

Development of a Low Cost Multi-Robot Autonomous Marine Surface Platform

A. Valada, P. Velagapudi, B. Kannan, C. Tomaszewski, G. Kantor, and P. Scerri

Abstract In this paper, we outline a low cost multi-robot autonomous platform for a broad set of applications including water quality monitoring, flood disaster mitigation and depth buoy verification. By working cooperatively, fleets of vessels can cover large areas that would otherwise be impractical, time consuming and prohibitively expensive to traverse by a single vessel. We describe the hardware design, control infrastructure, and software architecture of the system, while additionally presenting experimental results from several field trials. Further, we discuss our initial efforts towards developing our system for water quality monitoring, in which a team of watercraft equipped with specialized sensors autonomously samples the physical quantity being measured and provides online situational awareness to the operator regarding water quality in the observed area. From canals in New York to volcanic lakes in the Philippines, our vessels have been tested in diverse marine environments and the results obtained from initial experiments in these domains are also discussed.

1 Introduction

Understanding and monitoring complex natural systems is critical to our ability to use the resources provided by those systems while ensuring the systems stay healthy over long periods of time. Ecologically sensitive marine areas are large and remote, making it difficult to see impacts from threats such as climate change and pollution. Scientists and governments are interested in tracking contaminants in the water, assessing environmental change, monitoring sea levels, surveying the coastline and understanding biodiversity. Unfortunately, monitoring and understanding large spatially distributed systems is difficult due in part to partial observability, complex dynamics, high cost, and the diversity of data that must be collected. Radically new approaches to sampling and monitoring that enable collection of collecting large

A. Valada · P. Velagapudi · B. Kannan · C. Tomaszewski · G. Kantor · P. Scerri
The Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, USA
e-mail: {avalada, pkv, bkannan, ckt, kantor, pscerri}@cmu.edu

amounts of data over large areas and extended periods are required if we are to balance economic development with conservation for future generations [1].

Since the 1990's Unmanned Surface Vehicles (USV) or Autonomous Surface Craft (ASC) have been demonstrated to work in a broad set of applications such as bathymetric mapping, environmental sampling, defense and robotics research, yet there has been very little commercialization or adoption in these domains. Unlike ground or aerial vehicles, the extremely dynamic and frangible environment in which ASC's are deployed make the task of developing these machines far more complex. For example, in applications such as monitoring pollution around coral reefs, the system has to operate in a manner that does not damage the delicate ecosystem. For long range deployments, communication range and data transmission bandwidth are also important considerations. Although a few commercial ASC's have been developed in the past few years, the high cost of these systems has prevented many agencies from adopting them. Most of these vehicles are fitted with expensive navigation systems and sensors, thereby bringing the cost of the entire vehicle to several thousand dollars. These are only some of the challenges that the ASC community is facing today.

In this paper we present our work on developing a multi-robot ASC platform which addresses some of these challenges. We present the Cooperative Robotic Watercraft (CRW) platform, a commercial smartphone based solution that leverages the built-in inertial and visual sensors of modern mobile devices for navigation and localization. In recent years, smartphones have redefined the notion of mobile computing platforms. Unrelenting advances in mobile processors, low power sensors, and battery technology have made smartphones more powerful, affordable, and portable, opening up a number of interesting and novel applications. By using a combination of WiFi and 3G, we improve communication reliability, extend the range of system deployments beyond the limited range offered by standard 802.11 networks and take advantage of existing infrastructure while maintaining a low development cost for each craft - below \$800. The reduced vehicle cost allows for development of a fleet of boats, which can cooperate to perform tasks and achieve goals more efficiently and expediently than a single craft. We envision very large fleets of CRW, perhaps even numbering in the hundreds, autonomously exploring large bodies of water under the supervision of a small number of operators. Previous work has detailed the challenges involved in such coordination from a multi-agent perspective, including challenges in task allocation, information sharing, and adjustable autonomy [9].



Fig. 1: Team of CRW sampling at Taal lake in Philippines.

The overarching goal of our work is to develop a low-cost multi-robot marine surface platform that is easy to deploy and versatile enough to be used for a wide range of applications. This results in several unique design constraints that have not previously been explored. Section 3 outlines these considerations and describes the specific design choices that we made. We then discuss our validation experiments in cooperative adaptive water sampling, where a suite of sampling algorithms ranging from random search to highest uncertainty tracking were developed and system performance in different operating environments (lakes, rivers, and floods) as well as for varying team sizes were compared. The results from several field trials and subsequent analysis is presented in Section 4. Finally, Section 5 provides a brief conclusion along with a description of our plans for future research.

