Robot Boats as a Mobile Aquatic Sensor Network

Kian Hsiang Low[†], Gregg Podnar[‡], Stephen Stancliff[‡],
John M. Dolan[‡], and Alberto Elfes[§]
Dept. Electrical and Computer Engineering[†], Robotics Institute[‡], Carnegie Mellon University
5000 Forbes Avenue Pittsburgh PA 15213
NASA Jet Propulsion Lab[§]
MS:198-235 4800 Oak Grove Drive Pasadena CA 91109
{bryanlow, gwp, stancliff, jmd}@cs.cmu.edu, elfes@jpl.nasa.gov

ABSTRACT

This paper describes the Multilevel Autonomy Robot Telesupervision Architecture (MARTA), an architecture for supervisory control of a heterogeneous fleet of networked unmanned autonomous aquatic surface vessels carrying a payload of environmental science sensors. This architecture allows a land-based human scientist to effectively supervise data gathering by multiple robotic assets that implement a web of widely dispersed mobile sensors for in situ study of physical, chemical or biological processes in water or in the water/atmosphere interface.

Categories and Subject Descriptors

I.2.9 [**Robotics**]: autonomous vehicles, sensors; I.2.11 [**Distributed Artificial Intelligence**]: intelligent agents, multiagent systems

General Terms

Algorithms, Design, Experimentation, Human Factors

Keywords

telesupervision; autonomous surface vessels (ASV); sensor web; inference grids; log-Gaussian process.

1. INTRODUCTION

Earth and ocean science research use data obtained from space, the atmosphere, the land, surface water, and the ocean to improve understanding of the Earth and its natural processes. Developing better models of ocean processes is crucial for assessing global warming and for meteorological and ecological studies, while surface water quality data give us insight into the suitability of the water for human use. Both are necessary for assessing water's ability to support aquatic life. At least three in every four humans depend on "surface water as their primary source of drinking water"[1]. Water resources in a region are critical to economic growth as well as to quality of life.

Freshwater sensing is typically done with fixed sensors near tributaries and using periodic sampling. Ocean sensing is typically done with satellites, buoys, and crewed research vessels. Satellites are limited by cloud cover, temporal and geographical coverage, and resolution. Manned research vessels are expensive to deploy, and fixed sensors and anchored buoys cannot be easily moved to cover larger areas or to monitor regions of increased interest.

Deploying mobile sensor platforms to augment fixed sensors and sensor buoys will help us better understand and model how contaminants/pollutants move and affect our environment, from industrial chemical and oil spills, to Harmful Algal Blooms (HABs). These platforms will also enable in situ meteorological studies of the atmosphere/ocean interface for hurricane prediction.

This paper describes a multi-robot science exploration software architecture and system called Multilevel Autonomy Robot Telesupervision Architecture (MARTA) [11]. MARTA was initially developed for supervisory control of multiple robot assets exploring the lunar surface [4, 8, 12]. It has subsequently been expanded and instantiated in a system called TAOSF (Telesupervised Adaptive Ocean Sensor Fleet) for coordination and control of multiple robot boat exploration vehicles for oceanic [5, 6, 13] and freshwater research [10]. For the ocean boats, we have selected as an application the characterization of HABs [1], while for the freshwater vehicles, we are concentrating on their use in water quality studies.

[1]National Research Council. Earth Science and Applications from Space: National Imperatives for the Next Decade and Beyond, 2007.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
ESSA Workshop '09, April 16, 2009, San Francisco, California, USA
Copyright 2009 ACM 978-1-60558-533-8/09/04 ...$5.00.

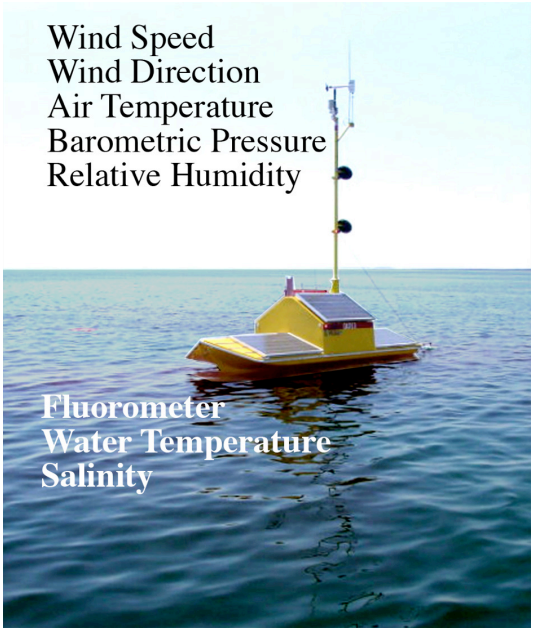

Figure 1: OASIS Platform with atmospheric and sea water sensors identified.

