

Multi-Robot Coordination through Dynamic Voronoi Partitioning for Informative Adaptive Sampling in Communication-Constrained Environments

Stephanie Kemna¹ John G. Rogers III² Carlos Nieto-Granda³ Stuart Young² Gaurav S. Sukhatme¹

Abstract—Autonomous underwater vehicles (AUVs) are cost- and time-efficient systems for environmental sampling. Informative adaptive sampling has been shown to be an effective method of sampling a lake or ocean for environmental modeling. In this paper, we focus on multi-robot coordination for informative adaptive sampling. We use a dynamic Voronoi partitioning approach whereby the vehicles, in a decentralized fashion, repeatedly calculate weighted Voronoi partitions for the space. Each vehicle then runs informative adaptive sampling within their partition. The vehicles can request surfacing events to share data between vehicles. Simulation results show that the addition of the coordination with dynamic Voronoi partitioning results in obtaining higher quality models faster. Thus we created a decentralized, multi-robot coordination approach for informative, adaptive sampling of unknown environments.

I. INTRODUCTION

Autonomous vehicles are cost- and time-efficient systems for exploration and mapping of unknown environments. We are interested in the problem of modeling these unknown environments efficiently. For example; modeling algae abundance in lake environments. Note that this is in essence a problem of learning a spatial field. Previous works have shown the benefits of informative and adaptive sampling over traditional methods [1]–[3]. Multi-robot approaches can reduce the time required to explore or map an area. Our focus is on improving the multi-robot system efficiency and modeling performance through the addition of decentralized multi-robot coordination approaches.

We investigate a multi-robot coordination approach for decentralized, informative, adaptive sampling with autonomous underwater vehicles. We want a decentralized approach for robustness of the system, such that there is no single point of failure, and to allow for anytime prediction; any robot at any point in time should have a reasonable model of the whole environment. Furthermore we want an approach that keeps the required amount of communication to a minimum. Especially in underwater environments, communication is severely limited, and we cannot assume continuous access to a reliable communication channel.

*This work was supported by the Army Research Laboratory, and the Office of Naval Research (N000141410536).

¹ Department of Computer Science, University of Southern California, Los Angeles, CA, USA {kemna, gaurav}@usc.edu

² United States Army Research Laboratory, Adelphi, MD, USA {john.g.rogers59.civ, stuart.h.young.civ}@mail.mil

³ Robotics and Intelligent Machines, Georgia Institute of Technology Atlanta, GA, USA cnieto6@cc.gatech.edu

To decentralize the approach, each robot maintains their own model of the environment. At certain points in time, robots can share measurements with each other, and add these to their own models. Each robot uses its own model, together with the model’s entropy, to decide which locations to sample next. We combine this with dynamic Voronoi partitioning, to decrease overlap in actions and decisions between vehicles, effectively coordinating actions between the vehicles, with minimal communication overhead. The Voronoi partitioning is re-calculated after every data sharing event, such that the partitioning changes with the uncertainty in the model, i.e. where the vehicles need to sample. In such, we create a decentralized, multi-robot coordination approach for informative, adaptive sampling of unknown environments.

A. Related Work

Single-vehicle informative sampling work for spatial field modeling has been pioneered by [1-2, 4]. They showed that the use of information-theoretic measures as part of the path planning approach, can greatly improve modeling performance. Low et al. [2] modified this initial off-line planning approach to run as adaptive sampling. In adaptive sampling, the vehicle plans the next steps during the mission, based on the (last) measurements that were taken. These works into informative and adaptive sampling were extended to multi-robot approaches. Singh et al. [1] used sequential allocation, where a centralized path planner planned paths for each robot in sequence, and all robots ran informative sampling in parallel. Ouyang et al. [4] used active coordination in a decentralized approach. Vehicles would coordinate their actions, but only when there were other vehicles within the planning neighborhood of the own vehicle.

Typical multi-robot coordination approaches used in other domains, e.g. exploration and mapping with ground vehicles, include auction-based methods [5]–[7] and spatial segregation, typically through Voronoi partitioning [8]–[10]. Nieto-Granda et al. [11] used three heuristics for multi-robot coordination; 1) additional robots wait until they are explicitly needed, 2) robots travel in groups and split when appropriate, 3) robots travel in teams of two and split when needed. The results indicated that a divide & conquer strategy is more effective than keeping robots waiting in reserve.

The multi-robot strategies used in adaptive sampling are mostly considering only local collaboration or requiring a central planner. We do not want to use any central com-

ponents in our system, to avoid having a single point of failure. Furthermore, we would prefer strategies that consider the global space, in order to be less susceptible to local minima, and to account for the exploration-exploitation trade-off. Some commonly used methods of coordination, such as auction-based methods, require a fair amount of communication between the vehicles. We want to keep communication between vehicles to an absolute minimum, given the constraints of the acoustic communication channel, and to limit time spent on communication, versus sampling.