2 Related Work

Although there has been considerably less work done on ASC's compared to AUV's, there are still some notable platforms that have made significant contributions to this field. One of the first ASC's that was developed was the ARTEMIS platform, which was used to perform bathymetry and test control algorithms [2]. After several iterations of the ARTEMIS platform, a new platform known as AutoCat [3] was developed to improve upon the shortcomings of its predecessors. Tele-supervised Adaptive Ocean Sensor Fleet [4] is a deep sea multi-robot science exploration system that combines a group of robotic boats to enable in situ study of phenomena in the ocean-atmosphere interface, on the ocean surface as well as the sub-surface beneath it. The NAMOS project [5] developed an airboat which works in conjunction with static nodes that are to be deployed prior to sampling. The static nodes provide temporal coverage while the airboat samples a good spatial resolution of data. The OASIS platform is a long-duration solar-powered autonomous surface vehicle, designed for autonomous global open-ocean operations [6]. Recently there has also been some work on developing a unique ASV which runs on the energy harnessed from waves. This ASV known as the Wave Glider [7] is also equipped with solar panels and has been demonstrated to work well for long distance as well as extended deployments. There has also been some exciting work in developing multi-robot systems with a wide variety of capabilities [8], but under tightly controlled conditions. The authors in [9] developed an ASV equipped with a LIDAR, camera, radar and an imaging sonar. The ASV is used for mapping above as well as below the surface of the water while accounting for GPS degradation near bridges and tree canopies. In the more recent work [10], the ASV developed at ETH was designed with a winch mechanism, that enables the vehicle to collect measurements at various depths, ranging from the surface to 130 m below.

Although these platforms are extremely capable and engineered specific to the requirements of the operating domain, the large associated cost with these platforms make them infeasible for large scale deployments.

3 Cooperative Robotic Watercraft

The primary objective of our work is to design, develop and test an application independent low cost, multi-robot platform that can be used for collecting scientific data

and to monitor dynamic spatial phenomena over large distributed areas with minimal human supervision. In this paper we discuss five challenging areas that impact the process of developing autonomous watercraft teams. These include: *hardware design, sensing and autonomy, user interface, communication and coordination, and exploration and coverage.*

As the system is equipped with a smartphone, the CRW platform has the capability to use either WiFi or 3G for communications. Application specific algorithms decide where the vessels need to go and determine which is best for each required location. In the future, we anticipate making this a distributed process, but it has not been required until now as good wireless connectivity has been available. The exploration and coordination aspects are handled according to the application specific tasks. The following sections describe the other key areas.

3.1 Hardware Design

For the design of individual craft we chose that of a traditional airboat, which is ideally suited for navigating shallow or ecologically sensitive areas such as reefs and estuaries without damaging the marine ecosystem due to its flat bottom and absence of operating components below the waterline. Another advantage of this design is that it can be used in urban floods where debris filled water often hinders rescue operations. Throughout the development process, several design iterations of various sizes ranging from 40 cm to 70 cm long were fabricated and tested to achieve our desired performance characteristics.

The main design criteria was to develop a system that is low-cost as well as easy to manufacture and repair. Unlike other expensive autonomous marine vehicles, the CRW platform was designed to work cooperatively in large teams to accomplish a task that would otherwise be impractical by a single vessel, which placed a few unusual constraints on the physical design of the vehicle. Simplicity and ease of fabrication were critical to keeping the cost of the craft as low as possible, however overall system robustness and modularity was also highly desired. Since even a relatively low failure rate for a large team would result in a cost-prohibitive amount of time for repairs, the individual airboats were designed with modular, interchangeable components, which can easily be swapped in the field.

The hull of the vessel is fabricated from multiple sheets of extruded polystyrene foam to be resilient and inexpensive. The sheets are bonded together into a single volume and carved using a heated nichrome wire setup. The hull is then sanded, painted, and coated with sealants for weatherproofing. An important feature of the platform is the self containment of the system components: the propulsion and steering system is bolted to a single deck plate which is mounted on top the vessel's hull, the electronics box slots into a dedicated compartment in the hull and is secured to the deck plate, and the smartphone is mounted to the deck plate at the bow of the craft. Any of these components can be easily detached and replaced within a couple of minutes.

