Persistent robot tasking for environmental monitoring through crowd-sourcing

Wesam Al-Sabban*, Jnaneshwar Das† and Ryan N. Smith‡
*School of Electrical Engineering and Computer Science, Queensland University of Technology, Brisbane, QLD 4001, Australia
Email: alsabban@student.qut.edu.au
†Robotic Embedded Systems Laboratory, University of Southern California, Los Angeles, CA 90089
Email: jnaneshd@usc.edu
‡Department of Engineering and Physics, Fort Lewis College, Durango, CO 81301, USA
Email: rnsmith@fortlewis.edu

***Abstract*—Persistent monitoring of the ocean is not optimally accomplished by repeatedly executing a fixed path in a fixed location. The ocean is dynamic, and so should the executed paths to monitor and observe it. An open question merging autonomy and optimal sampling is *how and when* to alter a path/decision, yet achieve desired science objectives. Additionally, many marine robotic deployments can last multiple weeks to months; making it very difficult for individuals to continuously monitor and retask them as needed. This problem becomes increasingly more complex when multiple platforms are operating simultaneously. There is a need for monitoring and adaptation of the robotic fleet via teams of scientists working in shifts; crowds are ideal for this task. In this paper, we present a novel application of crowd-sourcing to extend the autonomy of persistent-monitoring vehicles to enable non-repetitious sampling over long periods of time. We present a framework that enables the control of a marine robot by anybody with an internet-enabled device. Voters are provided current vehicle location, gathered science data and predicted ocean features through the associated decision support system. Results are included from a simulated implementation of our system on a Wave Glider operating in Monterey Bay with the science objective to maximize the sum of observed nitrate values collected.**

I. INTRODUCTION

Effective observation of ocean phenomena requires a persistent presence to understand the complex interactions occurring over multiple spatiotemporal scales (i.e., m to km and months to years). In specific regions, such observation can help address fundamental questions in differentiating processes that happen locally from those processes that happen remotely and are advected into the region. In recent years, autonomous vehicles have started playing a key role in effective, efficient and adaptive data collection, facilitating simultaneous and rapid measurements over long periods of time. Such sensor platforms provide fine-scale resolution far surpassing previous sampling methods, such as infrequent measurements from ships or static measurements from buoys. Existing vehicles used for such persistent monitoring campaigns are the ocean gliders; underwater gliders [1], [2] and the Liquid Robotics Wave Glider [3], [4].

With these vehicles providing such a prolonged endurance, the research question changes from *where* to sample to *how and when* to alter a path or decision to achieve desired science objectives; some of which may be competing. Specifically,

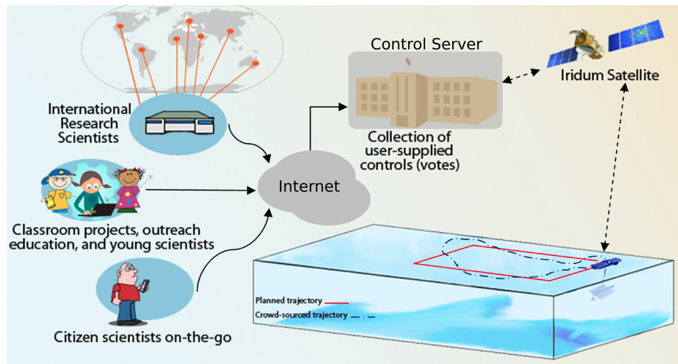

Fig. 1: Overview of the proposed crowd-sourced control system for a Liquid Robotics Wave Glider.

how do we obtain maximal information about the environment to help us better understand the dynamic processes that are occurring? Previous efforts have utilized a human-in-the-loop, decision-support approach for large campaigns to adapt paths and goals as data are obtained [5], [6]. This method has proven to be very successful, but it is not scalable when the mission duration for a fleet of vehicles is longer than a few days. Hence, we are interested in enabling input from multiple sources (crowdsourcing) to reduce the mission-planning burden on a single person or team, and learning from this collective decision making to extend the autonomous capabilities of persistent marine vehicles.