Previously, we have investigated the improvements that can be obtained from informative adaptive sampling with multiple robots sharing data through timed surfacing events or acoustic communications [12]. In this work, we show that the modeling performance for the multi-robot approach can further be improved through coordinating actions between vehicles. This prevents the scenario where vehicles choose the same waypoints, based on similar models of the world. We use an approach that combines Voronoi partitioning with a communication strategy for data sharing between vehicles, to achieve a decentralized multi-robot coordination approach.

Previous approaches using Voronoi partitioning include the works of Patten, Schwager and Soltero [10, 13, 14]. Patten et al. [10], and Schwager et al. [13] used a single initial Voronoi partitioning is used to divide the sample space across the robots. Patten et al. then had each robot run a travelling salesman problem (TSP) tour over sampling locations. These sampling locations have been greedily selected based on mutual information from within the Voronoi region. Schwager et al. and Soltero et al. [13, 14] discussed coverage control approaches. In both cases, the density function to be estimated is modeled using a linear combination of a set of basis functions. Schwager et al. [13] had each robot also run a TSP tour inside their partition, but over all sampling locations, to create the model. The estimated density function is then used to calculate the (weighted) Voronoi centroids, to which robots are deployed for sensor coverage. In Soltero et al. [14], robots either already have the density function, or run a lawnmower survey to obtain it. Robots then run paths whose waypoints are Voronoi generators. The waypoints move based on a control law, that takes into account the density function, towards the (weighted) centroid of the area. The paths thus change shape to monitor the interesting regions of the survey area only. Both of these approaches run two stages: model estimation, followed by static sensor deployment or persistent monitoring. Note that in this work, we focus only on the field estimation, and creating an informative model. We do not assume to know the density function beforehand, and are interested in using online and adaptive approaches for modeling it.

II. METHOD

In this paper, we develop a dynamic Voronoi partitioning approach for decentralized adaptive sampling. The vehicles create a model of the environment using Gaussian Process regression, and running adaptive sampling utilizing the pos-

terior map entropy of the model. To differentiate vehicle actions, each vehicle runs Voronoi partitioning over possible sampling locations. The posterior map entropy is used here as a density function for shifting the centroids of the Voronoi partitions towards interesting regions. While the vehicles are sampling, at depth, they can not share data with each other. However, they can request data sharing events, at which time data is shared, and the Voronoi partitions are re-calculated. This section will briefly explain each of these subparts of our dynamic Voronoi partitioning¹ for vehicle coordination in informative adaptive sampling.

A. Gaussian Process regression & path planning

A common technique for spatial modeling, known in geostatistics as Kriging, is Gaussian Process (GP) regression [16]. GP regression is a non-parametric modeling technique, where the GP is completely specified by its mean function and its covariance matrix. While the GP is a non-parametric model, its performance is affected by hyperparameters. These are typically estimated from the data using gradient-based optimization [16]. In our experiments, we use a zero-mean prior and an isotropic squared exponential covariance function for the GP [12]. Furthermore, we use resilient backpropagation (*Rprop*) for finding hyperparameters that maximize the likelihood of the model [17].

To model biological data, we use log-Gaussian Processes (ℓ GP), as in [2]: Let Y_x denote an ℓ GP, modeling the sensor value y_x at location $x \in \mathcal{X}$, where $\mathcal{X} \subset \mathbb{R}^2$. Let $Z_x = \log_e Y_x$, denote a GP. Then we can create the ℓ GP using GP regression by utilizing the fact that $z_x = \log_e y_x$. The GP's posterior mean and variance, $\mu_{Z_x|d_i}$ and $\sigma_{Z_x|d_i}^2$ (for sampled data d_i), are used to calculate the posterior mean and variance of the ℓ GP [2]:

$$\mu_{Y_x|d_i} = \exp\{\mu_{Z_x|d_i} + \sigma_{Z_x|d_i}^2/2\} \quad (1)$$

$$\sigma_{Y_x|d_i}^2 = \mu_{Y_x|d_i}^2 (\exp\{\sigma_{Z_x|d_i}^2\} - 1) \quad (2)$$

Each vehicle thus creates an ℓ GP model on-board, as it is sampling the environment.

The goal of the vehicle's actions is to reduce the uncertainty in the model. The vehicle executes movements using a waypoint behavior. We take a greedy approach to finding future sampling locations, based on the vehicle's ℓ GP model. As optimization criteria, we use posterior map entropy for the log-GP, as derived in [2]:

$$\mathbb{H}[\mathbf{Y}_{x_{i+1}}|d_i] = \log \sqrt{2\pi e \sigma_{Z_{x_{i+1}}|d_i}^2} + \mu_{Z_{x_{i+1}}|d_i} \quad (3)$$

where d_i is the already sampled data. Thus the vehicle chooses as its next waypoint an unvisited location with highest posterior map entropy. For the log-GP, as follows from Equation 3, this means locations with high posterior variance (to reduce it) and high expected sensor values

¹Note, we call our approach 'dynamic Voronoi partitioning' because we re-calculate the (weighted) Voronoi partitions after every data sharing event. Our approach is different from Bakolas and Tsiotras' 'dynamic Voronoi diagram' [15], where a Voronoi partition is re-calculated given expected time to reach Voronoi centers, or 'terminal positions'.

