2020

Multi-Robot Informative Path Planning in Unknown Environments Through Continuous Region Partitioning

Amitabh Bhattacharya
University of North Florida, a.bhattacharya@unf.edu

MULTI-ROBOT INFORMATIVE PATH PLANNING IN UNKNOWN ENVIRONMENTS THROUGH CONTINUOUS REGION PARTITIONING

by

Amitabh Bhattacharya

A thesis submitted to the School of Computing

in partial fulfillment of the requirements

for the degree of Master of Science in Computer and Information Sciences

UNIVERSITY OF NORTH FLORIDA

SCHOOL OF COMPUTING

April 2020

The thesis "Multi-robot Informative Path Planning in Unknown Environments through Continuous Region Partitioning" submitted by Amitabh Bhattacharya in partial fulfillment of the requirements for the degree of Master of Science in Computer and Information Sciences has been

Approved by the thesis committee: Date:

Ayan Dutta, Ph.D.

Thesis Advisor and Committee Chairperson

O Patrick Kreidl, Ph.D.

Thesis Advisor and Committee Chairperson

Anirban Ghosh, Ph.D.

Committee Member

ACKNOWLEDGEMENTS

First, I would like to offer my sincerest thanks and gratitude to my thesis advisors, Dr. Ayan Dutta and Dr. O Patrick Kreidl for their continued support and mentoring throughout the development of this thesis and for the duration of my time at the University of North Florida.

Additionally, I would like to thank Dr. Anirban Ghosh, the member of my thesis committee for his vital guidance.

I would also like to thank Drs. Sherif Elfayoumy, Asai Asaithambi and Swapnoneel Roy for providing much needed advice and support during my pursuit of the masters degree.

This research is supported by the UNF Foundation Board Grant received by Dr. Dutta and Dr. Kreidl and the UNF School of Computing.

CONTENTS

Figures

Tables

ABSTRACT

This research activity is primarily focused to obtain information from an environment with the help of a group of coordinated robots. Each robot is responsible to plan its path independently but the robots, as an overall system, have a common goal of maximum information collection. This domain of research is known as Multi-Robot Informative Path Planning (MIPP). MIPP is very motivating due to its challenging nature and numerous real-world applications. It has shown its presence from semiautomatic applications like robotic search and rescue to fully automatic applications like interplanetary missions.

We consider the NP-Hard problem of MIPP in an unknown environment having communication constraints and budget limitations in robots. We propose a novel approach that uses continuous region partitioning to efficiently divide an initially unknown environment among the robots based on the discovered obstacles. The research objective is to collect higher amount of information and target an uniform work distribution among robots. Simulation results show that our proposed approach is successful in reducing the initial imbalance in work distribution of the robots while ensuring close-to-reality spatial modeling within a reasonable amount of time.

CHAPTER 1

INTRODUCTION

Multi-robot informative path planning (MIPP)[SKG+07] is an important problem in robotic systems due to its practical relevance. Potential applications include interplanetary missions,[BSA+12] robotic search and rescue,[LN13] collecting samples of ambient phenomena like fire, nuclear radiation, soil moisture, or algal concentration, from difficult and hazardous environments such as post-disaster scenarios, combat environments, forests, or water bodies such as oceans and lakes.[Sol04] The main objective in MIPP is to coordinate a set of robots so that they can collect samples of the ambient phenomenon and report it back to a central location such as a base station. Robots need to perform this collection task in an energy efficient manner so that they can collect the maximal information from the environment with the limited battery energy available with them. The problem is challenging because the locations in the environment where the most information is available is not known *a priori* and the robots have to explore the environment and discover these locations in real time. Also, the robots have to coordinate their movements so that they do not end up collecting redundant samples and are also able to avoid collisions with each other. In this research, we consider another practical aspect of MIPP in the form of communication constraints - also not known *a priori* are the locations in the environment where the robots can form a partial or complete network via wireless connectivity to exchange local knowledge.

We initially divide the environment into non-overlapping regions using Voronoi partitioning[Vor07] and allocate each region to a robot where it is responsible for information collection We make an observation that in an unknown environment where the loca-

tions, numbers and shapes of the obstacles are unknown, it might happen that the areas of an initial partition are imbalanced. As the robots navigate the environment and discover obstacles, their perceptions about their allocated regions change. Therefore, we posit that their allocated partitions may benefit from dynamic adjustments as each robot discovers inaccessible sub-regions (e.g., due to obstacles) within its current partition, over time rendering the load of the information collection task better balanced across all robots. Robots partition the region based on current perceptions of the environment; so, if perceptions include newer information, then there is opportunity to refine partitions upon regaining connectivity with neighboring robots.

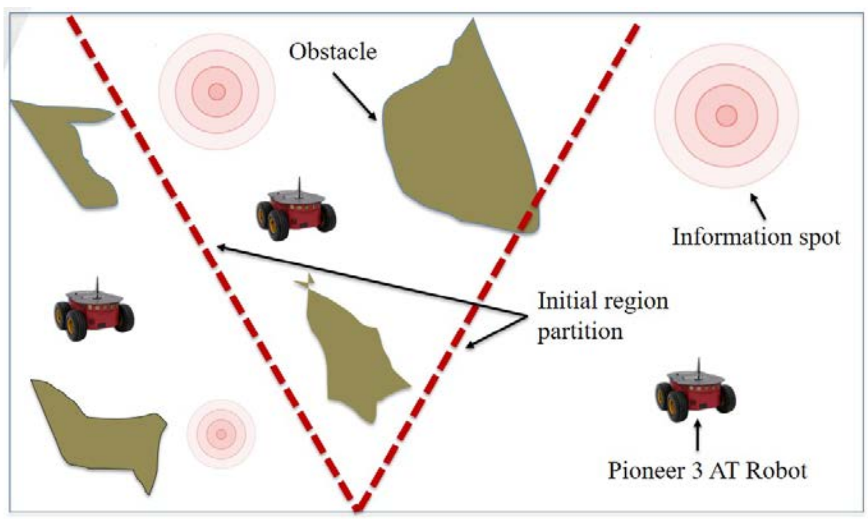

Figure 1.1: A view of an environment where robots are randomly dropped for information collection. The environment shows the presence of unknown obstacles and information spots. Each robot is allocated an initial non-overlapping region partitioned of the whole environment.

We have verified our proposed approach in simulation with up to eight robots having different communication and LIDAR (Light Detection and Ranging Device) ranges performing the task of information collection within an environment having different distributions of information samples. Our results show that the proposed approach is

able to successfully reduce the difference between the optimal and initially allocated obstacle-free area distribution among the robots, while incurring a reasonable time for up to 8 robots.

The contributions in this research are as follows:

- To the best of our knowledge, it presents the first formalization of the MIPP problem under communication constraints in an unknown environment.
- It proposed a greedy informative path planning strategy combined with a load-balancing algorithm to recursively repartition the multiple robots' allocated Voronoi components and achieve a better load-balanced information collection system.
- It extensively tests our framework with varying number of robots as well as varying communication and laser ranges within a high-fidelity robotics simulator, Webots.

This thesis document is organized as follows: First we discuss related literature in section 2. We next describe the problem under study in this research in section 3, including the initial Voronoi partitioning strategy and the greedy informative path planning technique our approach employs. In section 4, we discuss the recursive Voronoi repartitioning strategy and section 5 details the experimental evaluations and simulation results. Certain limitations of our approach, and their implications in our study, are also identified. Finally, in section 6, we conclude and suggest potential future directions.

CHAPTER 2

LITERATURE REVIEW

One of the very first studies on the MIPP problem is due to [SKG^{+}07]. They have proposed a branch-and-bound based budgeted path planning technique to solve the problem. This work uses Gaussian Processes (GP)[RW06] to model the environmental spatial phenomena. GP is a popular regression model, which has been used not only in MIPP studies but also for general information collection algorithms using static wireless sensors.[KSG08] The authors have used two metric to study the information collection quality: Entropy and Mutual Information. In our work, we use Entropy to be our quality metric. The MIPP problem with mobile robots having periodic connectivity has been studied in [HS10]. In this study, the robots have to come within each other's communication ranges after a certain time interval. Recently, Dutta et al. proposed a solution for the MIPP problem where the robots need to maintain a connected network throughout the information collection process.[DGK19] The authors have used bipartite matching to avoid inter-robot collisions while collecting maximal information and $u - v$ node separators to maintain connectivity among the robots. Our work in this study does not require the robots to maintain either continuous or periodic connectivity rather if they come within each other's contact, they intelligently share information and balance their workloads.

In [CLD13, LDK09], the authors have proposed two greedy strategies for information collection. The proposed strategies also provide provable near-optimal guarantees. However, the solution is centralized and does not scale well with number of robots. A poly-logarithmic approximation bound in a metric space is provided by the work in [LHL16]. A greedy information collection strategy has also been used in [DD16] where

a group of robotic modules form user-defined configurations while collecting maximal information from an environment. Our work in this research uses a similar greedy informative path planning strategy. Sampling-based path planners have also been used for information collection[HS14] where the authors have proposed a variant of the RRT* algorithm.[KWP+11] Distributed informative path planning algorithm has recently been proposed by Luo and Sycara,[LS18] where they have used a Gaussian mixture model to better prediction results. Hitz et al. has proposed a dynamic programming-based information collection framework in [HGP+14]. Partially Observable Markov Decision process (POMDP) has also been used for modeling the information collection phenomena.[MCdFB+09] Singh, Krause, and Kaiser studied a non-myopic version of the MIPP problem.[SKK09] A similar strategy in a spatio-temporal field is used in [MKGH07]. Environment monitoring with a team of underwater vehicles is studied in [KCS16, LLHJ+16]. However, most of these discussed studies on MIPP do not take the communication constraints of the robots, a highly practical consideration, into account. Moreover, the environment structure, i.e., the obstacle locations and shapes are supposed to be known. In this research, we relax these assumptions and the robots collect information from an initially unknown environment while gracefully handling their restrictive communication ranges.

Voronoi partitioning is a scheme of decomposing an area into disjoint cells where every point in each cell is closer to the center of that cell than any other cell.[Vor07] Voronoi partition has been used in solving several problems in robotics. Choset and Burdik[CB95] use a generalized Voronoi diagram for a robot path planning problem. Zhou et al.[?] has recently proposed an online collision avoidance mechanism using buffered Voronoi cells. In [HDG16], the region to be covered by a group of robots is decomposed into Voronoi cells and each robot is responsible for covering the corresponding Voronoi cell. Moreover, if part of one robot's allocated cell is inaccessible to that robot, other robot(s)

covers the inaccessible part of the cell. A similar area coverage strategy using a group of networked robots has been studied in [BSM+10], where the environment and the obstacles are assumed to be non-convex. Coverage path planning in an initially unknown environment using a multi-robot system has been studied in [GWD12] where the authors have proposed an online Voronoi partition based coverage algorithm. In [CMKB04, DCFB12], the authors have proposed a multi-robot coverage path planning strategy using Voronoi partitioning where each robot communicates its position while dynamically adapting the partitions with nearby robots complete coverage. Recently, a Manhattan distance-based Voronoi partitioning scheme is proposed in [**?**] for multi-robot area coverage. In a different work, Nair and Guruprasad has combined geodesic and Manhattan distance-based partitioning strategies for coverage in presence of obstacles.[?] The authors in [SMR06] propose a control mechanism that allows a swarm of robots to position themselves to optimize the measurement of sensory information in the environment. A Voronoi repartitioning strategy based on the unbalanced areas of the initial partitions has been studied in [TWWS08]. We use their proposed algorithm for Voronoi cell repartitioning among the communicating robots. Although not exactly using Voronoi cells, but decentralized area partitioning strategy has been previously used in [AAMO14] for area patrolling. However, the authors modeled the solution in a way that the robots periodically come into each other's contact for sharing information, which is not required in our approach. Moreover, the paper does not handle unknown obstacles in the environment.