In the following sections, we will discuss the overall MARTA architecture, describe the robot boats being developed and used, and describe field tests conducted under controlled conditions. We conclude with an outline of the next steps in the development of MARTA.

2. ROBOT BOAT PLATFORMS

2.1 The OASIS Ocean Platforms

The OASIS (Ocean-Atmosphere Sensor Integration System) vessels are long-duration solar-powered autonomous surface vehicles (ASVs), designed for global open-ocean operations (Figure 1). Their development is funded by the National Oceanic and Atmospheric Administration (NOAA). The NOAA-funded OASIS Platform Build Team, which consists of EG&G, Zinger Enterprises, and Emergent Space Technologies, provides vehicle development, payload integration and testing, operations, and maintenance of the OASIS fleet and ground systems.

The OASIS platform is approximately 18 feet long and weighs just over 3000 lbs. The vehicle has a payload capacity of 500 lbs, and is designed to be self-righting to ensure survivability in heavy seas. It supports a wide range of communication links including spread spectrum radio, a cellular phone link, and an Iridium satellite link.

Three platforms (named OASIS-1, OASIS-2, and OASIS-3) are operated by NASA WFF and support operations for the TAOSF project. OASIS shakedown operations have been performed since early 2005 in the waters of the DELMARVA region, including the Chincoteague Bay and Pocomoke Sound. The first open-ocean deployment of the OASIS system was performed in November 2006. During this operation, the OASIS-2 platform successfully navigated over 8 nautical miles on a transect line established in the Atlantic Ocean off the coast from WFF.

Sensors onboard the OASIS platforms enable the collection of water salinity and conductivity data, sea surface temperature, and chlorophyll measurements. A rhodamine fluorometer was integrated to support mapping operations during TAOSF dye deployment tests. Dye is deployed to simulate HABs to allow development and testing of navigation and mapping algorithms. The forward payload bay provides space for installation of additional sensors. This bay includes a water flow-through system with manifolds and a de-bubbling system that simplifies installation of new sensors.

A mast-mounted meteorological station allows acquisition of atmospheric measurements, including barometric pressure, air temperature, relative humidity, wind speed, and wind direction. OASIS is also equipped with a forward-looking digital imaging system providing remotely located scientists with images of atmospheric and sea state conditions.

The off-board infrastructure developed by EST is known as the OASIS Mission Operations Environment (MOE). The MOE resides in the Wallops Coastal Ocean Observation Laboratory (WaCOOL) control room and provides applications and services that enable the WFF engineering and science operations team to perform platform commanding and telemetry monitoring, as well as communications management. The MOE also provides a middleware interface to enable other customers, such as the MARTA/TAOSF project, to integrate new systems that further enhance OASIS science operations.

In the current deployment configuration, both engineering telemetry (e.g., GPS position, roll, pitch, yaw, battery voltage) and science sensor data are communicated between each robotic platform and NASA's Goddard Space Flight Center via the Internet. It is from this distribution point that NOAA weather researchers will receive ocean and ocean/atmosphere interface sensor data.

2.2 The Robotic Sensor Boats (RSBs)

For the mobile deployment of freshwater quality sensors, a fleet of small and relatively inexpensive Robotic Sensor Boats (RSBs) is being developed at CMU (Figure 2). The design is based on commercially available components such as the hull, drive system, and communications system. Customization is kept to a minimum with the primary focus being kept on the sensing and navigation requirements.

The hull is a relatively inexpensive roto-molded recreational kayak that is fitted with a ducted thruster on

Figure 2: Robot Sensor Boat (RSB) under development.

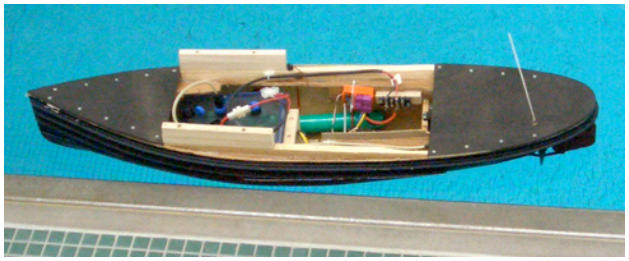

Figure 3: Initial testing of the telemetry and communications package on a small (1m length) platform.

each side. The thrusters are based on personal SCUBA propulsion units with screens to exclude foreign materials and to provide safety to humans operating in the same area. The dual thrusters allow simple differential steering and provide adequate speed when deploying each boat to its sampling area. A package of telemetry and water quality sensors is supported. A sealed spiral-cell deep-cycle marine battery provides power adequate for runs of many hours and also acts as ballast. Telemetry and science data, and well as control signals are all handled over a cellular modem. To develop and test the control and communication packages while the RSB is being built, a smaller platform was adapted (Figure 3).

