Doctoral Thesis Proposal:
Efficient, Self-deploying, and Reconfigurable Robotic Sensor Networks

Sandeep Manjanna

School of Computer Science

McGill University

msandeep@cim.mcgill.ca

Thursday 25th August, 2016

Abstract

Robotic Sensor Networks composed of robotic wireless sensing devices hold the potential to revolutionize environmental sciences by enabling researchers to collect data across expansive environments, over long, sustained periods of time. An efficient Robotic Sensor Network would be capable of physically reconfiguring itself to achieve efficient area coverage, in-depth examination of targets, reliable wireless connectivity, and dynamic protection against unexpected environmental developments. We aim to develop intelligent decision making algorithms for autonomous sensor networks to achieve generic tasks with utmost efficiency. Designing such algorithms requires tackling challenging problems that lie at the intersection of robotics, perception and communication. We address the problems related to automated configuration, adaptive sampling and reconstruction of the field, multi-robot coordination, and target behavior modeling.

We focus our research on probabilistic modeling of the states of the robotic network and also probabilistic representation of the spatial field. Research ideas for automated network configuration are proposed which aim at maximizing the utility of spatial coverage while minimizing the costs. To represent and reconstruct the spatial field, we propose to use Gaussian Process(GP) models. GP models can deal with noisy measurements, unevenly distributed observations, and fill small gaps in the data with high confidence while assigning higher predictive uncertainty in sparsely sampled areas. For persistent monitoring of a continuous field or discrete targets, we want to use probabilistic hypothesis density filters that provide an estimate of target density and use this information to reconfigure the robotic sensor network.

Thus we propose to develop a framework for a network of robots to autonomously plan and update their configuration such that they can efficiently observe and track spatially varying scalar field(s) that express environmentally important phenomena (e.g. temperature) and also discrete target groups (e.g. group of wifi-hotspot users). Our approach will be based on using a combination of topological representation of the robot network, and probabilistic model of the phenomena of interest. We provide details about the progress achieved till date and also set milestones to evaluate the progress towards our research goals.

Contents

1 Introduction

A robotic sensor network is a network of robots that are capable of sensing their environment, collaborating with other robots, and moving. We aim to develop intelligent decision making algorithms for autonomous sensor networks to achieve generic tasks with utmost efficiency. Designing algorithms for robotic sensor networks requires tackling challenging problems that lie at the intersection of robotics, perception and communication. We focus on addressing the problems related to automated configuration, adaptive sampling and reconstruction of the field, multi-robot coordination, and target behavior modeling.

One of the major applications of mobile robotics is exploring and mapping an unknown environmental field. Building a persistent map of an environmental field for a given region over a period of time requires continuous sensing. This problem has been tackled previously by deploying a network of static sensor nodes uniformly across the region of interest. But the environmental fields are usually non-uniform over a given region and dynamic over a time period. Persistent mapping of such dynamic fields requires the sensor nodes to estimate the changing field, then adapt their placements accordingly, and selectively sample the measurements to maximize the information gain while minimizing costs over time, distance traveled, and communication.

A robotic sensor node can be abstracted as three functional blocks: sensor to take measurements from the physical world, a decision making agent to model the environmental field and decide on an optimal policy, and an actuator to provide mobility to the sensor node. These functional units endow the robotic node with abilities to sense, assess and model the surroundings, decide on the optimal action, and execute the action.

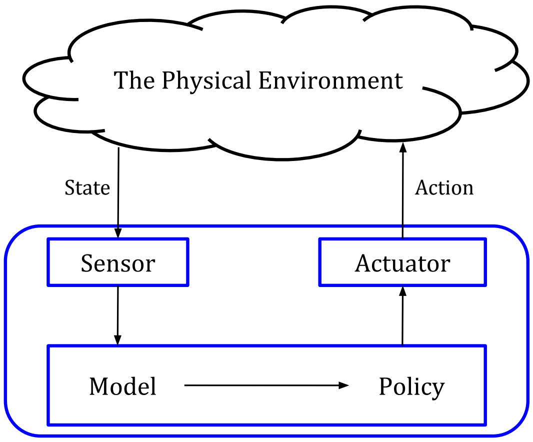

Figure 1: Block representation of a sensor node.

A team of such robotic sensor nodes creates an autonomous sensor network. Under ideal circumstances, a team of robotic sensor nodes collectively can be more efficient in sensing their environments. However, this comes with a penalty on communication and requirements to fuse the data captured by different sensor nodes.

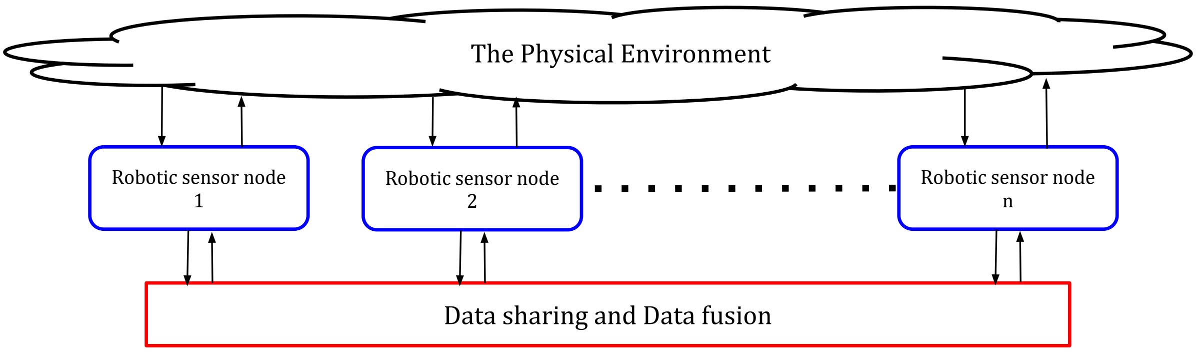

Figure 2: Block representation of a network of sensor nodes.

Mobility enables mobile sensors to flexibly reconfigure themselves to meet sensing requirements, improve the efficiency, and save sensing resources in area exploration, especially for large environments with limited number of sensors. Recently, adaptive sampling of sub-regions with high local-variance has become the focus of research in robotic exploration and mapping. The sub-regions of a spatial field which show high local-variance are referred to as hotspot-regions and such spatial fields with hotspot-regions are characterized by continuous spatially correlated measurements with the hotspot-regions exhibiting extreme measurements and much higher spatial variability than the rest of the field. Monitoring such environmental fields which include hotspot-regions will benefit by a team of network of robotic sensor nodes, instead of static sensor networks, as they can be more robust, more tolerant to faults, and capable of performing varied-resolution sampling.

Thus locomotion of sensor nodes facilitates a number of useful network capabilities, like the automated configuration; network re-configuration based on the target field; recover from harsh environment failures and partial breakages.

We propose to develop an adaptive network of robotic sensor nodes that is capable of - sampling the environment adaptively to produce a varied-resolution map of the spatial field; self deployment according to the spatial features of the field; tracking the changes in the field map and targets over time and building a transition model; respond to the changing field by reconfiguring the network topology.

1.1 Motivational Applications

Robotic sensing networks have many applications in exploration and mapping of an uncertain and dynamic environmental field. Mobile sensing platforms in the form of unmanned ground vehicles (UGV), unmanned aerial vehicles (UAV), autonomous surface vehicles (ASV) and autonomous underwater vehicles (AUV) can actively respond to dynamic changes in unknown environments and dynamically deploy themselves to enhance sensing capability in the area of interest. The mobile sensing nodes in the network can act collectively to mimic cooperative motions appearing in animal networks such as wolf packs, bird flocks and fish schools. Some of the examples that motivate us are as listed below,

- Environmental Sensing such as monitoring of ocean phenomena (plankton bloom, upwelling, contamination) [1] [2] [3], forest ecosystems [4], rare species [5], or pollution [6].
- Animal tracking and behavior estimation [7]; exploration and monitoring of migratory species and paths.
- Terrestrial Exploration such as Antarctic meteorite search [8]; geologic site survey and monitoring [9] and prospecting for mineral deposits [10]; or localized methane sources on Mars [11].
- Search and Rescue missions in forests and open-water environments.
- Urban Application such as surveillance at exhibitions, museums, public gatherings, and rallies; mobile wifi-hotspots in public spaces.

1.2 Challenges

A series of research issues that need to be addressed include measurement collection and evaluation, mobile network establishment, data fusion, sensor motion, and target tracking. Regardless of the applications, all these tasks aim at information acquisition. The kind of information to acquire and the way to acquire information can be considered the core problems in mobile sensing. Every new measurement should increase the knowledge base, and can be used to guide sensor motion, thus rebuilding network topology. The first challenge is to define a metric to evaluate potential measurement values.

Secondly, the individual sensor nodes should be able to communicate and share data with their neighboring nodes to enhance their understanding of the physical world by merging the data. Robotic sensor networks can decide about their new topology and trajectory based on the updated field map generated using data from all the sensor nodes in the network. Hence, data sharing and data fusion becomes an important challenge of such a system.