Our proposed method aims to bring the above-mentioned concepts together: how to use the idea of region partitioning for better load-balancing among the robots for information collection. The closest study to this work is due to Kemna et al.,[KRNG+17] which we complement and extend in the following way: the work in [KRNG+17] does not consider any obstacle in the environment whereas our repartitioning algorithm works

based on robots' explored regions and the unknown obstacles that they discover. Moreover, the work in [KRNG^{+}17] assumes the environment to be restrictive for communication whereas we consider a different but realistic constrained-communication model where the robots have limited communication ranges while the environment itself is not communication-restrictive.

CHAPTER 3

PROBLEM STATEMENT

Let $\mathcal{R} = \{r_1, \ldots, r_n\}$ denote a set of n robots dropped from an aircraft at a given set of locations $\mathcal{S} = \{s_1, \ldots, s_n\}$ within a convex polygonal environment $E \subset \mathbb{R}^2$ to be navigated. Environment E may include inaccessible sub-regions, due to static obstacles, also assumed to be convex polygons i.e., $E = E_{free} \cup E_{obst}$ with $E_{free} \cap E_{obst} = \emptyset$. All robots are initialized with common knowledge of $\mathcal{R}$, $\mathcal{S}$ and E and but with no knowledge about the location nor geometry of the obstacles, initially assuming $E = E_{free}$. Inaccessible regions are rather discovered during navigation, each robot equipped with a laser rangefinder by which to detect (and circumnavigate, when feasible) nearby obstacles. Robots are also each equipped with localized navigation (e.g., using GPS), an information collecting sensor (e.g., a camera) and wireless radio to share location, path, obstacle and sensor information with other robots when within communication range CR. Our problem setup assumes that both the radio range (for inter-robot communication) and the laser range (for obstacle detection) are greater than the sensing range (for information collection). We do not require the distances between initial locations in S to necessarily fall within communication range CR, i.e., the environment E itself is not communication-restrictive.

3.1 Initial Voronoi Partitioning

The Voronoi partition is a widely used mechanism for separating a space into non-overlapping components based on the "nearness" concept. Given a convex polygonal region E and n points in the set $\mathcal{S} = \{s_1, \ldots, s_n\}$, we can associate a polygonal Voronoi

component E_i with every point $s_i \in \mathcal{S}$ as the following:

$$E_i = \{q \in E; \| q - s_i \| \leq \| q - s_j \| \; \forall s_j \neq s_i \in S\} \tag{3.1}$$

where $\| q - s_i \|$ denotes the Euclidean distance between two points $q, s_i \in E$. Intuitively, each Voronoi component E_i is the collection of those points which are closer to s_i than any other site $s_j \in \mathcal{S}$. The intersection of any two Voronoi components is either null, a line segment or a single point. Each Voronoi component is a topologically connected non-null set. This standard partitioning of the environment into n Voronoi components for n sites can be done in $\mathcal{O}(n \log n)$ time.[Ste87]

Recall that all robots $\mathcal{R}$ are initialized with common knowledge of starting locations $\mathcal{S}$ and environment E, and none have initial knowledge of the obstacles and thus will initially assume that $E = E_{free}$. It follows that all robots will begin with a common understanding of the Voronoi partition $E_1, \ldots, E_n$ associated with their initial air-drop locations $\mathcal{S}$. Each such component E_i is then interpreted as a hard constraint on the portion of the environment that robot r_i is permitted to navigate within. Also recall that this initial partitioning makes no assumptions that the distances between locations in $\mathcal{S}$ are within communication range CR. Of course, whenever robots fall within communication range CR, information gathered during navigation (e.g., detected obstacles) can inform subsequent Voronoi repartitioning of the environment. See Figure 3.1.

3.2 Informative Path Planning

Consider a set of information collection points, or the Points of Interest (POI), that cover environment E in a grid structure. Each robot moves from one $poi \in POI$ to the other via straight lines in a deterministic manner and, upon reaching any poi, uses its sensor(s) to collect the information within that poi's grid-cell. Each robot's sensing process is assumed to have negligible noise, but only within the current grid-cell; specifi-

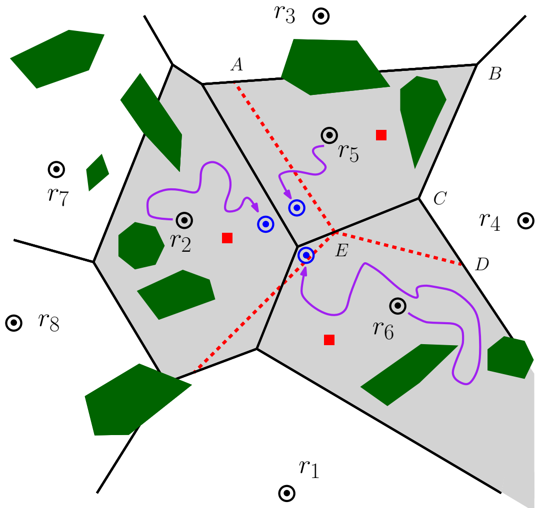

Figure 3.1: Eight robots (circles) navigate through their initial Voronoi components (separated by black-solid lines) of an environment with obstacles (green polygons). Robots r_2, r_5 and r_6 eventually move within communication range (blue circles), where the union of their respective components (shaded gray) is repartitioned (separated by red-dashed lines) to re-balance the remaining collection load based upon updated perceptions of the free versus inaccessible sub-regions.

cally, the information collected from the current poi will be interpreted as ”ground truth” for that poi, but no direct information from (unvisited) neighboring pois is revealed. Observe that, under these motion and sensing assumptions, at any point of time the POI set can be decomposed into two disjoint subsets, U and V, corresponding to the unvisited and visited collection points, respectively, with equivalence to the unobserved and observed grid-cells in E. Finally, each robot is assigned a budget B, which indicates the maximum number of poi's that can be visited. The objective of the robotic team is to follow maximally-informative paths, or equivalently minimizing redundancy in information collection, subject to the path-length budget B per robot.

We model the observable environmental phenomena using Gaussian Processes (GPs),[GKS05] which assume that all the collection locations in the environment (POI) generate information according to a Gaussian random vector $\mathbf{X}$ with known (prior) mean vector μ and covariance matrix Σ. By virtue of the above motion and sensing assumptions, it follows that for robot paths that render a set U of unvisited poi's, the Gaussian random

vector $\mathbf{X}_U$ characterizing the uncollected information has (posterior) mean vector and covariance matrix given by equations

$$\begin{aligned}\mu_{U|\mathbf{X}_V} &= \mu_U + \Sigma_{UV}\Sigma_{VV}^{-1}(\mathbf{X}_V - \mu_V)\\ \Sigma_{UU|\mathbf{X}_V} &= \Sigma_{UU} - \Sigma_{UV}\Sigma_{VV}^{-1}\Sigma_{UV}\end{aligned} \tag{3.2}$$

with $\mathbf{X}_V$ denoting the set of measurements sensed in the visited poi's and the prior statistics organized into the corresponding block forms with respect to U and V i.e.,

$$\mu = \begin{bmatrix} \mu_U \\ \mu_V \end{bmatrix} \quad \text{and} \quad \Sigma = \begin{bmatrix} \Sigma_{UU} & \Sigma_{UV} \\ \Sigma_{UV} & \Sigma_{VV} \end{bmatrix}.$$

The diagonal elements of covariance matrix $\Sigma_{UU|\mathbf{X}_V}$ in (3.2) express the posterior variances $\sigma^2_{u|\mathbf{X}_V}$ of every unvisited point-of-interest u in U. The associated posterior entropies are, in turn, given by

$$H(u|\mathbf{X}_V) = \begin{cases} \frac{1}{2}\log(2\pi e\sigma^2_{u|\mathbf{X}_V}), & u \in U \cap E_{free} \\ -PEN, & u \in U \cap E_{obst} \end{cases} \tag{3.3}$$

and objectively quantify which unvisited locations in the environment, when considered individually, contain the highest amount of uncertainty i.e., contribute most to the mean-square-error when predicting the future information measurements via the posterior mean vector $\mu_{U|\mathbf{X}_V}$. Observe that a negative penalty value is introduced for unvisited locations rendered inaccessible by discovered obstacles, promoting compatibility between obstacle avoidance capabilities of the robots and any subsequent path planning driven by uncertainty reduction.

Our multi-robot planning approach to uncertainty reduction will adopt the straightforward "greedy" informative path generation strategy: each robot r_i simply employs (3.3) to calculate the best neighboring location to visit (within its Voronoi component

E_i) based on the measurements from all previously visited locations V and up-to-date knowledge of discovered obstacles E_{obst}. More formally, given robot r_i is currently in location poi_{curr}, it will select informative location poi_{next} according to

$$poi_{next} = \arg \max_{u \in E_i \cap U \cap N(poi_{curr})} H(u|\mathbf{X}_V) \tag{3.4}$$

with $N(poi)$ denoting the set of all neighboring informative locations to any given poi in POI. The greedy path planning strategy has been proven to yield acceptable results in information models using the assumed spatially-localized Gaussian Processes.[CLD13, KRNG+17] Given the fact that each robot calculates only the next best poi to visit in every iteration (i.e., the planning horizon is 1) in a distributed manner (i.e., calculates only its own path) and poi_{next} is a neighbor of the robot's current location, the above-mentioned path planning strategy will scale polynomially.[CLD13, KSG08]

CHAPTER 4

ALGORITHM

As the robots do not have any prior knowledge about the locations and the shapes of the obstacles in the unknown environment, it may happen that the initial partitioning of the environment is imbalanced, i.e., some robots have more free informative locations in its Voronoi cell and the others have less. As the robots move in the environment and collect information, they may also discover obstacles. When they later meet other robots, they share their updated local knowledge of the environment and may repartition their currently allocated Voronoi cells.

4.1 Voronoi Repartitioning for Load Balancing

In this section, we present the load balancing algorithm used in our recursive region partitioning algorithm. This algorithm is an adaptation of the virtual-force based load-balancing algorithm proposed by Tewolde et al. in.[TWWS08]

Assume that during the task of information collection, a set $S \subseteq \mathcal{R}$ of k robots get positioned such that their induced communication graph is connected. In such a situation, the load balancing algorithm is triggered to balance the workload between these k robots. To achieve this, the algorithm performs Voronoi partitioning on the total area of the current k cells and possibly adjust the geometry of the cells along with the center of these cells. For a robot $r_i \in S$, its current workload w_i is calculated using the following formula:

$$w_i := \text{polygonal-area}(C) - \text{polygonal-area of the detected obstacle cells in } C. \quad (4.1)$$

The underlying idea is to allocate more cells to a robot whose workload is low thereby reducing workloads of the overloaded robots.

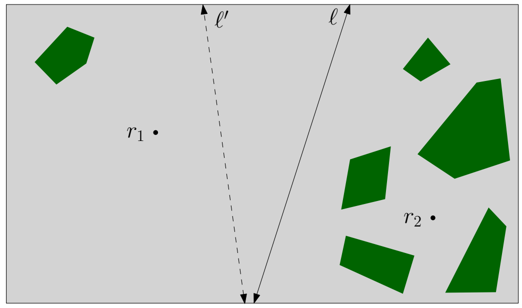

Figure 4.1: An illustration of the dynamic load-balancing is shown using two robots r_1, r_2. The line ℓ is the common boundary of the initial Voronoi partition of the two robots. After load-balancing, ℓ' is the new common boundary. Observe that after load-balancing, r_1 has lesser area to cover whereas r_2 has now more area to cover. The load-balancing happens based on the robots' currently detected obstacles in the environment (green polygons).