3.1.1 Propulsion and Steering

The propulsion and steering system is one of the most important components of the watercraft. Our initial design was based on that of a conventional airboat, in

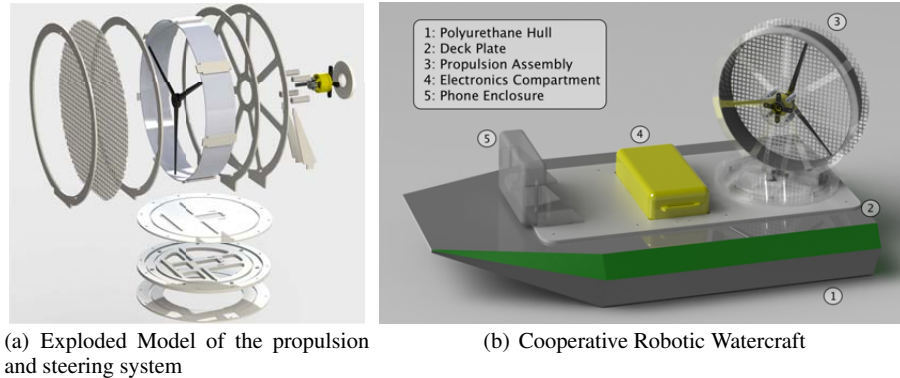


Fig. 2: CAD models of the CRW platform.

which a shrouded propeller powered by a brushless motor is mounted at the rear end of the vessel and a pair of rudders is used to direct thrust in order to steer the vehicle. Despite the tried and true nature of the traditional design, the production and maintenance of the rudder mechanism was cumbersome and its complexity did not support our vision of a modular system. After experimenting with several initial prototypes, we designed a rudder-less steering system where the entire propulsion assembly is actuated to steer the vehicle. In our design, the motor is mounted within a close-fitting PVC duct, which serves to decrease thrust losses and noise due to the tip vortices of the propeller. The duct is mounted to a circular acrylic base, which is free to rotate on bearings within a larger acrylic retaining ring. A servo is top-mounted to the circular base, with its horn protruding below the assembly where it can engage a matching cutout on the deck plate. Using the servo, the entire duct can pivot allowing for more efficient thrust vectoring than the traditional rudder system, particularly at rudder angles further from the neutral position. At full charge the propulsion system is capable of producing up to 8.2 Newtons of thrust, which can be entirely directed along any vector to the rear of the vessel, facilitating a small turning radius. An exploded model of the propulsion assembly is shown in Figure 2(a).

The primary material used in the propulsion and steering system is weather and UV resistant acrylic, which is cut from sheets using a laser cutter. The acrylic components are then assembled and bonded together with a special adhesive before installing the PVC ducting, motor, propeller, and protective mesh. Once templates were created for these acrylic components, manufacturing was seamless and many propulsion assemblies could be produced in a matter of hours. This new design proved not only to be more robust and efficient but also improved the overall maneuverability of the vehicle and could easily be swapped out in case of failure.

3.1.2 Electronics

The onboard electronics are decomposed into two major components: a phone assembly which houses the smartphone and orients it to best align its sensors with the vessel, and an electronics assembly that contains the main power source and actuation drivers for the vehicle as well as embedded electronics to interface with

additional external sensors. The heavier electronics assembly is placed close to the center of mass of the vessel and is partially recessed into the deck. The assembly is powered by a 10 Ah NiMh battery, which provides about 150 minutes of continuous operation. An Arduino Mega provides a fast and flexible array of digital and analog I/O for controlling the propulsion system and external sensor modules. Sensors such as fluorimeters and sonars are directly plugged into the Arduino through analog, digital, or serial channels. Bluetooth is used as an interface between the smartphone and the electronic components, enabling us to keep the two components decoupled physically and architecturally. Figure 3 shows the system architecture diagram.

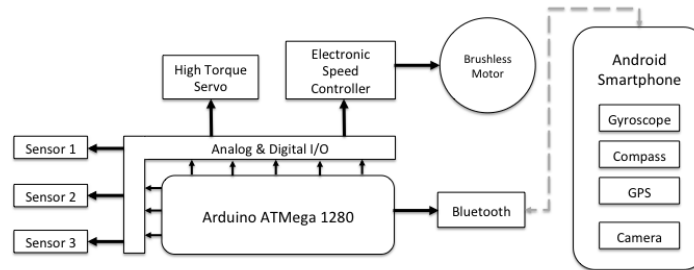


Fig. 3: System architecture diagram.

