Design of the USYD RowBot WAM-V System for the 2018 RobotX Challenge

Alexander Norris, Jackson Shields, John Sumskas

***Abstract*—The University of Sydney participated in the international Maritime RobotX challenge for the first time in December 2018. With a new and young team, USYD RowBot's objective was to participate in and complete as many of the competition tasks as possible. USYD RowBot's work with the Wave Adaptive Modular Vessel (WAM-V) has delivered a viable system for 2018 and laid the foundations for future competitions and research opportunities.**

I. Introduction

A surge of popularity in Unmanned Surface Vessels (USV) as well as development and interests in robotic and autonomous systems has inevitability led to desire to develop autonomous USVs (ASVs). This has been made easier with advancements in both hardware and software, such as improved Commercial Off-The-Shelf (COTS) sensors and comprehensive software architectures for automation (namely ROS).
The biennial Maritime RobotX Challenge serves to stimulate ASV technology development and practise by providing a space for competition and innovation in a complex and collaborative environment. USYD RowBot is a team from the University of Sydney which has entered the 2018 RobotX challenge for the purpose of learning, collaborating and contributing to the state of the art in marine autonomy.
USYD RowBot is a small team, comprised mostly of undergraduate engineering students and with the assistance of technical staff from the Australian Centre for Field Robotics (ACFR) and faculty staff from the University of Sydney.

II. Design Strategy

A. Development Goals

USYD RowBot's chief aim is to deliver a robust, capable ASV which can operate with many mission profiles in a wide variety of environments. From this perspective, the 2018 RobotX challenge is a critical step forward in this development process. This challenge is the first trial of USYD's WAM-V platform and it presents many of the tasks that a typical AMS could be expected to complete in future operations.
With this view of the competition, USYD RowBot aims to demonstrate an ASV which can acquit itself well in surface perception, navigation and propulsion while also being readily extensible to other missions in the future.

B. Team Capabilities and Limitations

USYD RowBot was founded in late 2017, with development of the WAM-V system beginning in the second quarter of 2018. At this time, ACFR staff had implemented an operational propulsion and localization solution on the vessel for their own research applications. It was the the student team's task to augment this existing platform for use in the RobotX Challenge. None of the participating students had experience with comparable systems and none had been involved in projects of such scale, extent and complexity. Furthermore, the team has operated from a relatively small budget pending additional sponsorship.
To compensate for these limitations, the team decided to pursue a rapid development process which made the most use of existing hardware and software solutions. This allowed the team to deliver an operational vessel without needing to spread its resources too thinly.

C. Task Selection

The tasks presented at the 2018 RobotX Challenge are complex, varied and thus require a complex and multi-functional vessel to complete them. In view of its limitations in development time and manpower, USYD RowBot elected to limit its mission scope for the 2018 challenge so that it could deliver effective solutions to the majority of tasks.
The team therefore decided to focus its resources on effective and reliable surface perception and navigation at the cost of subsurface perception and actuation. This philosophy is grounded in the understanding that observing and responding to the surface environment is the most critical aspect of an ASV's functionality.
With the scope of development limited to surface perception and response, the development of underwater acoustic perception and AUVs was postponed until later competitions. The development of a launcher system for detect and deliver was similarly postponed until time becomes available.
It is USYD RowBot's belief that the omission of these less essential capabilities was necessary so that the team could concentrate its resources on solving the problems it could reasonably contend within its small development period.

III. Vehicle Design

A. Power and Propulsion

The WAM-V's propulsion is delivered by 2 Torqueedo electric outboard thrusters in a differential arrangement. This provides the vessel with abundant thrust and heading control which is adequate for the team's requirements. These thrusters are connected to the WAM-V pontoons via custom flotation pods in place of the factory-standard pods delivered with the vessel. The new pods are more suitable for fine control of the vessel and allow for a mounting method which is better adapted for future ACFR usage of the vessel.

To power the AMS, 2 banks of lithium-ion batteries are attached to the WAM-V. These batteries provide for a mission duration in excess of 4 hours with high-power use.

B. Localization

Accurate localization is critical for the functioning of the WAM-V in all environments and usage cases, so the vessel is equipped with a Novatel INS system with differential GPS. This system provides a complete navigational solution to the vessel with accurate heading and a low-drift position fix.

Although SLAM could technically be utilized while on the RobotX 2018 challenge course, future applications for the WAM-V involve offshore activities beyond visual range of reference objects along with operations in built-up areas where GPS data is unreliable and large objects could move unexpectedly. The INS has been chosen for these situations and so serves as a dependable benchmark for the vessel's perception and navigation systems regardless of the availability of stationary landmarks or a consistent GPS signal.