Ocean science is a field that has traditionally required large research vessels, expensive equipment, significant research funding, and access to secluded regions of the earth. Thus, exploring large areas and enabling citizen scientists have been challenging, and one of the reasons that we know relatively little about the Earth's oceans. With all life dependent upon an active hydrosphere, the ocean is an important resource to study and observe. The more observers and scientists, the more we will learn. Thus, this research aims to provide access to ocean science through robotics to children, hobbyists and citizen-scientists, while also bringing together expert researchers across the globe.

The basic idea is relatively simple – utilise crowd-sourced input to control a Liquid Robotics Wave Glider (WG) in the ocean (see outline in Fig. 1). Users can *vote* on a path or

978-0-933957-40-4 ©2013 MTS

Fig. 2: The Liquid Robotics Wave Glider platform.

direction of travel to achieve their desired goal. In general, crowd-sourcing is the act of outsourcing tasks, traditionally performed by an employee or contractor, to an undefined, large group of people or community, *a crowd* through an open call. Here, the crowd-sourced input provides the available options, and the on-board autonomy system decides what is best to execute, given the current operational constraints, e.g., hazard mitigation, collision avoidance, vehicle speed, etc. The voters do not get total control, but only supply the vehicle with options for its next move. Here, we present the web-based architecture and a proof-of-concept experiment for crowd-controlling a WG.

We provide the details on the development and implementation of our decision and control algorithm and architecture, and present the results of a pilot study where a group of users were given access to vote within our framework to control a WG.

II. BACKGROUND AND MOTIVATION

The ocean is dynamic, and persistent monitoring in an ocean region is not optimally accomplished by repeatedly executing a fixed path. An open question merging autonomy and optimal sampling is emphhow and when to alter a sampling path/decision, while achieving the desired (and sometimes competing) science objectives. The persistence of ocean gliders (deployments lasting multiple weeks to months) makes it nearly difficult for individuals to continuously monitor and retask them; this becomes more complex with multiple platforms operating together. There is a need for monitoring and adaptation via teams in shifts, and crowdsourcing is a solution for this task. We present a novel application of crowdsourcing to extend the autonomy of persistent monitoring vehicles to enable non-repetitious sampling over long periods of time. Additionally, we aim to learn from *expert* inputs to increase information gain and democratize science objectives. We demonstrate our method with results from a simulation that asked users to control a Wave Glider in the presence of a dynamically changing field.

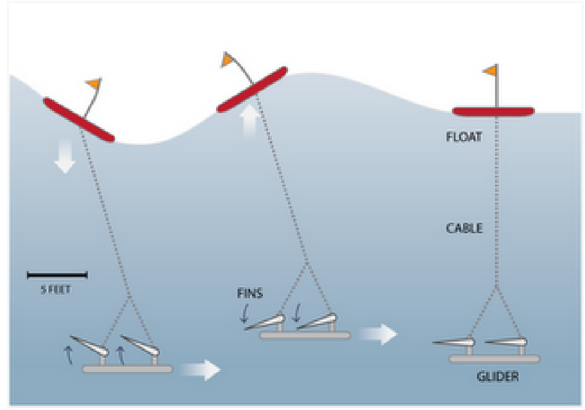

Fig. 3: The basic mechanics of locomotion for the Wave Glider.

The motivation for our proposed study is five-fold.

- Increase the ability of of ocean gliders to react to their dynamic environment and gather meaningful science data.
- Connect ocean researchers across the globe.
- Provide access to ocean sampling to a wide range of people.
- Outreach education and engagement for students of all ages.
- Enable citizen scientists to contribute to ocean research.

A. Liquid Robotics Wave Glider

We utilise the Liquid Robotics Wave Glider as the test-bed vehicle for demonstrating and implementing our techniques, see Fig 2. By harvesting abundant natural energy from ocean waves, the WG provides a platform for the study of persistent robotic control [3], [4]. The WG is a hybrid sea-surface and underwater vehicle that is composed of a submerged *glider* that is attached to a surface float via a tether. The WG is propelled by converting vertical wave motions into forward thrust, see Fig. 3. As waves pass by the surface float, the submerged glider tugs the it along a prescribed path. With a demonstrated endurance exceeding one year; during the PacX Challenge [7], a WG set a new world record for the longest distance traveled by an autonomous vehicle; a $9,000$ nm $(16,668$ km) scientific journey across the Pacific Ocean. The WG offers a unique platform for investigating the advantages of persistent monitoring for robotics research and ocean scientists alike. Persistent platforms like the WG are also ideal for education and outreach. They can teach children about the multiple elements of large science experiments (similar to controlling the Mars Rover). Also, there is interesting research in the democratization of large-scale experiments, e.g., the platforms have physical constraints with an objective to accomplish multiple (and potentially conflicting) science goals. How do the vehicles focus on the issues important to many?