Lastly, the energy consumption of the robotic sensor nodes need to be optimized to achieve an efficient coverage of the region. On-board energy is mainly consumed for sensing, data transfer, and sensor motion. High sampling frequency results in high energy consumption for sensing. Low Sampling frequency achieves higher energy efficiency, however, sensing performance is directly related to the number of measurements collected. Thus, an intelligent sampling strategy becomes important for saving on-board energy.

2 Background concepts

2.1 Sensor network deployment

Designing and deploying a sensor network aim to achieve two important qualities: network coverage and network connectivity. Network coverage refers to how well the sensor network covers the area of the phenomena being monitored, whereas connectivity refers to the ability of active nodes to stay connected. Wang et al. quantified the relationship between connectivity and coverage [12]. They define k-coverage as any location in the region of interest being monitored by at least k nodes, and k-connected if the network remains connected even when $k-1$ nodes fail. Planning a deployment a-priori is not always possible and it depends on the nature of monitored environment. To provide coverage and connectivity in a given area with the least number of sensors possible, two sampling based deployment approaches are proposed by Isler et al. [13]: (1) *concurrent deployment*, where the number and locations of the nodes are decided prior to deployment; and (2) *incremental deployment*, where the feedback about current coverage and connectivity after each sensor placement is used to decide the next placement. The authors propose using random sampling without replacement to achieve incremental deployment, i.e. the area already covered cannot be randomly picked again. The area covered by sensing range but not covered by the communication range of any sensor is added to the sampling domain.

The achievements in low power processors, wireless networking, and sensor technology gave rise to the field of wireless sensor networks. From a robotics perspective, we can view teams of robots with sensing, communication, and locomotion capabilities as mobile sensor networks. Locomotion facilitates a number of useful network capabilities, including the ability to self-deploy; that is, starting from some compact initial configuration, the nodes in the network can spread out such that the area 'covered' by the network is maximized. Howard et al. present a potential-field-based approach to deployment. In their work [14], the fields are constructed such that each node is repelled by both obstacles and by other nodes, thereby forcing the network to spread itself throughout the environment.

Recently, researchers are interested in leveraging the mobility of robotic sensor nodes to provide targeted coverage without resorting for exhaustive coverage. In [15], Stump et al. discuss the problems like finding the minimum number of robots and their locations so as to establish communication of a set of targets with a base station and solve for deployment of robots to maintain the communication links of moving targets with the base station. Another interesting problem arises when the network of robots is assigned to monitor a set of stationary targets. If the number of targets to be monitored is more than the available robotic sensor nodes, then it is possible to be in a situation where all the targets cannot be monitored all the times. In such scenarios few robots are required to actively visit more than one location to monitor multiple targets. Tokekar et al. present path planning techniques which compute paths for all robots such that every target is visible from at least one path [16].

2.2 Adaptive sampling

Generally, the problem of multi-robot exploration can be stated as follows [17]: n identical robots set out to explore an unknown area, each robot is equipped with sensing, localization, mapping, and limited-range communication capability. Designing a coordinated exploration algorithm to carry out the mission reliably and quickly has been the focus of many recent approaches. Adaptive sampling refers to strategies [18] in which the procedure for selecting locations to be included in robot paths depends on the sampling data observed during exploration. When the environmental phenomena are smoothly varying without any local-maxima peaks, non-adaptive strategies are known to perform well [19]. However, if the environment contains high local-variance, adaptive sampling can exploit the clustering phenomena to map the environmental field more accurately than non-adaptive sampling. Low et al. present an adaptive multi-robot exploration strategy that can perform both wide-area coverage and adaptive sampling of high local-variance regions using non-myopic path planning [18]. A key feature of this approach is in covering the entire adaptivity spectrum, thus allowing strategies of varying adaptivity to be formed and theoretically analyzed in their performance. They conclude that a more adaptive strategy improves mapping accuracy.

In contrast to random exploration of the environmental field [20], directed exploration selects robot paths to observe regions of high uncertainty. Directed exploration strategies that focus on feature sampling expect areas of high uncertainty to contain highly-varying measurements [10] [21]. Another approach to seek the trade-off between cost and information is to use down-sampling. Compressive sensing, a recently developed down sampling and reconstructing method yielding a sub-Nyquist sampling criterion, uses condensed linear measurements for reconstruction under a sparse domain without losing useful information [22] [23] [24].

2.3 Mobile Sensor Network Localization

Location awareness is important, as applications such as environment monitoring, target tracking and intrusion detection need to know the locations of mobile node precisely. Existing approaches for network localization include two categories: range based approaches [25] and range free approaches [26]. Range based localization can provide more accurate position estimates compared to range free localization, however, it costs more due to special hardware required. Localization methods for mobile sensor networks in general use a combination of beacon nodes, whose location information is known, and sensor nodes with unknown locations. Sensor nodes compute their locations on receiving broadcasting messages from the beacon nodes. The scenario with mobile beacon nodes and mobile sensor nodes is discussed by Zhang et al. in [27], where Monte Carlo Localization method is used to achieve high sampling efficiency and localization accuracy.

Multi robot SLAM has mostly been addressed in data fusion aspect characterized by two major sources of uncertainty due to the noise in sensing and motion without considering controlled mobility [28]. Extended Kalman filters and particle filters have been successfully implemented for data fusion of multi robot SLAM [29]. Active SLAM for multiple cooperating agents on trajectory control is developed, where each agent shares map information over a data fusion network [30].

2.4 Multi-robot coordination and communication

Good spatial distribution, good overall performance, robustness, reliability, flexibility, concurrency, and versatility are few of the advantages of a multi-robot system over a single robot listed by Yan et al. in their survey on multi-robot systems [31]. Even though they have numerous advantages, multi-robot systems do have costs. One such overhead is coordination and communication between the robots. The overall system performance can be directly affected by the quality of coordination and control. Static coordination [32] refers to an offline set of conventions that are decided before starting the task. Dynamic coordination [33] [34] occurs online during the task execution and is based on the analysis of the observations and current state of the system. Both these approaches have their own pros and cons. For example, static method can handle complex tasks, but its real-time controlling can be poor, whereas dynamic method can well meet the capability of real-time, but it has difficulty in dealing with more complex tasks. Hence, in practice it is common to see the combination of both static and dynamic coordination methods.

Communication, as a means of coordination, enables robots to share position information, state of the environment, sensor measurements, and also enable individual robot to learn about the intentions, goals, and actions of other robots. Yan et al. classify the communication structure based on the information transfer modes, namely explicit and implicit communication [31]. Explicit communication refers to the means for the direct exchange of information between the robots. Klavins [35] introduced a notion of communication complexity as a means to investigate the scalability of multi-robot algorithms in terms of how much coordination they require. Klavins also presented several communication schemes that cover several natural communication complexity classes from $O(n^2)$ communication to $O(1)$ communication. Rekleitis et al. propose a multi-robot coverage path planning for a team of robots with limited communication, where the environment is divided into strips and each strip is explored by a single robot, while the others remain stationary to observe the moving robot and estimate its position [36]. While this has the advantage of improving the overall accuracy of the map, it does nothing to speed the exploration process. On the contrary, the robots are forced to remain near each other in order to stay visible. Implicit communication refers to the way in which the robot gets information about other robots in the system through the environment. Pagello et al. [37] present an approach for coordinating a team of soccer playing robots through implicit communication where the cooperation between the robots is based on the form of the observed behavior of other robots. We plan to apply both these techniques together to complement each other, thus we can achieve both accuracy and stability provided by explicit and implicit communications respectively.

Decision-making can be regarded as a cognitive process resulting in the selection of a course of action among several alternative scenarios. Two popular classes of approach for decision making in multi-robot scenarios are centralized [38] [39] and decentralized [40] [41]. The centralized approach faces challenges in scalability, robustness and vulnerability. Decentralized approach can be further divided into distributed architectures and hierarchical architectures. There is no central control agent in distributed architectures [42], such that all the robots are equal with respect to control and are completely autonomous in the decision-making process. In hierarchical architectures, there

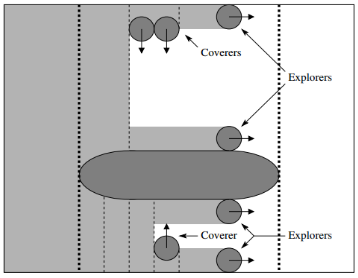

(a) Two sub-teams of robots encircling an obstacle [36].

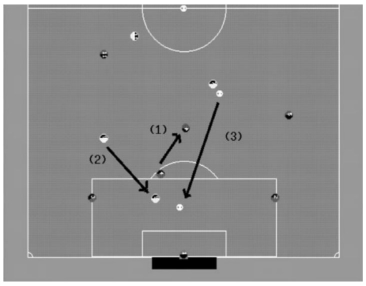

(b) Emergent pass behavior [37].

Figure 3: Multi-robot coordination.

exist one or more local central control agents which organize robots into clusters. The hierarchical architecture [43] is a hybrid architecture, intermediate between a centralized architecture and a distributed architecture. We will discuss our approach to distributed coordination in later sections.