C. Sensors

The team's focus on robust surface perception led to the selection of simple and proven technologies to accurately monitor the surrounding environment.

1) LIDAR: Scanning LIDAR is a critical component of the WAM-V's perception system. After an assessment of the team's requirements and means, a single RoboSense RS16 unit was selected to provide the key object detection functionality for the platform. This unit has proved to be capable, reliable and simple to integrate with the ROS system. The LIDAR is mounted on top of the vessel's perception tower, above the camera array.

2) Imaging: Visual perception is achieved with an array of 3 Logitech C270 USB webcams on a centralized perception tower. The cameras each have a $140°$ field of view and are mounted on $60°$ offsets to deliver a $260°$ total field of view with 10 degrees of overlap between the frames of adjacent devices.

The team considered the use of machine vision cameras, but these were abandoned due to cost and time constraints. The selected cameras are not expected to serve on the vessel for an extended period of time, and so their relative fragility without a protective housing was considered to be a viable trade-off.

3) Sensor Mounting: A custom mounting solution, termed the perception tower, was designed and 3D printed to ensure that the vessel's primary sensors can be positioned correctly. The perception tower is designed to be modular and easy to disassemble for transport and for the substitution of equipment. It it critical that the cameras can be readily swapped in case of a failure. The camera shells have been sealed by the team with liquid sealant, but the possibility of lens fogging under humidity could not be fully eliminated. To mitigate this issue, the cameras can be rapidly swapped out with minimal tooling in the unlikely event of a failure.

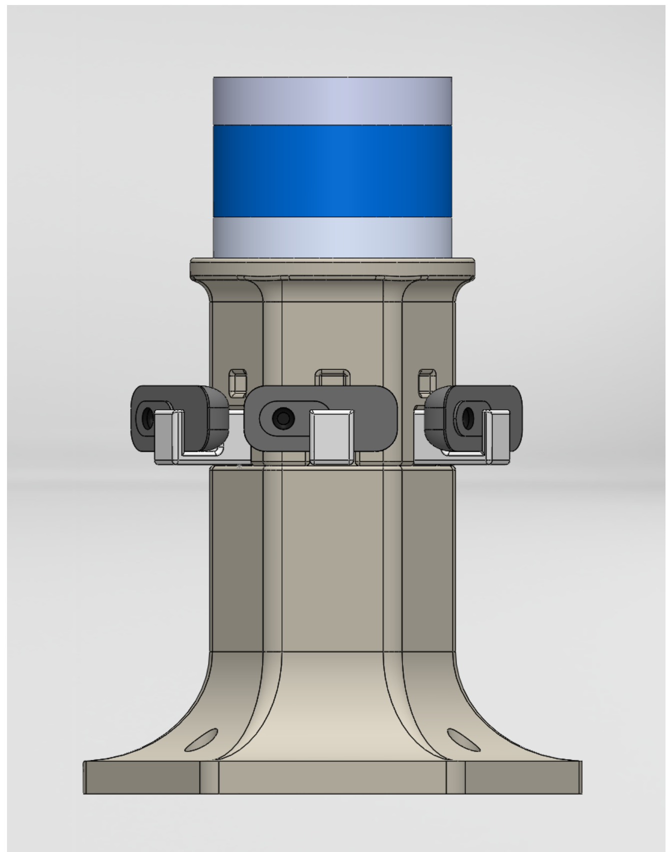

Fig. 1. CAD model of the perception tower, mounting a RoboSense LIDAR and 3 USB webcams

4) Future Sensor Hardware: There are many sensors which USYD RowBot is planning on integrating onto the WAM-V for use in both future competitions and other research and development projects. These sensors stand to both improve and extend the capabilities of the existing perception suite.

Plans have been proposed by the team to supplement the existing RoboSense RS16 LIDAR with new solid-state LIDAR units. These new devices would be installed to cover some of the blind spots in the current LIDAR system and to increase the probability of detecting small objects in close proximity to the vessel.

The imaging suite is planned to be comprehensively upgraded in the future, with a focus on delivering a rugged, all-around vision capability. Various machine vision cameras are being assessed for use as primary cameras for the system and sealed enclosures are planned to allow for webcams to cover side

and aft vision angles. Stereo vision systems, smart machine vision cameras and multi-spectral imaging options have also been considered for future implementations.
The WAM-V available to the team has already been modified by ACFR for use with various acoustic devices, namely multi-beam sonar and USBL units. Time and budget constraints prevented the team from integrating a hydrophone array onto the vessel for the 2018 competition, but this capability is a key part of development plans for future competitions.
In addition to the aforementioned sensor package upgrades, various new sensor technologies are being assessed for extending the capabilities of the WAM-V for research applications. The team is considering the addition of marine radar units and meteorological equipment to supplement the perception and data-collection capabilities of the vessel.

