VILLANOVA UNIVERSITY
College of Engineering

The Villanova University Autonomous Surface Vehicle Team
8th Annual AUVSI Foundation's International RoboBoat Competition

Journal Paper

June 15, 2015

1. Abstract

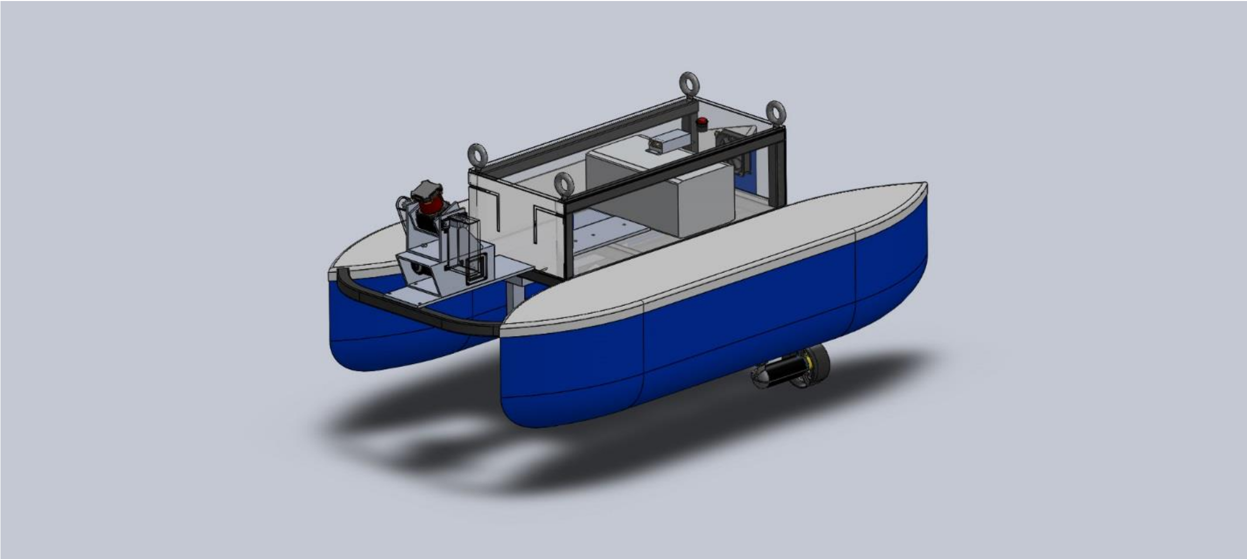

Figure 1: SolidWorks model of *SeaCat*.

SeaCat (Figure 1) is a twin-hulled catamaran-style autonomous surface vehicle designed and built entirely by Villanova University students. Making its fourth appearance in the eighth annual AUVSI Foundation's International RoboBoat Competition, *SeaCat* has proven to be a stable and adaptable platform upon which modifications can be made to meet the challenge requirements from year to year. This year's tasks are obstacle avoidance, automated docking, acoustic beacon positioning, and an interoperability challenge.

For the 2015 RoboBoat Competition, the Villanova University ASV Team looked to advance the system's hardware and software with the lessons learned from last year's RoboBoat and Maritime RobotX competitions. Some critical and significant additions to the boat include: a heat dissipation system to eliminate thermal concerns from years past, a deployable boom with an adjustable hydrophone baseline for better acoustic sensing, and a rotational invariance feature to the template matching algorithm in order to accomplish the interoperability task. Also, the vision assembly was changed to a belt system, its stepper motor was moved within the housing, and the potentiometer was switched from a three-turn to a signal turn variant. In combination with newly enhanced image fusion between the LIDAR and video camera, *SeaCat*'s vision system is now more robust, precise, and compact overall.

2. Introduction

Autonomous vehicle technology is a rapidly growing interdisciplinary field due to its great potential for advancing industries in both the military and civilian sectors. With the ability to operate independently of human input, autonomous vehicles offer unique solutions to a variety of problems facing society.