3. THE MARTA ARCHITECTURE

The MARTA system architecture (Figure 4) provides

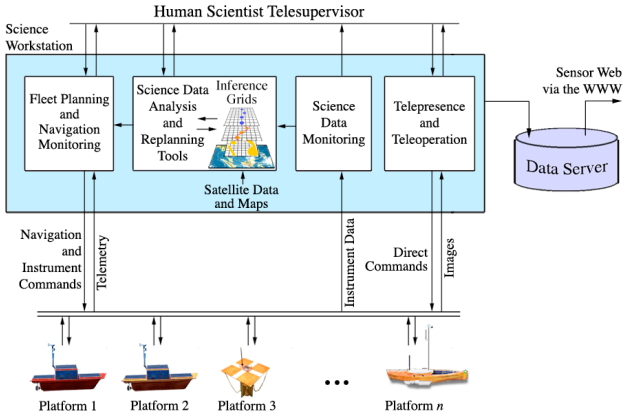

Figure 4: MARTA Telesupervision Architecture.

an integrated approach to multi-robot coordination and multi-level robot-human autonomy. It allows multiple robotic sensing assets (both mobile and fixed) to function in a cooperative fashion, and the operating mode of different robot platforms to vary from full autonomy to teleoperated control.

High-level planning and monitoring allows a human telesupervisor to assign to a fleet of robotic assets high-level goals, such as specifying an area of ocean to investigate. The steps to achieve the goals are planned, and operational commands are sent to each robot by the Fleet Planning and Navigation Monitoring. As the robots execute these plans, their operation is monitored both by the Fleet Planning and Navigation Monitoring module and by the human telesupervisor. Adaptive replanning of the robot assignments is based on sensor inputs (dynamic sensing) and coordination between multiple assets, thereby increasing data-gathering effectiveness while reducing the human effort required for tasking, control, and monitoring of the vehicles. The assets depicted in Figure 4 include two OASIS platforms, a drift buoy that reports its position as it moves with the current, and a Robot Sensor Boat for surface water operations.

Multi-level autonomy includes low-level autonomy on each independently operating robot; autonomous monitoring of the fleet; adaptive replanning; and when necessary, intervention by the human telesupervisor either with manual replanning, or by taking direct control of a robot via teleoperation.

Algorithms for science analysis of the acquired data can perform an initial assessment of the presence of specific science signatures of immediate interest. This can be done both onboard each robot and at the telesupervisor's workstation. Web-based communications support both control and communications of the robotic fleet over long distances, and the sharing of newly sensed and historical data with remote experts.

4. MAPPING HARMFUL ALGAL BLOOMS (HABS)

Interest in Harmful Algal Bloom (HAB) detection has grown in recent years for scientific, commercial public health reasons. Depending on the type of algae present, HABs have been shown to be dangerous to sea life and to human health.

There is a significant interest in identifying environmental factors that contribute to the occurrence of HABs, so that these may be incorporated in bloom prediction algorithms. These factors may include time of year, salinity, and sea-surface temperature to predict the abundance (low, medium, or high) and type of HABs.

MARTA will provide the following advantages over existing systems for observing and analyzing HABs:

- Dynamic tasking and adaptation;
- Higher in situ resolution and greater insensitivity to cloud cover in comparison to current satellite systems;
- Access to and greater agility in coastal waters than what is available through buoys;
- Real-time multipoint science data observations and generation of associated interpretations by remotely located oceanographers.

Since HABs are sporadic ocean phenomena, we are using rhodamine WT (water-tracing) dye as a HAB surrogate for initial experiments and tests.

For investigating an algal bloom, researchers at Carnegie Mellon University take control of the OASIS platforms and provide high-level planning and monitoring. The OASIS collection of weather-related data mentioned earlier is not interrupted, but the additional sensing for investigating the algal bloom is brought online to map the extent of a bloom, chlorophyll concentrations, and with additional sensors, eventually allow a shore-based biologist the ability to assess whether the algal species are harmful. If they are, local authorities can be notified at aquaculture businesses, fisheries, and beaches.

5. FIELD TESTS

5.1 Field Test 1 - Aug. 2007

For controlled experiments such as mapping while compensating for drift currents, geometric patches of rhodamine water-tracing dye are laid (Figures 5 and 6). An aerostat, tethered to a manned tender boat, carries an instrument package aloft including GPS, altimeter, compass, and video camera to provide validation images of data gathered by OASIS platforms that are mapping with fluorometers.

Methodology: Inference Grids. Observation of natural processes is limited by spatial and temporal sensor footprint, coverage, resolution, sampling rate, and measurement uncertainty. Markov Random Fields (MRFs) provide a natural formulation to represent spatially and temporally distributed observations. We use a discretized version of MRFs, called a spatio-temporal Markov Random Lattice (ST-MRL), to encode the data obtained by the different sensors and agents. Each cell in the lattice corresponds to a spatial volume and a time slice, and stores a stochastic vector with the probabilistic state description of the various processes that have been measured at the given location and time interval. Efficient updating methods are used to improve the probabilistic description encoded in the lattice as new observations flow in from the various sensors and platforms being used [2].