Fig. 2. Computer and Electronics Enclosure

D. Computation

The computational requirements of the current vessel are sufficiently low that a single industrial PC is capable of processing the collected information from the sensor arrays and control the vessel. Future development, particularly the addition of acoustic arrays or more cameras may require this PC to be supplemented with additional computers for pre-processing some or all of the collected data.
All electronic equipment which is not IP-rated is stored along with the computer in a central instruments box. Cooling of this box is handled via conduction to the outside air, which is sufficient for the small number of efficient devices contained within.

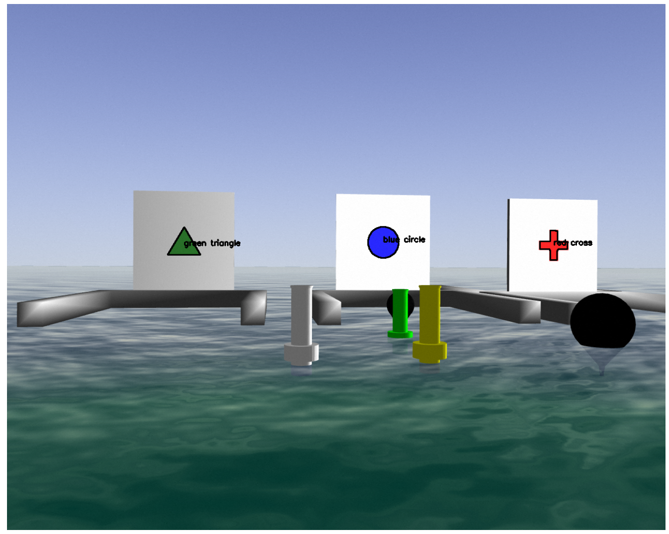

Fig. 3. Perception of Shapes in Simulation

E. Software

1) Operational Architecture: Our software is running on a combination of Robot operating System (ROS) and Lightweight communications and marshalling (LCM) [3]. Much of the low level control is executed with LCM to leverage various existing ACFR solutions while higher level function to carried out within ROS. We chose to use ROS because of the various tools already available to us, such as a WAM-V simulation package and many visualization and mission planning tools.
Our platform is running Debian Stretch with ROS Melodic and LCM installed.

2) Simulation: Simulation was a key part of the software development as testing opportunities were sparse. We used a developed simulation designed for the Virtual Maritime RobotX Challenge (VRMC) [1]. This simulation allowed us to prototype and test all of our high level functionality without the actual WAM-V. This include camera vision and lidar as well as our GPS and IMU.

F. Perception

The perception solution developed for the AMS is implemented with with a combination of existing ROS and OpenCV tools with custon-developed python scripts for fusion and operation.

1) LIDAR Processing: The LIDAR data is received as a point-cloud ROS message, courtesy of ROS drivers provided by RoboSense for use with the RS16 unit. This point cloud is then pre-processed with the pointcloud_to_laserscan ROS package to produce a 2D laser scan with points represented by bearing and range values.
The laser scan is then converted into an occupancy grid format, with cells storing the probability of occupancy, ranging from

0 to 100. Unknown cells are represented with -1. Successive hits on a particular cell increase that cell's probability value up to its maximum. The probability of cells measured to be empty is decreased over successive scans. Cells are considered empty when an object is detected behind the cell or no returns are observed from that cell when it is within the LIDAR's field of view and detection range envelope.

This mapping methodology was chosen for robust handling of spurious data, such as false positives or negatives as well as moving objects in the water. These false positives and negatives are anticipated in a marine environment due to random reflection and scattering of laser pulses off ripples in the water. By requiring multiple scans for LIDAR returns to significantly influence the occupancy grid, interference from these transient noise sources is mitigated.

2) Image Processing:

a) Data Collection: Image data is continually collected by ROS from all 3 webcams at their frame rate (30 fps) and stored as OpenCV image messages. These images are processed only on request from the various image processing functions in order to reduce the computation requirements for image-based perception. All images are converted into the HSV colour space before processing them for content.

b) Buoys: A mask-based approach is taken for buoy identification. Masks are generated for all possible colours of buoys which the vessel is expected to encounter during a run on the course. Contours are then generated from the image within a given Region of Interest (RoI). The contours are filtered by size and shape to classify the object within the RoI.

c) Docking Symbols: To begin symbol identification, a white mask is applied to the image to find the white background upon which the docking symbol is painted. Contours are then used to determine the shape of the symbol within the white background box. the shape is identified by counting the number of sides of the resultant contour. For circles, the contour is matched to an ellipse, with the areas compared to determine their similarity. To determine the colour of the shape, the internal colour of the shape is averaged and compared to reference colour values.

d) Light Buoy: Identifying the displayed colour of a light buoy is completed in a similar way to determining the colour of docking symbols. A white mask is used to detect the buoy's frame, and the shape and colour of the LED panel is confirmed by contours.