(interesting areas). When the AUVs start sampling, the model will initially be uniform. To this end, the vehicles' initial waypoint is a random location in the survey area. The vehicles make straight-line movements to waypoints. It is assumed that there are no obstacles within the survey area. Those could be handled via an obstacle avoidance behavior.

B. Dynamic Voronoi partitioning

For action coordination between vehicles, we use dynamic Voronoi partitioning. This allows us to coordinate vehicle movements without the need for constant communication. The vehicles independently estimate Voronoi regions, to limit their prospective sampling locations. Note that this is not used as a control law for vehicle movement. Assuming all vehicles have a fairly recent estimate of each other's positions, they can each calculate the Voronoi partitioning for the survey area: Each vehicle considers all (unvisited) sampling locations, and creates a subset with only those locations closest to itself. If the vehicles would not know the positions of all other vehicles, then one could use decentralized Voronoi partitioning, e.g. [18].

For these initial Voronoi partitions, we then consider the density function (i.e. the posterior map entropy), and calculate the weighted centroid for each partition. The weighted centroid is determined by considering all unvisited sampling locations in the partition and their posterior map entropy. Each location is weighted by its entropy (Equation 3) and the sums of all weighted latitude and longitude values are divided by the sum of all weights, to find the latitude and longitude of the partition's weighted centroid. These weighted centroids are used as the Voronoi generators for a second round of Voronoi partitioning. The resulting partitions are then used in the path planning algorithm (Section II-A).

C. Data sharing

Due to the limited bandwidth and potentially reduced throughput of underwater acoustic communications, it is not practical to have all vehicles broadcast their measurements underwater. We assume however that the underwater communication channel is stable enough to get some messages through, e.g. for inter-vehicle collision avoidance, and for requesting surfacing events.

Each vehicle, at any time, can request a surfacing event. It broadcasts a message to the other vehicle(s) to request surfacing, and meanwhile continues sampling. Vehicles receiving

a request send an acknowledgment and start surfacing. Upon receipt of acknowledgments from all other vehicles, the initiating vehicle also surfaces. Once on the surface, the vehicles initiate a handshake protocol, to make sure all other vehicles are at the surface and ready to start sharing data. After the handshake, all vehicles broadcast their (not yet transmitted) sensor measurements, through Wi-Fi communications. The received measurements from other vehicles are added to each vehicle's local ℓ GP. At this point, each vehicle will recalculate the Voronoi partitioning (as per Section II-B), and then re-commence with (at-depth) adaptive sampling.

There are many possible ways of determining when a surfacing event is required, e.g. the criterion could be time-based, or information-based. In this work, we take the approach of requesting surfacing events when the vehicles detect that they are sampling close to the border of their Voronoi region. In such case the interesting areas are near the border of their region, and thus it is considered a good time to reconsider the partitioning. Note that vehicles can not request a surfacing event within five minutes of a previous surfacing event.

III. EXPERIMENTAL SET-UP

We ran simulation experiments with two autonomous underwater vehicles, using the MOOS-IvP middleware [19]. The middleware includes a simple simulation of vehicle dynamics and PID control, as well as behavior-based autonomy. We use the following standard behaviors; loiter, waypoint, constant depth, and (inter-vehicle) collision avoidance. We use a constant speed of 1.5 m/s and constant depth of 5 m . Our path planning approach updates the waypoint behavior.

To simulate algae blooms in a lake, we generated a 2-D grid space ($400 \times 200\text{ m}$, 10 m spacing) with a 2-D Gaussian to represent a bloom, and additive Gaussian noise. For the data, we use a data value amplitude of 40, as a proxy for high Chlorophyll $\mu\text{g/L}$ values, with a noise amplitude of 10-20 percent. Within the MOOS-IvP simulation, each vehicle samples from this simulated data, and we add Gaussian noise ($\sigma = 1.5$) to model sensor noise [12]. Note that the sensors take point measurements, there is no range to the sensors, and all vehicles have the same sensor. We simulate two different scenarios for the grid space; a single bloom with low noise, and two (differently sized) blooms in opposite corners with higher noise, as shown in Fig. 1 and 2.