The purpose of the AUVSI Foundation's International RoboBoat Competition is to promote interest and research in this field and to help students recognize the real-world applications of autonomous systems in the context of surface vehicles. *SeaCat* is the Villanova University ASV Team's entry into this competition.

3. Physical Design Overview

SeaCat is a twin-hulled catamaran style vessel [1]. The two pontoon hulls are made from custom hand-laid fiberglass and have removable top covers which permit the placement of components within the pontoons. This results in a low center of mass which allows for the effective decoupling of yaw and roll motions of the boat. In combination, the twin hulled structure and low center of mass provide increased stability to the platform. Additionally, PVC tubing and Plexiglass were used for the construction of the frame, significantly reducing the weight of the vessel. This year's new hardware additions include a heat sink and thermally conductive silicon interface pad, and an actuated boom, detailed in Sections 4.3 and 4.4 respectively. A new vessel is currently being constructed for future autonomous surface vehicle competitions and will feature new fiberglass pontoons, more advanced thrusters, and a frame which will support a waterproof case for electronics. Complete waterproofing is the chief objective for *SeaCat*'s redesign.

4. Hardware

4.1 Real-Time Processor

SeaCat's main processor is a Speedgoat SN2415 real-time target machine. A real-time target machine is favorable for autonomous applications because of its ability to sample sensor data and output decisions rapidly rather than building up a queue of tasks. In context of the RoboBoat competition, it is imperative that state decisions are made quickly in order to effectively accomplish the desired tasks.

4.2 Integrated Vision System

SeaCat's LIDAR system was introduced in 2014 [2] to add depth to the data received from its video camera. A main focus for RoboBoat 2015 has been to improve this system in order to obtain more accurate and consistent measurements.

SeaCat must fuse the two images in order to effectively make vision-based decisions. To do this, first the measurements from the LIDAR are converted into three dimensional Cartesian coordinates. These coordinates are referenced to the center of the axis of rotation of the LIDAR, and must then be transformed to the origin and coordinate system used for the camera model (as seen in Figure 2). The X and Y measurements with reference to the camera are then normalized with respect to their Z components, and are then used in a pinhole camera model with radial lens distortion to identify the pixel values that the depth values are spatially paired with.

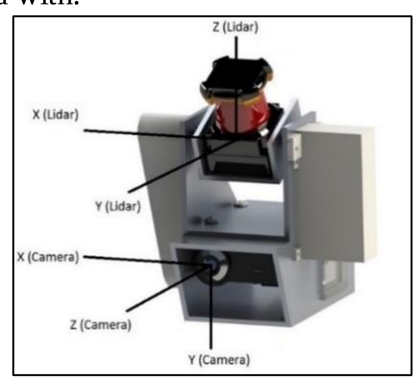

Figure 2: Vision System Assembly.

As there are many more data points gathered by the video camera within the field of view than the LIDAR, these transformed depth points are comparably very sparse. Morphological operations are used to connect neighboring sparse depth points to find regions in the video frame to match for color data. The image pairing can be seen in Figure 3.

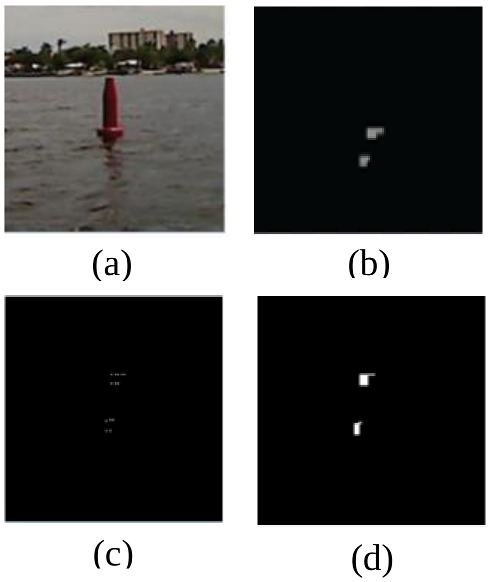

Figure 3: Example images showing video fusion. (a) The original RGB image, (b) raw LIDAR data, (c) LIDAR data plotted with respect to pixel indices, (d) joined LIDAR data creating a region of interest in the pixel frame.