B. Crowdsourcing

Crowdsourcing is the act of outsourcing tasks to an undefined, large group of people or community, *a crowd* through an open call. This concept has been applied across a broad set of actions and the concept has evolved into many different forms, see [8]. Crowdsourcing is most common in the processing and/or analysis of large data sets or the development of new products. Multiple past and present crowdsourcing projects can be found in [9]. One widely-used example of crowdsourcing is Amazon's Manual Turk, an online marketplace where users from around the world can *work* on projects posted by businesses and developers, providing an on-demand, scalable workforce [10]. Although crowdsourcing has been around for a long time, the word was formally coined in 2005 by Jeff Howe and Mark Robinson of Wired Magazine [11]. This area of research and applications has rapidly expanding with the existence of general-purpose platforms such as human-robot interaction [12], Environmental Monitoring [13]–[15], and data processing among others.

Crowdsourcing has recently been used in a scientific wildlife field study to examine migratory habits of birds [16]. In this research example, the authors developed and implemented multiple teleoperated robotic *observatories*, allowing birdwatchers to remotely see, track, and study animal activity and behaviors. In this study, a ranking system was developed to determine which single user was allowed to control the observatory camera at a given time. More than 600 users contributed to the online system, and requested more than 2.2 million images. These images were used by citizen scientists and experts to classify 74 unique species; eight of which where unknown to have breeding populations inside the survey region. This could not have been done by a single scientist or group, and utilized the crowd to assist in the data analysis process.

An application of crowdsourcing in robotics was the Mars exploration program implemented by NASA, [17]. Here, the crowd was allowed to vote on where the robot should land on the planet, and submit controls to the robot after successfully landing on Mars. More than 1,200 people signed up, contributed, asked questions, voted, and observed the development and implementation of Mars exploration project.

With the persistent presence in the ocean that is enabled through the use of the WG and other platforms, we begin to examine the utility of crowdsourcing in the ocean environment. As deployments become longer and longer, we can no longer rely on a single person or team to monitor and adapt a sampling mission. Mission plans and science goals should change as the ocean changes, and this necessitates decisions to be made continuously throughout a deployment. The most effective way to accomplish this is to distribute the decision-making task. We can use machine learning to determine patterns and trends in control strategies, weighting specific users based on experience or expertise. Through the use of this method, we can immediately adapt missions, and will enable the prediction of episodic events, potentially placing the assets in the right place at the right time to gather data with maximal information.

C. Persistent Monitoring

Persistent monitoring has been a focus in robotics recently, enabled by the increase in deployment duration and advances in autonomy. Some recent work on persistent monitoring can be found in [18], [19], where the frequency of visits to each region is adapted to the time scale on which that region changes. In the ocean environment, [20] presents an algorithm where the velocity of the vehicle is altered to change the sampling resolution within a given region. Persistent sampling is necessary in the ocean environment due to the multiple frequencies of spatiotemporal oscillation that need to be observed to study most phenomena.

D. Ocean Science Applications

Consider a field campaign involving multiple assets with the objective to track a dynamic ocean feature (e.g., ocean front or algal bloom) over an extended period. Scientist can use an Oceanographic Decision Support System (ODSS) to consult near real-time data, such as model outputs, satellite data, and weather conditions for each planning cycle [6]. The ODSS queries its database, incorporates ocean model outputs, and provides future projections to guide sampling into the spatiotemporal region of interest. With the guidance of expert scientists, an automated planner can then determine the most viable robotic option to be dispatched, given mission constraints. Plans can be continuously refined and updated throughout execution. The planner meanwhile will continuously conduct multi-criteria replanning and, after human validation, goals are communicated to deployed assets. Algorithms on-board the fleet ensure they operate safely and execute the prescribed mission. The above scenario works well when there is a single, distinct feature of interest, e.g., ocean front. However, what happens when there are competing science objectives (multiple features), and what does the fleet do when there is no distinct feature? Also, how can we teach the fleet to utilize tools like the ODSS to autonomously retask and adapt to the changing environment? These are questions that we aim to address through this research.