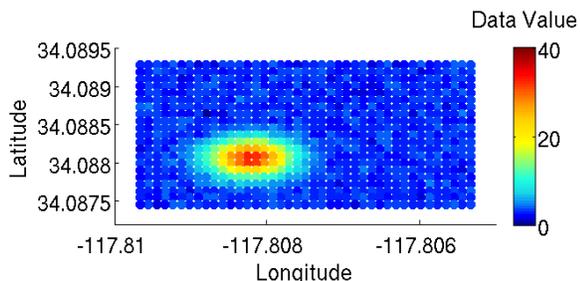


Fig. 1. Scenario 1: single bloom

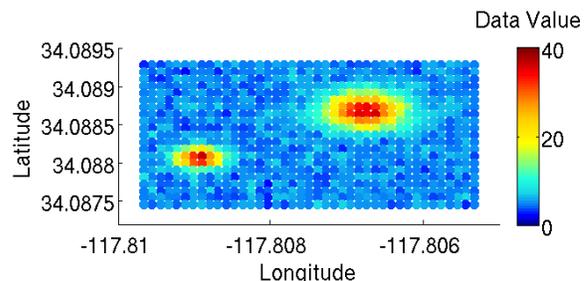


Fig. 2. Scenario 2: dual bloom

To simulate communications between the vehicles, we use two forms of simulated communications: acoustic and Wi-Fi. For acoustic communications (acomms), the *Goby* acomms suite enables us to simulate the whole acomms stack; modem driver, medium access control, priority-based message queuing, and encoding and decoding of messages [20]. We use a TDMA (Time Division Multiple Access) scheme for sharing the data channel between the vehicles. Each vehicle has a pre-assigned 3-second time slot, in which it can send one to two 32-byte messages. We limit the acomms range to 500 m, and are able to reduce the throughput of the channel to probabilistically drop messages, e.g. to drop 30% of the messages. For Wi-Fi communications, we simulate UDP communications, without restrictions on throughput or range².

1) *Experiments*: In our previous work [12], we showed the improvements in modeling performance when adding data sharing between vehicles. Vehicles sharing data performed better than those vehicles that did not share data and ran standard lawnmower surveys³ or adaptive sampling in parallel. In this paper, we compare the dynamic Voronoi partitioning approach for informative adaptive sampling with the data sharing approach, and run the following experiments:

- 2 AUVs, timed data sharing, no coordination,
- 2 AUVs, dynamic Voronoi partitioning,
- 3 AUVs, timed data sharing, no coordination,
- 3 AUVs, dynamic Voronoi partitioning.

In the timed data sharing approach, the AUVs surface every ten minutes⁴, to be able to share data. When all

²Assuming Wi-Fi via 2.4 GHz radio and good antennas, which we have empirically found capable of at least 500 m (unobstructed) range, which covers the current experimental area.

³A lawnmower survey is a typical coverage behavior where a vehicle travels back and forth across the survey area.

⁴This assumes that the vehicles have synchronized clocks, which is a given in our simulations, but should be paid attention to at field trials.

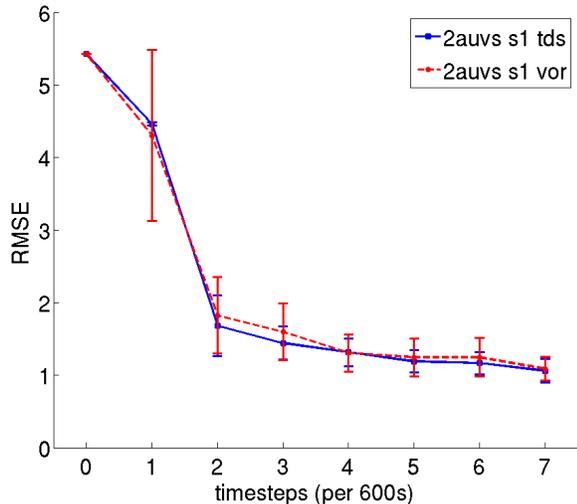


Fig. 3. Scenario 1 (s1), one bloom: Average RMSE and one standard deviation (SD) error bars, averaged over all AUVs in the 2-AUV mission, for both timed data sharing (*tds*) and dynamic Voronoi partitioning (*vor*).

vehicles are on the surface, they share data after executing the handshake protocol (Section II-C).

For every 2 AUV and 3 AUV experiments, we ran 10 and 15 simulations respectively, and we averaged over all results. Furthermore, we ran all these simulations over both scenarios (Fig. 1 and 2). The duration of each experiment was limited to the duration of simulations with the vehicles running lawnmower surveys. For example, to survey the area with three vehicles running high resolution, 20 m track spacing, lawnmower surveys takes on average approximately 3500 seconds. Therefore, after the vehicles have run adaptive sampling for 3500 seconds, they are requested to return to the start location, for the final data sharing and hyperparameter optimization.

IV. RESULTS

The results from all experiments are evaluated using root mean squared error (RMSE) between the posterior mean from the ℓ GP and the generated data, and using the negative log-likelihood (NLL). The RMSE captures the predictive mean performance of the ℓ GP, and the NLL incorporates also the predictive variances. We store the model predictions throughout the mission, at 10 minute intervals (600 s), for evaluation. Note that all figures in this section show results per these 10 minute time steps, and all changes therefore seem to occur at the same times. We run a first hyperparameter optimization at the first surfacing event. For the timed data sharing, this is at around 600 seconds into the mission. For the Voronoi mission, this is on average at 400-450 seconds into the mission. A second hyperparameter optimization is run at the end of the survey.