Associated with the ST-MRL we also maintain additional stochastic lattice-based layers for inference and decision, which are used to plan and control the activities of the robot platforms [3]. These layers include

Figure 5: OASIS-2 and OASIS-3 mapping a patch of rhodamine water tracing dye.

Figure 6: Aerial view from the aerostat observation platform. The chase boat (with aerostat tether) is in the upper part of the image, the OASIS boat in the lower part, and the rhodamine dye tracks laid by the chase boat to serve as surrogate HABs are towards the right part of the image.

Figure 7: A slanted view of the ocean area where the test was conducted. A spiral search pattern was executed by the boat. The measured rhodamine concentrations are shown as a vertical "ribbon" along the route taken by the OASIS boat.

vehicle navigation cost and risk to reach an area of interest; hypotheses of scientific events to be explored further; information metrics such as entropy to determine how the knowledge of a natural process is evolving, and where critical information is missing; and others. The augmented informational structure that incorporates both the ST-MRL and the information-theoretic inference and decision layers is called an Inference Grid (IG) [3, 7].

In the MARTA/TAOSF system, we use the Inference Grid (IG) model to represent multiple spatially- and temporally-varying properties, concerning both ocean processes and HABs. The rhodamine dye concentration measurements taken by the OASIS platforms during field tests are used as input to the Inference Grid mapping process. The fluorometer measurements are used to compute the presence or absence of dye for each cell in the area traversed by each OASIS boat. The probabilistic sensor model was derived from information on the sensitivity and performance of the sensors. The Inference Grid map of rhodamine dye for one of our mapping experiments (Figure 5) is shown in Figure 8.

Results and Analysis. Figure 6 shows an aerial view from the aerostat of the test area in the Chesapeake Bay; the OASIS robot boat platform is in the lower part of the image, close to the rhodamine dye tracks that serve as a surrogate for algal blooms during field testing. Figure 7 shows the search pattern executed by the OASIS boat to find the surrogate algal bloom, and Figure 8 displays an Inference Grid showing the areas with high probability of dye presence (in red), high probability of dye absence (in green), and high entropy or lack of information (in grey). Finally, in Figure 9, an

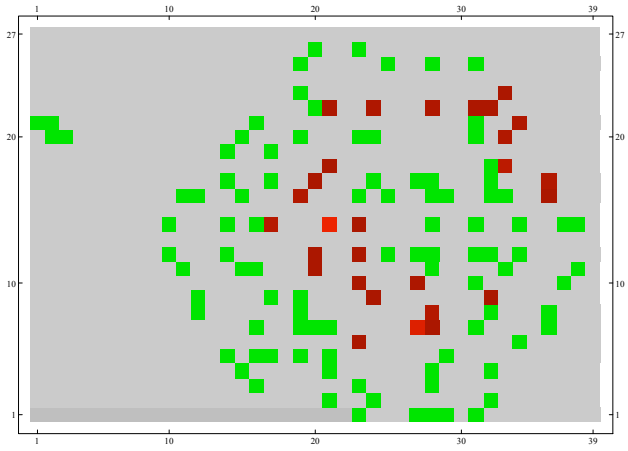

Figure 8: Inference Grid showing the areas with high probability of dye presence (in red), high probability of dye absence (in green), and high entropy or lack of information (in grey).

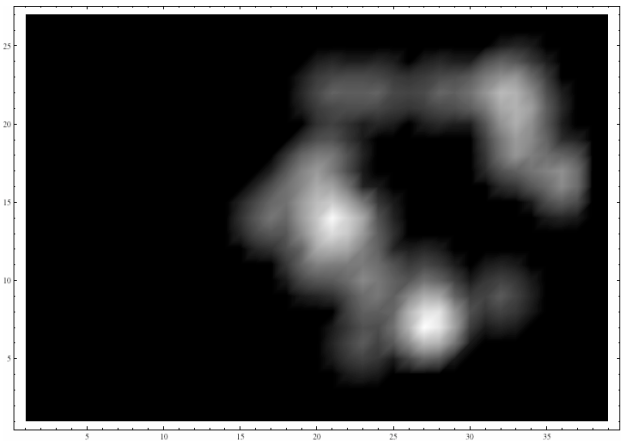

Figure 9: Inference Grid with inferred hypotheses of algal blooms (dye tracks).

Inference Grid with inferred hypotheses of algal blooms (dye tracks) is presented. The areas of high entropy can be used to replan the search pattern of the boat.

