Sensor Hysteresis Compensation using Adaptive Sampling

ABHINAV VALADA, The Robotics Institute, Carnegie Mellon University

Collecting samples from water bodies is not only important for understanding the physiology of aquatic life but also for understanding how these systems are affected by both natural changes in the environment as well as human activities. However collecting spatially distributed samples from the environment can be an exhaustive and complex task. First, for the data to be valid, the samples have to be collected in a short interval of time before the environment changes significantly, hence coordinated robot teams have to be used to accomplish this task. Secondly, many environmental sensors that are used have significant hysteresis in them, making sampling using them a very slow process. Rate dependent hysteresis can be observed in several sensors and poses a very serious problem to autonomous sampling where a robot continuously collects measurements while moving through a large field to create models of a physical parameter. This phenomenon cannot be ignored in dynamic sensing applications as the lag between the input and output causes a delay in the responsiveness of the sensor: for a change in input, the output of the sensor slowly and consistently approaches the actual value. In this research, I propose an adaptive sampling solution to this problem, where robots can be used to collect samples intelligently, while compensating for the hysteresis using a filter and a complementary RRT based planner that accounts for this effect.

Categories and Subject Descriptors: **I.2.9 [Robotics]**: Sensors

General Terms: Design, Algorithms, Sensing

Additional Key Words and Phrases: Adaptive sampling, autonomous mapping, field robotics, environmental data, hysteresis compensation

1. INTRODUCTION

Recent advances in Autonomous Surface Vehicle (ASV) technology have enabled these systems to be used in missions that involve sampling large bodies of water for extended periods in order to monitor dynamic spatial and temporal phenomena with little or no human supervision. Monitoring water bodies is not only important for understanding the physiology of aquatic life but also for understanding how these systems are a affected by both natural changes in the environment such as storms and volcanic eruptions as well as human activities such as surface run o from farms and industrial discharges. By collecting spatially distributed samples and analysing the data it may be possible to predict how some of these processes work and potentially prevent adverse ecological effects such as eutrophication, oxygen depletion, and accelerated aging. ASVs are a natural choice for this kind of application as they have the capability to sample large areas while providing real-time measurements. They have been successfully used for mapping applications both above and below the water surface [J. C. Leedekerken et al, 2010], even at varying depths [Hitz. G. et al, 2012]. Cooperative fleets of ASVs have advantages over a single ASV in reliability, coverage and fault tolerance. Moreover intelligent sampling techniques can greatly improve the efficiency and quality of sampling by adaptively determining the next sampling locations based on the previously measured data.

In this report, I describe my work on developing an algorithm to adaptively sample dissolved oxygen and temperature using autonomous boats. Initial experiments with autonomous sampling revealed that slow response of the polyethylene membrane used in dissolved oxygen sensors causes rate-dependent hysteresis, which significantly affects measurement accuracy. Similarly, with temperature sensors hysteresis is often caused by the introduction of some amount of strain or moisture penetrating inside the sensor. As these sensors are not specifically designed for dynamic measurements, a lag in the response of the sensor causes erroneous measurements if the vehicle travels at a rate that does not allow the sensor readings

to stabilize. Developing a sensor model that predicts the rate of change is one potential solution, but even small errors in this estimate can dramatically affect the final measurement value. In my work I adopt an approach in which I use the time derivative of measurements, rather than the measurements themselves, to alter a pre-defined set of bounds that converge to the true value over time. The vehicle then plans its path based on the expected sensor value, range of the bounds, and the frequency with which the cell has been previously visited. To validate this approach I compare my algorithm to a suite of other sampling algorithms including random walk and lawn-mower pattern.

The rest of the report is organized as follows. In Section 2, I formulate the problem in question and identify the associated related work in Section 3. I outline my bounding filter solution for modelling the environment in Section 5 and apply it to the problem of planning in Section 6. In Section 7, I describe the results obtained from comparison of my bounded filter approach with other standard algorithms and subsequently analyse the performance. Finally, in Section 8, I conclude the report with a summary of the results and an outline of future work directions.