A. Two AUVs

We ran simulation with two AUVs, running either timed data sharing or dynamic Voronoi partitioning. Results are averaged over all vehicles, for 10 simulations per experiment

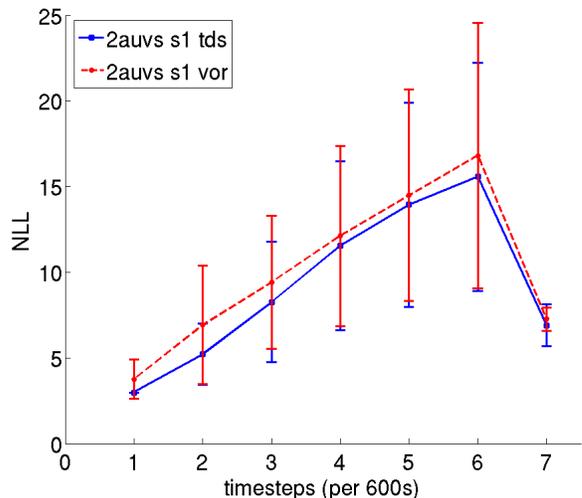


Fig. 4. Scenario 1 (s1), two blooms: Average NLL and one SD error bars, averaged over all AUVs in the 2-AUV mission, for *tds* and *vor*.

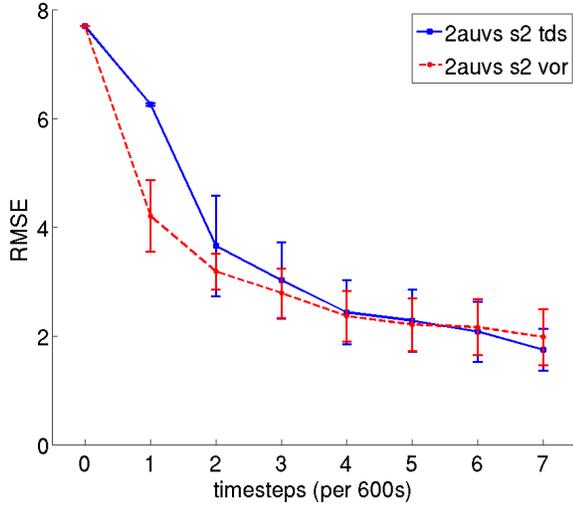


Fig. 5. Scenario 2 (s2), two blooms: Average RMSE and one SD error bars, averaged over all AUVs in the 2-AUV mission, for both *tds* and *vor*.

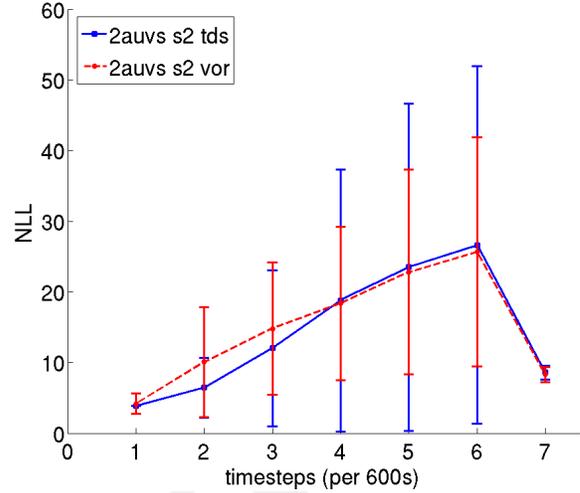


Fig. 6. Scenario 2 (s2), two blooms: Average NLL and one SD error bars, averaged over all AUVs in the 2-AUV mission, for both *tds* and *vor*.

and scenario. Figure 3-6 show the results, in terms of RMSE and NLL, for two AUVs running adaptive sampling with only data sharing, versus coordination through dynamic Voronoi partitioning. Note that the final time step is after the final hyperparameter optimization, when the survey is finished, and can hence lead to a bigger reduction of RMSE or NLL. The vehicles running timed data sharing surfaced 5 times during the mission. The number of surfacing events for the dynamic Voronoi approach was 7-9 times.

Figure 3 and 4 show the results for the first scenario, running with two AUVs. As can be seen, the performance for the two approaches is on average the same.

Figures 5 and 6 show the results for the second scenario, running with two AUVs. As can be seen, the RMSE initially drops more quickly for the dynamic Voronoi partitioning, showing improved performance. Both methods reach similar values at the end of the mission. The

performance is about the same in terms of NLL, but the dynamic Voronoi approach is in general more consistent in performance. Overall, for the experiments with 2 AUVs, we see that the modeling performance between the two methods is the same for the first scenario, but better with dynamic Voronoi partitioning in the second scenario.