Algorithm 1: The load-balancing algorithm.

Let w_i denote the workload of $r_i \in \mathcal{S}$.

Procedure `balanceWorkLoads()`

- iterations $\leftarrow 0$
- **while** *iterations* $\leq K$ **do**
  - Every robot $r_i \in S$ exchanges workload information with its neighbours $N(r_i)$.
  - We denote by $f_i := \sum_{r_j \in N(r_i)} \alpha(w_j - w_i)\vec{n_{ij}}$, the net virtual force acting on r_i where $\vec{n_{ij}}$ is the unit vector from r_i to r_j and α denotes a scale factor.
  - For every $r_i \in S$, $v_i \leftarrow v_i + (f_i - \lambda V_i)/\beta \cdot \Delta t$, where v_i is the current velocity of r_i, β is a constant for the robot mass, Δt is the time step size per iteration, and λV_i represents some viscous force in order to stabilize the deployment. Initially, every v_i is set to 0.
  - The new location of r_i is $x_{r_i} + v_i \cdot \Delta t$, where x_{r_i} is the current location of r_i. Compute the Voronoi diagram with the updated locations of the k robots and obtain the updated w_is.
  - **if** *the standard deviation of the* w_i*s* $\leq \tau$ **then**
    - Terminate.
  - iterations $\leftarrow$ iterations $+1$

The algorithm with all its technical details is presented in Algorithm 1. Refer to

Fig. 4.1 for an illustration using two robots. Note that in the original algorithm,[TWWS08] the number of iterations is not specified. However, in our experiments, we have found 1000 to be a reasonable upper-bound on the number of iterations required to achieve an equilibrium in the load-balancing process; we denote this upper bound by K. Also, in our experiments we terminate the load-balancing process when the standard-deviation of the w_i's reaches a value of τ or less. In our implementation, we have set $\alpha = 0.04$, $\beta = 1$, $\lambda = 1$, $\Delta t = 0.5$, $K = 1000$, and $\tau = 5$.

4.2 Recursive Region Partitioning for Information Collection

Each robot, r_i, will continue to generate a new path within its designated Voronoi cell, E_i, and will follow that path until one of the following two things happen:

1. The path length covered so far by r_i exceeds the assigned budget, B,

2. r_i *meets* (i.e., comes within the communication range of) another robot r_j.

While visiting new poi cells within its designated region, r_i will continuously broadcast messages containing its unique ID, current location, and its currently allocated region, i.e., E_i, a polygonal region. If r_i meets another robot(s), then depending on these received information from the other robot(s), they will decide whether they should repartition the union of their current Voronoi cells to achieve load-balancing or they should ignore repartitioning and continue to stay within their currently designated cells.

To keep track of contacted robots and their corresponding Voronoi cells along with their discovered obstacles so far, each robot, r_i, maintains a data-structure, named Perception of World (PoW_i). The PoW_i can be imagined as an 2D array of time-stamped objects indexed by poi's and contains information about the following: E_i's, discovered obstacle locations, and the sensed information measurements. It might happen that robot r_i discovered obstacles in robot r_j's current Voronoi cell, which r_j might or might not

Algorithm 2: Recursive region partitioning for information collection

// Each robot r_i executes the following procedures.

Procedure `pathPlanning()`

  if *Covered budget-limit distance* B **then**

    Broadcast the latest PoW_i along with a FINAL message to indicate that the budget-limit is reached.

  else

    Plan the next poi to visit in its own partition using Equation 3.4.

    Move to this poi location to collect information.

Procedure `RecursivePartitioning()`

  Create a data-structure, PoW_i, to keep track of the current perception about all the poi's.

  if *meets a robot* r_j *with which* r_i *has never met OR* PoW_i *and/or* PoW_j *have changed from the last time these robots met* **then**

    Update the corresponding Voronoi partitions.

  Plan the next poi to visit for information: *pathPlanning()*.

have any knowledge about. But when they meet, both of them will have the knowledge about these obstacles.

As the obstacles are assumed to be static, if PoW's are to be exchanged between r_i and r_j, this perception will be shared by them and corresponding repartitioning strategy will be implemented. Note that, PoW_i is a local data-structure, i.e., different robots' PoW_i's might be different in terms of stored content depending on which robot(s) they came into contact with. This data-structure will help the robots to partition their respective regions if required in a more intelligent way as we will see next.

When to repartition? Let $\bar{R}_i$ denote the set of robots that are within robot r_i's communication range. If r_i meets with $r_j \in \bar{R}_i$, one of the following situations can arise:

1. r_i has never communicated with r_j before.

2. r_i has communicated with r_j before.

    (a) r_i and/or r_j's perceptions about the other robots' partitions (PoW_i and/or PoW_j) have changed after r_i last communicated with r_j.

(b) PoW_i and PoW_j are still the same from r_i and r_j's last communication.

We will handle these cases one by one. We will start off with the simplest case $(2.b)$ – r_i and r_j's perceptions about the other robots' partitions (PoW_i, PoW_j) have not changed from the last time these two robots have met. In this case, they don't have to update their plans and/or their respective Voronoi cells and they can plan their next best poi's to visit within their current Voronoi cells (line 12 in Algorithm 2).

If r_i and r_j are meeting each other for the first time (case 1), they might or might not have current partition information about each other. In any case, they will repartition their current Voronoi cells based on their latest knowledge about the obstacles in their respective PoW's. As the environment is unknown to the robots in the beginning and they are discovering shapes and locations of the obstacles in the environment as they explore more, the initial Voronoi cells might not be '*balanced*' in terms of the free space available to explore. Consequently, one robot might be tasked with exploring significantly larger/smaller amount of region in the environment compared to the others. In order to minimize this potential high level of variance and bring balance to the obstacle-free areas of robots' exploration regions, we use a Voronoi cell repartitioning technique described in Algorithm 1.

As a result of this repartitioning strategy, the participating robots are allocated to a unique and new polygonal region (possibly concave) in the environment for information collection. The algorithm virtually shifts the sites for the Voronoi partitions on the plane such that the new cells are balanced in terms of number of free poi cells in them and it terminates when the standard deviation of the load distribution is sufficiently low enough (see Algorithm 1 for details). Note that the repartitioning algorithm is only executed on the area bounded by the union of the Voronoi cells of the coordinating robots ($\bar{R}_i$). As the robots discover more obstacles in the environment, their perception about the environment will change and this will be reflected in their continuously repartitioned regions.

In case r_i and/or $r_j \in \bar{R}_i$'s PoW's have changed after they last met (Case $2.a$), they will follow similar strategy to repartition their current Voronoi cells as prescribed for case 1. When a robot reaches its path budget limit, i.e., covered B informative locations, it will stop exploring the environment any further. But it will continue to broadcast its PoW along with a ***FINAL*** message so that the other robots can plan their partitions accordingly whereas the stopped robot will not repartition its latest Voronoi cell (lines 3-4 in Algorithm 2).

Proposition 1 *The region partitioning approach produces a complete coverage of the environment E.*

Coverage of E is complete when every $poi \in POI$ belongs to some Voronoi cell E_i of some robot r_i. After the initial partitioning, $\cup_n E_i = E$ and therefore, $POI \subseteq E$. When a set of robots $\bar{R}$ repartition their current cells, they do the repartitioning of the union of their current cell area, i.e., on $\mathcal{V} = \cup_{r_k \in \bar{R}} E_k$. Let $\mathcal{V}' = E \setminus \mathcal{V}$. As a result of the repartitioning, $\mathcal{V}$ does not change[TWWS08] and thus $\mathcal{V}'$ does not change. Therefore, $\mathcal{V}' \cup \mathcal{V} = E$. Hence, proved.

CHAPTER 5

EVALUATION

In this section we discuss the simulation settings and present the results found via rigorous simulation experiments within a 3D robot simulator Webots.[Mic04]

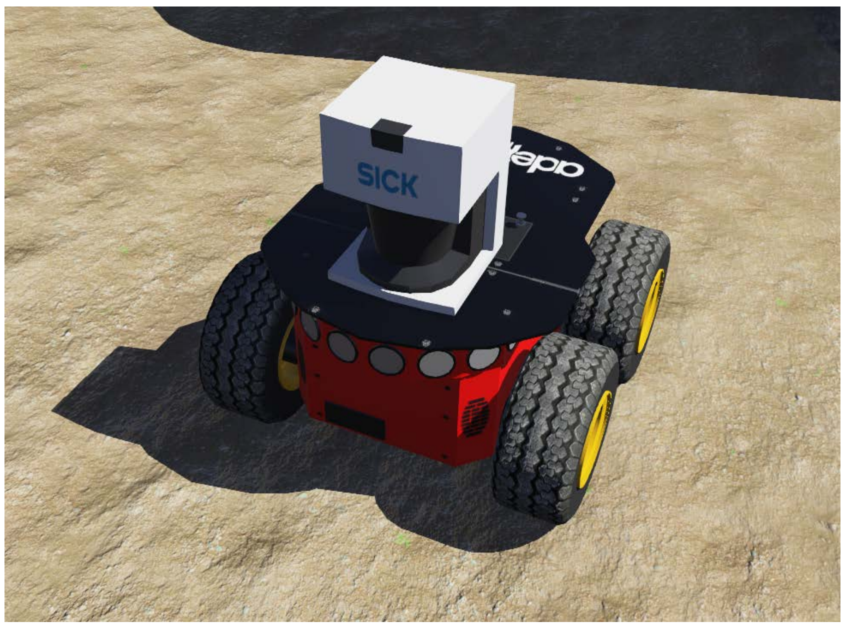

Figure 5.1: Simulated model of Pioneer 3AT robot equipped with a Sick LIDAR within Webots.

5.1 Settings and Environment Model

We have tested our proposed informative path planning approach in simulation, which is performed on a desktop computer with a 4.6 GHz. Intel core i7-8700 Processor, 16 GB RAM, and an AMD Radeon Rx 550 4 GB graphics card. We have used a simulated model of the Pioneer 3AT robot within the Webots simulator that is equipped with a GPS and a Sick LMS 291 LIDAR sensor (Fig. 5.1). Two different 49 meter $\times$ 49meter environments

(shown in Figure 5.2) are considered, choosing initial robot locations as well as rectangular obstacles at random.

We have varied the obstacle amount in the environment between $5, 10,$ and 15%. LIDAR ranges are varied between $\{5, 10, 15\}$ meters. The communication ranges are varied between the same distances. Each robot has the same communication and LIDAR ranges in any particular test setting. The budget is fixed to 50 grid cells in all the tests. Each robot starts off with the ground truth information measurement of 10% of the grid cells for calculating the initial values of the GP hyper-parameters. The main metrics that we evaluate in this research are: 1) entropy, 2) variance, 3) root mean square error (RMSE), 4) time, 5) load, and 6) amount of obstacle detected. As this is the first MIPP solution with communication-constrained robots that handles unknown environments having polynomial obstacles, we could not compare against any previous approach. In simulation, each test is run 10 times and the average results are presented here. The error bars in the plots indicate the standard deviations of the y-axes metrics. Table 5.1 summarizes the different parameters used in our experiments and the different values used for these parameters.

