine the pose of the LiDAR sensor relative to the parent coordinate frame. The resulting frame transform is then published on the /tf topic. The Laser Assembler Service subscribes to the /scan and /tf channel, projects the scans into Cartesian space, and then transforms them into the parent frame. The resulting point clouds are stored in a rolling buffer until the service is called. Clients that call the Laser Assembler must first pass two parameters: a begin time and an end time. The Laser Assembler then examines the time stamps of all the Point Clouds in the rolling buffer and combines those clouds which fall within the begin/end time into a single Point Cloud which is then returned to the client.

The Master Node publishes servo goal position commands on the /tilt_controller/command topic. It also places calls to the Laser Assembler service such that all the scans from an individual

forward or backward sweep of the servo are assembled into a single Point Cloud. Servo commands and assembly calls are timed to produce a new Point Cloud twice per second. Finally, the returned Point Clouds are published on the /laser_sweeps channel for visualization using Rviz.
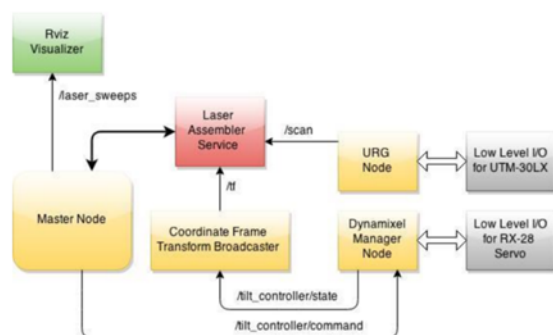

Figure 12 – LIDAR Flow Chart

VII.   Secondary Vehicle
i.   The Vehicle

The secondary vehicle was designed with a focus on simplicity and size. The vehicle consists of two side panels that were cut from a sheet of .25" thick high-density polyethylene (HDPE) and one bottom plate made of Delrin. The combination of these two plastics makes the submersible slightly negatively buoyant and a small foam float is attached to compensate for this. One possibility being discussed is the feasibility of having the vehicle be 3D printed as a single part. Powered by one vertical and one horizontal thruster, the vehicle has a single, upward facing USB camera. The thrusters used are created from two bilge pump motors commonly used for the MATE competition. Control and power are handled through a tether attached to VBASS, eliminating the need for a pressure housing on the submersible. The vehicle is designed to simply move down, then travel forwards and take a picture before being retrieved.

ii.   The Launch and Recovery System

To retrieve the secondary vehicle, a launch and recovery system (LARS) was made to winch up a connecting tether. Using a depleted 3D printer filament spool as the pulley, the mechanism pulls on a reinforced Ethernet cable. Generic Ethernet cables have a low tensile strength and can break under load so a reinforcement cable was added to increase the strength when load is applied.

The spool spins on two separate 3D printed polylactic acid (PLA) axle parts leaving a gap in the middle for the Ethernet cable to continue through one of the parts out to an electronics box. An issue came up where this would cause the wire to twist on itself in the middle of the spool as the mechanism turned. This was solved by using a slip ring press fitted in one of the axles to allow the Ethernet cable on the spool to turn freely while having a static Ethernet cable out.

Also made of 3D printed PLA parts, two side plates sandwich the spool and axle sub assembly to the boat. Both plates allow the axles to spin freely through the use of non-corrosive ceramic bearings. The non-slip ring side attaches a motor in direct drive to the axle to power the LARS. This system is fastened together using the same 6-32 screw and brass heat pressed inserts. A cover for the motor and

the tight press fits of the axles on the spool make the electronics of the system water resistant.
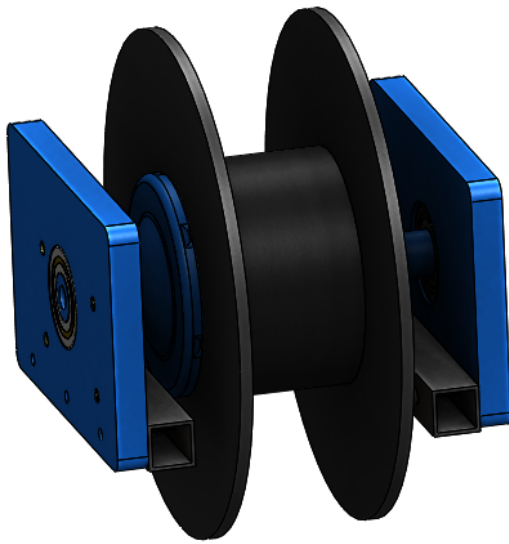
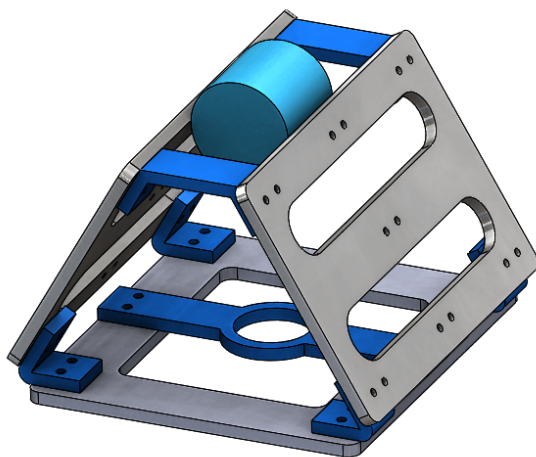Figure 13 – Spool and bearing housings of launch and recovery system (LARS)