To determine the colour sequence being displayed by the Light Buoy is handled by a progressively-filled buffer of size 3. Each change in panel colour triggers a new entry in the buffer, and an extended blank period (i.e. end of sequence or missed readings) resets the buffer to ensure the code is collected properly.

In order to ensure that subsequent colour detections are correct, after the first reading only a restricted RoI around the panel's position is processed for colour identification.

3) Sensor Fusion: Sensor fusion begins from the occupancy grid, where K-means clustering with a distance threshold is used to group occupied cells into representations of objects. These objects are filtered by size to provide an initial guess of their classification. Small objects (radius < 1.2 meters) are classified as buoys, larger objects are classified as docks, and very large clusters are classified as land. A list of object ID, position, type and confidence values is maintained by an object server for future reference.

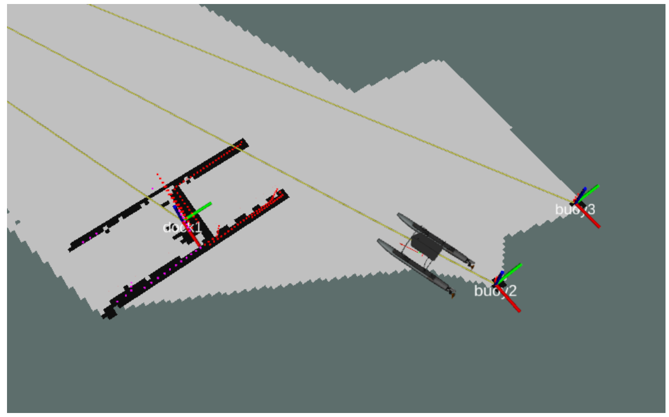

Fig. 4. Perception of Dock and buoys

For buoys, each detection triggers the image-based buoy classification routines, providing a bearing which the routine converts into the required RoI. With the classification returned, the confidence value attributed to it is range-dependant, with closer values resulting in a higher confidence. If the new confidence value is greater than the previous, the classification will be overwritten to reflect the better result.

The light buoy is a special case, where a process is continuously running when classification is required. This process is activated and deactivated as necessary to conserve computing resources. Whilst activated, the process will actively seek to detect the buoy and report its state. Once the state is determined, it and buoy classification are returned to the object server.

For docks, which have already been detected by their size, a rectangular rotated bounding box is fitted to the dock to determine its orientation. From this, the dock entry locations can be determined along with the locations from which the docking symbols should be captured and determined. Once the vessel arrives at these positions, the symbol detection routines are called to produce as classification.

4) Resultant Data: The object server periodically publishes the object list for most other processes (chiefly the mission planner) to leverage for their operation. High-level decisions are made based on the vessel's state and the object list.

The occupancy grid is used by the path planning software as its primary reference information. This allows for very fine detail to be provided to the path planner while streamlining data processing for other functions which do not require this resolution.

G. Control and Guidance

The WAM-V is a differential drive vehicle. The platform has an option to use either a Pure Pursuit Controller [2] or a non-linear guidance law controller, depending on the application. Each one uses a list of waypoint to go to as it target. With pure pursuit, the AMS will chase each waypoint to within a specified margin of error before chasing its next waypoint.

Another method was a non-linear guidance law controller [4]. The algorithm works by setting a virtual target point at the intersection of a user-defined circle around the USV, and the straight line connecting the two way-points.

Each of these methods output a course message containing a desired speed and angle. A PID controller is used to match the angle and speed to the given course command.

H. Planning

Mission planning is done primarily using a state machine of a ROS package called SMACH [5]. For each task there will be a series of states that the robot must traverse to complete its task. SMACH was used because it has a variety of useful features that extend a normal state machine. Firstly SMACH has a Visualization tool implemented, where the current state can be visualized and observed. This is helpful when debugging and inquiring about the actions of the platform. Secondly it allows for nesting of tasks and modification, such as allowing the implementation of state machines inside state machines. This is especially useful for the finals and semifinals where the course is the combination of tasks rather than just individual tasks.

I. Safety

Safety was a key concern for the team throughout development and care was taken to ensure that each subsystem included fail-safe features and would not interfere with the safety features of other subsystems on the vessel.