2. PROBLEM FORMULATION

For modeling purposes, the dissolved oxygen sensor by Atlas Scientific and the temperature sensor by Decagon Devices were chosen. Several samples were collected and the response of the sensor was recorded for different water samples. The response of the sensor was similar to a trend such as an exponential moving average. Hence we can model the sensor with hysteresis as a process with internal state, based on a exponential moving average.

$$s(t+1) = \alpha^t \psi(x) + (1 - \alpha^t)\, s(t)$$

In continuous time, this corresponds to a first order differential process.

$$\frac{\partial s}{\partial t} = \log(\alpha)\,(\psi(x) - s(t))$$

The objective is to estimate the value of $\psi(x)$ at some set points of time. An easy way of doing this is to simply visit each point and wait for a certain amount of time for the sensor output s to approach $\psi(x)$. However as α is small, $\frac{\partial s}{\partial t}$ will also be small, therefore it can take a very long time for s to approach $\psi(x)$. A sample trend from the dissolved oxygen sensor is shown in Figure 1.

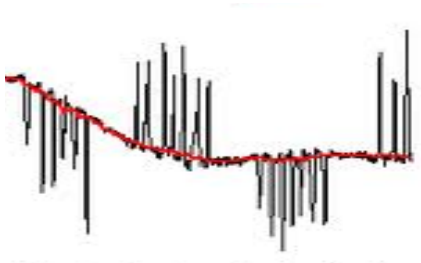

Fig. 1. Sample trend obtained from the dissolved oxygen sensor.

3. RELATED WORK

Several intelligent sampling strategies have been developed for autonomous vehicles that aim to identify hotspots, reduce resource costs, optimize sampling coverage or more accurately measure environmental phenomena. The authors in [Zhang. Bin et al, 2009] explored a sampling technique, using both a team of robotic boats and static sensor nodes, in which the sensing field is partitioned into sub regions either according to equal gains or equal area and each boat is assigned a specific subarea for sampling. The readings are gathered from the static nodes and paths are computed for the boat such that they reduce the integral mean squared error. In [Dan O. Popa et al, 2004], the authors explore an approach based on model parameter estimation of a variable in which the physical parameter being measured is assumed to be linearly distributed across the field and the algorithm aims to minimize the measured uncertainty in the field distribution. The algorithm also has multiple secondary objectives such as to minimize the network utility of multiple AUVs by controlling the sampling location and sampling rate using a potential function that encapsulates the network model and minimizing the energy consumption by varying the speed of the vehicle according to the energy available.

Thermoclines are believed to be an important breeding zone for marine microorganisms and hence a considerable amount of work has been done on thermocline detection and monitoring using sensor networks, gliders and other AUVs [Cruz, N.A. et al, 2010 and S. Petillo et al, 2010]. Zhang in [Bin Zhang et al, 2004] used a wireless sensor actuator network and a robot mule to detect thermoclines using distributed binary search. In this algorithm the nodes were assigned regions to sample and could move vertically by altering their buoyancy. Each node first localized the temperature variation in its own region, then combined this data with that of children nodes, forwarded it to the parent node and so on, until the final bulk data was transferred to the user. They further improved the performance of the algorithm by using a mobile robot to collect data from an active node and communicate it to another.

Sampling of Phytoplankton has also gained popularity in recent years as it plays a very important role in ocean ecology. In [Yanwu Zhang et al, 2009] the authors used a Dorado AUV to describe a method to detect and collect water samples at peak chlorophyll fluorescence, taking into consideration the delay in measurement while detecting the peak. The AUV followed a yo-yo pattern and used gradient following to detect a peak in the ascend stage and successfully collected the peak chlorophyll fluorescence sample at the same depth in the descend cycle.