3.2 Sensing and Autonomy

At the core of the CRW platform design is an Android smartphone which provides computing and inertial sensors in an integrated package. This design decision allows us to avoid the cost-prohibitive process of integrating sensory, computing, and communication components in a comparable form factor. Modern smartphones also provide access to multiple modes of communication such as 3G, WiFi and bluetooth. The phone is enclosed in a watertight case that is mounted upright on the deck of the vessel to best utilize the camera's field of view. The Google Nexus S makes an ideal candidate due to its powerful open source Android development environment.

Computationally intensive tasks and high-level controls are performed on the phone, while low-level sensor interfacing and high-speed control loops execute on the Arduino. The primary issue to overcome pertains to the quality of obtained data from the embedded sensors required for control, specifically the GPS, gyroscope, and compass. Layers of filters are required to smooth the data to extract sufficiently clean information to effectively control the boat.

At times throughout sampling it becomes desirable to collect a physical sample to verify sensor readings or for more detailed analysis than the onboard sensors can provide. A water sampling mechanism provides this capability, allowing for up to six unique samples to be collected before its storage must be emptied. The sampler is loaded with evacuated glass vials, which are punctured by needles affixed to individual rubber intake hoses leading beneath the boat. The vacuum in the vials is held by pinching the hoses against rounded pins from a tumbler lock, which ride against a central cam. A single actuator is then used to rotate the cam, leveraging the elasticity of the hoses to displace the tumbler lock pins and selectively draw samples

into each vial. By using individually sealed vials and a unique hose for each sample, cross contamination between samples is minimized without requiring any additional actuators.

As all the computation is performed on the Android smartphone, local intelligence for each boat resides on the phone. Layers of functionality separate general modules from application specific modules. An end-user interface provides a single operator with an overview of the state of the boats and provides high and low level commands for interacting with them. The boat executes the core functionality via the boat server. Client applications and additional modules running on the phone provide the domain specific functionality. This design allows us to make subtle changes for specific domains without modifying previously tested and reliable code. For example, the behavior to be exhibited by the vessel when it loses communication with the base station varies significantly depending on the domain. When doing testing, the vessel should immediately attempt to go back to home base; in a flood response, it may only return at regular intervals to provide data to first responders. This domain specific logic is captured in the client applications without adversely affecting the core functionality that implements the actions.

The top level intelligence of the boat, the reasoning about where and what the boat should do, is encapsulated in a proxy. Currently the proxy runs on an operator's machine and has relatively low-overhead in terms of communication with the boat. The proxy is responsible for path planning to implement high-level operator directives about areas to visit or search.

3.3 Operator Interface

A centralized operator interface provides an enhanced situational awareness about the multi-robot teams and the operating domain. The interface displays information about the locations of all the vessels, overlaid on a map of the environment. Using the interface the operator can specify high-level objectives either as waypoints, paths or areas to search, or low-level direction commands to the boats. The watercraft also transmit images captured from the on-board camera at approximately 1

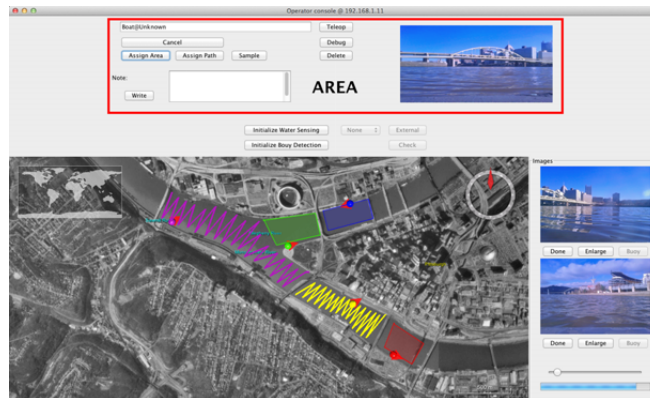


Fig. 4: Main panel of the operator interface.

Hz. An image queue on the operator's side receives and reorders the images, allowing the operator to observe, discard, or save images for later use. The operator interface emphasizes simplicity and reliability over complex functionality. A snapshot of the main panel of the interface is shown in Figure 4. The interface is also composed of subpanels that display application specific information. For example, the water sensing subpanel displays the sampling grid cells with the current mean of observations within each cell as well as buttons to modify the number cells and sensor channels. The debug subpanel displays additional vital information about the vehicle such as the yaw rate, coordinates of the current position and waypoints positions, controller gains and tuning sliders.

4 Experimental Results

The CRW platform is designed as an application independent fleet of ASV's. In order to evaluate the designed system, we performed experiments in different domains such as water quality monitoring, depth buoy verification, flood disaster mitigation and bathymetry. In this paper we will focus on water quality monitoring as this application has attracted much attention in the recent years.