III. PROBLEM STATEMENT

The primary objective for the presented tool is to engage more participants in the study of ocean science; specifically connecting experts from around the globe and enabling ocean science exploration by children. Additionally, we are interested in learning sampling behaviors from ocean experts to extend robotic autonomy and understand general science objectives within a given region. Given the unprecedented amount of data that are collected by marine robots, we require a method to rapidly process, analyze, and make decisions to adapt to the dynamically changing environment. This tool provides utility in three areas of research merging artificial intelligence and ocean science.

1) Enables the ability to democratize science goals.
2) Allows for modelling of collective control behaviour in response to continuously updated science data from various sources.
3) Demonstrates the capability and utility of decision support tools.

IV. SYSTEM DESCRIPTION

In the presented example, we are interested in maximizing the sum of observed nitrate values (henceforth called *bloom score*, a dimensionless measure) collected by a WG in

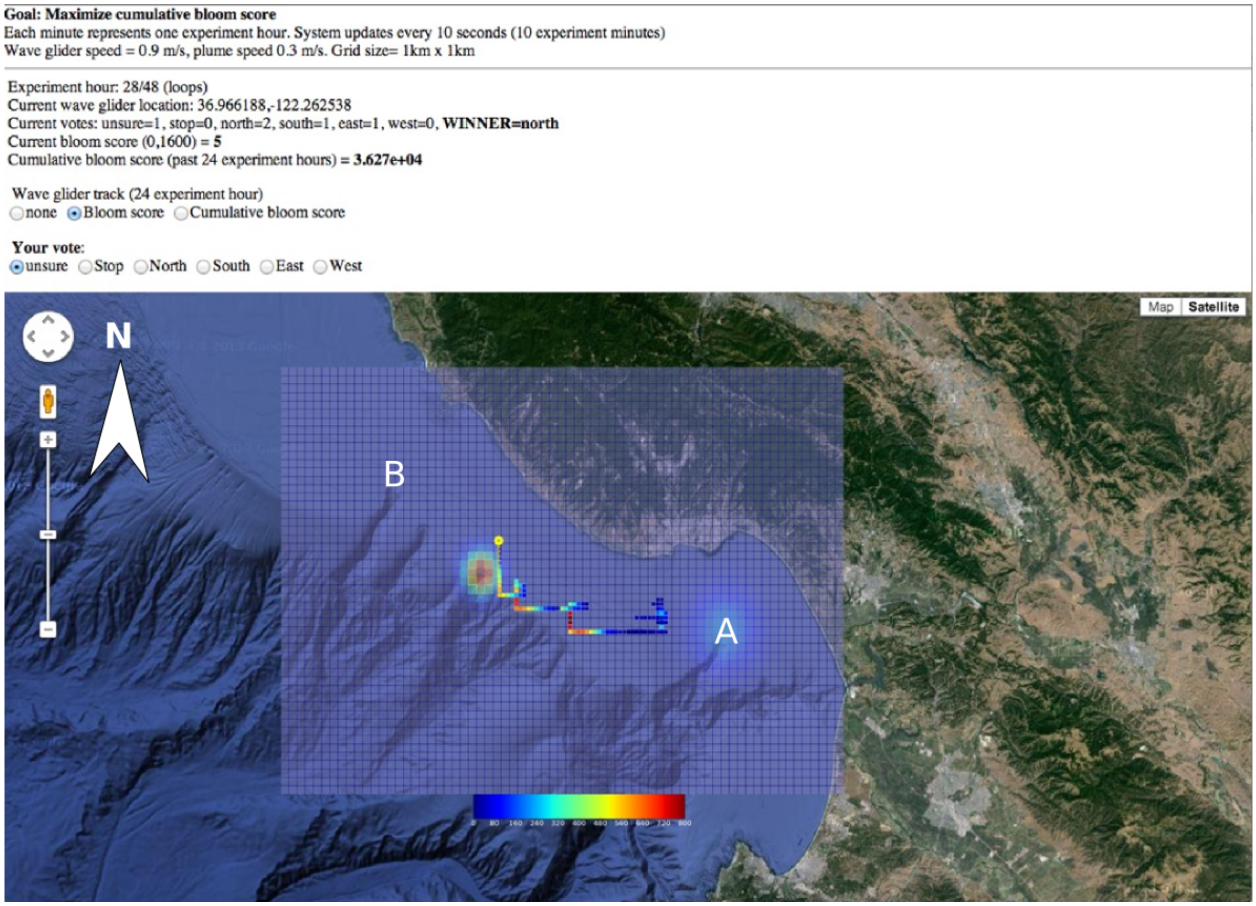

Fig. 5: Screenshot from the experiment when the wave glider was being controlled by five users, with two users concurring with the command direction *north*. Each grid cell represents a 1km x 1km area. The yellow dot marks the current location of the wave glider. In the simulation, the nitrate plume upwells at the location marked 'A', in the north Monterey bay. The plume then advects north-west towards location 'B' over a period of 2 days (48 minutes in simulation).

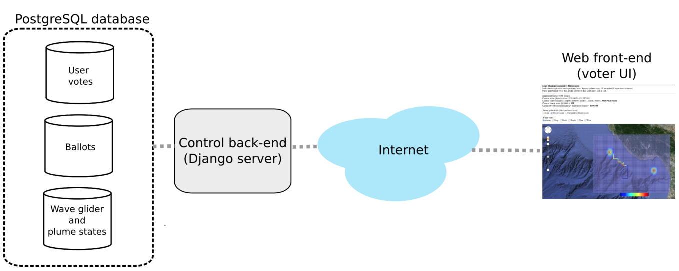

Fig. 4: System architecture for the crowd-control tool.