4. TEST PLATFORM

The algorithm discussed in this report was designed and implemented on the Cooperative Robotic Watercraft (CRW) platform [A. Valada et al, 2012]. CRW is a multi-robot autonomous surface vehicle, equipped with an Android smartphone that provides the inertial sensors and computing platform for the system. The CRW's design is similar to that of an airboat with a modified steering mechanism in which the entire propulsion assembly is actuated using servo motors, allowing for improved thrust vectoring, which enables sharper turns. The drive system and other electronics are interfaced to an Arduino microcontroller that communicates with the

smartphone via Bluetooth. Most of the autonomy software resides on the phone while some of the application specific intelligence, such as the sampling algorithms, are implemented on a centralized operator interface that interacts with the individual vehicles via 3G or WiFi.

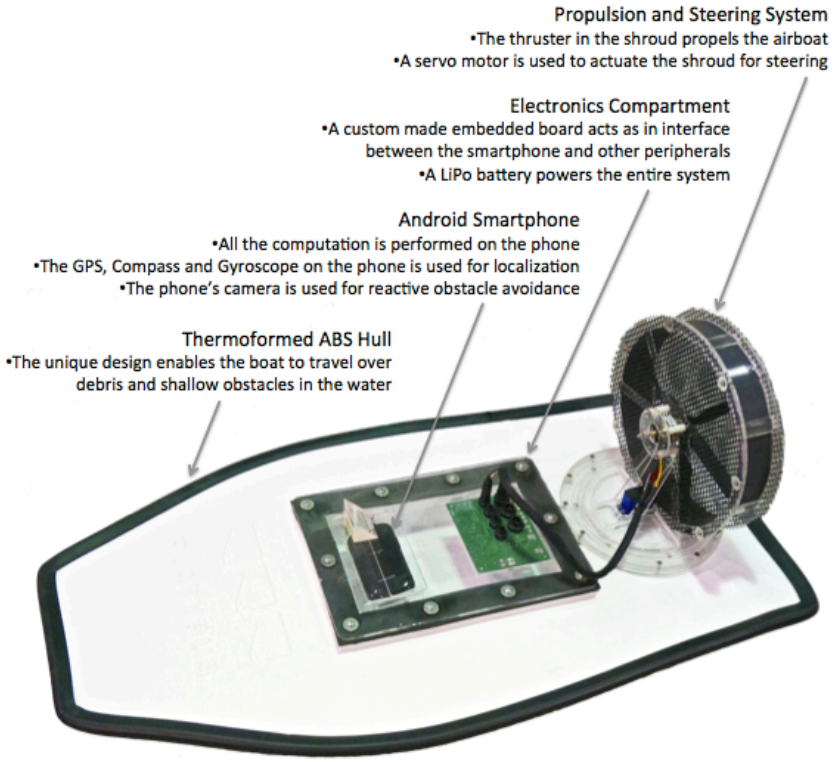

Fig. 2. Lutra 1.0, Cooperative Robotic Watercraft

Water quality sensors such as dissolved oxygen, temperature, specific conductivity and pH, are mounted on the vehicle and interfaced to the system through the Arduino. The camera on the smartphone provides real-time situational awareness about the operating environment using a steady stream of images that are processed through an image queue and displayed on the operator interface. A water sampling mechanism on board each vehicle has the ability to collect physical samples on demand for more detailed analysis in the laboratory. A diagram depicting the CRW is shown in Figure 2.

5. BOUNDED FILTER

Rate dependent hysteresis can be observed in several sensors and poses a very serious problem to autonomous sampling where a robot continuously collects measurements while moving through a large field to create models of a physical parameter. This phenomenon cannot be ignored in dynamic sensing applications as the lag between the input and output causes a delay in the responsiveness of the sensor: for a change in input, the output of the sensor slowly and consistently approaches the actual value. This effect is not as important in static sensing applications, where the rate of change of the physical parameter in the field is much slower than the hysteresis in the sensor. However, in our case the watercraft

traverses through the water while simultaneously taking measurements, making compensating for the hysteresis effect critical. Within the suite of sensors on the CRW platform, I observed this effect significantly in both temperature and dissolved oxygen.