Figure 14 – Minisub

C. Experimental Results
I.   Optics

A key component when approaching a computer vision challenge is selecting an appropriate color space. Common options include red green blue (rgb), hue saturation value (HSV), chromatic space (YCrCb), grayscale, grey world, and grey edge.

An in depth conversation about each paradigm is outside the scope of this paper, but each color space has its own intrinsic properties and alleged benefits when it comes to filtering out ambient illumination.

The color spaces of most interest are grey world and grey edge, and thus will be covered briefly.

The main difference between these two spaces is that grey world assumes the average color across an image is an estimate at the background illumination while grey edge assumes the gradient of an image holds the estimation. Both of these algorithms will scale each pixel in each frame by this factor in an attempt at color constancy. To compare these spaces to each other, as well as build in OpenCV color spaces (RGB, HSV, YCrCb, Greyscale), a program was developed that detects a buoy based solely on the output of a LiDAR.

This program will detect any and all buoys inside of the camera frame based on the method of moments and return the center of the object. A 20x20 pixel region of interest (ROI) is then outline around the centroid and transformed into the desired color spaces. The average value in this box is assumed to be the same color and is written to a text document. These values are then read by a Matlab program which plots each array vs time as well as finding the standard deviation.

The principle behind this experiment is to find the color space with the lowest deviation with respect to different colored objects, different lighting conditions, and different angles of attach between the boat, the buoy, and the sun.

Unfortunately, experimental data is not available at the time of this paper. However, data will be presented at the Marine Robotics Club's presentation at the 2016 RoboBoat Competition later this year.

D.   Acknowledgements

E.   References

[1] Austin, Thomas. The Application of Spread Spectrum Signaling Techniques to Underwater Acoustic Navigation. Woods Hole, MA. Print.

[2] Burdic, William. Underwater Acoustic System Analysis. 2nd ed. Los Altos Hills, California: Peninsula Publishing, 2002. 361-380. Print.

[3] M. Miranda, P.-P. Beaujean, E. An, M. Dhanak, "Homing an Unmanned Underwater Vehicle Equipped with a DUSBL to an Unmanned Surface Platform: A Feasibility Study", Proc. of the MTS/IEEE Oceans'2013, San Diego, CA, Sept. 2013, pp. 1-10.

[4] M. Miranda, "Mobile Docking of Remus-100 Equipped with USBL-APS To An Unmanned Surface Vehicle: A Performance Feasibility Study", Master's Thesis, Florida Atlantic University, Boca Raton, FL, May 2014.

[5] J. van de Weijer, Th. Gevers, "Color Constancy Base on the Grey Edge Hypothesis", Image Processing (ICIP), 2005 IEEE International Conference on, On page(s): 722-5, September 2005.

F.   Appendix-Outreach Activities

This season the Marine Robotics Club at FAU began working with the AD Henderson Middle School and FAU High School. Each Friday, several students would go to the school and assist in the development of four teams for the MATE competition: one high school team and three middle school teams. These volunteers were responsible for building the obstacles, guiding the students through the design process, and assisting the students in implementing their designs. The Marine Robotics Club, with help from FAU High, hosted a Regional Sea Perch Competition for middle and high school teams. The club hosted free bi-weekly seminars at the college level on many topics including 3D printing, OpenCV, and ROS. The club continued its bi-weekly meetings at Ramblewood Middle School to assist with their developing STEM program. In total, the club recorded over 120 volunteer hours.

This season also saw the introduction the Marine Robotics Club Seminar Series. Over the course of the school year team members shared skills and knowledge with other team members, the greater FAU student body, and in fact the entire internet through a dedicated youtube page (Owltube). Owltube has showcased over a dozen seminars and has amassed over 3000 minutes of view time. Topics covered included control systems, GUI building, 3d priting, and OpenCV.