Monterey Bay, CA USA. A dynamic field emulates upwelling events in the northern Monterey Bay bloom incubator. Within our simulation, nitrate gets upwelled (modeled as a Gaussian hotspot), near the Soquel canyon, and starts drifting westward, fading in intensity during transport. In the meantime, a new upwelled nitrate hotspot starts surfacing at Soquel Canyon. This process repeatedly loops. We examine the crowdsourced data to see if there is a trend or set of trends that users demonstrate when confronted with the aforementioned situation. If so, can we learn a policy that the WG can use for persistent monitoring when not controlled by an individual or the crowd?

Our web-based system was was designed using the Django web framework with a PostgreSQL database at the backend (see Fig. 4). Users have the option to cast a vote on a browser-based, web user-interface. The options are *unsure*, *stop*, *north*, *south*, *east* and *west*. Every 10 seconds (10 minutes in experiment time), the browser sends the vote selected by the user to the backend Django server. The server tallies the votes, and updates the state of wave glider based on the winning ballot. Fig. 5 shows the web-based voting interface which was used during the experiment described in Section V.

The system has two features to ensure that the platform stays within the domain of the experiment, and away from the shoreline. First, a boundary is maintained by the server as a polygon, and a check is carried our to ensure the wave glider is never commanded to move out of this region. Second, if a winning vote results in the platform hitting the boundary, all users' winning votes are redacted to *unsure* (see Fig. 6).

V. EXPERIMENT DESIGN

The experiment website designed for this project contains a realistic simulation of a nitrate plume that wells up to the ocean surface in north Monterey Bay. The plume then starts drifting westward due to ocean currents, while simultaneously losing strength and dissipating. We chose a nominal speed for the nitrate bloom to be 0.3 m/s. This process,