I propose an intelligent sampling solution to this hysteresis compensation problem in the form of a filter that accounts for this effect. Rather than recording the sensed value at a location, I maintain an upper and lower bound on the predicted value in each area and use the direction of change in the sensor measurement to adjust the bounds. For example, if the gradient is trending downwards (as shown in Figure 3), the actual value must be lower than the value reported by the sensor, hence the upper bound can be adjusted. The inverse holds when the values are trending upwards, allowing the lower bound to be increased.

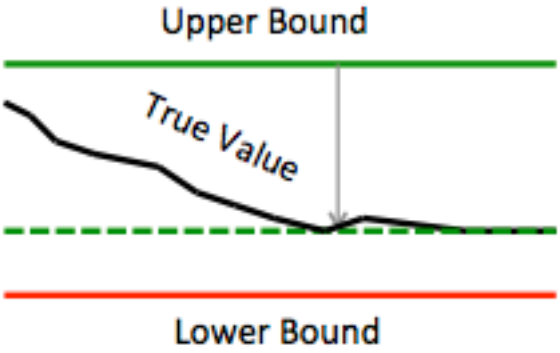

Fig. 3. Illustration of the bounding technique.

While the gradients tend to be consistent and reasonably noiseless, I use a median filter to remove occasional noise in sensor measurements. The median value over a window of readings is computed and linear regression is used to find a gradient across a larger window over the median filtered values. The gradient is then used to change either the upper or lower bound in the area if it has an absolute value above a constant defined as ϵ. Based on my initial experiments with water quality sensors exhibiting hysteresis I made two practical design decisions on the filter. Firstly, a zero gradient is the most useful gradient as it could be used to bring the upper and lower bound to the current value, since the sensor must be at the true value. However, in practice I found that this was misleading in view of discretization in the sensor output, as the gradient might appear to be zero even when it is not. Therefore I choose to ignore gradients of zero, though these will be considered in future work. Conversely, at times a sensor measurement can oscillate between two discrete levels causing an apparent gradient even when there is none. To avoid the impact of these erroneous readings on the filter, I made another design decision that requires the gradient be above ϵ, which filters out very small gradients due to oscillations in the sensor output. ϵ can be determined by analyzing a small data set obtained from a sample run. The pseudo code for the filtering process is shown below.

ALGORITHM 1. Bounded Filter

Function PROCESS(v)

```
        Data ← v_data
        Windows ← MedianFilter(Data)
        gradient ← LinearRegression(Windows)
        cell ← CellFor(v_position)
        if gradient < ϵ then
                if cell_lower < v_data then
                        cell_lower ← v_data
                end if
        else if gradient < ϵ then
                if cell_upper > v_data then
                        cell_upeer ← v_data
                end if
        end if
end function
```

6. BOUNDED PLANNER

Planning for information collection is a intriguing problem that has been extensively studied in recent years [Kian Hsiang Low et al, 2011]. Adopting the bounding filter opens up new challenges and opportunities for developing a planning algorithm. The sensor measurement with which the watercraft enters a cell is important as a change in the measurement will in turn affect the bounds of that cell. Closer the measurement with which the watercraft enters a cell is to the mid-point of bounds of the current cell, the more valuable the collected data is likely to be, since a change in the measurement will lead to the biggest expected change in the filter. Planning using the bounding filter is a challenge as it needs to take into account what measurement the sensor might output along the path, even though it can only be estimated based on the current value of the upper and lower bounds of a cell.

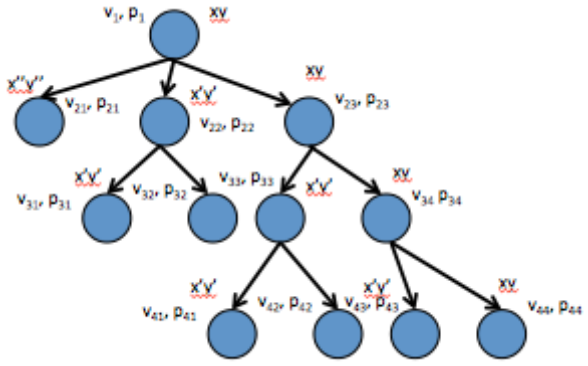

Fig. 4. Illustration of the planning tree.