5.2 Field Test 2 - Jul. 2008

In this field test, the rhodamine WT dye is laid down in two stripes. The OASIS-2 boat follows a raster scan, which progresses along the dye stripes. The OASIS-2 boat's trajectory in global (GPS) coordinates (i.e., with added drift) is *drift-corrected* to produce a corresponding trajectory in (water) surface coordinates (i.e., without drift). From Figure 10, we can observe that the latter drift-corrected trajectory stretches over a shorter distance (i.e., 233m) as compared to that of the former trajectory (i.e., 380.9m). With drift correction, the dye stripes do not "run away" from the boat.

Our aim is to reconstruct the (drift-corrected) rhodamine WT dye map (i.e., in surface coordinates) using the sparse dye data collected along OASIS-2 boat's tra-

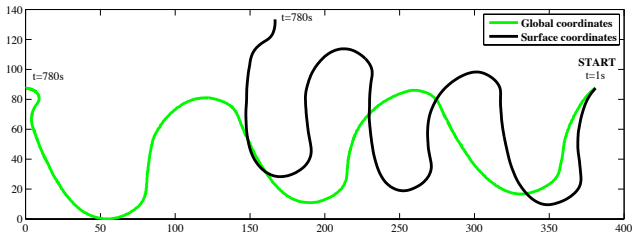

Figure 10: Comparison of OASIS-2 boat's trajectories with and without added drift (respectively, green and black paths). The units for the axes are in meters.

jectory over a 380.9m×87.5m sampling region (i.e., 780 fluorometer readings over a period of 13 minutes). In particular, we want to investigate if the dye-stripe phenomenon can be inferred from the reconstructed map.

Methodology: Log-Gaussian Process. The dye data (see Figure 11) collected by OASIS-2 boat is characterized by continuous, spatially correlated, positively skewed (Figure 12) fluorometer measurements. Using this sparse dye data, we will reconstruct the dye map in surface coordinates, which spans 233m by 102.4m. The properties of the fluorometer measurements suggest the use of the non-parametric probabilistic inference model called the *log-Gaussian process* (ℓGP). The advantage of this modeling technique is that it does not require any assumptions on the distribution underlying the sampling data. Furthermore, the prediction of fluorometer measurements at unobserved locations can be performed at any desired map resolution. We will now describe the log-Gaussian process briefly and refer the interested reader to [9] for more details.

Let $\mathcal{X}$ be the domain of the dye map corresponding to a finite, discretized set of locations. Each location $x \in \mathcal{X}$ is associated with an observed fluorometer measurement y_x and its corresponding random counterpart Y_x. Let $\{Y_x\}_{x\in\mathcal{X}}$ denote a ℓGP defined on the domain $\mathcal{X}$. That is, if we let $Z_x = \log Y_x$ (i.e., $Y_x = \exp\{Z_x\}$), then $\{Z_x\}_{x\in\mathcal{X}}$ is a *Gaussian process* (GP). This means the joint distribution over any finite subset of $\{Z_x\}_{x\in\mathcal{X}}$ is Gaussian. The GP can be completely specified by its (prior) mean and covariance functions

$$\mu_{Z_x} \stackrel{\triangle}{=} \mathbb{E}[Z_x] \ ,$$

$$\sigma_{Z_x Z_u} \stackrel{\triangle}{=} \mathrm{cov}[Z_x, Z_u]$$

for $x, u \in \mathcal{X}$. Let the data d denote n pairs of sampled locations and their corresponding observed measurements. Let $\mathbf{x}$ and $\mathbf{z_x}$ denote vectors comprising the location and measurement components of the data d. If the data d are available, the distribution of Z_x remains a Gaussian with the posterior mean and variance

$$\mu_{Z_x|d} \stackrel{\triangle}{=} \mathbb{E}[Z_x \mid d] = \mu_{Z_x} + \Sigma_{x\mathbf{x}}\Sigma_{\mathbf{xx}}^{-1}\{\mathbf{z_x}^\top - \boldsymbol{\mu}_{\mathbf{Z_x}}\} \quad (1)$$

