Distributed Multi-Robot Information Gathering using Path-Based Sensors in Entropy-Weighted Voronoi Regions

Alkesh K. Srivastava[1], George P. Kontoudis[1],
Donald Sofge[2], and Michael Otte[1]

[1] University of Maryland, College Park, MD 20742, USA,
[alkesh,kont,otte]@umd.edu,
[2] U.S. Naval Research Laboratory, Washington, DC 20375, USA
donald.a.sofge.civ@us.navy.mil

Abstract. In this paper, we present a distributed information-gathering algorithm for multi-robot systems that use multiple path-based sensors to infer the locations of hazards within the environment. Path-based sensors output binary observations, reporting whether or not an event (like robot destruction) has occurred somewhere along a path, but without the ability to discern where along a path an event has occurred. Prior work has shown that path-based sensors can be used for search and rescue in hazardous communication-denied environments—sending robots into the environment one-at-a-time. We extend this idea to enable multiple robots to search the environment simultaneously. The search space contains targets (human survivors) amidst hazards that can destroy robots (triggering a path-based hazard sensor). We consider a case where communication from the unknown field is prohibited due to communication loss, jamming, or stealth. The search effort is distributed among multiple robots using an entropy-weighted Voronoi partitioning of the environment, such that during each search round all regions have approximately equal information entropy. In each round, every robot is assigned a region in which its search path is calculated. Numerical Monte Carlo simulations are used to compare this idea to other ways of using path-based sensors on multiple robots. The experiments show that dividing search effort using entropy-weighted Voronoi partitioning outperforms the other methods in terms of the information gathered and computational cost.

Keywords: path planning, multi-robot exploration, distributed decision making, Shannon Information Theory, information-driven partitioning, Bayesian Probability

1 Introduction

Autonomous agents are increasingly used in scientific, commercial, and military applications. In critical missions of locating human survivors after disasters, the deployment of multi-robot systems may be more efficient than deploying a single

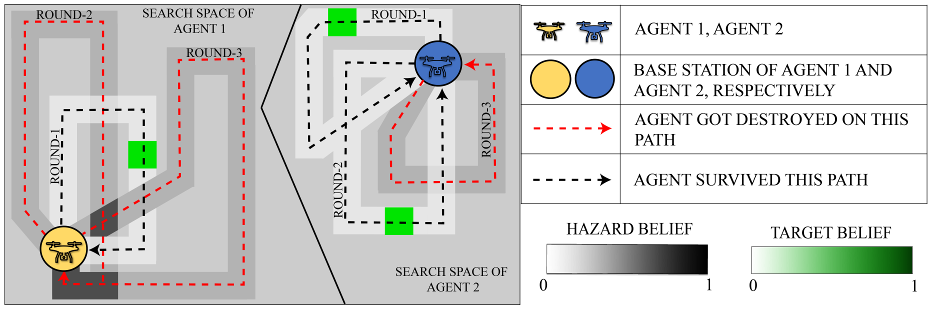

Fig. 1. Information gathering with two robots where the search space is decomposed in two entropy-weighted Voronoi partitions. The dotted lines represent the path that the agent traverses in the three search rounds. Robot 1 efficiently traverses the path of round 1, reporting a potential target at the location shown in green. However, it gets destroyed on round 2 and 3, and thus we only know that a hazard exist along these two paths. The overlapping areas of the round 2 and 3 for robot 1 have higher probability of hazard existence (darker gray). Robot 2 gets destroyed at round 1, and survives round 2 and 3. Since robot 2 reports a target at the same location at round 2 and 3, the probability of target existence is higher (darker green).

agent. In this work, we are primarily motivated by search and rescue missions of human survivors in lethally hostile environments. We consider a general case where information exchange between agents in the field is prohibited. Yet, the locations of *base stations*—where robots start and end each search round's journey through the environment—provide connectivity to a shared centralized server (i.e., before and after each journey). The environment is completely unknown and lethally hostile such that agents may be destroyed by a stationary hazards during each search round, and the locations of hazards are initially unknown. Due to the lack of communication in the field, agent destruction happens at an unknown location along the agent's path and leads to loss of any information collected about targets along that path. However, agent loss also *provides* information about the location of hazards within the environment (see Fig. 1). Even though we do not know where along the path the event of robot destruction has occurred, the shape of the path constrains the set of possible locations at which the hazard may have destroyed the agent.

The idea that a path's shape provides information that can be used to update a belief map (e.g., of hazard presence) given that an event (e.g., robot destruction) has occurred somewhere along a path is called a *path-based sensor* [1]. The objective of the work described in the current paper is to extend the use of path-based sensors to the case where multiple robots explore the environment simultaneously and in a distributed fashion, to improve the computational efficiency and accelerate the convergence speed of the hazard and target belief maps.

The proposed distributed methodology decomposes this problem into two distinct sub-problems – decomposition of search space, and path planning in the decomposed space. Thus, this technique is a decoupled approach consisting of two steps: (i) centralized Shannon entropy-driven partitioning of the environment into a set of disjoint regions (one region per agent); and (ii) distributed information theoretic path planning of each agent in its region. The first step assigns each agent a particular partitioned region in the search space. To ensure that each point in the partitioned space is given to the nearest base station and affects the entropy of the partitioned space the least, we use entropy-weighted Voronoi partitioning. This helps distribute the workload among the agents for the next step, where each agent simultaneously and locally plans an information-driven path in its corresponding search space, as described in Fig. 1. Thereby helping the multi-robot system to efficiently and robustly infer the hazard and target belief map of the environment.