2.5 Multi-target tracking and modeling

Knowledge about the position of moving objects can be used to improve the behavior of the system. While Kalman filters have been shown to provide highly efficient state estimates, they are restricted to Gaussian distributions over the state to be estimated. Particle filters have been introduced to estimate non-Gaussian, non-linear dynamic processes. Schulz et al. introduce a sample-based variant of joint probabilistic data association filters (JPDAFs) to track features originating from individual objects and to solve the correspondence problem between the detected features and the filters [44]. They use particle filters to track the states of the objects and applies JPDAFs to assign the measurements to the individual objects. Khan et al. present a Markov chain Monte Carlo based particle filter that effectively deals with interacting targets, i.e. targets that are influenced by the proximity and/or behavior of other targets [45].

In a multi-target environment, the number of targets also changes with time along with the states of the targets. This is because of the targets appearing in and disappearing off the region of interest. Often, the sensors are not able to detect all the existing targets. Moreover, the sensor also receives a set of spurious measurements (clutter) not originating from any target. As a result, the observation set at each time step is a collection of indistinguishable partial observations, only some of which are generated by targets. Out objective is to jointly estimate the time-varying number of target states from a sequence of observation sets in the presence of data association uncertainty, detection uncertainty, noise, and false alarms.

Uncertainties in association of measurements with appropriate targets is an intrinsic problem in multi-target tracking. Even if the sensor observes all targets and receives no clutter, single-target filtering methods are not applicable since there is no information about which target generated which observation. Most traditional multiple-target tracking formulations involve explicit associations between measurements and targets [44] [46]. The data association problem has been tackled by various formulations [47]. We propose to use random finite set (RFS) formulation as it is an emerging and promising approach as compared to the traditional association-based methods. In the RFS formulation, the collection of individual target states is treated as a set-valued state, and the collection of individual observations is treated as a set-valued observation. The problem of estimating states of multiple targets in the presence of clutter and association Uncertainties can be cast in a Bayesian filtering framework by modeling set-valued states and set-valued observations as RFSs. This approach to multiple-target tracking is an elegant generalization of the single-target Bayes filter.

3 Topic I : Efficient sensor network configuration

In this section, we discuss the problem of configuring the networks of robotic sensors in an efficient topology such that the information gain is maximized and the energy spent by the network is minimized. We propose an information driven approach for covering the region of interest by a network of robots. In this problem, we assume that the map of the physical world is known and the goal is to cover the region and efficiently reconstruct the spatial phenomenon in that region.

3.1 Efficient coverage

Coverage in mobile robotics is the task of determining a path that passes over all points of an area or volume of interest while avoiding obstacles. Our goal is to sense a particular region of interest in the environment and be able to reconstruct the measured environmental field using a network of robotic sensors. Since there are autonomous agents involved in coverage there is a need to consider the cost involved in terms of energy consumed and time required to finish the task. A good map of a scalar field requires complete coverage of the region or a good sparse coverage strategy along with an efficient interpolation technique. We propose to optimize the trade off between the environmental field mapping and the cost associated with the sensing.

3.1.1 Adaptive / Selective coverage

Adaptive coverage refers to covering a particular region of interest one step at a time, making the decision about the next step based on the observations made at the current step. This can (in some cases) be formalized as a reinforcement learning task. The agent has to make a sequence of decisions that depend on its observations. The environment is stochastic (because the underlying spatial phenomena are unknown). There is also a well defined reward function, and actions influence the rewards to be collected in the future. The state space in our case will constitute the locations ($x \in X$) of the agent in the world, the actions will make the agent transit from one location ($x \in X$) to another ($x' \in X$). The representation of rewards or costs can be achieved in multiple ways. We propose to model the rewards such that two different kinds of coverages are achieved.

One of the ways to approach adaptive sampling is to sample the world adaptively such that the uncertainty in spatial phenomenon is reduced with time. This encourages the agent to come up with an optimal policy to explore the region such that the spatial phenomenon map is complete. A random variable X_v is associated with each location $x \in X$. The joint distribution $P(Z_X)$ can then be used to quantify uncertainty in the prediction $P(Z_{X\setminus A}|Z_A = z_A)$ of phenomena at unobserved locations $Z_{X\setminus A}$, after making observations $Z_A = z_A$ at a small subset A of locations. To quantify this uncertainty the mutual information (MI) criterion can be used. For a set of locations, X, the MI criterion is defined as:

$$MI(A) = H(Z_{X\backslash A}) - H(Z_{X\backslash A}|Z_A), \tag{1}$$

where $H(Z_{X\backslash A})$ is the entropy of unobserved locations $X\backslash A$, and $H(Z_{X\backslash A}|X_A)$ is the conditional entropy of locations $X\backslash A$ after sensing at locations A. Hence mutual information measures the reduction in uncertainty at the unobserved locations. Therefore, in our approach, we would like to select locations that most reduce the uncertainty in spatial phenomenon prediction for the environment. Considering mutual information ($MI(A)$) as the reward function, we can apply Markov decision process (MDP) solvers, like Q-learning and Value iteration, to solve for an optimal policy that provides a sequence of locations to visit such that the uncertainty in the spatial map is reduced over time.

In our recent work [48], we presented an anytime algorithm that selectively covers a region with varying resolution. This approach is very useful when there is a partial prior knowledge about the environmental field. The region of interest is discretized into grid cells and each grid-cell is assigned a utility value equivalent to the integral of the underlying prior over that cell. The rewards are such that higher rewards imply regions with high interest, i.e. a hotspot-region. The goal is the plan a path for an agent to cover hotspot-regions with high resolution and remaining region with lower resolution. We used value iteration to compute the best action to be taken at a given state. Value iteration is a method of computing an optimal MDP policy. It computes the optimal value of a state $V^*(x)$, i.e. the expected discounted sum of rewards that the agent will achieve if it starts at that state and executes the optimal policy $\pi^*(x)$.

The optimal value function $V^*(x)$, $\forall x \in X$, is defined by the Bellman equation [49],

$$V^*(x) = \max_a \left(R(s,a) + \gamma \sum_{x' \in X} P(x'|x,a) V^*(x') \right), \tag{2}$$

where γ is a discount factor. Thus according to Eq.2, the value of a state x is the sum of instantaneous reward $R(s,a)$ and the expected discounted value of the next state $V^*(x')$, when the best available action ($a \in A$) is used. Optimal policy ($pi^*(s)$) defines an action for every state that achieves the optimal value. Given the optimal value function for all states, optimal policy is defined by,

$$\pi^*(s) = arg \max_a \left(R(s,a) + \gamma \sum_{s' \in S} P(s'|s,a) V^*(s') \right) \tag{3}$$

In this approach, we clear the rewards as and when the corresponding state is visited. Thus the reward function is changing over time as the agent clears the rewards. Even though this violates the Markov assumption over the entire coverage task, every state transition still holds good to be formulated as a one step MDP, where every state transition of the agent is modeled as MDP in a new world and the value function is computed over the updated rewards of the world. Thus the convergence of the value iteration technique still holds good for every state transition. This method of re-evaluating the utilities of all states iteratively for every step becomes computationally expensive and unrealistic to run on-line. Hence, we propose an approximation of updating the utilities of the states. The value at any state is influenced more by its neighborhood than by a state that is very distant from it. Hence, for a one-step transition we can achieve good utility updates by just considering a neighborhood subset of the state space ($S_n \subset S$) instead of complete state space (S). We have also extended this algorithm to adapt the actions based on the operational challenges posed by the environmental conditions. In [50] we extended this work to plan paths for a team of robots to adaptively explore the region. Fig.4 shows the robot trajectories generated by our algorithm to selectively cover the region of interest.

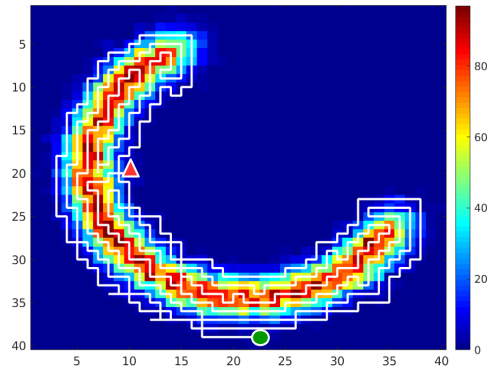

(a) Single robot trajectory [48]. The red triangle indicates the start and green circle indicates the end of the trajectory.

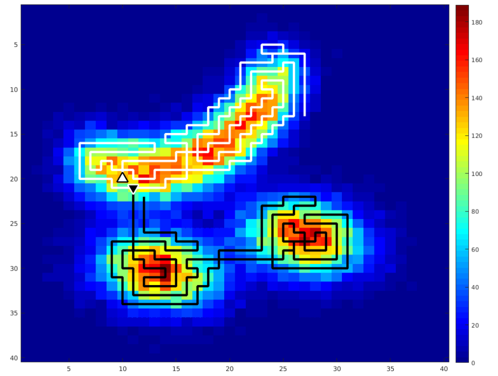

(b) Trajectories for two robots (black and white) [50].