4.1 Water Quality Monitoring

Investigating surface pollution is a difficult and complex task that can be challenging even for an autonomous watercraft. Pollutants such as petroleum hydrocarbons, heavy metals, polycyclic aromatic hydrocarbons, pesticides, and some toxic organic compounds are the common contaminants of concern. Even though there are several protocols for safely disposing such compounds, the lack of efficient monitoring techniques has led to illegal dumping in water bodies. Using cooperative ASC's we can drastically improve the management of water resources by using intelligent techniques such as adaptive sampling to analyze the trend of contamination and identify the sources of pollution. Analysis of surface water also gives us information on the interaction between the atmosphere and hydrosphere, which is of scientific interest. While performing sampling with an autonomous watercraft, our main objective is to collect measurements of interest in a way that best explains the hydrological processes. Different sampling applications require either spatial or temporal resolution of data; while sampling to identify the sources of pollution, it is desirable to sample at locations having greater variance, whereas the main objective while sampling to map a hydrological process is to obtain the highest resolution possible. In either of these applications the vehicles paths have to be planned and replanned according to the measurements obtained. Sampling of the entire field has to be performed before the field varies significantly, hence the use of a single vehicle to monitor large areas is impractical.

During our initial sampling efforts using the CRW system, we implemented a highest uncertainty tracking cooperative sampling algorithm and compared the results to those obtained using random search and lawnmower patterns. We chose temperature and specific conductivity as the quantities to be sampled, as conductivity gives us a measure of the total dissolved salts (TDS) in the water and temperature is an important factor that influences biological activity in the water. The sample

space is then discretized into grid cells where the number of cells are configurable. Determining the number of cells to use is a critical task as there will be an optimal number of cells for the selected sampling space for which an algorithm will work most efficiently, and choosing more or less cells than necessary might lead to over or under-sampling.

$$v_{i,j} = \sigma^2 \times \alpha^c \quad (1)$$

$$\alpha < 1.0$$

In random search sampling, the vehicle is given random cells to sample one after the other until the sampling time expires, while in lawnmower sampling, the vehicle sweeps the area in a lawnmower fashion. These two algorithms give us a good reference for comparison. As we use multiple vehicles for sampling, it was necessary to divide the sample space into sub regions and assign each vehicle to a specific subregion for the lawnmower sampling approach. For the random search approach, dividing the sampling region was not required since each vehicle is assigned a random cell to sample. In the maximum uncertainty tracking algorithm, each cell is assigned a value based on the variance and the number of times that cell has been visited. The standard deviation of all the observations in each cell is calculated and the new value of the cell is given by Equation 1, where c is the number of times that cell has been visited, σ^2 is the variance and α is a constant. A path is then planned to the cell which has the greatest v . Initially when the count of all the cells are zero, v is assigned a high sensor value so that the vehicle visits all the cells at least once, this gives an opportunity for the vehicle to explore the area before sampling for the maximum uncertainty.

For this algorithm to work efficiently the standard deviation at each cell should not vary very significantly from another. The time taken for the algorithm to converge directly correlates with the difference in the standard deviation between the cells. Other factors such as number of cells also influences the performance of the algorithm. After sampling the entire field, the surface is reconstructed using interpolation. The fidelity of the reconstruction is then determined using the estimation error obtained at different periods of time.

4.1.1 Experimental Domain I — Taal Lake

The Taal lake in Philippines is famous for many reasons. It is the largest fish farm in the Philippines and rather uniquely, the world's largest lake on an island on a lake on an island. In May 2011 however, it became infamous for another reason: it was the site of one of the largest fish kills in the world. About 956 metric tons of fish appeared dead on the surface of the lake, costing the fish farming industry nearly 70 million PHP (1.6 million USD) in lost revenue. One of the major reasons cited for this disaster is the sudden rise in temperature of the lake followed by rainfall that prevented the adequate exchange of oxygen on the surface. A team consisting of five CRW was deployed in Taal lake to observe the changes in the surface temperature before and after a tropical rainfall. Plots shown in Figure 5 displays an overlay of the surface temperatures obtained by the CRW team and an image captured during deployment is shown in Figure 1. It was observed that the average surface temperature varied from 31.81 degrees Celsius before the rainfall to 29.96 degrees Celsius after the rainfall.