Parameters	**Notations**	**Values**
Environment size	$\textsc{Env}_{size}$	49×49 m.^{2}
Grid cell size	GC_{size}	1×1 m.^{2}
Robot count	n	$\{2, 4, 6, 8\}$
Budget	B	50 grid cells
Communication range	CR	$\{5, 10, 15\}$ m.
Obstacle percentage	$O\%$	$\{5, 10, 15\}$
Initial training set size	TR	10% of all cells
LIDAR range	LR	$\{5, 10, 15\}$ m.

Table 5.1: List of parameters used in our experiments.

Recall from Section 3.2 the use of Gaussian Processes (GP) for information modeling. Our experiments assume a 49m. $\times$ 49m. grid where the corner points comprise the 50-by-50 poi's underlying a homogeneous, isotropic Gaussian Markov random field using exponential pairwise covariance functions: for any pair of poi's i and j at spatial

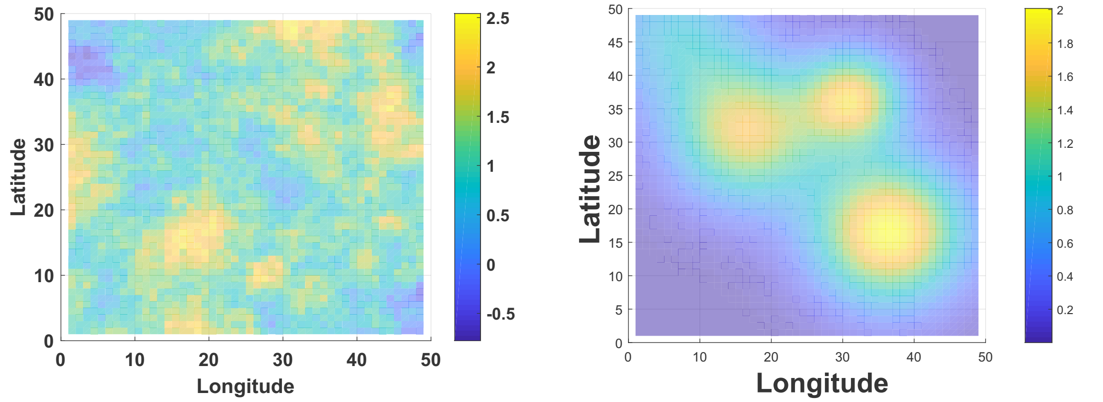

(a) (b)

Figure 5.2: The 2-D environmental information models used in this research with pair-wise covariance parameters (a) ENV1 with $\beta = 1$ and $\ell = 25\text{m}.$ and (b) ENV2 with $\beta = 0$ and $\ell = 25\text{m}$. The variable pattern of ENV1 in (a) is obtained as a sample from the Gaussian process with the associated covariance matrix but assuming a zero mean field, while the zero covariance matrix underlying ENV2 degenerates every sample to equal the mean field, crafted as a set of circular rises above zero in (b).

locations $\mathbf{p}_i$ and $\mathbf{p}_j$, respectively, let

$$\sigma_{ij}^2 = \beta^2 \exp\left(-||\mathbf{p}_i - \mathbf{p}_j||/\ell\right)$$

where $\beta > 0$ is the local standard deviation and ℓ (in meters) is the exponential rate of diminishing covariance between increasingly-distant poi's. Fig. 5.2(a) illustrates a sample from such a process with zero mean and using parameters $\beta = 1$ and $\ell = 25\text{m}.$ to define the covariance matrix. Observe that such a process, when $\beta = 0$, will deterministically render only the mean field: Fig. 5.2(b) illustrates such a mean field that is proportional to a five-component Gaussian mixture over the 2-D spatial region, each component located at its length-2 mean vector with spherical contours determined by a 2-by-2 diagonal covariance matrix. Note that ENV1 and ENV2 are comparable to the test environments used in [CLD13] and [HS14] respectively.

5.2 Results

In this section, we present the results obtained by our proposed approach in different environmental settings. This section is divided into six sub-sections, each of which discusses the results for one particular performance metric.

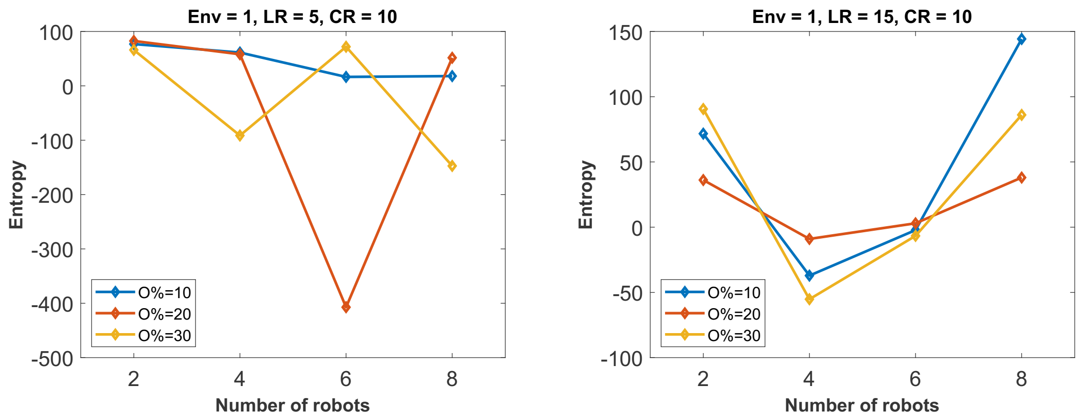

(a) (b)

Figure 5.3: Amount of entropy collected by the robots in ENV1: a) $LR = 5$, and b) $LR = 15$.

5.2.1 Entropy

First, we present the results for our optimization metric, entropy. The entropy, in information theory, is a metric demonstrating the randomness of a variable and is associated with average amount of information present in the variable's outcomes.[ent] A higher amount of information can be obtained by knowing about an entity possessing a greater entropy by value. As shown in Eq. 3.4, the objective of the robots is to collect maximal possible information from the environment through maximizing the entropy in every iteration. In order to decide the next 'best' location, each robot finds the neighbor poi that yields the maximum entropy. In our experiments, the robots use all the visited poi's to predict the information measurements in the unvisited poi's. The results of the entropy metric are shown in Fig. 5.3. Our results are consistent with the findings of [CLD13] and it can be observed that the greedy method cannot exploit the low spatial correlation

in the environment very well. Therefore, the entropy plots do not show a consistent trend. The myopic nature of the used planning strategy reinforces it further. Our proposed framework is generic – a more sophisticated informative path planning strategy (e.g., [CLD13, LS18, MKGH07]) can be used in place of this greedy strategy without impacting the proposed novel coordination mechanism.

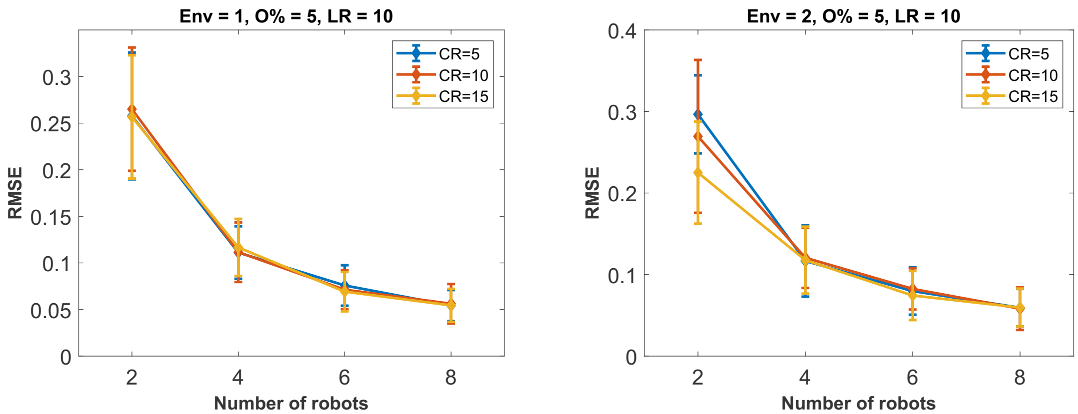

(a) (b)

Figure 5.4: RMSE in prediction for $O\% = 5$: a) ENV1 and b) ENV2.

5.2.2 RMSE

Root mean square error (RMSE) is a standard metric to evaluate the quality of a regression model.[RW06] The equation to compute the RMSE is

$$
\begin{gathered}
RMSE = \sqrt{\sum_{i=1}^{n} \frac{(\hat{y}_i - y_i)^2}{n}} \\
\hat{y}_1, \hat{y}_2, \ldots, \hat{y}_n \text{are predicted values} \\
y_1, y_2, \ldots, y_n \text{are observed values} \\
n \text{is the number of observations}
\end{gathered}
\tag{5.1}
$$

In our case, lower RMSE indicates that the GP prediction is able to closely model the underlying spatial phenomena shown in Fig. 5.2. The results are shown in Figs. 5.4 and 5.5. LIDAR range is fixed to 10 meters in these plots. We can observe that varying communication ranges of the robots do not have significant effect on the RMSE

metric. On the hand, with increasing number of robots in the environments, RMSE values go down consistently. This result is consistent with the results found by the previous information modeling strategies.[CLD13, SKG+07]

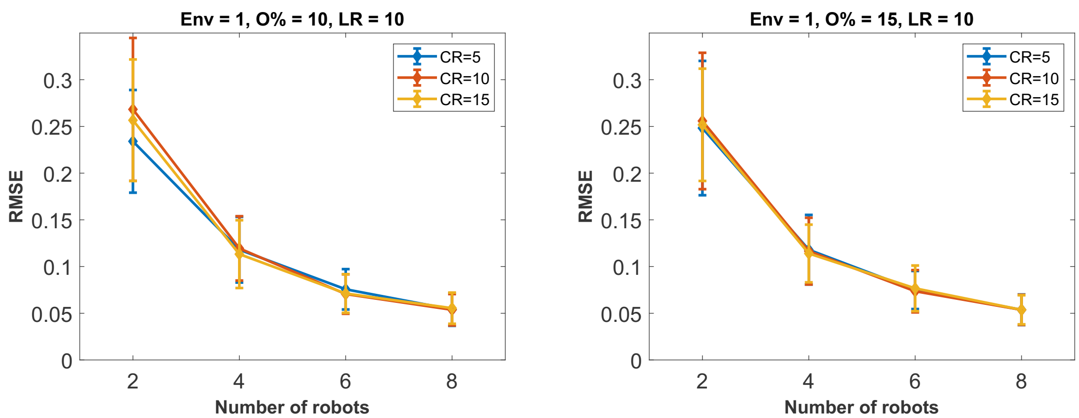

(a) (b)

Figure 5.5: RMSE in prediction for ENV1: a) $O\% = 10$, and b) $O\% = 15$.

We only present the results for the laser range of 10 meters as we believe it to be the best representative. For $LR = 10$, the maximum average RMSE value found in ENV1 is 0.26 with 2 robots and $O\% = 5, 10$ and $CR = 10$. On the other hand, the lowest RMSE value found in ENV1 is 0.05 with 8 robots and $O\% = 15$ and $CR = 10, 15$.