I implemented an RRT based planner with horizon bounds. In this approach a tree is used to expand to the most promising nodes. As paths in the tree are expanded, an estimate of the sensor value is maintained as the mid-point between the upper and lower bounds of the filter, which is then used as the expected value of the sensor as the watercraft leaves the cell. The value of going into a cell is estimated as a function of how far the expected sensor measurement with which the watercraft enters the cell is from the mid-point between the upper and lower bounds of the cell, current difference between the upper and lower bounds and the number of times that cell has

been previously visited on the path. The value of expanding a particular node in the search tree is a heuristic based on the current expansion depth and the rate that value has accumulated on the path so far. The aim of the heuristic is to encourage exploration of the paths that have the highest value. The pseudo code for my planner is given below.

ALGORITHM 2. Bounded Planner

```
Function BOUNDED PLANNER
        n.loc ← currentLocation
        n.value ← 0
        n.sensor ← currentSensorValue
        quwuw.add(n)
        while expansions < maxExpansions do
                n ← queue.poll
                for e ← getExpansions(n) do
                        e.value ← n.value + max(0.0, (e.cell_upper – e.cell_lower)) - |((e.cell_upper +
                        e.cell_lower)/2) - n.sensor)
                        expectedSensorChange ← min((e.cell_upper + e.cell_lower)/2 – n.sensor,
                        maxChange)
                        e.sensor ← n.sensor + expectedSensorChange queue.add(e)
                end for
        end while
end function
```

To further detail the functioning of the planner, lets take the following example. Three cells in the field are shown in Figure 5. Assume currently the robot is in the $cell_{12}$ i.e, the second cell in the first row. Now the current sensor reading in this cell is shown as a continuous black line. The upper bounds in all the cells are shown as a red line and the lower bound in all the cells are shown as a green line. The corresponding representation of the sensor reading measured in $cell_{12}$ in other cells (predicted value) is shown as a black dashed line. Now the robot will move to a neighboring cell that will give the maximum bounds contraction, in other words the cell where the predicted value will be closest to the midpoints of the bounds of that cell. In our example, the corresponds to $cell_{22}$ i.e, the second cell in the second row. Hence the robot will now move to $cell_{22}$ from $cell_{12}$, in order to efficiently use the filter.

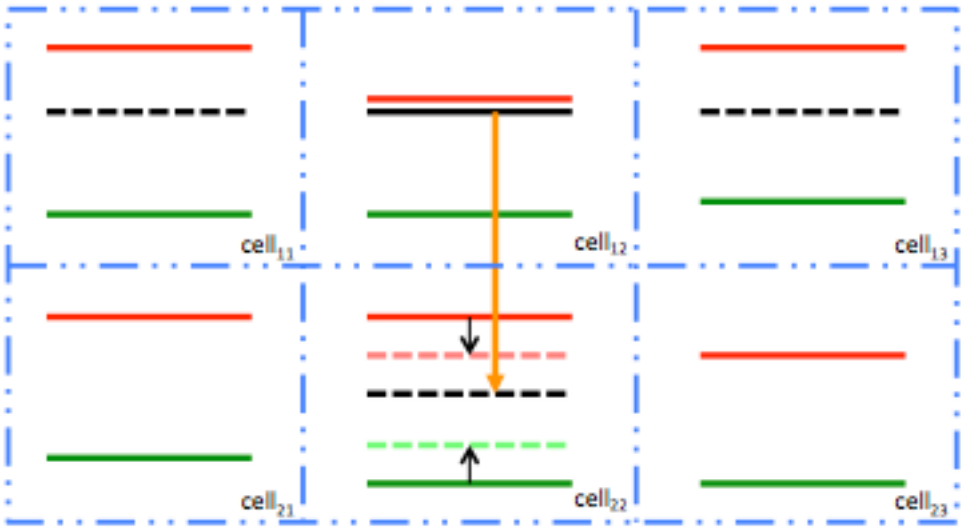

Fig. 5. Example of a cell selection step. Shown three cells in the field and the action taken by the planner based on the bounds