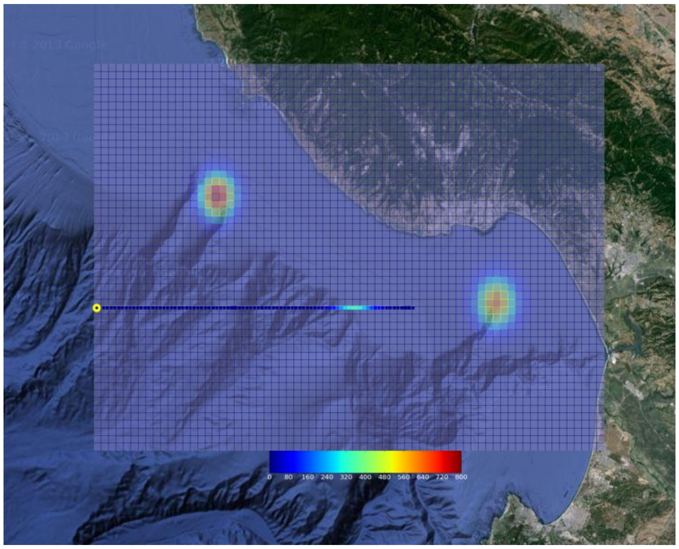

Fig. 6: Here, the WG has reached the boundary of the allowable survey region. Boundary monitoring ensures that the robot cannot be commanded to hit land or head outside the experiment domain. Winning votes that result in the boundary being hit are redacted, until the user recasts votes.

known as advection, takes approximately 48 hours, and is represented by 48 mins in the simulation. This phenomenon is repeated continuously throughout the experiment.

Participants were sourced from around the world through social media, and were instructed to steer the WG within Monterey Bay to accumulate the maximum bloom score (cumulative nitrate value). Votes from all active participants were counted every 10 seconds (equivalent to 10 experiment minutes) with the majority vote deciding the next action. The browser map in the interactive website was also refreshed every 10 seconds.

Safety measures were implemented to simulate an actual field trial. For example, we implemented a bounding region to prohibit the WG from hitting land or heading away to open ocean. If the vehicle hits this built-in boundary, any votes directing it into the boundary are changed to 'unsure' and the majority was retabulated. We assume a constant speed of 0.9 m/s for the WG throughout the experiment. The experiment started with a cumulative bloom score of 0, with the potential to achieve a maximum bloom score of 1600.

After the experiment, we analyzed the crowd commands and the reward (cumulative bloom score) achieved to determine how well the crowd was able to control the the robot to achieve the given science objective. A summary of results is presented in the following section.

VI. RESULTS

A four-hour, crowd-control experiment was carried out with multiple participants. None of the participants were informed about the detailed dynamics of the simulated nitrate plume, nor the dynamics of the WG. This trial represented over 9 hours of real experiment time. At all times, at least one person was controlling the WG, with a maximum of five participants voting for more than one continuous hour.

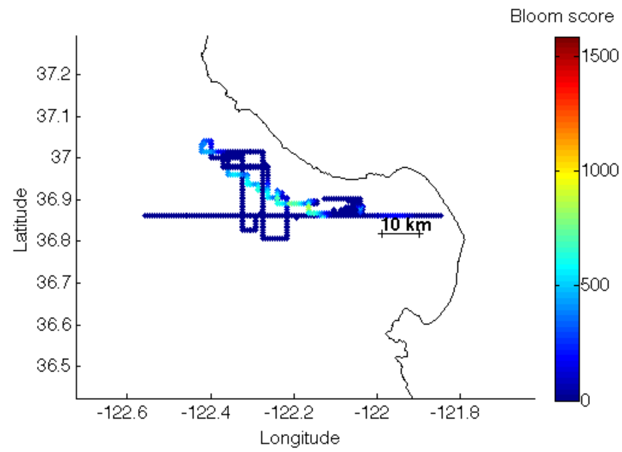

Fig. 7: Wave glider data over approximately 9 days of experiment time (224 minutes or ~ 3.7 hours in simulation time) in Monterey bay

Figure 5 displays a screenshot from when the simulation was taking place.

The participants managed to accumulate bloom scores consistently, following the plume, and returning to the location where it got upwelled at north Monterey Bay (Fig. 7). Additionally, as shown in Fig. 8, the participants demonstrated an improvement in the control of the platform during the second half of the experiment when they acquired a cumulative bloom score 1.6 times that from the first half.

VII. CONCLUSIONS

We presented the design and implementation of an online tool to control a persistent platform through crowdsourcing over long time durations. We analyzed how well a crowd can control this robot when science goals are specified. Results of a four-hour trial with multiple participants demonstrate the ability of our system to democratize the control of a robotic sampling platform to obtain scientific data.

Future work on this project is to learn wave glider control policies from crowd commands, where part of the crowd will consist of expert oceanographers. A learned policy can be used when the robot is operating without any human input (autopilot mode). Over time we aim to implement a system that weighs votes against a learned policy and known expert input to autonomously decide what action to perform.