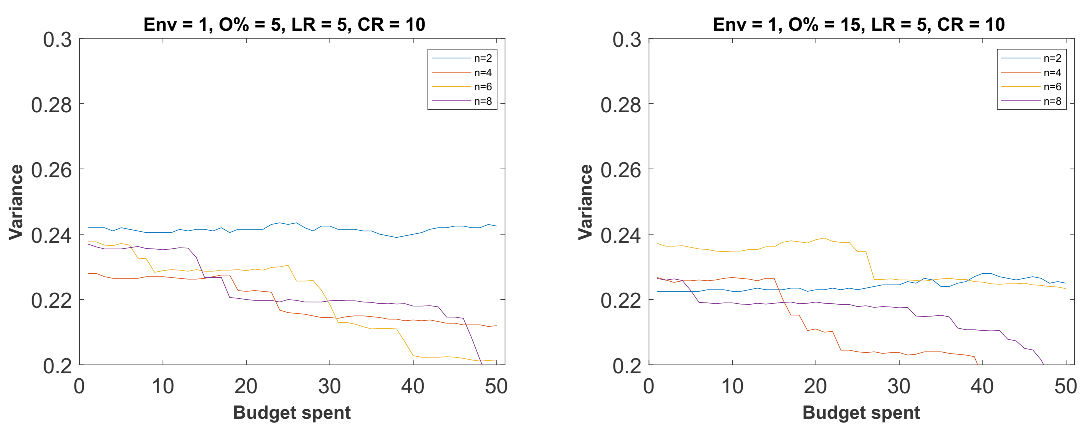

(a) (b)

Figure 5.6: Posterior variance for ENV1 with $LR = 5, CR = 10$: a) $O\% = 5$, and b) $O\% = 15$.

5.2.3 Variance

Following [KSG08], we are also interested in studying the posterior variance metric. The results for this metric are presented in Figs. 5.6 and 5.7. We can see that with time, the posterior variance values generally go down. The slope is usually steeper with more number of robots in the environment.

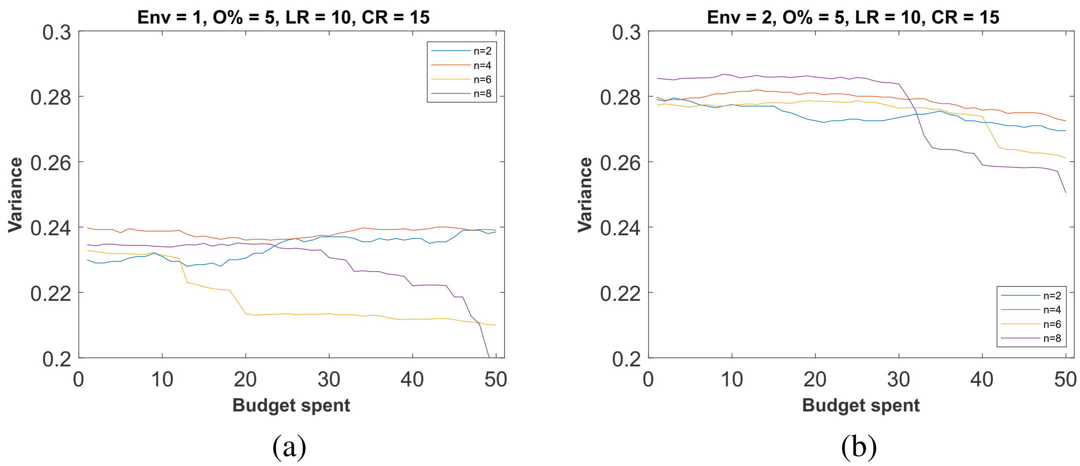

Figure 5.7: Posterior variance with $O\% = 5, LR = 10, CR = 10$ in a) ENV1, and b) ENV2.

5.2.4 Time

Next, we show the execution time of our proposed approach with different test settings. The result is presented in 5.8. From this plot, we can observe that with increasing robot count, the run time does not increase significantly. The average run time with 5% obstacles in ENV1 and 10 meter laser range across all CR and n is 136.89 sec. whereas with 5% obstacles it increases slightly to 138.43 sec. The maximum time taken by any simulation was 479.39 sec. with 8 robots and $O\% = 5$, $CR = 5$, which we believe to be fairly reasonable given the intractable nature of the handled problem.

Execution time of the proposed approach have not varied significantly in ENV2. The maximum run time for ENV2 is found to be 456.48 sec. with 8 robots and $O\% = 5$, $CR = 10$.

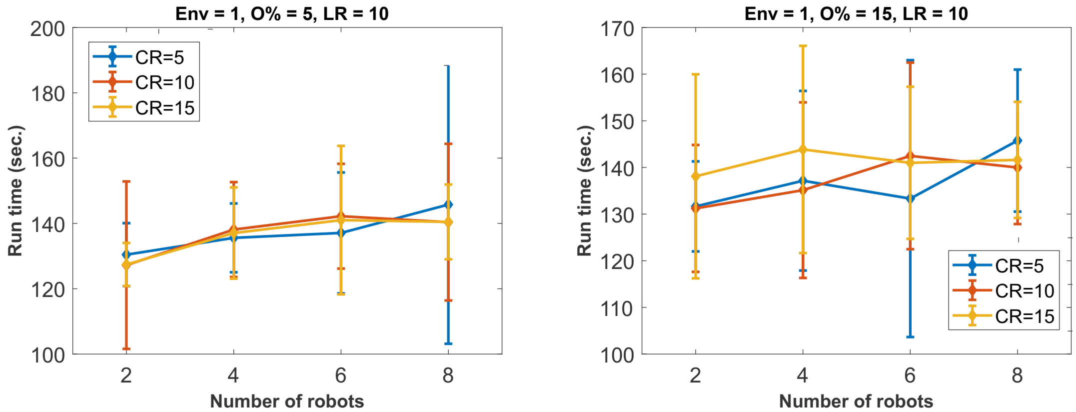

(a) (b)

Figure 5.8: Run time of the proposed approach for ENV1 with $LR = 10$: a) $O\% = 5$, and b) $O\% = 15$.

5.2.5 Load

To demonstrate the effectiveness of the use of the Voronoi repartitioning strategy, we use the following terminologies. Let OPT_l denote the optimal average workload, $Init_l$ denote the initial average workload of the robots based on the initial Voronoi partitioning without considering the obstacles in the environment, and let Fin_l denote the final average load of the robots. OPT_l is calculated using the following formulae:

$$OPT_l = (ENV_{size} - \frac{O\% \times ENV_{size}}{100})/n$$

To show the effectiveness of the repartitioning approach, we calculate the differences between OPT_l and $Init_l$, Fin_l. As the initial partitioning does not take the load into account, it can be highly imbalanced. We are interested in observing how the difference with the optimal has reduced over time. The initial and final differences with the optimal load are calculated as follows.

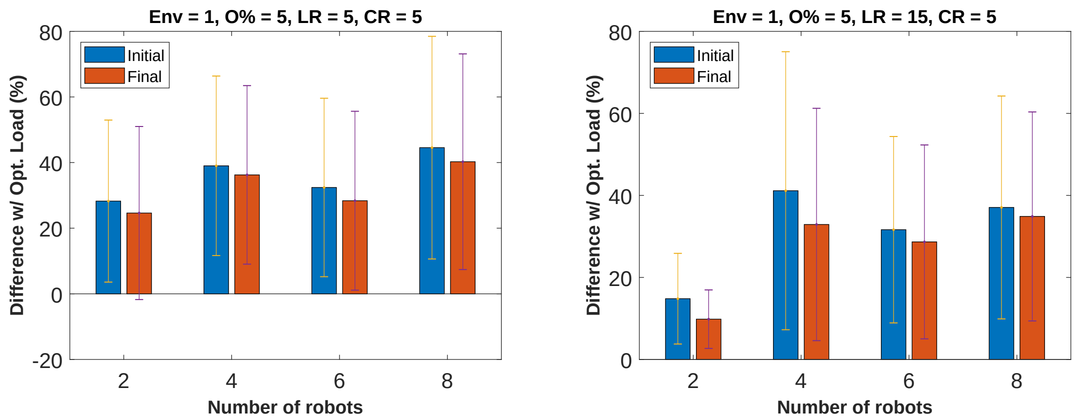

(a) (b)

Figure 5.9: Average load of the robots in ENV1 with $CR = 5$ and $O\% = 5$: a) $LR = 5$, and b) $LR = 15$.

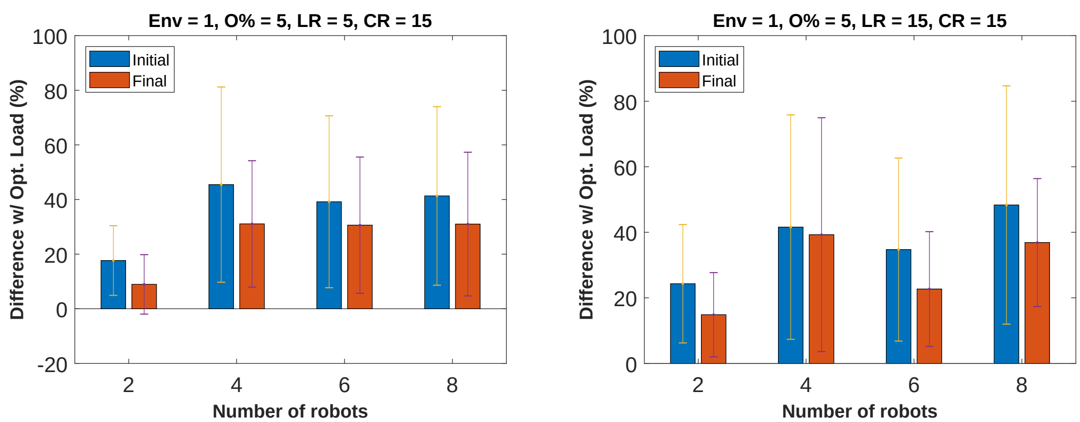

(a) (b)

Figure 5.10: Average load of the robots in ENV1 with $CR = 15$ and $O\% = 5$: a) $LR = 5$, and b) $LR = 15$.

$$
\begin{aligned}
Init_{diff} &= \frac{|Init_l - OPT_l|}{OPT_l} \times 100;\\
Fin_{diff} &= \frac{|Fin_l - OPT_l|}{OPT_l} \times 100;\\
\Delta &= Init_{diff} - Fin_{diff};
\end{aligned}
$$

The results are shown in Figs. 5.9-5.14. The y-axes in these plots show the differences of $Init_l$ and Fin_l with the OPT_l, i.e., $Init_{diff}$ and Fin_{diff}. The maximum differ-

ence (Δ) between $Init_{diff}$ and Fin_{diff} has been found to be 180.93% with 8 robots and $O\% = 15, CR = 15, LR = 15$. Generally, we have noticed that with higher communication range, the value Δ goes higher, i.e., the robots are better able to reduce the initial load imbalance. This is highly intuitive because of the fact that with higher CR, they can communicate with farther robots and the repartitioning process can happen between potentially more number of robots, which in turn bring lowers the load imbalance.

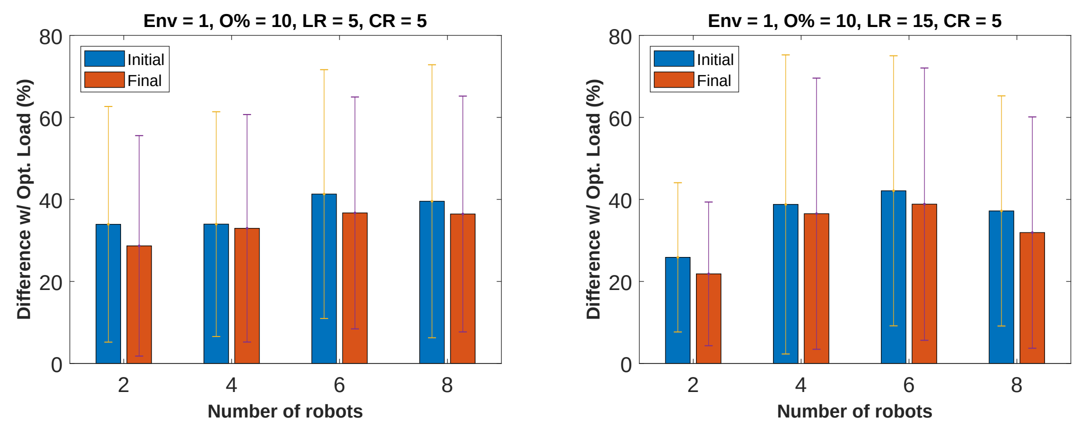

(a) (b)

Figure 5.11: Average load of the robots in ENV1 with $CR = 5$ and $O\% = 10$: a) $LR = 5$, and b) $LR = 15$.