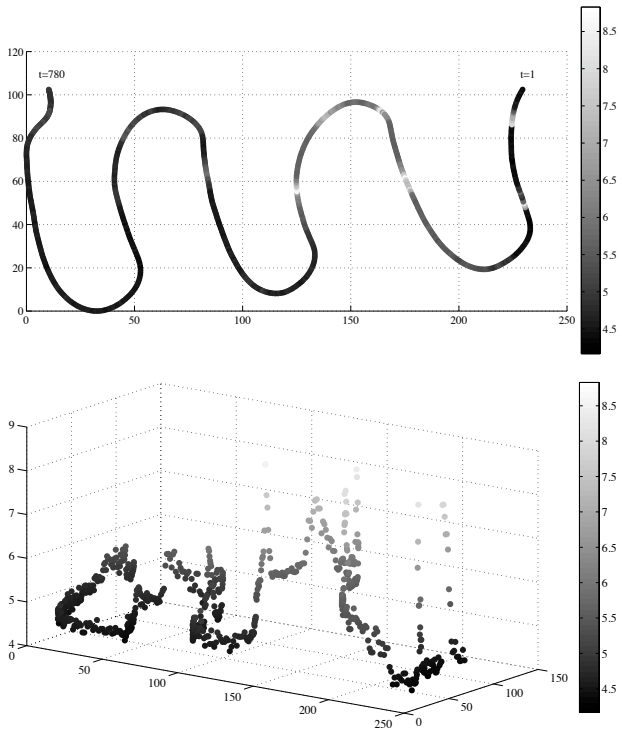

Figure 11: Continuous, spatially correlated fluorometer readings (measured in mg/m^3) of the drift-corrected trajectory (measured in meters).

$$\sigma^2_{Z_x|d} \stackrel{\triangle}{=} \mathrm{var}[Z_x \mid d] = \sigma_{Z_x Z_x} - \Sigma_{x\mathbf{x}}\Sigma_{\mathbf{xx}}^{-1}\Sigma_{\mathbf{x}x} \quad (2)$$

where, for the location components v, w in $\mathbf{x}$, $\boldsymbol{\mu}_{\mathbf{Z_x}}$ is a column vector with mean components μ_{z_v}, $\Sigma_{x\mathbf{x}}$ is a covariance vector with components $\sigma_{z_x z_v}$, $\Sigma_{\mathbf{x}x}$ is the transpose of $\Sigma_{x\mathbf{x}}$, and $\Sigma_{\mathbf{xx}}$ is a covariance matrix with components $\sigma_{z_v z_w}$. Note that the posterior mean $\mu_{Z_x|d}$ (1) is the best unbiased predictor for predicting the log-measurement z_x at the unobserved location x. An important property of the Gaussian posterior variance $\sigma^2_{Z_x|d}$ (2) is its independence of $\mathbf{z_x}$.

The ℓGP has the (prior) mean and covariance function

$$\mu_{Y_x} \stackrel{\triangle}{=} \mathbb{E}[Y_x] = \exp\{\mu_{Z_x} + \sigma_{Z_x Z_x}/2\} \ ,$$

$$\sigma_{Y_x Y_u} \stackrel{\triangle}{=} \mathrm{cov}[Y_x, Y_u] = \mu_{Y_x}\mu_{Y_u}(\exp\{\sigma_{Z_x Z_u}\} - 1)$$

for $x, u \in \mathcal{X}$. We know that the distribution of Z_x given the data d is Gaussian. Since the transformation from $\mathbf{z_x}$ to $\mathbf{y_x}$ is invertible, the distribution of Y_x given the data d is log-Gaussian with the posterior mean and variance

$$\mu_{Y_x|d} \stackrel{\triangle}{=} \mathbb{E}[Y_x \mid d] = \mathbb{E}[\exp\{Z_x\} \mid d] = \exp\{\mu_{Z_x|d} + \sigma^2_{Z_x|d}/2\} \quad (3)$$

$$\sigma^2_{Y_x|d} \stackrel{\triangle}{=} \mathrm{var}[Y_x \mid d] = \mu^2_{Y_x|d}(\exp\{\sigma^2_{Z_x|d}\} - 1) \quad (4)$$

where $\mu_{Z_x|d}$ and $\sigma^2_{Z_x|d}$ are the Gaussian posterior mean (1) and variance (2) respectively. Note that the log-

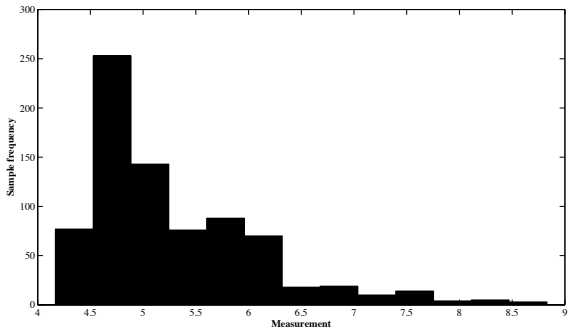

Figure 12: Positively-skewed fluorometer readings (mg/m^3): The positive skew results from a small number of high measurements and a huge pool of low measurements.

Gaussian posterior mean $\mu_{Y_x|d}$ (3) is the best unbiased predictor for predicting the actual measurement $y_x = \exp\{z_x\}$ at the unobserved location x. On the other hand, $\exp\{\mu_{Z_x|d}\}$ is a biased predictor of the actual measurement $y_x = \exp\{z_x\}$. This can be observed from (3) that $\mu_{Y_x|d} \geq \exp\{\mu_{Z_x|d}\}$.