ACKNOWLEDGMENTS

J. Das was supported in part by NOAA under award NA11NOS4780052 and NSF under award CCF-0120778. R.N. Smith was partially supported by both the Queensland Smart State Program and the Queensland University of Technology Early Career Academic Recruitment and Development (ECARD) program. We would like to thank members of the Computational Learning and Motor Control Lab and the Robotic Embedded Systems Lab for participating in the experiment.

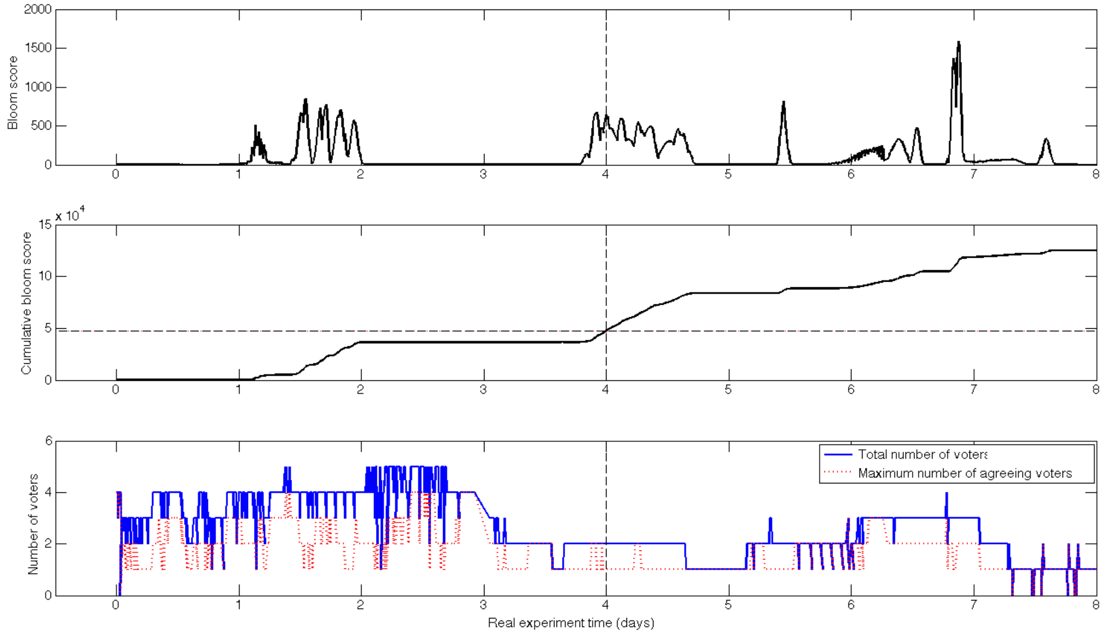

Fig. 8: Wave glider trajectory data over approximately 8 days of simulated experiment time (192 minutes or ~ 3.2 hours in simulation time). The participants performed better during the last half of the experiment when they acquired a cumulative bloom score 1.6 times that from first half.

REFERENCES

[1] Webb Research Corporation, http://www.webbresearch.com/slocum.htm, Slocum Glider, 2008, viewed July 2013.

[2] O. Schofield, J. Kohut, D. Aragon, E. Creed, J. Graver, C. Haldman, J. Kerfoot, H. Roarty, C. Jones, D. Webb, and S. Glenn, "Slocum gliders: Robust and ready," *Journal of Field Robotics*, vol. 24, no. 6, pp. 473–485, 2007.

[3] R. Hine, S. Willcox, G. Hine, and T. Richardson, "The wave glider: A wave-powered autonomous marine vehicle," in *OCEANS 2009, MTS/IEEE Biloxi - Marine Technology for Our Future: Global and Local Challenges*, oct. 2009, pp. 1 –6.