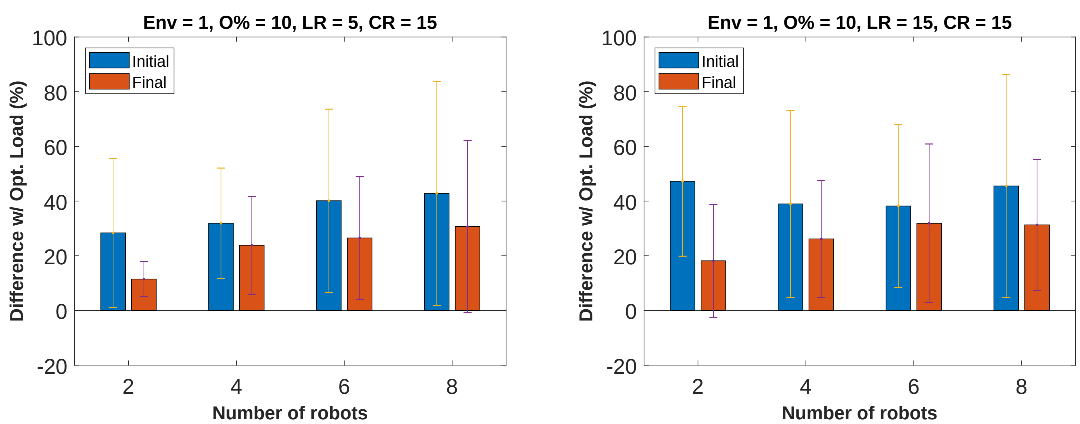

(a) (b)

Figure 5.12: Average load of the robots in ENV1 with $CR = 15$ and $O\% = 10$: a) $LR = 5$, and b) $LR = 15$.

Similarly, longer laser range helps the robots to detect farther obstacles. Along with higher communication range, this enables to robot to better reduce the Δ metric. This

is also evident from our results. For example, with $LR = 5$ and $CR = 5$, the average value of Δ across all n and $O\%$ is 3.51% while the maximum being 98.25%. On the other hand, with LR and CR both set to 15, this average increases to 10.12%. When we have fixed the obstacle locations and only robot locations are varied, a better load-balancing can be noticed (Figure 5.15.(a)). However, if only obstacles are randomly generated and the robot locations are fixed (Figure 5.15.(b)), the difference with the optimal load metric is similar to when both are randomly generated (Figure 5.12.(b)).

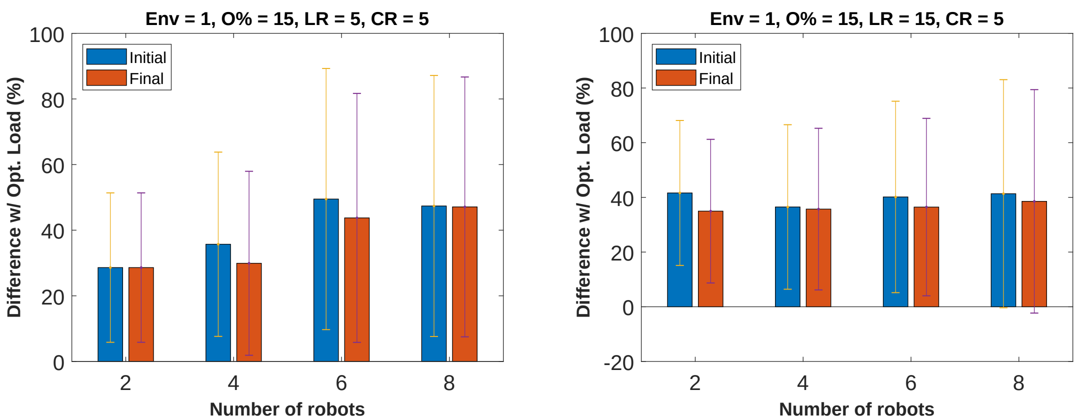

(a) (b)

Figure 5.13: Average load of the robots in ENV1 with $CR = 5$ and $O\% = 15$: a) $LR = 5$, and b) $LR = 15$.

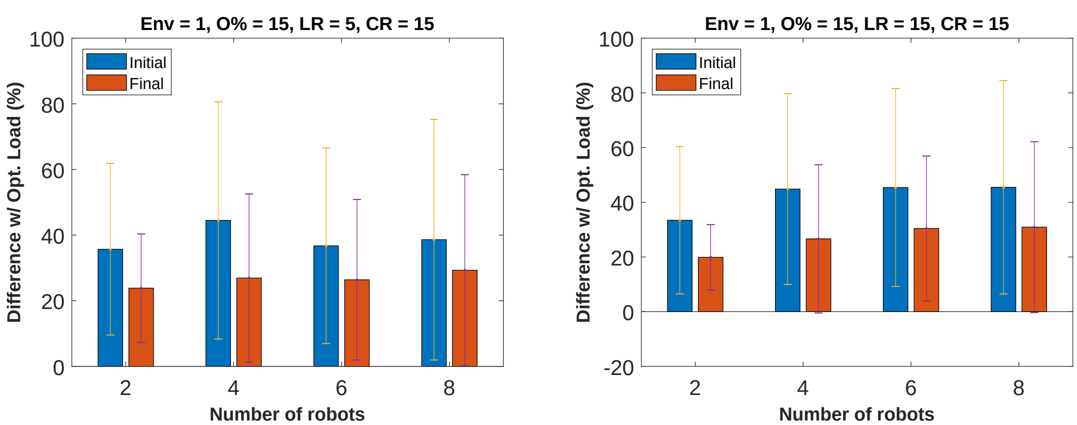

(a) (b)

Figure 5.14: Average load of the robots in ENV1 with $CR = 15$ and $O\% = 15$: a) $LR = 5$, and b) $LR = 15$.

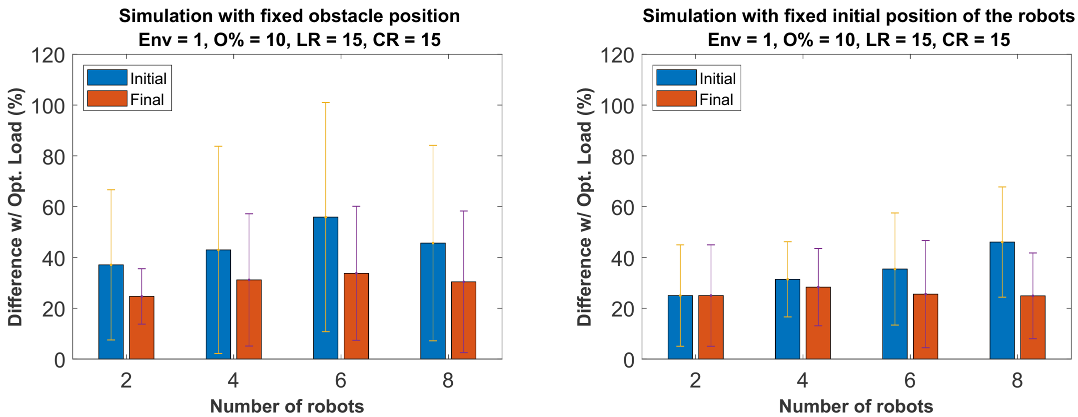

(a) (b)

Figure 5.15: Average load of the robots in ENV1 with $CR = 15$, $O\% = 10$, and $LR = 15$.

On the other hand, similar trend can be noticed for different obstacle percentages in the environment. For example, with $O\% = 5$, the average value of Δ across all $n, CR,$ and LR is 7.34%. This average value of Δ increases to 7.62% with $O\% = 15$. It is evident from the results that in a more cluttered environment, the repartitioning approach is more effective.

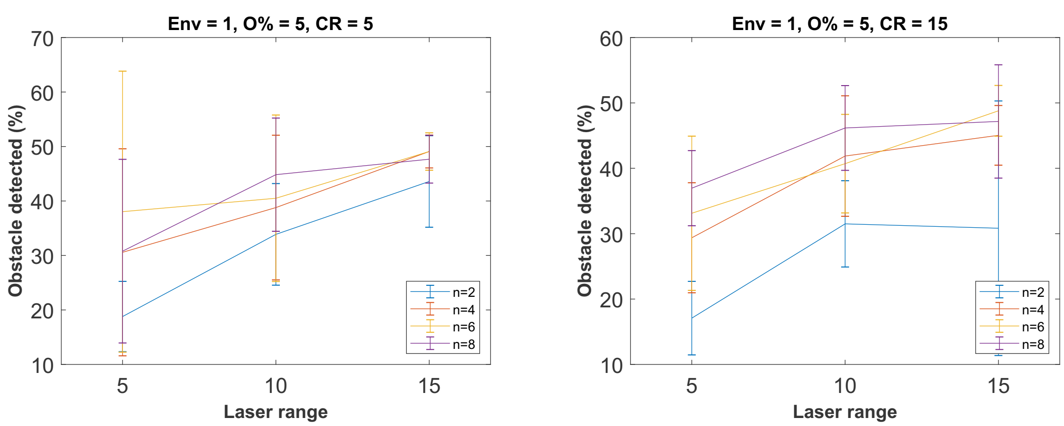

(a) (b)

Figure 5.16: Detected obstacle percentage for ENV1 with $O\% = 5$: a) $CR = 5$, and b) $CR = 15$.

5.2.6 Detected Obstacles

We are also interested in finding how many obstacles are detected by the robots. This is important for applications like environment mapping. The results are presented in Figs. 5.16-5.18. The detected obstacle percentages are plotted against the laser ranges of the robots' on-board LIDARs. Higher the range, farther obstacles are detected by the robots. One should note that, the results presented here represent the union of all discovered obstacles by all the robots, i.e., if an obstacle is observed by more than one robot, it is counted only once.

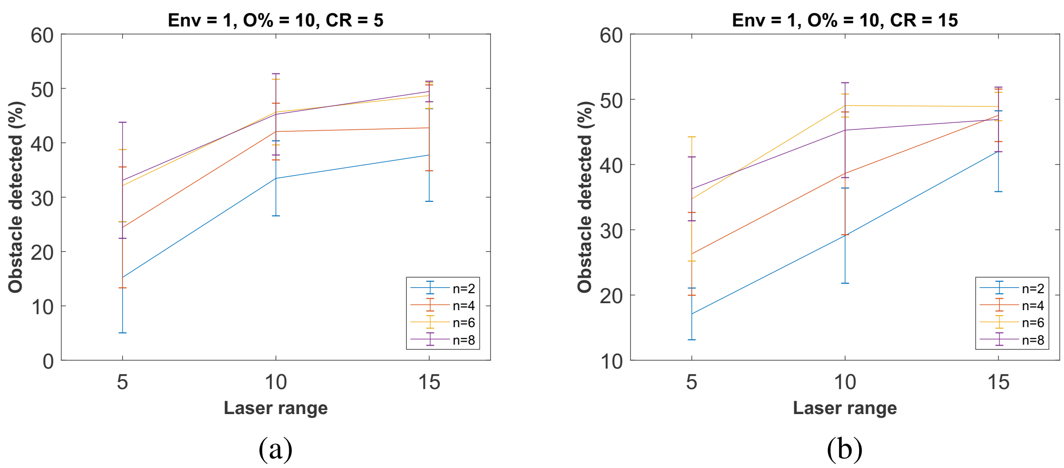

Figure 5.17: Detected obstacle percentage for ENV1 with $O\% = 10$: a) $CR = 5$, and b) $CR = 15$.