4.3 Design for Thermal Management

In past years, *SeaCat* has had problems with the Speedgoat overheating and causing other components of the system to fail. To solve this issue, a heat sink made of unpolished aluminum 6061 and a 3M Thermally Conductive Silicon Interface Pad 5519 were installed on the boat.

Selected for its low contact and thermal resistances, the aluminum was manufactured into a T-shape and positioned with the top of the "T" lying beneath the Speedgoat and the length of it extending to the bow of the boat. To minimize air gaps and ultimately reduce resistance to conduction at the interface, the silicon pad was placed between the Speedgoat and the sink. This particular thermally conductive interface pad was chosen because of its high compressibility and thermal conduction value of 4.1 W/m-K.

To study the temperature profile of the boat and the effectiveness of the heat sink, tests using a solar emulator were conducted before and after the installation of the heat dissipation equipment. With the solar emulator on and no heat sink on the boat, there was a continuous increase in temperature until the test had to be terminated in order to prevent damage to the Speedgoat. This test was then repeated with the heat sink installed (Figure 4) and results showed a steady state temperature at 60 °C. A DAQ 9213 and type E thermocouple wires applied at five locations on the Speedgoat (Figure 5) were used to gather the temperature data. From this data, it was concluded that the heat sink and thermally conductive silicon interface pad effectively dissipate the heat generated by the real-time processor, and will alleviate the thermal problems.

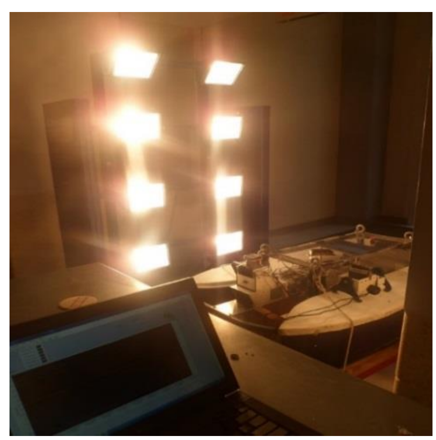

Figure 4: Experimental set up, solar emmulator on with heat sink installed.

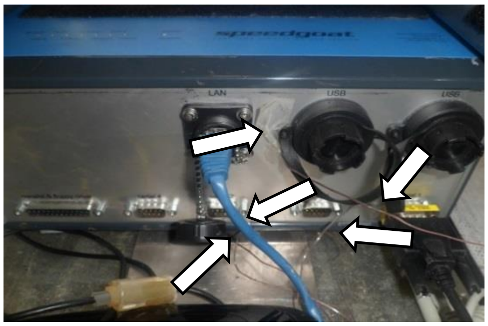

Figure 5: Thermocouples attached to the Speedgoat.

4.4 Boom System

For this year's competition, an actuated boom system was implemented on *SeaCat* with the purpose of having the hydrophones be selectively deployable for the acoustics task. This design is optimal because it (a) allows the hydrophones to be submerged deeper in the water than the bottom of the pontoons, (b) does not affect the dynamics of the boat during normal navigation, and (c) reduces the risk of damage to the hydrophones when not in use because they are nested between the pontoons when the boom is retracted.

The system's main components are a two inch stroke length Firgelli Automations electric linear actuator, which is managed by the Speedgoat via voltage relays, and a custom built hydrophone assembly with variable position arrays. This assembly is composed of a plate fastened to the boom and two 3-D printed mounts which secure the hydrophones (Figure 6). This design is critical because it optimizes the signal from the underwater acoustic beacon, allowing for better distance and bearing estimates.