[4] J. Manley and S. Willcox, "The wave glider: A persistent platform for ocean science," in *Proceedings of the IEEE/MTS OCEANS Conference*, Sydney, Australia, May 2010, pp. 1 – 5. [Online]. Available: http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=5603614

[5] J. Das, T. Maughan, M. McCann, M. Godin, T. O'Reilly, M. Messié, F. Bahr, K. Gomes, F. Py, J. Bellingham, G. S. Sukhatme, and K. Rajan, "Towards mixed-initiative, multi-robot field experiments: Design, deployment, and lessons learned," in *International Conference on Intelligent Robots and Systems*, 2011, pp. 3132–3139. [Online]. Available: http://robotics.usc.edu/publications/719/

[6] K. Gomes, D. Cline, D. Edgington, M. Godin, T. Maughan, M. McCann, T. O'Reilly, F. Bahr, F. Chavez, M. Messié, J. Das, and K. Rajan, "Odss: A decision support system for ocean exploration," in *IEEE International Conference on Data Engineering*, 2013. [Online]. Available: http://robotics.usc.edu/publications/790/

[7] Liquid Robotics, "Pacx – the pacific crossing," 2013, [Online; accessed 25-July-2013]. [Online]. Available: \url{http://liquidr.com/pacx/pacific-crossing.html}

[8] J. Howe, *Crowdsourcing: How the power of the crowd is driving the future of business*. Random House, 2008.

[9] Wikipedia, "List of crowdsourcing projects — wikipedia, the free encyclopedia," 2013, [Online; accessed 25-July-2013]. [Online]. Available: http://en.wikipedia.org/w/index.php?title=List_of_crowdsourcing_projects&oldid=564636368

[10] Amazon, "Amazon manual turk – artificial artificial intelligence," 2013, [Online; accessed 25-July-2013]. [Online]. Available: https://www.mturk.com/mturk/welcome

[11] Wikipedia, "Crowdsourcing — wikipedia, the free encyclopedia," 2013, [Online; accessed 25-July-2013]. [Online]. Available: http://en.wikipedia.org/w/index.php?title=Crowdsourcing&oldid=565540383]

[12] S. Chernova, N. DePalma, E. Morant, and C. Breazeal, "Crowdsourcing human-robot interaction: Application from virtual to physical worlds," in *RO-MAN, 2011 IEEE*. IEEE, 2011, pp. 21–26.

[13] G. Little, L. B. Chilton, M. Goldman, and R. C. Miller, "Turkit: tools for iterative tasks on mechanical turk," in *Proceedings of the ACM SIGKDD workshop on human computation*. ACM, 2009, pp. 29–30.

[14] M. S. Bernstein, G. Little, R. C. Miller, B. Hartmann, M. S. Ackerman, D. R. Karger, D. Crowell, and K. Panovich, "Soylent: a word processor with a crowd inside," in *Proceedings of the 23nd annual ACM symposium on User interface software and technology*. ACM, 2010, pp. 313–322.

[15] J. P. Bigham, C. Jayant, H. Ji, G. Little, A. Miller, R. C. Miller, R. Miller, A. Tatarowicz, B. White, S. White *et al.*, "Vizwiz: nearly real-time answers to visual questions," in *Proceedings of the 23nd annual ACM symposium on User interface software and technology*. ACM, 2010, pp. 333–342.

[16] S. Faridani, B. Lee, S. Glasscock, J. Rappole, D. Song, and K. Goldberg, "A networked telerobotic observatory for collaborative remote observation of avian activity and range change," in *Proceedings of the 2009 IFAC Workshop on Networked Robotics (Moore KL, Ed.). International Federation of Automatic Control. Elsevier, Oxford, United Kingdom. Citeseer*, 2009.

[17] NASA, "Mars exploration program," July 2013. [Online]. Available: http://mars.jpl.nasa.gov/index.cfm

[18] S. L. Smith and D. Rus, "Multi-robot monitoring in dynamic environments with guaranteed currency of observations," in *IEEE Conference on Decision and Control*, Atlanta, GA, December 2010.

[19] S. L. Smith, M. Schwager, and D. Rus, "Persistent robotic tasks: Monitoring and sweeping in changing environments," *IEEE Transactions on Robotics*, vol. 28, no. 2, pp. 410 – 426, 2012.

[20] R. N. Smith, M. Schwager, S. L. Smith, B. H. Jones, D. Rus, and G. S. Sukhatme, "Persistent ocean monitoring with underwater gliders: Adapting sampling resolution," *Journal of Field Robotics*, vol. 28, no. 5, pp. 714 – 741, September/October 2011.