Figure 4: Trajectories generated by selective coverage algorithm [48] overlaid on underlying reward distributions. The color-bars indicate the rewards.

3.1.2 Efficient reconstruction

Efficient coverage needs an efficient way to reconstruct the spatial field with as sparse samples as possible. In this section we will discuss some of the techniques that we plan to use to address the problem of field representation and reconstruction. Field reconstruction in a sensor network setup needs to deal with noisy measurements, unevenly distributed observations, and fill small gaps in the data with high confidence while assigning higher predictive uncertainty in sparsely sampled areas. One of the techniques to naturally handle these necessities is Gaussian Process (GP) models. GP models also have the advantage of not assuming a fixed discretization of the space and of additionally providing predictive uncertainties. The explicit model of uncertainty that a GP provides has led to their successfully application in a wide range of other robotic applications. Hence, we propose to model the distribution of valuable/interesting locations in the world as a Gaussian Process: a probabilistic representation of a scalar field [51]. In this modeling of the spatial phenomena, the value of the field $f(x)$ at each location x is estimated by a Gaussian distribution with mean $m(x)$ and a covariance function or the kernel $k(x, x')$:

$$f(x) \sim \mathcal{GP}(m(x), k(x, x')) \tag{4}$$

The assumptions about the smoothness of the estimated field are incorporated in the covariance function. Input to the model is a set of noisy observations $\{(x_i^*, f_i^*)|i = 1, 2, ..., n\}$, where $f_i^* = f(x_i^*) + \epsilon_i$ and $\epsilon_i \sim \mathcal{N}(0, \sigma_n^2)$. A Gaussian measurement noise with zero mean is assumed over the observations. The predicted value for f at new location x is estimated as a Gaussian distribution:

$$f'(x) = \mathbf{k}(x)^T (K + \sigma_n^2 I)^{-1} \mathbf{f}^* \tag{5}$$

$$Var[f'] = k(x, x) - \mathbf{k}(x)^T (K + \sigma_n^2 I)^{-1} \mathbf{k}(x) \tag{6}$$

where $\mathbf{k}(x) = [k(x_1, x), .., k(x_n, x)]^T$, $K = [k(x, x')]_{x,x' \in \{x_i^*\}}$, and $\mathbf{f}^* = [f_1^*, f_2^*, .., f_n^*]^T$. The form of the kernel function $k(x, x')$ needs to be chosen such that it reflects the assumptions about the scalar field. We plan to use the *radial basis kernel function* with the form,

$$k(x, x') = exp\left(-\frac{||x - x'||^2}{2l^2}\right). \tag{7}$$

The characteristic length-scale parameter l specifies how quickly the value of the scalar field becomes uncorrelated with distance. The global noise variance σ_n^2, and the length-scale l are known as the hyperparameters of the process and represented as $\Theta = (\sigma_n, l)$. To best represent the underlying data, the hyperparameters Θ need to be adapted. One of the ways to achieve this is by maximizing the marginal log likelihood of the training data w.r.t. the hyperparameters. In our recent work [48], we used Gaussian processes to create depth profile of the sea-floor using sparse altitude measurements from an autonomous surface vehicle. We also analyzed the variation in the accuracy of the depth map in accordance with the sparsity of measurements (Fig.5). This gives us an idea about coming up with a trade-off between building an accurate map and minimizing the cost associated with sampling. We also used similar approach with other water quality sensor measurements. These maps are very useful in representing the environment field and figuring out the local-maxima peaks for further planning of the robotic network.

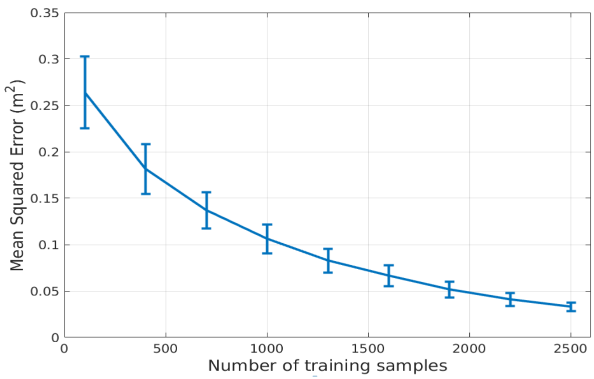

Figure 5: Efficiency of GP model over no. of samples [48].

One of the limitations of standard GP framework is the assumption of constant length-scale over the whole input space. Intuitively, the length-scales describe the area in which observations strongly influence each other [52]. For environmental fields with or without local-maxima peaks, we propose to use locally varying length-scales to account for different situations. In our case, where we want to sample spatial fields with varying resolutions, high-resolution at high local-variance regions and lower resolution at low local-variance regions, we plan to use low length-scale for regions with high sampling density and higher length-scale when the sampling density is low. We will use an extension of the squared exponential covariance function [53] to address this problem of non-stationarity. The squared exponential covariance function takes the form,

$$k(x_i, x_j) = |\Sigma_i|^{\frac{1}{4}}|\Sigma_j|^{\frac{1}{4}}\left|\frac{\Sigma_i + \Sigma_j}{2}\right|^{-\frac{1}{2}} \times exp\left[-(x_i - x_j)^T\left(\frac{\Sigma_i + \Sigma_j}{2}\right)^{-1}(x_i - x_j)\right] \tag{8}$$

Each input location x' is assigned a local Gaussian kernel matrix Σ' and the covariance between two targets f_i and f_j is calculated by averaging between the two local kernels at the input locations x_i and x_j. Thus the modeled covariance of the corresponding target values is influenced by the local characteristics at both locations (x_i and x_j). In addition, to achieve the smoothing behavior, we plan to calculate the local length-scale following the techniques proposed by Lang et al. [54]. For an input location x_i, Lang et al. estimate the gradient $(\bigtriangledown f)_i$ from the observations in the local neighborhood. Thus we can adapt the the local kernels Σ_i using gradient information. But for a network of robots to achieve reconstruction of the map of a region in a distributed way we cannot directly apply GPs because of their unfavorable scaling: $\mathcal{O}(N^3)$ in time and $\mathcal{O}(N^2)$ in space, where N is the number of training samples. Various approximation methods have been proposed in the literature to overcome these limitations [55] [56]. Most of these methods fall within sparse approximation where a low-rank approximation is applied to the covariance matrix of the GP prior using a smaller subset of M ($<< N$) inducing variables [57]. In this case complexity in time and memory space are reduced to $\mathcal{O}(M^2N)$ and $\mathcal{O}(MN)$, respectively.

We plan to build on tiling approximation strategy proposed by Plagemann et al. [52], in which they split the input space into overlapping rectangular segments and assign an individual GP model to each of the segments. This sub-model is then only provided with observations from within its segment. For a prediction at input location x, the GP segment that is most likely to have best approximation for x is used to predict the function value (f) at x. In our work, the approach is to assign segments to different robotic sensor nodes based on their locations and use these segments to generate the observation maps around the nodes helping in decision about the next potential location to visit.

3.2 Multi-agent coordination

In this section we will discuss the coordination techniques that we plan to consider for the optimal behavior of the network of sensing robots. As discussed in Section 2.4, centralized decision making is capable of producing optimal plans, but it suffers from the cons like non-scalability, vulnerability, and non-robustness. In our multi-robot setup we would like the robots to be capable of making decisions for themselves and still come up with an optimal plan for the group task. In our recent work [50], we used a hybrid decision making architecture with an assumption of full communication between the robots. We present an algorithm to plan paths for multiple robots to selectively cover the region of interest such that the information gain per unit time is maximized. We used a shared memory between the robots where the robots could see the complete current state of the environment and also their neighbor's position on the map. But, each robot decided its own optimal policy for its next action based on the current state of the world and its neighbors. Thus we tried a hybrid method between centralized and distributed control architectures. We achieved this by assigning a reward function for robot r_k proportional to the distance of the grid-cell location l_i from other $n-1$ neighboring robots and inversely proportional to the distance of the grid-cell location l_i from its own location l_{r_k}, thus encouraging robot k to visit locations that are close to the l_{r_k} and farther from other neighboring agents,

$$R(l_{r_k}, l_i) \propto \frac{\sum\limits_{\substack{j \in (1,2,..n) \\ j \neq k}} distance(l_{r_j}, l_i)}{distance(l_k, l_i)}. \tag{9}$$

There has been many approaches to multiple robot exploration problem: closest frontier cell approach [58], rendezvous-based strategies [59], sub-region based exploration [36], but with less focus on coverage efficiency and overall mission efficiency. Recently, negotiation-based coordination algorithms [60] [61] [40] have shown a reduction in time and travel distance in multi-robot missions because of their capabilities to address both coordination and task distribution at once and present task-independent strategies for multi-robot coordination. We plan to approach the distributed coordination problem by using distributed bidding algorithm [17]. Sheng et al. propose a nearness measure,

$$\lambda_i = e^{-\frac{d_1}{r_c}} + \alpha e^{-\frac{d_2}{r_c}} + \cdots + \alpha^{n-2} e^{-\frac{d_{n-1}}{r_c}} \tag{10}$$

where, λ_i is the nearness measure of the $i^t h$ robot, $d_1 \leq d_2 \leq \cdots \leq d_{n-1}$ are the distances $d(R_i, R_j)$ with $j \in (1, 2, ..., n)$, n is the number of robots within the same subnetwork, and r_c is the communication range. $\alpha(< 1)$ is a positive fading factor that will help to disperse the robots and maximize the coverage area.