(a) (b)

Figure 6: Hydrophone assembly, (a) front view and (b) side view.

4.5 Hydrophones and Amplifiers

For the acoustic beacon positioning task, *SeaCat* is equipped with two Teledyne Reson TC4013 omnidirectional hydrophones mounted to the boat's boom. The hydrophones are oriented outward with an adjustable ultra-short baseline set to .75 times the wavelength of the signal to be detected [3]. Each hydrophone is connected to a Teledyne Reson EC6081 voltage preamplifier located in the boat's right pontoon. The preamplifiers have adjustable high and low pass filters and gain in order to help isolate and amplify the desired signal which will come from the underwater acoustic beacon. An Arduino Due is used to receive and interpret relative data from the hydrophones and then communicate it to the Speedgoat. This process is explained in detail in Section 5.4.

4.6 Unmanned Aerial Vehicle (UAV)

One application of autonomous vehicles that is currently being explored is the integration of multiple vehicles to collaboratively accomplish a mission. To highlight this challenge, this year's RoboBoat Competition features an interoperability task. While this task was modified to eliminate the physical UAV portion, advancements were still made with *SeaCat* towards accomplishing it.

A UDI U818A UFO quadcopter was used as a prototype to plan for the interoperability task. For the competition, the team intended to use a FPV-Factory Splash Drone AUTO because of its waterproof design and features conducive to autonomy, including GPS navigation and an open source platform.

The team found the most challenging aspect of this task to be the successful landing of the UAV on the boat. A proposed solution to this is an actuated spooling system on board *SeaCat*. Upon completion of the challenge, the drone would enter a hovering mode and be retracted until an on board limit switch is flipped, indicating that the quadcopter has landed.

The image recognition software for this task is a modified version of the image recognition software currently used for the automated docking task but with more templates and a rotational invariance feature. This is discussed in Section 5.5.

5. Software

SeaCat has three processors running simultaneously for distributed control and parallel development. The primary processor for higher logic is the Speedgoat, which is programmed using Simulink's Real-Time Workshop, and compiled to the target using XPC Target Turnkey. The vision system is run by a PandaBoard, which is also programmed using Simulink's Real-Time Workshop, but compiled through the RLNX package developed at Villanova University. The acoustic system is run by an Arduino Due, which is programmed using Arduino's IDE. These three systems communicate through RS-232 serial connections, using JSON protocol for messaging.

5.1 Finite State Machine (FSM)

The Finite State Machine (FSM) is a structure (Figure 7) which allows the system to approach challenges in a logical manner [1]. The FSM organizes challenges of the competition into individual states which are enabled upon completion of the previous state. Completion is achieved by either a preset timeout in case of hardware or software malfunction, or a successful run. The FSM also extends into the individual tasks where a lower level structure governs the execution of subtasks and routines. This allows for a modular programming process where each action can be written and debugged on a distinct and unique basis.

5.2 Color Recognition with Bayesian Filtering

Perception of the color of an object is directly related to the lighting conditions that object is subject to. In outdoor applications, the lighting conditions are constantly changing; for instance, – as a cloud blows past, the lighting conditions change, causing the sensor to read different color information. Even orientation makes a large difference, as sunlight is highly directional. An example of the variance in perception of color can be seen in Figure 8. Both of these images were taken within ten minutes of each other, facing the buoys from opposite sides.