7. RESULTS

I evaluated the potential of my algorithm in a simulated environment partitioned into a ten by ten grid and created a model of the dissolved oxygen sensor with mild hysteresis. The more number of grids, the better is the resolution but the longer it takes, therefore there is a trade off between the sensing time and resolution of sensing. Two experiments were performed to determine the utility of the filter and the planning approach. In the first experiment, the value of each cell in the environment was draw from a uniform random distribution and in the second experiment, the value of each cell in the environment was drawn from a mixture of gaussians. The results obtained using a simple averaging filter and the bounded filter along with three different path planning algorithms: random walk, lawnmower pattern and the bounded planner, are shown in Figures 6 and 7. In the figures the solid lines represent the bounded filters result with the appropriate planner and the dashed lines represent the averaging filters result with the appropriate planner. In the bounded filter, I measure the error from the mid-point of the filter which may not always be a good measure of the information in the filter. For example, one of the bounds may get changed much earlier than the other due to gradients in the environment. The graphs represent the average result of 100 randomly generated environments and one simulated boat. In both cases, random movement with the bounded filter eventually leads to the lowest error. The difference is dramatically higher in the random environment as the random variation between the cells makes the hysteresis effect more crucial and therefore makes the averaging approach perform more poorly. The lawnmower pattern performs well initially, in part because measurements are collected uniformly all over the field, in turn significantly reducing the overall error. However, with the bounded filter, the lawnmower approach asymptotes towards some non-zero error as the watercraft enters cells from the same direction with the same hysteresis trend each time.

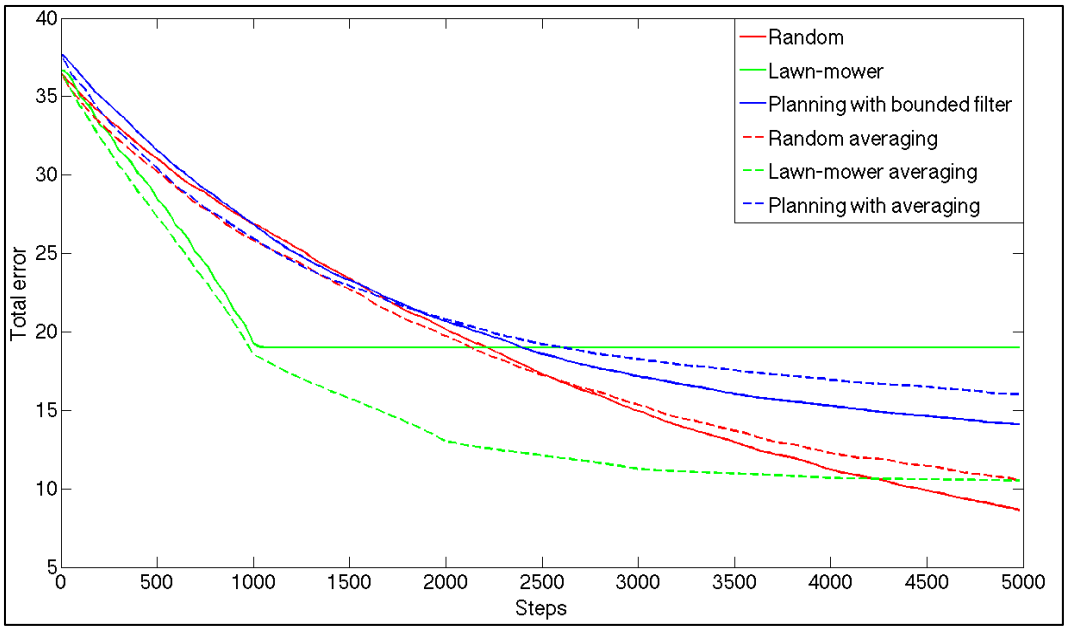

Fig. 6. Comparison of different planning algorithms using the bounded filter and the averaging filter in a simulated environment having a random distribution.