Results and Analysis. The log-Gaussian process is used to predict the fluorometer measurements at $233 \times 103 = 23999$ unobserved locations that are evenly spaced throughout the sampling region in surface coordinates (Figure 13). Figure 13a shows the predicted map (3) while Figure 13b shows the variance (4) associated with the predicted locations. From Figure 13a, we can observe two roughly parallel dye stripes running from the right with higher concentration to the left with lower concentration. The left part of the dye stripes appears to be more diffused, thus resulting in a smaller inter-stripe gap. From the lengthscale hyperparameters of the log-Gaussian process (obtained using maximum likelihood estimation), we also learn that there is a much higher degree of spatial correlation between fluorometer measurements along the horizontal axis than along the vertical axis. This allows the two dye stripes to be inferred relatively well.

Using the predicted dye map (Figure 13a), we further perform binary classification (Figure 13c) via column-wise thresholding to determine if each map location is in a dye stripe or not. From Figure 13, we can see that the top dye stripe is better identified than the bottom one. Figure 14 shows the row-wise profile of the predicted dye map (Figure 13a); each circle denotes the mean predicted fluorometer measurement for a particular row (≈ 0.9945m wide). If we set the threshold to the mean fluorometer measurement of the predicted dye map (i.e., ≈ 5.3188mg/m^3), Figure 14 then shows that the top and bottom dye stripes in Figure 13 are approximately 15.9m and 22.9m wide, respectively, with an inter-stripe gap of roughly 9.9m.

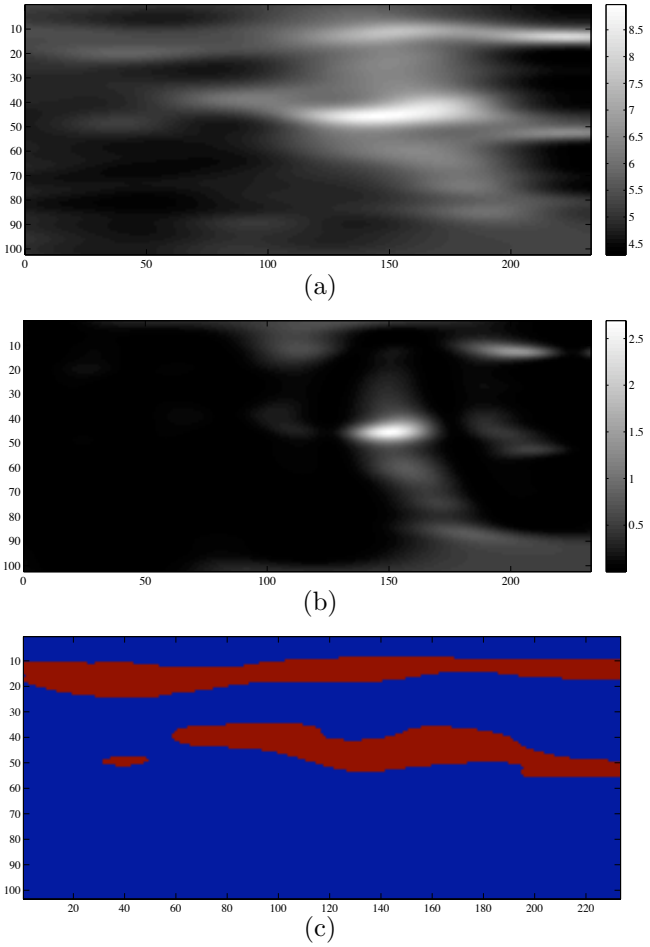

Figure 13: Map inference of rhodamine dye concentration: (a) Prediction, (b) variance, and (c) binary classification via thresholding. The units for the axes are in meters.

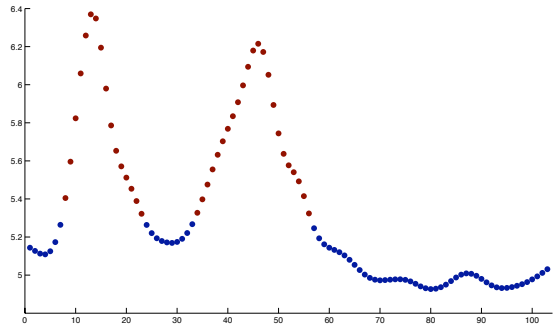

Figure 14: Row-wise profile showing the plot of mean predicted fluorometer measurement (mg/m^3) vs. row number. The red circles are above the threshold, which is set to the mean fluorometer measurement of the predicted dye map. The blue circles are below the threshold.

6. CONCLUSIONS

The MARTA-based TAOSF multi-level autonomy control architecture provides many advantages over existing systems for observing and analyzing HABs including: dynamic tasking and adaptation; higher in situ resolution and greater insensitivity to cloud cover (as compared with satellite systems); access to and greater agility in coastal waters than that available through buoys; real-time multipoint science data observations and generation of associated interpretations by remotely located oceanographers.