In a given task, every robot can compute a measure for task-success (S_a) with the next action a based on the current state of the world, current position and state of other task-mate robots, and target actions chosen by its task-mates. For example, in the case of exploration the robot calculates the information gain I_i for each frontier cell i. The cost (C_a) associated with each action of the robot is computed. In our example problem, cost would be the distance of the frontier cell i from the robot's current location. Lastly robot's nearness (λ) is computed using the Eq.10. The total gain of the robot if action a is chosen is given by,

$$g_a = \omega_1 S_a - \omega_2 C_a + \omega_3 \lambda \tag{11}$$

where $\omega_i, i \in (1, 2, 3)$ represent positive weights. Now the maximum gain over all possible actions will be placed as a bid between all the robots. The bid B of a bidding robot is given by,

$$B = \max_a (g_a). \tag{12}$$

The bidding robot broadcasts this bid value to all of its neighbors and waits for a fixed time t_{bid} to hear a response back. If there is no other robot bidding during this period or no other robot provide a better bid, the robot goes ahead with the decided action. If there are any new changes to the state of the world during the bidding time period, then the bidding robot recalculates its bid value. And if there is any other robot who bids a better bid value, then the robot waits for the winner's identification and includes the destination state of the winner robot to calculate its next bid values. Obviously the bidding wait period t_{bid} is an important parameter. If t_{bid} is very small, it is highly possible that fewer robots will participate in the bidding. If t_{bid} is very big, it is very likely that many robots will participate in the bidding, which implies a better bid will be generated but the waiting time will be longer. Hence, the choice of the t_{bid} has to be made very carefully and should be backed by proper analytical backing.

4 Topic II : Reconfiguration of sensor networks to monitor spatio-temporal fields

Persistent monitoring of a spatial phenomenon that is varying temporally (Fig.6) requires on-line adaptation of the topology of the network of robots because the measurements by stationary robots cannot capture the dynamics of the spatio-temporal fields. To achieve this, the autonomous agents should be capable of estimating the temporal evolution of the phenomenon. In this section we discuss techniques to model the changing field by tracking the targets that move around in the region of interest. We are not interested in tracking each target separately. Instead we want to build a continuous field based on the density of the targets in the sub-regions and model the changes in this field as the targets move over time.

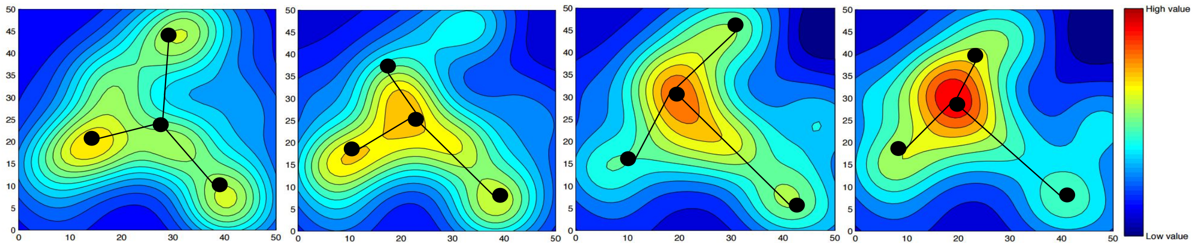

Figure 6: An example of a dynamic spatio-temporal field modeled using Gaussian radial basis functions is shown at different time instants. The black dots and lines represent the expected sensor network topologies.

4.1 Multi-target motion modeling

Consider a state space $\mathcal{X}$ with transition density $f_{t|t-1}(x_t|x_{t-1})$ from time $t-1$ to t and observation space $\mathcal{Z}$ with a probability density $g_t(z_t|x_t)$ of receiving the observation z_t. For multiple-target tracking, we represent the collection of target states and measurements at time t as finite sets,

$$X_t = \{x_{t,1}, x_{t,2}, ..., x_{t,m(t)}\} \in \mathcal{F}(\mathcal{X}) \tag{13}$$

$$Z_t = \{z_{t,1}, z_{t,2}, ..., z_{t,n(t)}\} \in \mathcal{F}(\mathcal{Z}) \tag{14}$$

where $m(t)$ is the number of targets in the region og interest at time t, $n(t)$ is the number of observations made at time t, and $\mathcal{F}(\mathcal{X})$ and $\mathcal{F}(\mathcal{Z})$ are the collections of all finite subsets of $\mathcal{X}$ and $\mathcal{Z}$. The multiple-target tracking problem can then be posed as a filtering problem with (multiple-target) state space $\mathcal{F}(\mathcal{X})$ and observation space $\mathcal{F}(\mathcal{Z})$. Analogous to single target system, uncertainty is characterized by modeling multiple-target state X_t and measurement Z_t as random finite sets (RFSs). The Bayes filter recursion to propagate the multiple-target posterior in time is given by,

$$p_{t|t-1}(X_t|Z_{1:t-1}) = \int f_{t|t-1}(X_t|X)p_{t-1}(X|Z_{1:t-1})\mu_s(dX) \tag{15}$$

$$p_t(X_t|Z_{1:t}) = \frac{g_t(Z_t|X_t)p_{t|t-1}(X_t|Z_{1:t-1})}{\int g_t(Z_t|X)p_{t|t-1}(X|Z_{1:t-1}\mu_s(dX)} \tag{16}$$

where μ_s is an appropriate reference measure of $\mathcal{F}(\mathcal{X})$ [62]. The recursion Eqs. 15 and 16 involve multiple integrals on the space $\mathcal{F}(\mathcal{X})$, which are computationally intractable. The combinatorial nature of the multiple-target densities and the multiple integrations on the (infinite dimensional) multiple-target state space make the multi-target Bayes filter computationally intractable. *Probability hypothesis density filter* (PHD filter) is an approximation developed to alleviate the computational intractability in the multiple-target Bayes filter. PHD filter propagates the first-order statistical moment, or intensity, of the RFS of states in time rather than propagating the multiple-target posterior density in time [47].

The first order moment (*intensity*) for an RFS X on $\mathcal{X}$ with probability distribution P is a nonnegative function v on $\mathcal{X}$ such that for each region $S \subseteq \mathcal{X}$,

$$\int |X \cap S| P(dX) = \int_S v(x)dx, \tag{17}$$

i.e. the integral of v over any region S gives the expected number of elements of X that are in S. Given a multiple-target state X_{t-1} at time $t-1$, the state X_t is given by the union of surviving targets, the spawned targets, and the spontaneous births.

$$X_t = \left[\bigcup_{\zeta \in X_{t-1}} S_{t|t-1}(\zeta)\right] \cup \left[\bigcup_{\zeta \in X_{t-1}} B_{t|t-1}(\zeta)\right] \cup \Gamma_t \tag{18}$$

where, Γ_t is the RFS of spontaneous birth at time t, and $B_{t|t-1}(\zeta)$ is the RFS of targets spawned at time t from a target with previous state ζ. Using the intensities of these RFSs, i.e. $\gamma_t(\cdot)$ is the intensity of RFS Γ_t, $\beta_{t|t-1}(\cdot|\zeta)$ is the intensity of RFS $B_{t|t-1}(\zeta)$, $P_{S,t}(\zeta)$ is the probability that a target still exists at time t given that its previous state is ζ, $P_{D,t}(x)$ is the probability of detection, and $\kappa_t(\cdot)$ is the intensity of clutter RFS K_t, it can be shown that the posterior intensity can be propagated in time under PHD recursion [47],

$$v_{t|t-1}(x) = \int P_{S,t} f_{t|t-1}(x|\zeta) v_{t-1}(\zeta) d\zeta + \int \beta_{t|t-1}(x|\zeta) v_{t-1}(\zeta) d\zeta + \gamma_t(x) \tag{19}$$

$$v_t(x) = [1 - P_{D,t}(x)] v_{t|t-1}(x) + \sum_{z \in Z_t} \frac{P_{D,t}(x) g_t(z|x) v_{t|t-1}(x)}{\kappa_t(z) + \int P_{D,t}(\xi) g_t(z|\xi) v_{t|t-1}(\xi)}. \tag{20}$$

The posterior intensity as seen in Eq.19 and Eq.20 is a function on the single-target state space $\mathcal{X}$ unlike the multiple target recursion that operates on $\mathcal{F}(\mathcal{X})$. Hence, the PHD recursion requires much less computational power. Particle-PHD filter has been a popular implementation due to its ability to deal with time-varying number of nonlinear targets with relatively low complexity. However, the large number of particles and unreliable clustering techniques for extracting state estimates leave this approach in disadvantage. Instead we plan to use Gaussian mixture PHD filters [63] that propagate posterior intensity in time as measurements arrive. The PHD recursion does not admit closed-form solutions in general, and numerical integration suffers from the *curse of dimensionality*. Vo et al. in their work [63], successfully demonstrated that for a certain class of multiple-target models, i.e. linear Gaussian multiple-target models, the PHD recursion in Eq.19 and Eq.20 admits a closed-form solution. They provide proof showing that if the posterior intensity at time $t-1$ is a Gaussian mixture, then the predicted intensity for time t is also a Gaussian mixture and this further implies the posterior intensity at time t is also a Gaussian mixture. Given the Gaussian mixture intensities $v_{t|t-1}$ and v_t, the corresponding expected number of targets $\hat{N}_{t|t-1}$ and $\hat{N}_t$ can be obtained by summing up the appropriate weights.