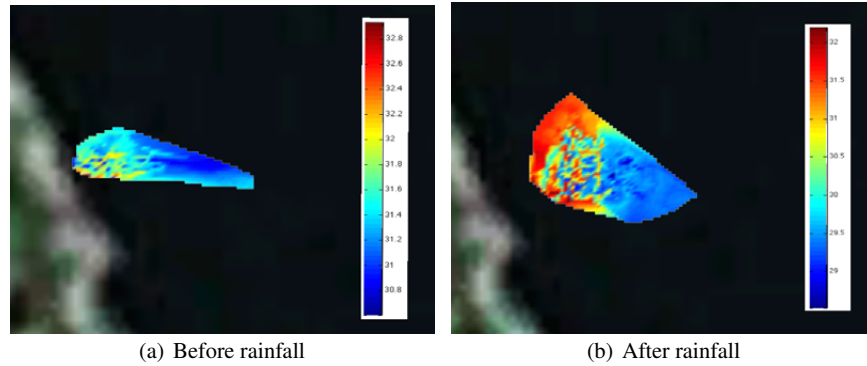


Fig. 5: Surface temperature variation before and after rainfall in Taal lake.

4.1.2 Experimental Domain II — Gowanus Canal

The Gowanus Canal is one of the most polluted bodies of water in the United States and was recently named a superfund site by the Environmental Protection Agency. Effluents from chemical and manufacturing plants along the canal are the major source of contamination of the waterbody. A team of CRW was used to autonomously map the specific conductivity distribution in the area to analyze the spatial variation in the field. Significant spatial variation suggests manual sampling is not viable as several hundred points would have to be sampled to get a low estimation error. Figure 6(c) shows a plot of the sample space discretized into 10×10 cells with the standard deviation of each cell and the number measurements obtained in them indicated by red circular markers. It can be seen that even though a large number of samples are taken in most of the cells, there is a high variation in the measured value. The mean of the observations in the cells varied from 7.7050 dS/m to 7.5607 dS/m across the sensing field and the variance was calculated to be 0.0023. The spatial distribution of the specific conductivity of a section of the Gowanus Canal is

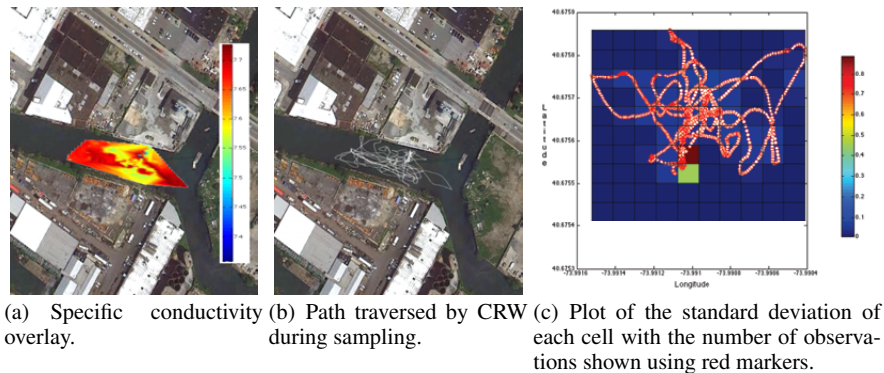


Fig. 6: Specific conductivity variation in a section of the Gowanus canal.

shown in Figure 6(a). The average measured specific conductivity was 7.6387 dS/m, which is well above safe limits.

4.1.3 Experimental Domain III — Allegheny River

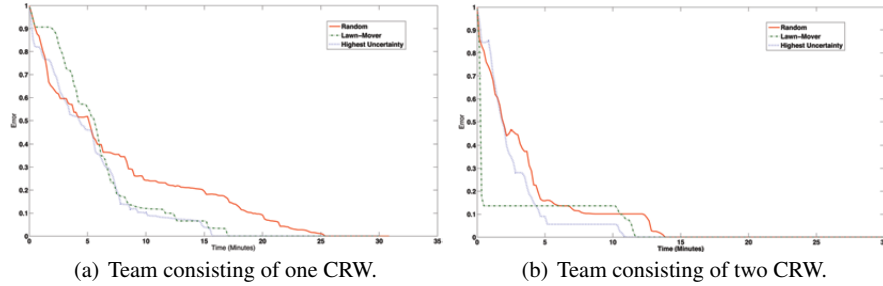


Fig. 7: Error comparison between random search, lawnmower pattern and highest uncertainty sampling.