Most common methods used for color analysis use a heuristic approach to color detection, meaning threshold values are chosen, and if the color values for an object fall within those thresholds, the object is considered to be that color. This is often used in conjunction with a color-space conversion

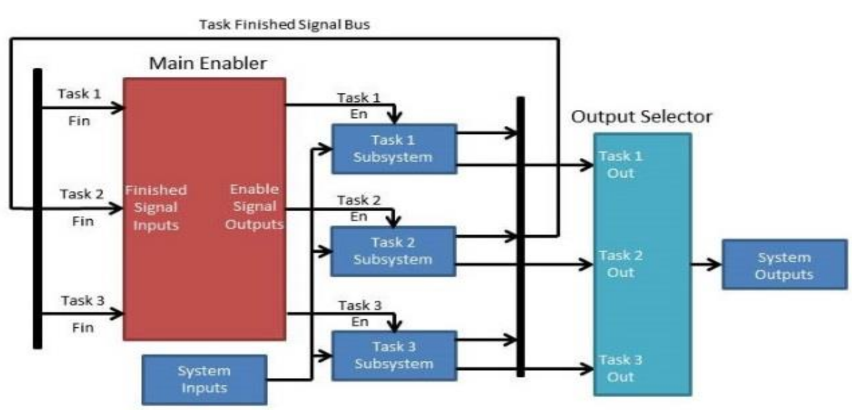

Figure 7: Structure of the Finite State Machine.

[4] [5], such as the hue/saturation/value (HSV) color space. By converting the RGB pixel values to HSV, the color data does not vary as much with slight variations in lighting conditions, but it is not completely decoupled from the lighting. Notice in Figure 9, how a red, white, and green buoy were each observed and a histogram was taken of the hue data over the region it occupied. Each of the three different colored buoys had a spike at the right end of the hue spectrum region, and are hence not discriminated.

Figure 8: Variance in outdoor lighting conditions.

Rather than approaching the problem by trying to find the best way to determine color in an image, a probabilistic approach of recursive color state estimation was developed to gain certainty in color estimation over several samples [6]. This is done by using the LIDAR to first identify objects in the image frame, and determining the mean HSV values over that region. The color estimation algorithm was developed from Bayes theorem. Bayes theorem states that the probability of a quantity (the state, in this case the color of the object) can be inferred and updated from data (the sensor measurement, in this case the HSV values) through use of a generative model (the probability distribution of sensor readings over known states).

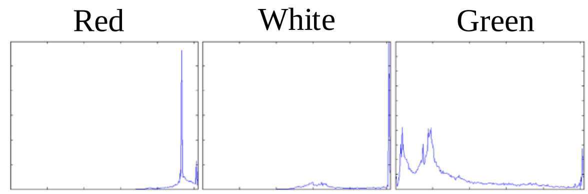

Figure 9: Normalized hue distributions over red, white, and green buoys.

For this particular application, there are a known, finite number of colors the buoys can be, and three different color channels. The recursive model is the distribution of HSV values from previously collected data of the known buoy colors. Figure 10 shows an example run of the implemented algorithm converging upon the correct identification of a red buoy over time.

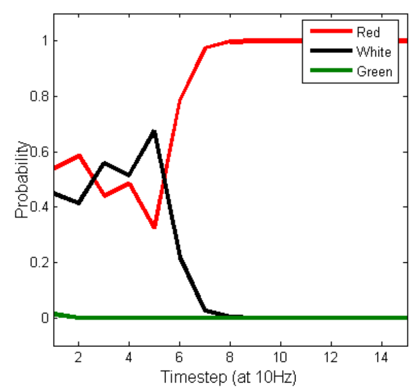

Figure 10: Recursive color state estimation of a red buoy under sunny

5.3 Simultaneous Localization and Mapping

The development of a list of landmarks that have been observed by the vessel over time is mapping, and the relative location of the vessel within a map is localization. Without both, the vessel can only react to sensor readings from the current time step, having no memory of what had been previously observed. Localization and mapping is done on board the vessel by leveraging past and current sensor readings together using a Simultaneous Localization and Mapping (SLAM) algorithm [6]. This allows the vessel to do active path planning, as well as refine estimations of vehicle position and orientation.