In the multiple-agent case, we will use a similar model where each of the agent is assigned a sub-region within the region of interest and each agent tracks the number of targets present in its particular region at any given time. Based on this number of targets in its region, every agent decides to explore neighboring regions or continue to stay in their own region. We also plan to merge the field map and get overall trajectories and distribution of the mobile targets. Once we have the targets tracked, decision about the topology of the network is made such that the target coverage is maximized. We will discuss more on the reconfiguration of the network in the next section.

4.2 Data merging and network reconfiguration

Each robot in the network is assigned with a sub-region around it for which the robot is responsible to generate a target density map. The region-based approach using the robot and target densities does not always maximize the total coverage in a region; for example, if the sensor ranges of robots in a region overlap, the total coverage within the region is often small even though there are enough robots in the region. In order to obtain a good robot spread within a region, a visibility maximization method can be used. The visibility of a region R is defined as [64]:

$$Visibility(R) = \frac{The covered area of region R}{Total area of region R}. \tag{21}$$

The movement of robots within a region is modified in order to pursue two goals in parallel: each robot tries to maximize the number of targets tracked, and to maximize the coverage of regions. As discussed in Section 4.1 we plan to use PHD filters to generate a target density map for every sub region around each robot. To obtain the density map for entire region of interest, the maps from sub regions need to be merged. But we need to consider the fact that there are regions in between the covered sub-regions where there is no data and also there will be regions where there is an overlap between the sub-regions. We propose to use Gaussian Processes with radial basis kernels (Section 3.1.2) to interpolate and merge the maps from all sub-regions into a global target-density map.

Once the global target density map is built, the aim is to optimally distribute the robotic sensor nodes such that the number of targets tracked is maximized. This problem has similar flavor as the ones explained in Section 3.1.1, where there is an underlying distribution and the agent needs to maximize the coverage for higher density regions. But one change is that the robot to target ratio needs to be considered in current problem. We cannot have all the robotic nodes concentrated in one sub-region. One way to handle this is by including the target-robot ratio of the sub-region into the reward function. The reward gained by robot r_k at location l_{r_k} by taking an action to go to location l_i can be represented according to,

$$R(l_{r_k}, l_i) \propto \frac{N_{t_i}}{N_{r_i}}, \tag{22}$$

where N_{t_i} and N_{r_i} represent the number of targets and the number of robots in the sub-region around location l_i respectively.

5 Proposed research plan

In this section we list the progress we have made to this date, and the proposed milestones to evaluate progress in our research plan.

5.1 Progress to date

We have discussed some of the recent work done so far in Section 3.1. Here we present a list that summarizes our current progress.

- **Sparse Sampling and Field Reconstruction:** Building a map of a field map of an environmental phenomenon requires an enormous amount of data and achieving this with an autonomous vehicle becomes challenging because of limited sensing range and battery life. Having an efficient reconstruction technique which can construct the field with sparse data points becomes very important. In one of our recent studies [48], we illustrate the performance of the Gaussian process regression technique over the number of training samples used for field reconstruction (Fig.5). We constructed a bathymetric map from sparse sonar point-data using a Gaussian process model. The results show that a decent reconstruction could be achieved even when the sampling rate is reduced to half.

- **Selective Coverage by Multiple-robots:** One of the essential features of a good coverage algorithm is to examine salient regions in the increasing order of the value-gain associated with them. This becomes significant when the task is to cover a given region in limited time. We present a value iteration based reward-driven finite-horizon algorithm to extract the maximum amount of valuable data in the least amount of time [48]. We extended this approach to multi-robot scenario [50] where we explored some of the multi-robot coordination strategies as discusses in Section 3.2.

- **Search and Exploration strategies:** In a recent collaboration [65], we analyzed three search strategies to efficiently search passively moving targets by modeling their distribution and motion. We provided a comparison of performance in terms of their success rate and the cost of search. In this work, we also presented performance bounds for guaranteed capture of the target and minimum time of capture. This study will be of use when we have heterogeneous platforms in our network and have the active agents looking for passive nodes to exchange information.

- **Data Correlation in Heterogeneous Robotic setup:** Another collaboration with Li et al. [66] introduced an array of data mapping tools like Zernike-moment image matching, and Fabmap image classification. Through this work we compared and correlated the data collected by an ensemble of heterogeneous autonomous sensor systems to produce a comprehensive view of the health of the coral reef.

- **Frontier based Exploration and Task distribution:** Recently in a summer school we implemented a frontier based exploration algorithm which aims to minimize the map-uncertainty by predicting the actions based on Gaussian Process predictions. We also tried a task coordination experiment with two robots sharing the tasks of water-surface exploration and water-sampling. We modeled the water sampling task as a secretary hiring problem where the water from particular spot is sampled only when the required water-quality threshold is crossed. We are still working on implementation for a multi-robot scenario and expect to publish this work in Fall 2016.

5.2 Proposed Milestones

In this thesis we plan to develop a framework for a network of robots to autonomously plan and update their configuration such that they can efficiently observe and track spatially varying scalar field(s) that express environmentally important phenomena (e.g. temperature) and also discrete target groups (e.g. group of wifi-hotspot users). Our approach will be based on using a combination of topological representation of the robot network, and probabilistic model of the phenomena of interest. Following is the list of problems that lie within the interest of this thesis,

- Sensor network monitoring and tracking 1 target.
- Sensor network monitoring and tracking 2 targets.
- $N_t \leq N_r$ (Number of targets (N_t) is less than or equal to the number of robots(N_r)).
- $N_r < N_t < 2N_r$ (Number of targets is greater than the number of robots, but not too big).
- $N_t >> N_r$ (Continuous field).
- $N_h \leq N_r$ (Number of local-maxima peaks (N_h) is less than or equal to the number of robots(N_r)).

Another dimension can be added to the list of problems by considering the static and dynamic targets. Thus there is a lot of problems in the robotic sensor network problem space that the thesis aims to address. We created milestones in order to evaluate our progress towards the thesis goals presented in Sections 3 and 4. In this section, we propose a tentative list of milestones that will guide us towards our research goals. This list will evolve according to any new developments we make.

- **Multi-robot Selective Coverage Experiments on Real Robots:** The value-iteration based coverage algorithm [48] discussed in Section 3.1.1 has been tested on a single Kingfisher platform (Fig.7(a)). We plan to validate this approach in a multiple-robot [50] setup on real-robots. The plan is to use autonomous surface vehicles (Fig.7(a)) to map the water-surface phenomena and terrestrial robots (Husky platforms Fig.7(b)) to map the spatial fields.

(a) Autonomous Surface Vehicle (Kingfisher robot) (b) Terrestrial Vehicle (Husky robot) (c) Indoor platform (Turtlebot robot)

Figure 7: Robots planned to be used in our experiments.

- **Evaluation of Multi-robot Selective Coverage:** As discussed in Section 3.1.2, we plan to implement varying length-scale Gaussian Processes based on the sampling resolution and rugosity in spatial fields. We want to use this reconstruction technique as an evaluation for our online selective/adaptive coverage algorithms.

- **Coordination between Multiple Robotic Sensor Nodes:** In our recent work [50], we presented a distributed decision making system with continuous communication. We want to explore three different models of communication in multi-robot coordination: Continuous communication, Range limited communication, and No communication. We proposed an auction based approach for coordination between multiple-robots in Section 3.2. In this approach the robots use a range-limited communication model. We plan to implement a frontier based multi-robot exploration algorithm presented in [17] with a auctioning strategy for coordination between the robots. For no-communication model of coordination, we plan to decompose the region of interest into multiple sub-regions and using bidding mechanism to assign robots to these sub-regions. Then we use a distributed decision making mechanism to decide actions for each of these robots within their sub-region.