B. Three AUVs

We also ran simulation with three AUVs, running timed data sharing and dynamic Voronoi partitioning. Results are averaged over all vehicles, for 15 simulations per experiment and scenario. Figures 7-10 show the averaged results for the simulations with three AUVs. Note that the final time step is after the final hyperparameter optimization, when the survey is finished, and can hence lead to a bigger reduction of RMSE or NLL. The vehicles running timed data sharing surfaced 4 times during the mission. The average number of surfacing

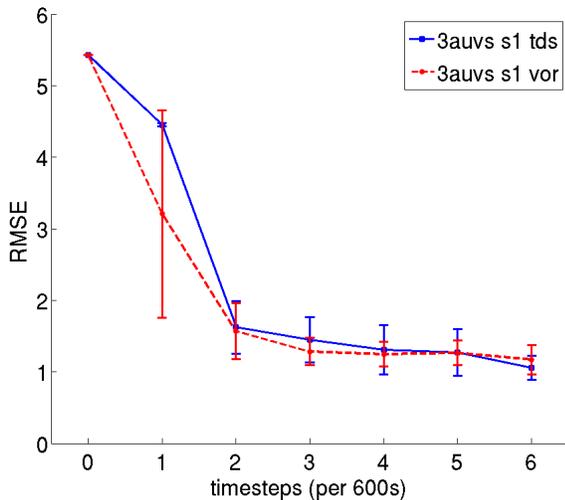


Fig. 7. Scenario 1 (s1), one bloom: Average RMSE and one SD error bars, averaged over all AUVs in the 3-AUV mission, for both *tds* and *vor*.

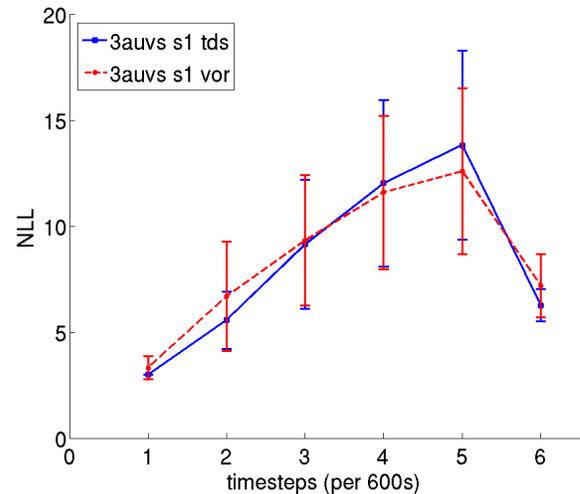


Fig. 8. Scenario 1 (s1), one bloom: Average NLL and one SD error bars, averaged over all AUVs in the 3-AUV mission, for both *tds* and *vor*.

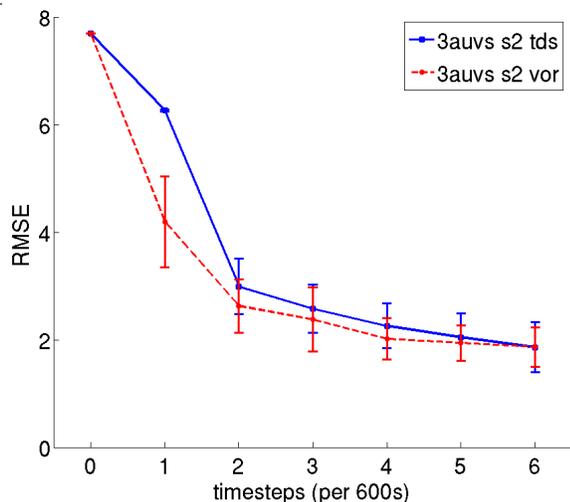


Fig. 9. Scenario 2 (s2), two blooms: Average RMSE and one SD error bars, averaged over all AUVs in the 3-AUV mission, for both *tds* and *vor*.

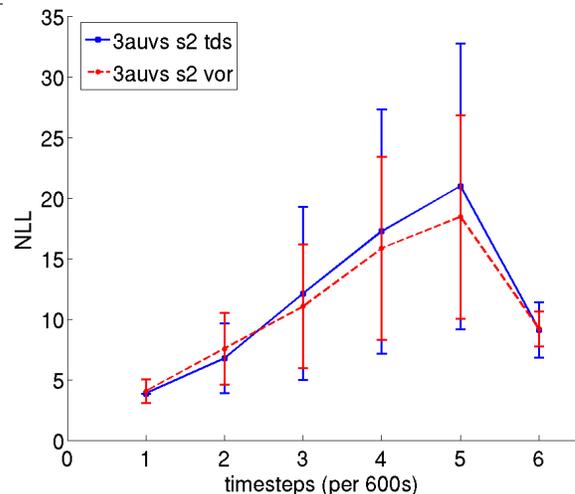


Fig. 10. Scenario 2 (s2), two blooms: Average NLL and one SD error bars, averaged over all AUVs in the 3-AUV mission, for both *tds* and *vor*.

events for the dynamic Voronoi approach was 6-7 times.