As expected, with higher laser range, generally the robots were able to detect more obstacles in the environment. On the other hand, with more robots present in the environment, more obstacles are detected in general. This phenomena is evident in Figs. 5.16-5.18. For example, with 8 robots, $O\% = 5$ and $CR = 5$, on an average 30.79% obstacles are detected with a laser range of 5 meters. This average value increases to 47.67% with a laser range of 15 meters. With 2 robots in the environments, these numbers drop to 18.79% and 43.58% respectively. On the other hand, with 15% grid cells being obstacles in the environment, these averages are 33.40%, 45.47% with 8 robots and 16.13%, 50.60% with 2 robots respectively. Although the average detected obstacle

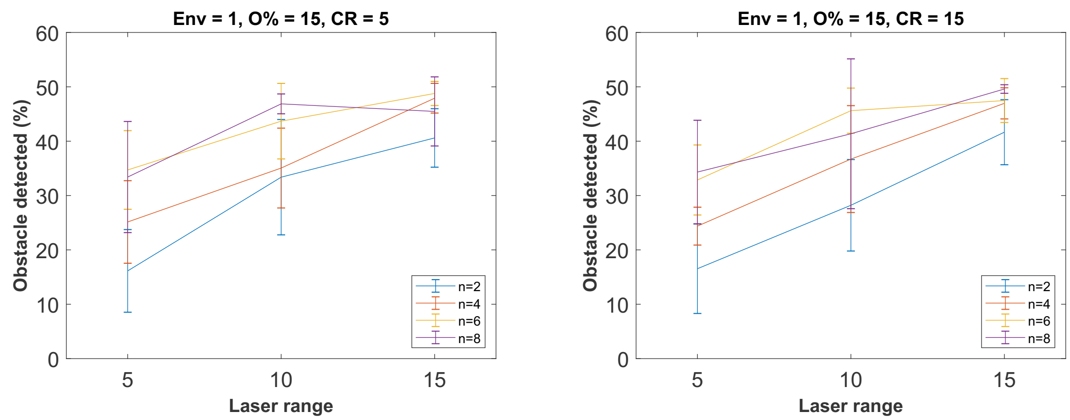

(a) (b)

Figure 5.18: Detected obstacle percentage for ENV1 with $O\% = 15$: a) $CR = 5$, and b) $CR = 15$.

percentage numbers with $O\% = 15$ are slightly lower than that with $O\% = 10$, but in reality the robots detected more obstacle grid cells because the percentage numbers are relative to the total obstacle count in the environment. E.g., with 8 robots and LR set to 15, 114 grid cells are detected as obstacles with $O\% = 5$ and 328 cells as obstacles with $O\% = 15$. Communication range does not have a significant impact on this metric.

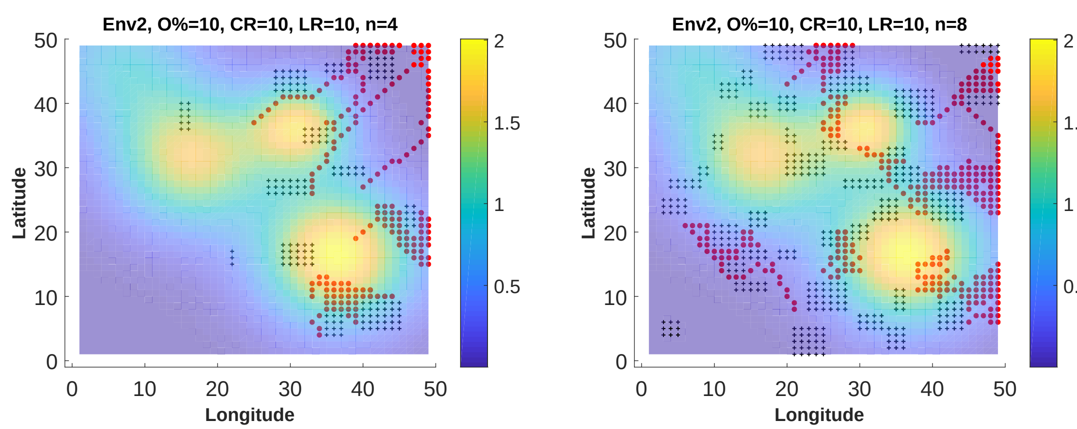

(a) (b)

Figure 5.19: Paths followed by n robots in ENV2 with $O\% = 10, CR = 10, LR = 10$: a) $n = 4$, and b) $n = 8$.

Sample Paths. Finally, we show (Figs. 5.19-5.21) a set of sample informative paths planned by different number of robots in the two tested environments. The red circles represent the sequence of cells visited by the robots, i.e., paths and the black + signs

indicate the detected obstacles in the environments. As each robot's local GP model is trying to reduce the information measurement uncertainty, we can notice that the robots are usually scattering in the environment while collecting information. We also observe that with higher laser range, more obstacles in the environment are detected by the robots as expected.

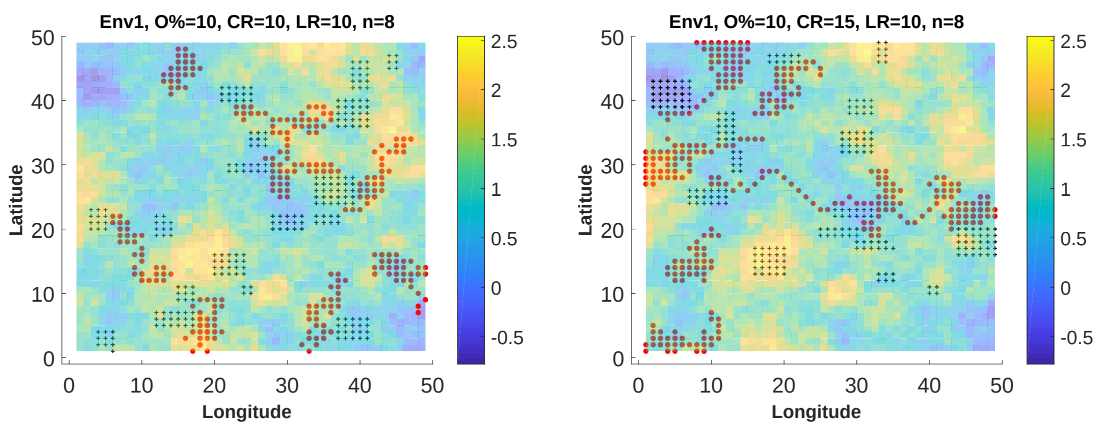

(a) (b)

Figure 5.20: Paths followed by 8 robots in ENV1 with $O\% = 10, LR = 10$, and: a) $CR = 10$, and b) $CR = 15$.

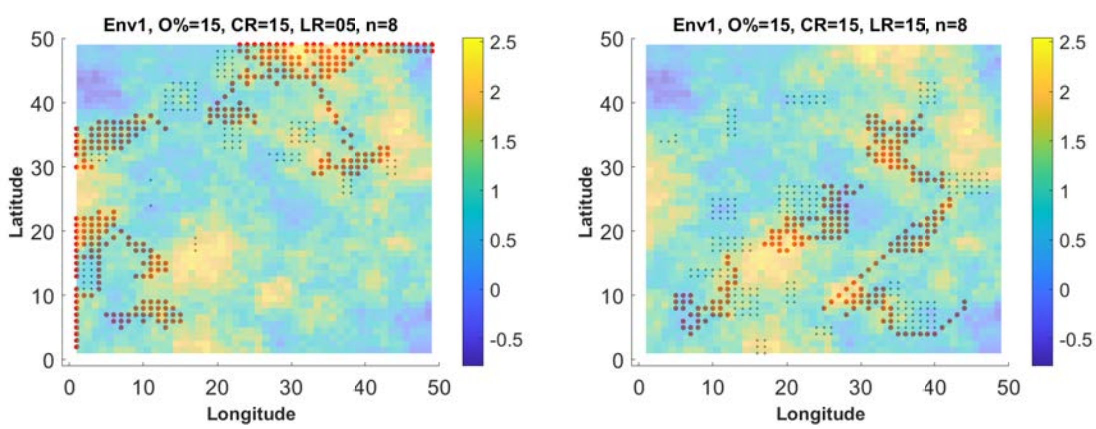

(a) (b)

Figure 5.21: Paths followed by 8 robots in ENV1 with $O\% = 15, CR = 15$, and: a) $LR = 5$, and b) $LR = 15$.

Key Findings. The key findings from the results can be summarized as follows: 1) Repartitioning is more effective in a cluttered environment; 2) with higher values of CR

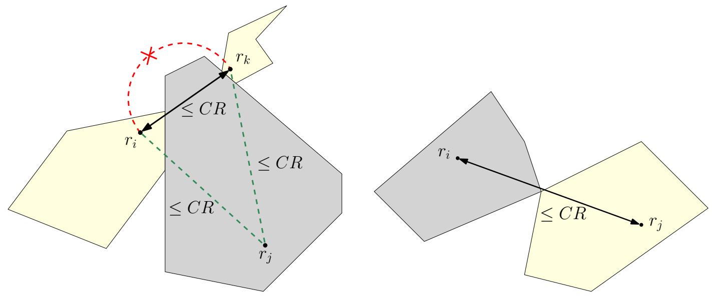

(a) (b)

Figure 5.22: Two limiting scenarios are illustrated: (a) The robots r_i, r_k will not participate in repartitioning since their Voronoi cells do not share boundary; (b) The robots r_i, r_j will not participate in repartitioning since their boundaries touch only at a point.

and LR (i.e., better sensors), load balancing is more effective; 3) with higher number of robots present in the environment, prediction model is better; 4) the percentage of obstacles identified in the environment is proportional to the number of robots used.

5.3 Limitations and Discussion

Our results show that when two or more robots come within communication range, they are successfully able to reduce the imbalance in their workloads by repartitioning their current Voronoi cells. Although our approach can handle online repartitioning successfully with convex obstacles, it does not work with concave obstacles. Online repartitioning with concave obstacles will require further sophisticated communication mechanism. As our model does not require the robots to form a connected network at any point of time, we cannot guarantee a complete allocation of all regions with concave obstacles present in the environment with our current approach. Our proposed approach also does not handle a couple of corner cases as shown in Fig. 5.22.

The first scenario is depicted in Fig. 5.22.(a) where two robots r_i and r_k are within their communication ranges but their current Voronoi cells do not intersect with each other. As these cells are disjoint, if the robots were to do repartition their cells, the union of them would produce multiple disjoint polygons. Our algorithm does not par-

ticularly handle this scenario. Thus, r_i and r_k continue their information collection work in their current cells without executing the repartitioning procedure.

Secondly, if two robots r_i and r_j are within their communication ranges but their respective Voronoi cells intersect on a single point, the robots do not participate in the repartitioning process. As the Voronoi cells intersect on a point only, if one of the robots were to go to the other side of that intersecting point after the repartitioning, it would need to go through that single point. We do not explicitly handle this situation and the robots continue to collect information in their current cells. The scenario is shown in Fig. 5.22.(b).