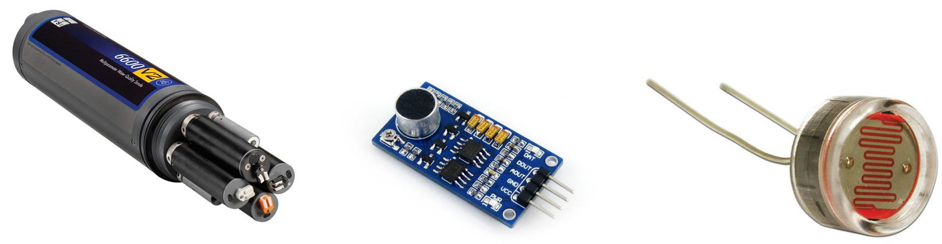

(a) Water quality measurement sensors

(b) Audio sensor

(c) Photo sensor

Figure 8: Sensors to be used in exploration field experiments.

  For experimental validation, we plan to use autonomous surface vehicles (Fig.7(a)) and terrestrial robots (Husky platforms Fig.7(b)). The goal of exploration experiments will be to map the surface water quality measures collected using water-quality sensors (Fig.8(a)). With the terrestrial exploration we can think of applications like mapping the audio field or the light field in the given region of interest using audio and photo sensors (Fig.8(b) and (c)).

- **Region-based Probability Hypothesis Density filters for Distributed Tracking of multiple targets:** As proposed in Section 4.1, one of our milestones is to achieve a distributed tracking of evolution in multi-target density distribution using PHD filters. To achieve this we plan to implement a region-based target density map where each of the robot in the network is assigned a particular sub-region to track. Then the task is to merge the density distribution from all the robots to get a target-density map of the whole region. We plan to use Gaussian mixture PHD filter implementation presented in [63]. The plan for evaluation of this technique is to use simulations on ROS-Stage and ROS-Gazebo. We plan to use real robots in an extension of this methodology explained in the next milestone.

- **Reorganization of Robotic Network based on the Target Density Distribution:** In this milestone we plan to implement and validate the algorithms to adapt according to the temporally changing target density distribution. We plan to use an adaptive sampling approach similar to the one discussed in Section 3.1.1 to follow the target density instead of the underlying spatial field. Here, we are tracking the targets as continuous density distribution rather than discrete targets.

Hence, we can use similar adaptive sampling techniques used for continuous spatial fields. We propose an experiment involving mobile wifi-hotspots which change their configuration according to the motion of potential wifi users. We plan to use indoor setup for these experiments with robots like Turtlebots (Fig.7(c)) equipped with wifi-hotspots.

The following item is our "long shot goal", for which the progress on the prior tasks is critical, and could possibly not be a part of this thesis.

- **Network of Heterogeneous Robots:** We are interested in exploring similar robotic sensor network problems when the sensor robots are all not the same (heterogeneous). Operating with heterogeneous network of robots opens a plethora of interesting problems related to coordination, data fusion, task distribution, communication, and data sampling. We plan to solve problems involving robots with varying capabilities. For example, having a network of surface vehicles and underwater vehicles coordinating to monitor a coral reef by mapping the reef at different resolutions.

6 Conclusion

We present a summary of our planned contributions to Robotic Sensor Network framework. We discuss the challenges in achieving an autonomous network of robotic sensor nodes. The problems in robotic sensor network that have been of interest are related to tracking of either continuous spatial field or discrete targets. We present a background for these problems and also proposed our ideas to address both of these tracking problems simultaneously.

We propose to develop a framework for a network of robots to autonomously plan and update their configuration such that they can efficiently observe and track spatially varying scalar field(s) that express environmentally important phenomena (e.g. temperature) and also discrete target groups (e.g. group of wifi-hotspot users). Our approach will be based on using a combination of topological representation of the robot network, and probabilistic model of the phenomena of interest. We described the progress we have made upto this date and proposed milestones which will help us evaluate the progress towards our research goal.

References

[1] E. Fiorelli, N. E. Leonard, P. Bhatta, D. A. Paley, R. Bachmayer, and D. M. Fratantoni, "Multi-auv control and adaptive sampling in monterey bay," *IEEE Journal of Oceanic Engineering*, vol. 31, no. 4, pp. 935–948, 2006.

[2] N. E. Leonard, D. A. Paley, F. Lekien, R. Sepulchre, D. M. Fratantoni, and R. E. Davis, "Collective motion, sensor networks, and ocean sampling," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 48–74, 2007.

[3] E. J. Englund and N. Heravi, "Phased sampling for soil remediation," *Environmental and ecological statistics*, vol. 1, no. 3, pp. 247–263, 1994.

[4] M. A. Batalin, M. Rahimi, Y. Yu, D. Liu, A. Kansal, G. S. Sukhatme, W. J. Kaiser, M. Hansen, G. J. Pottie, M. Srivastava *et al.*, "Call and response: experiments in sampling the environment," in *Proceedings of the 2nd international conference on Embedded networked sensor systems.* ACM, 2004, pp. 25–38.

[5] W. Thompson, *Sampling rare or elusive species: concepts, designs, and techniques for estimating population parameters.* Island Press, 2013.

[6] H. Chang, A. Q. Fu, N. D. Le, and J. V. Zidek, "Designing environmental monitoring networks to measure extremes," *Environmental and Ecological Statistics*, vol. 14, no. 3, pp. 301–321, 2007.

[7] C. Veibäck, G. Hendeby, and F. Gustafsson, "Tracking of dolphins in a basin using a constrained motion model," in *Information Fusion (Fusion), 2015 18th International Conference on.* IEEE, 2015, pp. 1330–1337.

[8] D. S. Apostolopoulos, M. D. Wagner, B. N. Shamah, L. Pedersen, K. Shillcutt, and W. L. Whittaker, "Technology and field demonstration of robotic search for antarctic meteorites," *The International Journal of Robotics Research*, vol. 19, no. 11, pp. 1015–1032, 2000.

[9] R. Castano, R. C. Anderson, T. Estlin, D. DeCoste, F. Fisher, D. Gaines, D. Mazzoni, and M. Judd, *Rover traverse science for increased mission science return.* Pasadena, CA: Jet Propulsion Laboratory, National Aeronautics and Space Administration, 2003.

[10] K. H. Low, G. J. Gordon, J. M. Dolan, and P. Khosla, "Adaptive sampling for multi-robot wide-area exploration," in *Proceedings 2007 IEEE International Conference on Robotics and Automation.* IEEE, 2007, pp. 755–760.

[11] V. Formisano, S. Atreya, T. Encrenaz, N. Ignatiev, and M. Giuranna, "Detection of methane in the atmosphere of mars," *Science*, vol. 306, no. 5702, pp. 1758–1761, 2004.

[12] X. Wang, G. Xing, Y. Zhang, C. Lu, R. Pless, and C. Gill, "Integrated coverage and connectivity configuration in wireless sensor networks," in *Proceedings of the 1st international conference on Embedded networked sensor systems.* ACM, 2003, pp. 28–39.

[13] V. Isler, S. Kannan, and K. Daniilidis, "Sampling based sensor-network deployment," in *Intelligent Robots and Systems, 2004.(IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on*, vol. 2. IEEE, 2004, pp. 1780–1785.

[14] A. Howard, M. J. Matarić, and G. S. Sukhatme, "Mobile sensor network deployment using potential fields: A distributed, scalable solution to the area coverage problem," in *Distributed Autonomous Robotic Systems 5.* Springer, 2002, pp. 299–308.

[15] E. Stump, N. Michael, V. Kumar, and V. Isler, "Visibility-based deployment of robot formations for communication maintenance," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on.* IEEE, 2011, pp. 4498–4505.

[16] P. Tokekar and V. Kumar, "Visibility-based persistent monitoring with robot teams," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on.* IEEE, 2015, pp. 3387–3394.

[17] W. Sheng, Q. Yang, J. Tan, and N. Xi, "Distributed multi-robot coordination in area exploration," *Robotics and Autonomous Systems*, vol. 54, no. 12, pp. 945–955, 2006.

[18] K. H. Low, J. M. Dolan, and P. Khosla, "Adaptive multi-robot wide-area exploration and mapping," in *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1.* International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 23–30.

[19] A. Singh, R. Nowak, and P. Ramanathan, "Active learning for adaptive mobile sensing networks," in *Proceedings of the 5th international conference on Information processing in sensor networks.* ACM, 2006, pp. 60–68.

[20] R. McCartney and H. Sun, "Sampling and estimation by multiple robots," in *MultiAgent Systems, 2000. Proceedings. Fourth International Conference on.* IEEE, 2000, pp. 415–416.

[21] A. Singh, A. Krause, C. Guestrin, W. J. Kaiser, and M. A. Batalin, "Efficient planning of informative paths for multiple robots." in *IJCAI*, vol. 7, 2007, pp. 2204–2211.

[22] E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *IEEE signal processing magazine*, vol. 25, no. 2, pp. 21–30, 2008.

[23] D. L. Donoho, M. Elad, and V. N. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," *IEEE Transactions on information theory*, vol. 52, no. 1, pp. 6–18, 2006.

[24] S. Huang, "Adaptive sampling with mobile sensor networks," 2012.