Figures 7 and 8 show the RMSE and NLL for the first scenario with a single bloom. The figures show that with the coordination through Voronoi partitioning, the modeling performance improves both in terms of RMSE and NLL, in particular over the first 3-4 time steps (1800-2400 seconds) of adaptive sampling.

Figures 9 and 10 show the RMSE and NLL for the second scenario with two blooms. Note that, due to the higher noise in the dual bloom scenario (see Figure 2), the RMSE and NLL are higher than for the single bloom scenario. Again we see that the performance is better when using coordination through Voronoi partitioning, instead of only data sharing. Overall, we see for the experiments with 3 AUVs that the use of dynamic Voronoi partitioning for vehicle coordination results in higher quality models, in particular for the second scenario with two blooms.

V. DISCUSSION

The results show that the modeling performance improves with addition of a coordination approach, in multi-robot informative adaptive sampling. For the 2-AUV experiments, the performance is the same for the experiments on the first scenario. For the 3-AUV experiments, the performance improvement is also not as well defined for this scenario. This is likely because the single bloom scenario, with low noise, is a simple scenario, where the addition of vehicles, or coordination in the multi-robot approach, does not pay off.

For the second scenario, the performance improvement is more pronounced. We see that in particular the RMSE improvement is bigger at the start of the mission, and overall the NLL performance is more consistent. The improvement is bigger at the beginning, because this is where the biggest improvement can be obtained. After ca. 3-4 time steps (1800-2400 seconds) of sampling, the performance starts to level off for all methods. This indicates that the model after this

time span is about as good as it gets, and reaching the limit in possible performance given model and sensor noise. The addition of another hyperparameter optimization step at this time may aid in improving modeling performance. However, it is also an indicator that the mission could be ended earlier for all these adaptive sampling approaches.

We have not evaluated the overall time taken for the missions, because time was used to regulate mission duration, and no other stopping criteria were tested. All missions took exactly the same amount of time, and vehicles traveled approximately the same distance, given that the AUVs never stopped and ran at the same speed. If we were to let all robots run for infinity, every approach would end up with the same model performance. This is also evidenced by the final performance being similar, if not the same, for all methods. We are thus looking for improvements in modeling performance at early stages of the mission. In [12], we showed that the modeling performance for adaptive sampling is greatly improved in the early stages of the mission, after which the performance levels out, suggesting the adaptive mission could have been ended earlier. For the 3-AUV Voronoi approaches, we see also see improvements in the early stages, which mean that (a) we could potentially end the mission earlier, and (b) if we had to stop the mission at an earlier stage, we would have a better model. This type of ‘anytime prediction’ capability is very promising.

As mentioned in Section IV-A and Section IV-B, the numbers of surfacing events for the dynamic Voronoi approach was a little higher for both the 2-AUV and 3-AUV experiments (7-9 versus 5, and 6-7 versus 4, respectively). For timed data sharing, the vehicles surfaced every 10 minutes (600 s). For dynamic Voronoi partitioning, the trigger for surfacing events existed of a vehicle nearing the Voronoi border of their partition. The number of surfacing events thus depends on the size of the Voronoi partition, the location of the vehicle, and the location of informative unsampled areas in this partition. The number of surfacing

events therefore is not optimized (i.e. minimized). It is limited however by a ‘blackout’ period of five minutes after each previous surfacing event, to prevent immediate re-surfacing, for example if a new Voronoi region border happens to be near a vehicle. It would be interesting to investigate different triggers for requesting surfacing events, and to develop a strategy for minimizing surfacing, while still maintaining the benefits of regular data sharing and anytime prediction capability. While our current research efforts focus on shallow waters and relatively shallow depths of features of interest, the minimization of surfacing events will be of greater importance for deep water operations.

Asynchronous surfacing strategies, via a central sharing vehicle/system, were not evaluated in this work either, because we wanted to avoid having a single point of failure on the surface or on shore, and we wanted to avoid the addition of a surfaced data mule for both cost and vehicle safety reasons. Another approach for asynchronous surfacing would be to have vehicles share information in subsets or teams, and to share with those who do not have the data yet, whenever possible. This would require a lot more internal accounting and the capability of vehicles to share or negotiate what has (not) been transferred. While asynchronous sharing would remove the current system’s ability to provide a good model from any vehicle at any time, it is an interesting avenue for future work. In particular when considering the application of adaptive sampling systems to environments where any form of communication can only be achieved in certain areas or at certain times.

In the future, we intend to explore different coordination approaches, as well as the effect on performance of the use of different path planning algorithms, and the application to larger areas. Furthermore, efforts are ongoing to demonstrate the adaptive sampling performance through field trials. Thus we also aim to create new scenarios for future simulations based on field data.