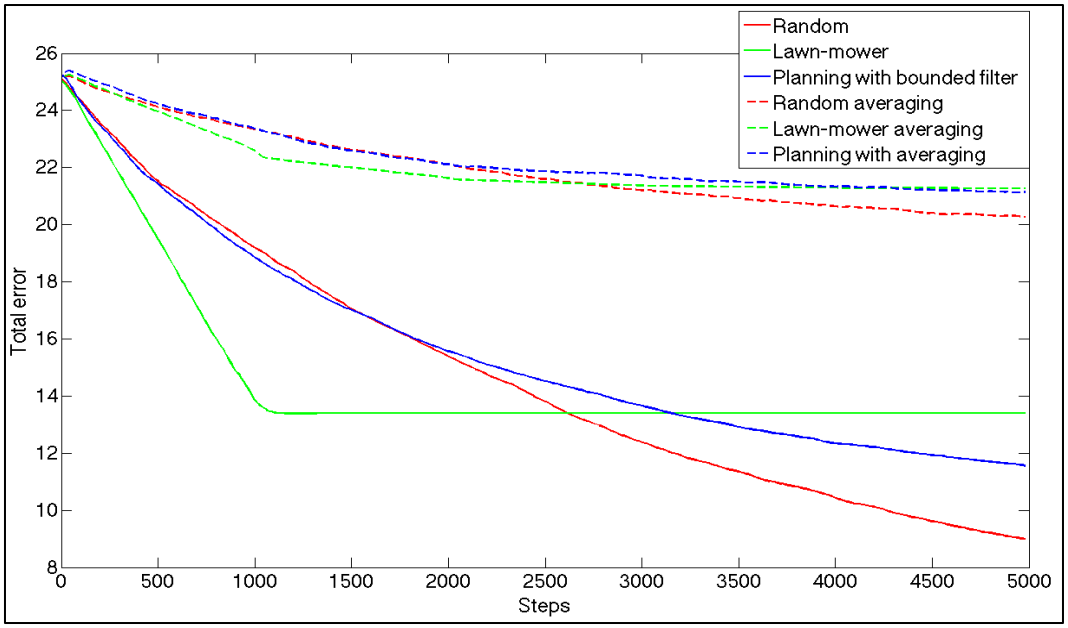

Fig. 7. Comparison of different planning algorithms using the bounded filter and the averaging filter in a simulated environment having a Gaussian distribution.

The bounded planner surprisingly does not perform better than random walk over the long run, although it has an advantage for a short period. In early stages, the good performance appears due to visiting a wider number of cells and the reason for the poorer performance later on is not completely clear and will be investigated in future work. As the robot collects samples and tags each sample with the current

GPS position, I plotted the samples collected using a dissolved oxygen sensor during the initial data collection experiments. The plotted graph is shown in Figure 8.

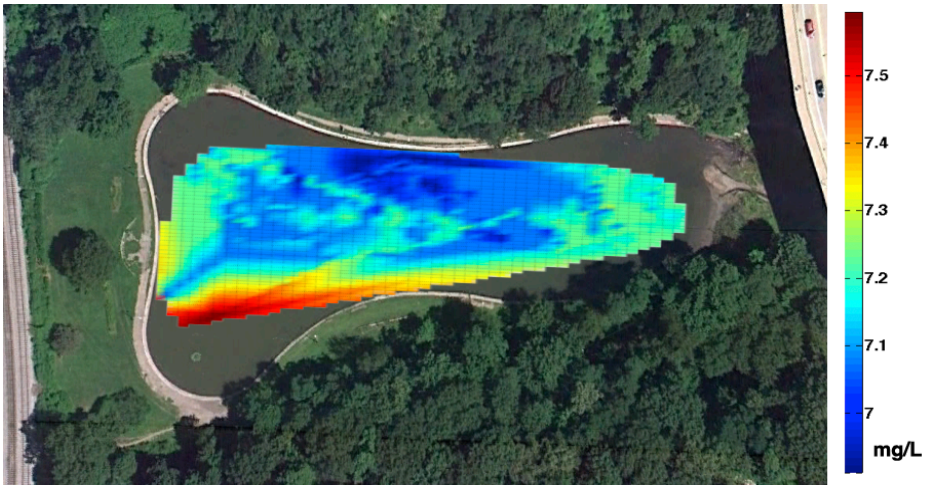

Fig. 8. Variation observed in dissolved oxygen during initial data collection experiments.