On the other hand, if r_i and r_j are within each other's communication range and r_j and r_k are within each other's communication range but r_i and r_k are outside of each other's communication range, then the repartitioning is possible through r_j (see illustration in Fig. 5.22.(a)).

Our work in this research is the first to study the MIPP problem under communication constraints in unknown environments. Given that there is no guarantee that two or more robots will meet each other during the information collection process, it is difficult to provide optimal load-balancing. However, we have modeled the solution in a way that whenever robots meet with each other, they locally balance their loads. Although we have used a greedy informative path planning strategy that has been proved to yield reasonably good solutions in the literature,[DGK19,CLD13] any other sophisticated information collection strategies, such as proposed in [LS18], or mutual information-based optimization technique instead of entropy-based, such as proposed in [KSG08], can be used in its place. This will not change the overall coordination strategy for the communication constrained robots in an initially unknown environment proposed in this

research. This will help the future researchers in this topic to better reason about impact of the highly practical yet less studied communication constraints of the robots, especially in environments where the obstacle locations and shapes are initially unknown. Moreover, our proposed framework can easily be adapted for coverage path planning, i.e., to cover all the free locations in the environment while handling the robot's limited communication ranges in an unknown environment.

CHAPTER 6

Conclusion and Future Work

In this research, we have proposed a continuous Voronoi partitioning based informative path planning algorithm to collect a maximal amount of information from an unknown environment using a set of mobile robots. In our model, the robots are randomly placed in the environment and are allocated a non-overlapping initial region, which is a part of the bigger environment, to start the information collection. At this moment, each robot only knows about its starting position and the allocated region but do not have any knowledge of the distribution and amount of obstacles. As the robots move to collect more information, they possibly discover new obstacles and their corresponding perceptions about their allocated workloads, i.e., the amount of area they need to cover for information collection, change. When two or more robots come within each other's communication ranges, they share their current perceptions about the environment and based on their local knowledge bases, they repartition their initial Voronoi cells, which have been allocated to them without taking the obstacles of this unknown environment into account. This helps the robots to achieve a better workload balance. Our used path planning approach is distributed in nature and each robot plans its local path only for one future informative way-point at a time. Results show that our proposed strategy helps the robots to model the underlying spatial phenomena that are close to reality while successfully reducing the imbalance in the amount of allocated free areas to the robots.

In the future, we plan to observe the enhancement in the information collection and environment modeling by switching from greedy to a novel path planning algorithm, e.g., dynamic programming. A dynamic programming approach considers path plan-

ning by looking ahead a certain number of steps into the future with a maximum look-a-head limit to the horizon. This approach is analogous to planning a move in a game that involves decision making by considering a series of moves ahead.

Furthermore, we plan to test the proposed approach using a team of unmanned aerial vehicles, incorporate time-varying obstacles and various connectivity constraints into our proposed framework and perform experiments with real robots.

BIBLIOGRAPHY

[AAMO14] Jose J Acevedo, Begona C Arrue, Ivan Maza, and Anibal Ollero. A decentralized algorithm for area surveillance missions using a team of aerial robots with different sensing capabilities. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 4735–4740. IEEE, 2014.

[BSA+12] WB Banerdt, S Smrekar, L Alkalai, T Hoffman, R Warwick, K Hurst, W Folkner, P Lognonné, T Spohn, S Asmar, et al. Insight: an integrated exploration of the interior of mars. *LPI*, (1659):2838, 2012.

[BSM+10] Andreas Breitenmoser, Mac Schwager, Jean-Claude Metzger, Roland Siegwart, and Daniela Rus. Voronoi coverage of non-convex environments with a group of networked robots. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 4982–4989. IEEE, 2010.

[CB95] H Choset and J.W Burdick. Sensor based planning, part i: The generalized voronoi graph. In *Proc. IEEE Int. Conf. on Robotics and Automation*, 1995.

[CLD13] Nannan Cao, Kian Hsiang Low, and John M Dolan. Multi-robot informative path planning for active sensing of environmental phenomena: A tale of two algorithms. In *Proceedings of the 2013 International Conference on Autonomous agents and Multi-agent Systems*, pages 7–14, 2013.

[CMKB04] Jorge Cortes, Sonia Martinez, Timur Karatas, and Francesco Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on robotics and Automation*, 20(2):243–255, 2004.

[DCFB12] Joseph W Durham, Ruggero Carli, Paolo Frasca, and Francesco Bullo. Discrete partitioning and coverage control for gossiping robots. *IEEE Transactions on Robotics*, 28(2):364–378, 2012.

[DD16] Ayan Dutta and Prithviraj Dasgupta. Simultaneous configuration formation and information collection by modular robotic systems. In *2016 IEEE international conference on robotics and automation (ICRA)*, pages 5216–5221. IEEE, 2016.

[DGK19] Ayan Dutta, Anirban Ghosh, and O Patrick Kreidl. Multi-robot informative path planning with continuous connectivity constraints. In *Robotics and Automation (ICRA), 2019 IEEE International Conference on*. IEEE, 2019.

[ent] Entropy (information theory) - wikipedia. `https://en.wikipedia.org/wiki/Entropy_(information_theory)`.

[GKS05] Carlos Guestrin, Andreas Krause, and Ajit Paul Singh. Near-optimal sensor placements in gaussian processes. In *Proceedings of the 22nd international conference on Machine learning*, pages 265–272. ACM, 2005.

[GWD12] K Guruprasad, Zachary Wilson, and Prithviraj Dasgupta. Complete coverage of an initially unknown environment by multiple robots using voronoi partition. In *International Conference on Advances in Control and Optimization in Dynamical Systems*, 2012.

[HDG16] Kurt Hungerford, Prithviraj Dasgupta, and KR Guruprasad. A repartitioning algorithm to guarantee complete, non-overlapping planar coverage with multiple robots. In *Distributed Autonomous Robotic Systems*, pages 33–48. Springer, 2016.

[HGP+14] Gregory Hitz, Alkis Gotovos, Franccois Pomerleau, Marie-Eve Garneau, Cedric Pradalier, Andreas Krause, and Roland Y Siegwart. Fully autonomous focused exploration for robotic environmental monitoring. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 2658–2664. IEEE, 2014.

[HS10] Geoffrey Hollinger and Sanjiv Singh. Multi-robot coordination with periodic connectivity. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 4457–4462. IEEE, 2010.

[HS14] Geoffrey A Hollinger and Gaurav S Sukhatme. Sampling-based robotic information gathering algorithms. *The International Journal of Robotics Research*, 33(9):1271–1287, 2014.

[KCS16] Stephanie Kemna, David A Caron, and Gaurav S Sukhatme. Adaptive informative sampling with autonomous underwater vehicles: Acoustic versus surface communications. In *OCEANS 2016 MTS/IEEE Monterey*, pages 1–8. IEEE, 2016.

[KRNG+17] Stephanie Kemna, John G Rogers, Carlos Nieto-Granda, Stuart Young, and Gaurav S Sukhatme. Multi-robot coordination through dynamic voronoi partitioning for informative adaptive sampling in communication-constrained environments. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pages 2124–2130. IEEE, 2017.

[KSG08] Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9(Feb):235–284, 2008.

[KWP+11] Sertac Karaman, Matthew R Walter, Alejandro Perez, Emilio Frazzoli, and Seth Teller. Anytime motion planning using the rrt. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pages 1478–1483. IEEE, 2011.

[LDK09] Kian Hsiang Low, John M Dolan, and Pradeep Khosla. Information-theoretic approach to efficient adaptive path planning for mobile robotic environmental sensing. In *Nineteenth International Conference on Automated Planning and Scheduling*, 2009.

[LHL16] Zhan Wei Lim, David Hsu, and Wee Sun Lee. Adaptive informative path planning in metric spaces. *The International Journal of Robotics Research*, 35(5):585–598, 2016.

[LLHJ+16] Pierre FJ Lermusiaux, Tapovan Lolla, Patrick J Haley Jr, Konuralp Yigit, Mattheus P Ueckermann, Thomas Sondergaard, and Wayne G Leslie. Science of autonomy: Time-optimal path planning and adaptive sampling for swarms of ocean vehicles. In *Springer Handbook of Ocean Engineering*, pages 481–498. Springer, 2016.

[LN13] Yugang Liu and Goldie Nejat. Robotic urban search and rescue: A survey from the control perspective. *Journal of Intelligent & Robotic Systems*, 72(2):147–165, 2013.

[LS18] Wenhao Luo and Katia Sycara. Adaptive sampling and online learning in multi-robot sensor coverage with mixture of gaussian processes.

In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6359–6364. IEEE, 2018.

[MCdFB+09] Ruben Martinez-Cantin, Nando de Freitas, Eric Brochu, José Castellanos, and Arnaud Doucet. A bayesian exploration-exploitation approach for optimal online sensing and planning with a visually guided mobile robot. *Autonomous Robots*, 27(2):93–103, 2009.

[Mic04] Olivier Michel. Cyberbotics ltd. webots™: professional mobile robot simulation. *International Journal of Advanced Robotic Systems*, 1(1):5, 2004.

[MKGH07] Alexandra Meliou, Andreas Krause, Carlos Guestrin, and Joseph M Hellerstein. Nonmyopic informative path planning in spatio-temporal models. In *AAAI*, volume 10, pages 16–7, 2007.

[RW06] Carl Edward Rasmussen and Christopher KI Williams. *Gaussian processes for machine learning*, volume 1. MIT press Cambridge, 2006.

[SKG+07] Amarjeet Singh, Andreas Krause, Carlos Guestrin, William J Kaiser, and Maxim A Batalin. Efficient planning of informative paths for multiple robots. In *IJCAI*, volume 7, pages 2204–2211, 2007.

[SKK09] Amarjeet Singh, Andreas Krause, and William J Kaiser. Nonmyopic adaptive informative path planning for multiple robots. In *Twenty-First International Joint Conference on Artificial Intelligence*, 2009.

[SMR06] Mac Schwager, James McLurkin, and Daniela Rus. Distributed coverage control with sensory feedback for networked robots. In *robotics: science and systems*, pages 49–56, 2006.

[Sol04] Neal Solomon. Methods and apparatus for multi robotic system involving coordination of weaponized unmanned underwater vehicles, February 12 2004. US Patent App. 10/421,983.

[Ste87] Fortune Steven. A sweepline algorithm for voronoi diagrams. *Algorithmica*, 2(1-4):153, 1987.

[TWWS08] Girma S Tewolde, Changhua Wu, Yu Wang, and Weihua Sheng. Distributed multi-robot work load partition in manufacturing automation. In *Automation Science and Engineering, 2008. CASE 2008. IEEE International Conference on*, pages 504–509. IEEE, 2008.

[Vor07] G Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. *Journal für die Reine und Angewandte Mathematik*, 133:97–178, 1907.

VITA

Amitabh Bhattacharya has Bachelor degree in Computer Science and Engineering from Rajiv Gandhi Proudyogiki Vishwavidyalaya, INDIA. Amitabh expects to receive a Master of Science in Computer and Information Sciences with a concentration in Computer Science from the University of North Florida in May 2020. Dr. Ayan Dutta and Dr. Patrick O Kreidl are serving as Amitabh's thesis advisors. Previous to the graduation, Amitabh has spent thirteen years within the Space industry located in Bengaluru, INDIA working as a Data Scientist.

His research work focuses on multi-robot informative path planning with the value of distributed coordination and multi-stage look-ahead. He has co-authored the paper *Multi-robot Informative Path Planning in Unknown Environments Through Continuous Region Partitioning* which is accepted at the 32nd International FLAIRS Conference, 2019. The research also won the Real-World Poster Award at the UNF Symposium in November 2019.