SLAM works by combining the relative measurements from the vision system with a positional estimate formed by a dynamic model of the vessel to determine the location of a landmark [7]. The landmarks currently identified are then compared with previously identified landmarks to determine whether or not the new landmark has previously been observed, or if a landmark should have been observed (as in, the previous landmark was within the sensor's field of view with relation to the position and orientation of the vessel) but was not. If a previously identified landmark is not observed when it is within the field of view over several samples, the landmark is removed (considered a poor initial measurement). As sensor readings are imperfect, no two time steps will measure the position of the landmark precisely the same. To account for this, the comparison of the landmarks is done through a method known as clustering [8], where geometric distances between the landmarks are compared, and a heuristic value is chosen to determine if the two landmarks are considered close enough to be identified as the same object.

The use of a heuristic method can lead to data association issues, where two separate objects may be identified as the same landmark, or the same object over two different time steps may be identified as different landmarks. To address this, the implementation on the vessel is a development of the FastSLAM method [9]. FastSLAM deals with the data association uncertainty through the use of a particle filter. The particle filter creates several isolated map estimates (for *SeaCat*, 50 particles) each with a small amount of Gaussian random noise added to the sensor readings. These maps are compared at the end of each time step to validate the clustering. This also makes the algorithm more robust to spurious measurements.

To summarize, the method works by first reading in the sensor data for position and orientation of the vessel, as well as objects being observed by the vision system. This data is then used for each particle. The next step is prediction, where the dynamic model is used to estimate the motion of the vessel using the last position, the control inputs to the thrust system, and applying a movement variance to account for noise. The positional estimate from the dynamic model and the sensor readings are then leveraged together in the measurement update, where they are clustered together with old landmark data. During this measurement update, old landmarks are checked for validity to see if they are continuing to be observed if they are within the field of view of the vision system. This develops the map, and each particle's map is then compared to check for data association issues. A flowchart showing the SLAM approach for a single scan be seen in Figure 11.

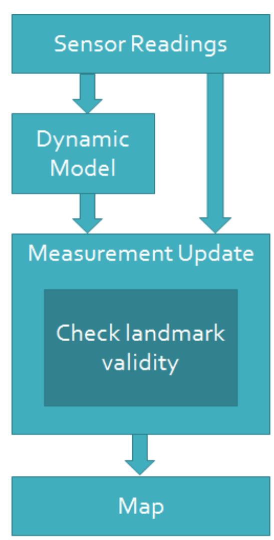

Figure 11: Flowchart of SLAM functionality.

5.4 Acoustic Localization

The hydrophones are spaced apart to optimize the incoming frequency from the underwater acoustic beacon. The signal output from the hydrophones, however, is minimal and must be amplified in order for an Arduino Due to read it. The Arduino Due pulls the peak values from each signal and calculates the difference between them. Then, because sound intensity from a point source obeys the inverse square law, the inverse square law is used to determine the distances from the hydrophones to the point source, or in this context, the underwater acoustic beacon. From these values and the known hydrophone baseline, the location of the transponder can be estimated.

Approximations using the inverse square law are most accurate when there are no disturbances, i.e. reflections or noise. To mimic ideal conditions, the Arduino Due also runs a finite impulse response filter (FIR) and a Kaiser window function. In combination, the input signal is normalized, producing more accurate heading and bearing estimates.

5.5 Rotational Invariant Template Matching

For the interoperability task, an algorithm is used to identify a seven-segment display with a hexadecimal character representation. The picture returned to *SeaCat* is first filtered for the expected range in the hue/saturation/value (HSV) channels. Then, the resulting binary image is evaluated by image analysis tools to orient the pattern's primary axis vertically. In order to determine the resulting picture's orientation and whether it requires rotation or not, the picture is scaled and analyzed for the location of an orientation reference point. Finally, the adjusted binary image is compared pixel by pixel to templates in order to identify the correct hexadecimal character. Figure 12 is an example of a test image before and after the algorithm was applied.