The vessel is equipped with 4 emergency stop buttons around its perimeter, which immediately cut power to all the actuators on the vessel. This emergency stop system includes a relay which can be triggered from software in ROS and a custom-built wireless emergency stop unit developed by the team. The wireless emergency stop operates using the LoRa protocol and allows for the vessel to be shut down safely from over 300m away. This quoted range is the furthest separation tested by the team while beyond line of sight, and is far below the true capabilities of the equipment.

IV. Experimental Results

A. Lab Testing

Lab testing was critical in validating the low-level function of the various components of the AMS. Each software subsystem and hardware device was function tested before integration with any other system.

Lab testing proved particularly important for the Torqueedo thrusters, which plagued the team with teething problems as multiple bugs were identified and fixed.

B. Simulation

With the limited field tests available to the team, it proved imperative to retain an accurate and comprehensive simulation environment. In order to improve the simulation's accuracy and to tune it for the real-world operation of the vessel, data was recorded from the field tests to be for configuring the simulator.

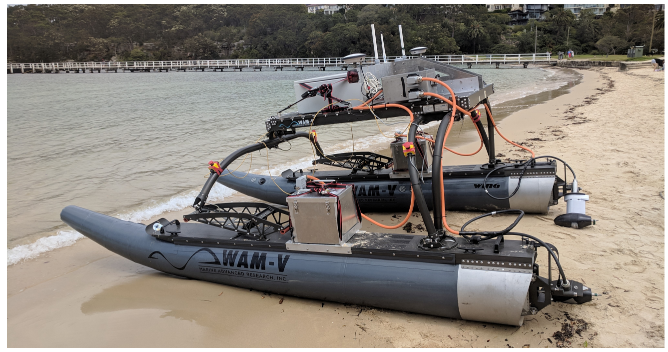

Fig. 5. The WAM-V on field test day 2, flotation pods detached

C. Field Testing

Field testing of the AMS was completed in various stages throughout development. This testing was primarily to inform and validate the outcomes of simulation testing.

The vessel's propulsion and navigation systems were tested on 3 days in the latter months of the development period. This testing occurred in Chowder Bay in Sydney's north and proved that the vessel could follow a commanded course and deliver the required level of control and reliability.

Delays in the manufacturing of the upgraded flotation pods and sensor mounting tower prevented the testing of the complete system before shipping. This would have been a major issue if not for the benefits of testing within the Gazebo simulation environment.

The perception systems were field tested separately using test targets built by the team. These field tests were chiefly designed to validate the performance of the cameras in varied lighting conditions. The ability for the cameras to reliably identify an object in frame by colour is a lynch-pin for the system so field tests for this were prioritized.

V. Platform Assessment and Future Development

1) Current Capabilities: USYD RowBot believes that the current AMS system to be delivered for the 2018 RobotX challenge is capable surface vessel which can acquit itself well in Hawaii and provides a suitable platform for future development. Its perception, propulsion and autonomy all stand to be upgraded and supplemented for research and competition in the coming years. The team also looks forward to using the vessel with AUVs, UAVs and deck-based actuators in the future.

The vessel's reliability is judged to be adequate for competition but not for a long-duration mission. Upgrades will need to be made to the vision and propulsion systems to sufficiently ruggedise them for such an application.

2) Development Paths: Future competitions and research will require many upgrades and alterations to the current vessel to improve and extend its capabilities. The vessel's perception suites stand to be upgraded in ways previously discussed under Vehicle Design.

Acknowledgments

USYD RowBot would like to thank Robonation for organizing the RobotX 2018 challenge and for providing the WAM-V vessel to the University of Sydney.

We thank the Australian Centre for Field Robotics for its crucial financial, technical and organizational assistance to the team.

We would also like to thank our generous sponsors; Thales Group Australia and Novatel for their assistance in this endeavour, we couldn't have done this without them.

References

[1] C. Aguero and B. Bingham, *Virtual Maritime RobotX Challenge (VMRC)*, 2018, https://bitbucket.org/osrf/vmrc

[2] R. Craig Coulter, *Implementation of the Pure Pursuit Path Tracking Algorithm*, Carnegie Mellon University, USA, 1992

[3] A. S. Huang, E. Olsen and D. C. Moore, *LCM: Lightweight Communications and Marshalling*, MIT, USA, 2010

[4] H. Niu, Y. Lu, A. Savvaris and A. Tsourdos, *Efficient path following algorithm for unmanned surface vehicle*, Cranfield University, UK, 2016

[5] J. Bohren and I. I. Y. Saito, *SMACH*, 2018, http://wiki.ros.org/smach