By using the TAOSF architecture, it increases data-gathering effectiveness and science return while reducing demands on scientists for robotic asset tasking, control, and monitoring. The data are also made available to other scientists via the world-wide-web both as they are collected, and from an historical data archive server.

Acknowledgements

This work was supported by NASA award NNX06AF27G, "Telesupervised Adaptive Ocean Sensor Fleet", granted under the Advanced Information Systems Technology program of NASA's Earth Systems Technology Office (ESTO). The TAOSF project is a collaboration among Carnegie Mellon University (CMU), NASA Goddard Space Flight Center (GSFC), NASA Goddard's Wallops Flight Facility (WFF), Emergent Space Technologies, Inc. (EST), and the Jet Propulsion Laboratory (JPL). Work on the OASIS platforms is conducted by Emergent Space Technologies, Inc., EG&G, and Zinger Enterprises under award NA03NOS4730220 from the National Oceanic and Atmospheric Administration (NOAA), U.S. Department of Commerce.

7. REFERENCES

[1] J. M. Dolan, G. W. Podnar, A. Elfes, S. Stancliff, E. Lin, J. Higinbotham, J. C. Hosler, J. Moisan, and T. A. Moisan. Smart ocean sensing using the telesupervised adaptive ocean sensor fleet system. In *Proc. ESTC*, 2008.

[2] A. Elfes. Multi-source spatial fusion using Bayesian reasoning. In M. A. Abidi and R. C. Gonzalez, editors, *Data Fusion in Robotics and Machine Intelligence*. Acad. Press, 1992.

[3] A. Elfes. Robot navigation: Integrating perception, environmental constraints and task execution within a probabilistic framework. In L. Dorst, M. van Lambalgen, and F. Voorbraak, editors, *Reasoning with Uncertainty in Robotics*, pages 91–130, LNCS 1093, Springer, Berlin, 1996.

[4] A. Elfes, J. M. Dolan, G. W. Podnar, S. Mau, and M. Bergerman. Safe and efficient robotic space exploration with tele-supervised autonomous robots. In *Proc. AAAI Spring Symposium,* Technical Report SS-06-07, pages 104–113, 2006.

[5] 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, and E. A. Kulczycki. The telesupervised adaptive ocean sensor fleet (TAOSF) architecture: Coordination of multiple oceanic robot boats. In *Proc. IEEE Aerospace Conference*, 2008.

[6] A. Elfes, G. W. Podnar, J. M. Dolan, S. Stancliff, E. Lin, J. C. Hosler, T. J. Ames, J. R. Moisan, T. A. Moisan, J. Higinbotham, and E. A. Kulczycki. The telesupervised adaptive ocean sensor fleet. In *Proc. SPIE Conference on Atmospheric and Environmental Remote Sensing Data Processing and Utilization III: Readiness for GEOSS*, volume 6684, 2007.

[7] A. Elfes, G. W. Podnar, R. F. Tavares Filho, and A. Pavani Filho. Inference grids for environmental mapping and mission planning of autonomous mobile environmental robots. In *Proc. Workshop Sensing a Changing World*, 2008.

[8] E. Halberstam, L. Navarro-Serment, R. Conescu, S. Mau, G. W. Podnar, A. D. Guisewite, H. B. Brown, A. Elfes, J. M. Dolan, and M. Bergerman. A robot supervision architecture for safe and efficient space exploration and operation. In *Proc. 10th Biennial Int'l Conf. on Engineering, Construction, and Operations in Challenging Environments (Earth & Space 2006)*, 2006.

[9] K. H. Low, J. M. Dolan, and P. Khosla. Adaptive multi-robot wide-area exploration and mapping. In *Proc. AAMAS*, pages 23–30, 2008.

[10] G. W. Podnar, J. M. Dolan, and A. Elfes. Networked architecture for robotic environmental ocean science sensors. In *Proc. Workshop Sensing a Changing World*, 2008.

[11] G. W. Podnar, J. M. Dolan, A. Elfes, and M. Bergerman. Multi-level autonomy robot telesupervision. In *Proc. ICRA 2008 Workshop on New Vistas and Challenges in Telerobotics*, 2008.

[12] G. W. Podnar, J. M. Dolan, A. Elfes, M. Bergerman, H. B. Brown, and A. D. Guisewite. Human telesupervision of a fleet of autonomous robots for safe and efficient space exploration. In *Proc. HRI*, pages 325–326, 2006.

[13] G. W. Podnar, J. M. Dolan, A. Elfes, S. Stancliff, E. Lin, J. C. Hosler, T. J. Ames, J. R. Moisan, T. A. Moisan, J. Higinbotham, and E. A. Kulczycki. Operation of robotic science boats using the telesupervised adaptive ocean sensor fleet system. In *Proc. ICRA*, pages 1061–1068, 2008.