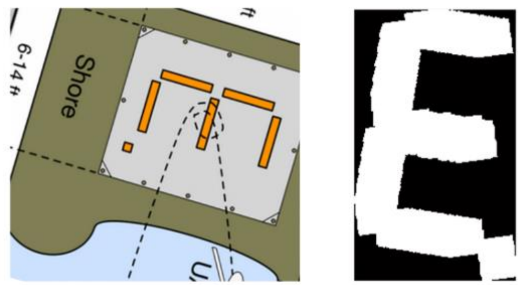

(a) (b)

Figure 12: Image before (a) and after (b) filtering and rotation.

6. Conclusion

Building upon the experience of previous years, *SeaCat* has been altered for this year's competition with the modification of the vision system and the addition of both an automated boom and a rotation invariance feature to the template matching algorithm. *SeaCat* has undergone countless hours of both laboratory and in-water testing and has proven to be a stable and dependable platform for the 2015 competition. The Villanova University ASV Team is proud of what it has accomplished and is confident going into the competition.

7. References

[1] E. Zhu, D. DeGaetano, G. Cheng, M. Benson, G. Jenkins, F. Ferrese and C. Nataraj, "An Experimental Autonomous Surface Vehicle With Vision-Based Navigation, Obstacle Avoidance & SLAM," Naval Engineers Journal, vol. 126, no. 4, pp. 173-177, December 2014.

[2] J. Anderson, A. Lebbad, C. Nataraj and G. Clayton, "An Intergrated LIDAR/Video Vision System for Autonomous Surface Vehicles," Naval Engineers Journal, vol. 126, no. 4, pp. 169-172, December 2014.

[3] F. Jaffres, T. Austin, B. Allen, R. Stokey and C. Von Alt, "Ultra Short BaseLine Acoustic Receiver/Processor," in Oceans - Europe 2005, 2005.

[4] J. Akita, "Real-time Color Detection System using Custom LSI for High-Speed Machine Vision," Lecture Notes in Computer Science, vol. 1856, pp. 128-135, 2000.

[5] Y. HuaYun, "A method of color detection base on xy-H," Applied Mechanics and Materials, vol. 263, pp. 525-528, 2013.

[6] A. Lebbad, N. DeLeo, C. Matada Shivananda, F. Ferrese and C. Nataraj, "Probablistic Vision Based Guidance & SLAM For Autonomous Surface Vehicles," Naval Engineers Journal, vol. 126, no. 4, pp. 95-99, December 2014.

[7] S. Thrun, W. Burgard and D. Fox, Probabilistic Robotics, Cambridge: The MIT Press, 2006.

[8] R. Clason, "Finding Clusters: An Application of the Distance Concept," The Mathematics Teacher, 1990.

[9] M. Montemerlo, S. Thrun, D. Koller and B. Wegbreit, "FastSLAM: A Factored Solution to the Simultaneous Localization and Mapping Problem," in Eighteenth Nation Conference on Artificial Intelligence, 2002.

8. Acknowledgements

Funding from Naval Engineering Education Center is gratefully acknowledged. In addition, NREIP support is deeply appreciated for DeAngelo, Hsu, Jenkins, and Leighow. The Villanova University ASV Team would like to acknowledge hardware donation from Speedgoat GmBH, hardware accommodations from The Modal Shop, Inc., software support from MathWorks, Inc., and unfettered access to testing facilities at VideoRay, LLC. Additional technical assistance was provided by J. Wesley Anderson, Zack Nowosad, Chris Townsend, and Edward Zhu. Dr. Garett Clayton's direction for LIDAR algorithms is very much appreciated.

9. Team Members

Students

Anderson Lebbad	MSME ’15
Adam DeAngelo	BSME ’16
Gus Jenkins	BSME ’16
Michael Trigo	BEME ’16
Christopher Hsu	BSME ’17
Scott Leighow	BSME ’17
Samantha Tropeano	BSME ’17

Advisor

Dr. C. Nataraj