A more rigorous experimental validation of the random search sampling, lawnmower pattern, and the highest uncertainty algorithms were performed on a branch of the Allegheny river shielded from the main flow, located by Washington's Landing in Pittsburgh, PA. All three sampling algorithms were evaluated using a team of CRW consisting of three vessels equipped with sonars as well as conductivity and temperature sensors. The sampling area was about 6510 square meters, discretized into 10×10 cells and the sampling time for each algorithm was 30 minutes. To quantify the performance of the highest uncertainty tracking adaptive sampling algorithm, we calculated the normalized error and compared it with that of the random search sampling and lawnmower patterns. The error was calculated by subtracting the value of each cell at the end of 30 minutes to that at each sampling interval. This process also gives us the rate of convergence to the final measurements. It is assumed that the field does not vary significantly for the period of sampling and the values of the cells at the end of the sampling time are an approximate measure of the ground truth. To analyze the final field estimation error, the sum of errors of all the cells at discrete sampling intervals was plotted. The graph in Figure 7a shows the trend obtained from testing using one CRW. With the random search sampling, more measurements end up being taken in the cells that are in the center of the grid, as the vessel traverses to random waypoints from one end of the sample space to another. In the case of the lawnmower pattern, a more uniform trajectory is followed, hence exploration is sequential and slow. With the highest uncertainty algorithm, priority is first given to exploration of the sample space and then to tracking the cells with highest variance. It can be seen in Figure 7a that this algorithm has more of a uniform downward trend and the final value converges faster than both random search and lawnmower patterns.

The time taken to converge to the final value can be improved by cooperatively sampling the area, thereby reducing the time required for exploration and maximizing coverage. To demonstrate the performance of the algorithms using multiple




Fig. 8: Error comparison for lawnmower pattern sampling with CRW teams of varying sizes.

vessels, experiments on the algorithms were performed using teams of two and three CRW. In Figure 7b, the lawnmower pattern sampling has the least initial error, this can be attributed to the fact that each vessel samples only the subregion assigned to it, hence the initial explorative sampling is fast and even, but the convergence to the final value slows down as the vessels follow a fixed pattern. With the highest uncertainty algorithm, fast convergence is achieved as the observations from both the vessels are used to calculate the uncertainty and plan the next waypoint accordingly. With a single CRW, convergence was achieved in 15.6 minutes, whereas with two CRW, convergence was achieved in 10.4 minutes. A comparison of the error for the lawnmower pattern with varying team sizes is shown in Figure 8. With a team consisting of three vessels, convergence to the final value was achieved within 10.5 minutes, which is still higher than the result achieved using two vessels with the highest uncertainty algorithm.




Fig. 9: Terrain model of a section of the Allegheny riverbed.

Bathymetric mapping is one of the most successful applications for ASV's and AUV's as bathymetric maps are an important navigational aid for all marine vehicles. A team of CRW equipped with sonars were used to map the river bed near Washington's Landing. To create a three dimensional surface model, a grid is first constructed using the range of the obtained values and interpolation is performed on the measured data. Gaussian smoothing is applied to smooth any noise in the measured data and a surface is fit on it. A map developed from sampling portraying the terrain model and isobaths for the same area is shown in Figures 9 and 10.

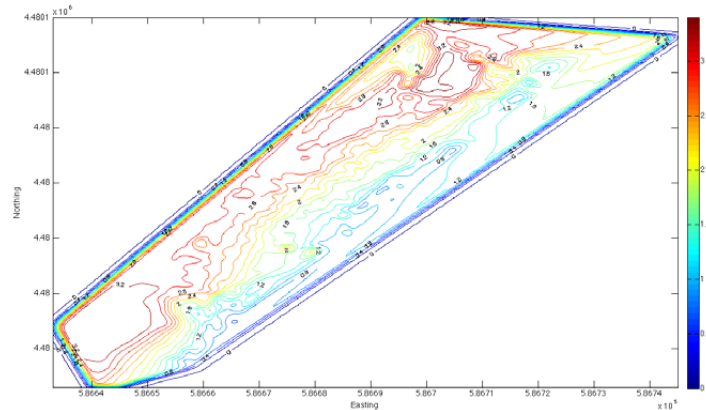


Fig. 10: Countour map of a section of the Allegheny river bed.