VI. CONCLUSION

Overall, we have shown that informative adaptive sampling with multiple robots benefits from the addition of action coordination between vehicles. Previous works have focused mostly on coordinating within close range for local planning, or coordinating between vehicles by planning sequentially. We developed an approach where the vehicles can coordinate their actions, without limiting the coordination approach or the path planning to a local neighborhood. Furthermore, our approach is completely decentralized, and creates a robust multi-robot system for adaptive sampling. In combination with the communications strategy, this makes the approach applicable for deployments in hazardous or communication-constrained environments. Our dynamic Voronoi partitioning technique for multi-robot coordination is thus an effective method in running decentralized, informative adaptive sampling with multiple robots, in unknown environments.

ACKNOWLEDGMENTS

The authors would like to thank David A. Caron for providing the biological perspective to our robotic sampling efforts. Furthermore we would like to thank members of the Robotic Embedded Systems Laboratory at the University of Southern California for their constructive feedback on this work, and the anonymous reviewers for their valuable feedback.

REFERENCES

- [1] A. Singh, A. Krause, C. Guestrin, W. Kaiser, and M. Batalin, “Efficient Planning of Informative Paths for Multiple Robots,” in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2007.
- [2] K. H. Low, J. M. Dolan, J. Schneider, and A. Elfes, “Multi-Robot Adaptive Exploration and Mapping for Environmental Sensing Applications,” PhD thesis, CMU, 2009.
- [3] D. R. Thompson, D. S. Wettergreen, and F. J. C. Peralta, “Autonomous Science during Large-Scale Robotic Survey,” *Journal of Field Robotics*, vol. 28, no. 4, pp. 542–564, 2011.
- [4] R. Ouyang, K. H. Low, J. Chen, and P. Jaillet, “Multi-Robot Active Sensing of Non-Stationary Gaussian Process-Based Environmental Phenomena,” in *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2014.
- [5] R. Zlot, A. Stenz, M. B. Dias, and S. Thayer, “Multi-Robot Exploration Controller by a Market Economy,” *IEEE International Conference on Intelligent Robotics and Systems (IROS)*, pp. 3016–3023, 2002.
- [6] R. Simmons, D. Apfelbaum, W. Burgard, and D. Fox, “Coordination for multi-robot exploration and mapping,” *AAAI Conference on Artificial Intelligence (AAAI)*, 2000.
- [7] W. Sheng and N. Xi, “Multi-robot area exploration with limited-range communications,” in *IROS*, 2004, pp. 1414–1419.
- [8] D. E. Soltero, M. Schwager, and D. Rus, “Generating informative paths for persistent sensing in unknown environments,” *IROS*, no. 2010, pp. 2172–2179, 2012.
- [9] A. Marino, G. Antonelli, A. P. Aguiar, A. Pascoal, and S. Chiaverini, “A decentralized strategy for multirobot sampling/patrolling: Theory and experiments,” *IEEE Transactions on Control Systems Technology*, vol. 23, no. 1, pp. 313–322, 2015.
- [10] T. Patten, R. Fitch, and S. Sukkarieh, “Large-Scale Near-Optimal Decentralised Information Gathering with Multiple Mobile Robots,” in *Proceedings of Australasian Conference on Robotics and Automation*, December 2013.
- [11] C. Nieto-Granda, J. G. Rogers, and H. I. Christensen, “Coordination strategies for multi-robot exploration and mapping,” *The International Journal of Robotics Research (IJRR)*, vol. 33, no. 4, pp. 519–533, 2014.
- [12] S. Kemna, D. A. Caron, and S. G. S., “Adaptive Informative Sampling with Autonomous Underwater Vehicles: Acoustic versus Surface Communications,” in *Oceans*, September 2016.
- [13] M. Schwager, M. P. Vitus, S. Powers, D. Rus, and C. J. Tomlin, “Robust adaptive coverage control for robotic sensor networks,” *IEEE Transactions on Control of Network Systems*, 2014.
- [14] D. E. Soltero, M. Schwager, and D. Rus, “Decentralized path planning for coverage tasks using gradient descent adaptive control,” *IJRR*, vol. 33, pp. 401–425, March 2014.
- [15] E. Bakolas and P. Tsiotras, “The Zermelo–Voronoi diagram: A dynamic partition problem,” *Automatica*, vol. 46, no. 12, pp. 2059–2067, 2010.
- [16] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. MIT press, 2006.
- [17] M. Blum and M. Riedmiller, “Optimization of Gaussian Process Hyperparameters using Rprop,” in *European Symposium on Artificial Neural Networks, CML*, 2013.
- [18] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, “Coverage control for mobile sensing networks,” in *IEEE International Conference on Robotics and Automation (ICRA)*, vol. 2, no. May, 2002, pp. 1327–1332.
- [19] M. R. Benjamin, P. Newman, J. J. Leonard, and H. Schmidt, “An Overview of MOOS-IvP and a Users Guide to the IvP Helm Autonomy Software,” MIT, Tech. Rep. MIT-CSAIL-TR-2010-041, 2010.
- [20] T. Schneider, “Goby Underwater Autonomy Project - User Manual,” 2014. [Online]. Available: <http://gobysoft.org/doc/2.0/>