[25] B. Dil, S. Dulman, and P. Havinga, "Range-based localization in mobile sensor networks," in *European Workshop on Wireless Sensor Networks.* Springer, 2006, pp. 164–179.

[26] T. He, C. Huang, B. M. Blum, J. A. Stankovic, and T. Abdelzaher, "Range-free localization schemes for large scale sensor networks," in *Proceedings of the 9th annual international conference on Mobile computing and networking.* ACM, 2003, pp. 81–95.

[27] S. Zhang, J. Cao, C. Li-Jun, and D. Chen, "Accurate and energy-efficient range-free localization for mobile sensor networks," *IEEE Transactions on Mobile Computing*, vol. 9, no. 6, pp. 897–910, 2010.

[28] S. Thrun and Y. Liu, "Multi-robot slam with sparse extended information filers," in *Robotics Research. The Eleventh International Symposium.* Springer, 2005, pp. 254–266.

[29] A. Howard, "Multi-robot simultaneous localization and mapping using particle filters," *The International Journal of Robotics Research*, vol. 25, no. 12, pp. 1243–1256, 2006.

[30] M. Bryson and S. Sukkarieh, "Architectures for cooperative airborne simultaneous localisation and mapping," *Journal of Intelligent and Robotic Systems*, vol. 55, no. 4-5, pp. 267–297, 2009.

[31] Z. Yan, N. Jouandeau, and A. A. Cherif, "A survey and analysis of multi-robot coordination," *International Journal of Advanced Robotic Systems*, vol. 10, 2013.

[32] S. Kato, S. Nishiyama, and J. Takeno, "Coordinating mobile robots by applying traffic rules." in *IROS*, vol. 92, 1992, pp. 1535–1541.

[33] L. Iocchi, D. Nardi, and M. Salerno, "Reactivity and deliberation: a survey on multi-robot systems," in *Workshop on Balancing Reactivity and Social Deliberation in Multi-Agent Systems.* Springer, 2000, pp. 9–32.

[34] E. Todt, G. Rausch, and R. Suárez, "Analysis and classification of multiple robot coordination methods," in *Robotics and Automation, 2000. Proceedings. ICRA'00. IEEE International Conference on*, vol. 4. IEEE, 2000, pp. 3158–3163.

[35] E. Klavins, "Communication complexity of multi-robot systems," in *Algorithmic Foundations of Robotics V.* Springer, 2004, pp. 275–291.

[36] I. Rekleitis, V. Lee-Shue, A. P. New, and H. Choset, "Limited communication, multi-robot team based coverage," in *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, vol. 4. IEEE, 2004, pp. 3462–3468.

[37] E. Pagello, A. D'Angelo, F. Montesello, F. Garelli, and C. Ferrari, "Cooperative behaviors in multi-robot systems through implicit communication," *Robotics and Autonomous Systems*, vol. 29, no. 1, pp. 65–77, 1999.

[38] F. Tang and L. E. Parker, "Asymtre: Automated synthesis of multi-robot task solutions through software reconfiguration," in *Proceedings of the 2005 IEEE international conference on robotics and automation.* IEEE, 2005, pp. 1501–1508.

[39] K. M. Wurm, C. Stachniss, and W. Burgard, "Coordinated multi-robot exploration using a segmentation of the environment," in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems.* IEEE, 2008, pp. 1160–1165.

[40] B. P. Gerkey and M. J. Mataric, "Sold!: Auction methods for multirobot coordination," *IEEE transactions on robotics and automation*, vol. 18, no. 5, pp. 758–768, 2002.

[41] R. T. Vaughan, K. Stoy, G. S. Sukhatme, and M. J. Mataric, "Lost: Localization-space trails for robot teams," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 796–812, 2002.

[42] S. C. Botelho and R. Alami, "M+: a scheme for multi-robot cooperation through negotiated task allocation and achievement," in *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*, vol. 2. IEEE, 1999, pp. 1234–1239.

[43] T. Fukuda, T. Ueyama, Y. Kawauchi, and F. Arai, "Concept of cellular robotic system (cebot) and basic strategies for its realization," *Computers & electrical engineering*, vol. 18, no. 1, pp. 11–39, 1992.

[44] D. Schulz, W. Burgard, D. Fox, and A. B. Cremers, "Tracking multiple moving targets with a mobile robot using particle filters and statistical data association," in *Robotics and Automation, 2001. Proceedings 2001 ICRA. IEEE International Conference on*, vol. 2. IEEE, 2001, pp. 1665–1670.

[45] Z. Khan, T. Balch, and F. Dellaert, "An mcmc-based particle filter for tracking multiple interacting targets," in *European Conference on Computer Vision*. Springer, 2004, pp. 279–290.

[46] S. S. Blackman, "Multiple hypothesis tracking for multiple target tracking," *IEEE Aerospace and Electronic Systems Magazine*, vol. 19, no. 1, pp. 5–18, 2004.

[47] R. P. Mahler, "Multitarget bayes filtering via first-order multitarget moments," *IEEE Transactions on Aerospace and Electronic systems*, vol. 39, no. 4, pp. 1152–1178, 2003.

[48] S. Manjanna, N. Kakodkar, M. Meghjani, and G. Dudek, "Efficient terrain driven coral coverage using gaussian processes for mosaic synthesis," in *CRV '16: Proceedings of the 2016 International Conference on Computer and Robot Vision*. IEEE Computer Society, June 2016.

[49] R. Bellman, "Dynamic programming princeton university press," *Princeton, NJ*, 1957.

[50] S. Manjanna, N. Kakodkar, and G. Dudek, "Multi-Robot Path Planning for Selective Coverage," in *ICRA 2016 Workshop on Fielded Multi-robot Systems Operating on Land, Sea, and Air*, Stockholm, Sweden, May 2016.

[51] C. E. Rasmussen, "Gaussian processes for machine learning," 2006.

[52] C. Plagemann, S. Mischke, S. Prentice, K. Kersting, N. Roy, and W. Burgard, "Learning predictive terrain models for legged robot locomotion," in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2008, pp. 3545–3552.

[53] C. Paciorek and M. Schervish, "Nonstationary covariance functions for gaussian process regression," *Advances in neural information processing systems*, vol. 16, pp. 273–280, 2004.

[54] T. Lang, C. Plagemann, and W. Burgard, "Adaptive non-stationary kernel regression for terrain modeling." in *Robotics: Science and Systems*, 2007.

[55] C. Walder, K. I. Kim, and B. Schölkopf, "Sparse multiscale gaussian process regression," in *Proceedings of the 25th international conference on Machine learning*. ACM, 2008, pp. 1112–1119.

[56] M. Lázaro-Gredilla and A. Figueiras-Vidal, "Inter-domain gaussian processes for sparse inference using inducing features," in *Advances in Neural Information Processing Systems*, 2009, pp. 1087–1095.

[57] J. Quiñonero-Candela and C. E. Rasmussen, "A unifying view of sparse approximate gaussian process regression," *Journal of Machine Learning Research*, vol. 6, no. Dec, pp. 1939–1959, 2005.

[58] B. Yamauchi, "Frontier-based exploration using multiple robots," in *Proceedings of the second international conference on Autonomous agents*. ACM, 1998, pp. 47–53.

[59] N. Roy and G. Dudek, "Collaborative robot exploration and rendezvous: Algorithms, performance bounds and observations," *Autonomous Robots*, vol. 11, no. 2, pp. 117–136, 2001.

[60] R. Zlot, A. Stentz, M. B. Dias, and S. Thayer, "Multi-robot exploration controlled by a market economy," 2002.

[61] R. Simmons, D. Apfelbaum, W. Burgard, D. Fox, M. Moors, S. Thrun, and H. Younes, "Coordination for multi-robot exploration and mapping," in *AAAI/IAAI*, 2000, pp. 852–858.

[62] B.-N. Vo, S. Singh, and A. Doucet, "Sequential monte carlo methods for multitarget filtering with random finite sets," *IEEE Transactions on Aerospace and electronic systems*, vol. 41, no. 4, pp. 1224–1245, 2005.

[63] B.-N. Vo and W.-K. Ma, "The gaussian mixture probability hypothesis density filter," *IEEE Transactions on signal processing*, vol. 54, no. 11, pp. 4091–4104, 2006.

[64] B. Jung and G. S. Sukhatme, “Tracking targets using multiple robots: The effect of environment occlusion,” *Autonomous robots*, vol. 13, no. 3, pp. 191–205, 2002.

[65] M. Meghjani, S. Manjanna, and G. Dudek, “Multi-target rendezvous search,” in *IROS ’16: Proceedings of the 2016 International Conference on Intelligent Robots and Systems*. IEEE, October 2016.

[66] A. Q. Li, I. Rekleitis, S. Manjanna, N. Kakodkar, J. Hansen, G. Dudek, L. Bobadilla, J. Anderson, and R. N. Smith, “Data correlation and comparison from multiple sensors over a coral reef with a team of heterogeneous aquatic robots,” in *International Symposium of Experimental Robotics (ISER)*, Tokyo, Japan, Mar. 2016.