5 Conclusion and Future Work

In this paper we described the design and architecture of a team of CRW and analyzed its performance in water quality monitoring applications. The integrated smartphone solves many inherent problems associated with ASV's relating to long range communication, situational awareness, system integration, and cost effectiveness. The custom designed propulsion assembly not only increased the efficiency of steering of CRW but also remarkably improves the fault tolerance of the entire system while compared to the traditional design. Additionally, by incorporating multiple layers of communication and a centralized coordination protocol, we ensure efficient and robust operation of the autonomous system. We compared the estimation error of the highest uncertainty tracking algorithm with random search sampling and lawnmower patterns for teams of varying sizes. The highest uncertainty algorithm not only has lower average error but also converges faster than the random search sampling and lawnmower patterns. Our experiments totaling to more than 180 boat hours in different operational and environmental conditions and with with teams of up to five CRW have confirmed the versatility and the serviceability of our platform. A table detailing the approximate testing time in water at each location along with the team size is shown in Table 1.

Ongoing work is focused on system modeling and development of a nonlinear controller with an integrated planner for this underactuated system. Unlike con-

Table 1: Deployment statistics

Experimental Domain	Team Size	Time in Water (Hours)
Taal Lake	5	95
Gowanus Canal	2	15
Allegheny River	3	70

ventional airboats, the actuated propulsion assembly enables the CRW to perform extremely tight turns, thereby allowing for very accurate tracking by planning paths using the system dynamics. Future work includes developing more intelligent multi-robot adaptive sampling algorithms that reduce the estimation error and improve the spatial resolution of surface reconstruction. Another promising application for our system is depth buoy verification, where a team of CRW is used to visually identify the depth buoys which are placed in river channels to mark the safe passage of ships and barges, a process that is currently performed manually at a considerable expense.

References

1. Niemi G, Wardrop D, Brooks R, Anderson S, Brady V, et al.: Rationale for a New Generation of Indicators for Coastal Waters. *Environment Health Perspective* 112(9): doi: 10.1289/ehp.6903, (2004)
2. T. Vaneck, J. Manley, C. Rodriguez, and M. Schmid.: Automated Bathymetry using an Autonomous Surface Craft. *NAVIGATION, Journal of the Institute of Navigation*. Vol. 43 No 4, Winter (1996-1997)
3. J. Manley, A. Marsh, W. Cornforth, and C. Wisema.: Evolution of the Autonomous Surface Craft AutoCat, *Proc. of Oceans 2000, MTS/IEEE Providence, RI, October, (2000)*
4. A. Elfes, G.W. Podnar, J.M. Dolan, S. Stancliff, E. Lin, J.C. Hosler, T.J. Ames, J. Higinbotham, J.R. Moisan, T.A. Moisan, et al.: The telesupervised adaptive ocean sensor fleet architecture: Coordination of multiple oceanic robot boats. *Proc. of the IEEE Aerospace Conference, (2008)*
5. G.S. Sukhatme, A. Dhariwal, B. Zhang, C. Oberg, B. Stauffer and D.A. Caron.: The Design and Development of a Wireless Robotic Networked Aquatic Microbial Observing System.” *In Environmental Engineering Science, Vol. 24, No. 2, pp. 205-215, Mar (2007)*
6. G.W. Podnar, J.M. Dolan, A. Elfes, S. Stancliff, E. Lin, JC Hosier, T.J. Ames, J. Moisan, T.A. Moisan, J. Higinbotham, et al.: Operation of robotic science boats using the telesupervised adaptive ocean sensor fleet system, *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA '07), (2007)*
7. R. Hine and P. McGillivray.: Wave powered autonomous surface vessels as components of ocean observing systems, *Proc. of PACON 2007, Honolulu, HI, June (2007)*
8. K.H. Low, G. Podnar, S. Stancliff, J.M. Dolan, and A. Elfes.: Robot boats as a mobile aquatic sensor network., *In Proc. W. on Sensor Networks for Earth and Space Science Applications at the Int. Conf. on Information Processing in Sensor Networks, (2009)*
9. J. C. Leedekerken, M. F. Fallon, J. J. Leonard.: Mapping Complex Marine Environments with Autonomous Surface Craft, *12th International Symposium on Experimental Robotics 2010, New Delhi & Agra, India, December (2010)*
10. Hitz G, Pomerleau F, Garneau M-, Pradalier C, Posch T, Pernthaler J, Siegwart RY.: Design and application of a surface vessel for autonomous inland water monitoring, *IEEE Robotics and Automation Magazine, (2012)*
11. P. Scerri, B. Kannan, P. Velagapudi, K. Macarthur, P. Stone, M.E. Taylor, J. Dolan, A. Farinelli, A. Chapman, B. Dias, et al.: Flood disaster mitigation: A real-world challenge problem for multi-agent unmanned surface vehicles, *AAMAS '11 Workshop on Autonomous Robots and Multi-robot Systems, (2011)*