8. CONCLUSIONS

In this paper, I address the problem of using a robot to accurately map different water quality parameters using an intelligent sampling algorithm. I describe an adaptive autonomous sampling approach to compensate for the hysteresis that is observed in some water quality sensors using a bounded filter and a complementary RRT based planner. The performance of this algorithm was analyzed using results from simulation. The results show that the bounded filter has the least total error and converges the fastest when compared to a lawn-mower pattern and random walk.

Future work on the bounded filter includes developing a method for using zero gradients to bring the bounds to the current measured value and creating a new metric that can be used to assign values to cells while evaluating the bounded filter, as currently I only take the midpoint of the bounds as the cell value at each time-step.

ACKNOWLEDGMENTS

The author would like to thank Dr. Paul Scerri for his ideas and support throughout the project

REFERENCES

J. C. Leedekerken, M. F. Fallon, J. J. Leonard, Mapping Complex Marine Environments with Autonomous Surface Craft, 12th International Symposium on Experimental Robotics 2010, December 2010, New Delhi & Agra, India.

Hitz. G, Pomerleau. F, Garneau. M.-E and Pradalier. C et al., Autonomous Inland Water Monitoring: Design and Application of a Surface Vessel, IEEE Robotics & Automation Magazine, Vol. 19, Issue. 1, March 2012, pp. 62-72.

Zhang, Bin and Sukhatme, Gaurav S., Adaptive Sampling for Field Reconstruction With Multiple Moblie Robots, The Path to Autonomous Robots, Springer, 2009, pp. 1-3.

Dan O. Popa and Arthur C. Sanderson and Rick Komerska and Richard Blidberg and et al., Adaptive Sampling Algorithms for Multiple Autonomous Underwater Vehicles, IEEE/OES Autonomous Underwater Vehicles, June, 2004.

Cruz, N.A. and Matos, A.C., Adaptive sampling of thermoclines with Autonomous Underwater Vehicles, Oceans 2010, September, 2010, pp. 1-6.

S. Petillo, A. Balasuriya, and H. Schmidt, Autonomous adaptive environmental assessment and feature tracking via autonomous underwater vehicles, in Proc. IEEE Int. Conf. Oceans 2010, Sydney, Australia, May, 2010.

Bin Zhang and Gaurav S. Sukhatme and Aristides A. G. Requicha, Adaptive Sampling for Marine Microorganism Monitoring, IEEE/RSJ International Conference on Intelligent Robots and Systems, 2004, pp. 1115-1122.

Yanwu Zhang, Robert S. McEwen, John P. Ryan and James G. Bellingham, An Adaptive Triggering Method for Capturing Peak Samples in a Thin Phytoplankton Layer by an Autonomous Underwater Vehicle, Oceans, October, 2009, pp. 1-5.

A. Valada, P. Velagapudi, B. Kannan, C. Tomaszewski, G. Kantor and P. Scerri, Development of a Low Cost Multi-Robot Autonomous Marine Surface Platform, The 8th International Conference on Field and Service Robotics, Japan, July, 2012.

Kian Hsiang Low, John M. Dolan, and Pradeep Khosla, Active Markov Information-Theoretic Path Planning for Robotic Environmental Sensing, AAMAS-11, May, 2011, pp. 753-760.

Bryan, Jeff Schneider, Actively Learning Level-Sets of Composite Functions, Proceedings of the 25th International Conference on Machine Learning, July, 2008.