Related Work: The concept of *path-based sensors* was introduced in [1], which focused on the scenario where multiple agents explored the environment *sequentially.* The current paper compares three different extensions of the path-based sensor idea in which m agents explore the environment in parallel during each search round.

We now discuss two different bodies of related work. We begin by reviewing the literature that studied the passive cooperation problem for coverage path planning (CPP) algorithms, and then we review work on information gathering. CPP algorithms use a combination of environmental partitioning and passive communication [3]. In [4], Nair and Guruprasad use a partition-and-cover approach based on Voronoi partitioning to achieve a passive cooperation between agents to avoid task duplicity. Distributed strategies are used in [5–12] to keep track of regions covered. The use of geodesic or Manhattan distance as a distance metric to improve exploration is proposed in [13]. Cooperative sweeping of the environment by multiple robots is studied in [14] and the use of entropy as a weight for Voronoi tessellation was discussed in [15]. In our work, we employ a divide-and-conquer approach to environmental partitioning using entropy-weighted Voronoi decomposition. Our work assumes that base stations have a centralized communication topology, and we focus on using path-based sensors to gather information about hazards when communication is denied in the field.

Recursive Bayesian filters are used in the literature to update belief maps of the environment [16, 17]. However, [17] assumed knowledge of exact location of malfunctioning agents, i.e., false positive location is known. A formal derivation of mutual information gradient for information gathering is presented in [18]. The latter introduced information surfing, the idea of using the gradient of mutual information for information gathering, and the work is extended in [19, 20], but with the assumption that a hazard could not destroy the agents. The focus of this work is on combining path-based sensors and information gathering in environments with lethal hazards and targets, using a team of imperfect robots that may malfunction (false positive) or report false negative observations.

Contribution: The contribution of this paper is twofold. First, we present a distributed methodology for information gathering in a multi-robot system using path-based sensors. The proposed technique is a decoupled two-step scheme that uses a distributed information-theoretic mechanism for partitioning the environment and local information-theoretic planner to maximize information gathering. Second, we also synthesize a centralized global planner for information gathering in a multi-robot system that uses sequential Bayesian filter to calculate estimated belief maps of the environment.

2 Problem Formulation

Consider a spatial environment which is divided into $a \times b$ discrete cells composing the search space $\mathcal{S} \in \mathbb{R}^2$. A team of M agents is deployed in the environment, where each agent is denoted by $i = 1, \dots, M$. The search space includes n base stations $D = \{d_1, d_2, \dots, d_n\} \subset \mathcal{S}$, where $n \leq M$. Note that in this paper we consider the case of $n = M$. Each agent i is located at its corresponding base station d_i to explore a specific bounded region $S_i \subset \mathcal{S}$. The union of the regions of exploration S_i constitute the search space $\bigcup_{i=1}^{m} S_i = \mathcal{S}$, with no overlap $\forall i \neq j$, $S_i \cap S_j = \emptyset$, where $i, j \in \{1, 2, \dots, m\}$. Each agent i can visit up to l_i cells c in its search space forming the path $\zeta_i = \langle c_1, c_2, \dots, c_{l_i} \rangle \subseteq S_i$ such that the exploration starts and ends at the agent's corresponding base station d_i, i.e., $c_1 = c_{l_i} = d_i$. When all the surviving agents reach their respective base stations at the end of their exploration, we terminate one search round $r \in \mathbb{Z}_{>0}$. We define the time taken by an agent i to move from cell c_j to c_{j+1} as one timestep t. Thus, the duration of a search round can be defined as $t_r \coloneqq \max_{i \in \{1,\dots,M\}} l_i t$.

Definition 1. *Path-based sensor [1]*
A sensor that reports whether or not an event has occurred somewhere along a path, but has no conception of where along the path that event has taken place.

The search space includes hazardous elements $\mathcal{Z} = \{0, 1\}$ that may destroy the agent, where $Z = 1$ indicates the presence of a hazard, while $Z = 0$ denotes the absence of a hazard in a particular cell. If the agent i is destroyed *anywhere* along the path ζ_i then the path-based sensor (Definition 1) is triggered ($\Theta = 1$), and if the agent i survives the path ζ_i then we consider that the path-based sensor is not triggered ($\Theta = 0$). We assume that the path-based sensor may report false positive and false negative triggering. False-positive accounts for faulty or malfunctioning robots that get destroyed regardless of the presence of a hazard, whereas false-negative accounts for the chance of the robot surviving a cell despite having a hazard. The search space also includes some elements of interest, hereafter referred to as targets, $\mathcal{X} = \{0, 1\}$, where $X = 1$ indicates the presence of a target, while $X = 0$ denotes the absence of a target in a particular cell. The presence of a target is recorded by a noisy sensor that may also report a false positive or a false negative observation of the target. The sensor used for the detection of targets is not a path-based sensor. In other words, if a robot survives a path, the target sensor reports the exact whereabouts of the target observation along the survived path.

All agents start their exploration simultaneously at the beginning of a search round. Inter-agent communication is prohibited at all times. However, the successful traverse of a path by an agent indicates information about absence of hazards, and each survived agent also transmits information about targets along their path to its base station. In other words, each agent i reports its observation only to its base station d_i. A central server σ can obtain this information from each base station to update the global belief map.

Problem 1. Given a team of M agents, with inter-agent communication prohibited at all times, in an environment with multiple available base stations D that communicate with a central server σ, the task of the agents is to explore a completely unknown environment efficiently and gather information about targets X and hazards Z.

Remark 1. The use of the centralized server to partition the environment and update shared belief maps is convenient, but not required. If a completely decentralized and distributed algorithm is desired, then the use of the centralized server can be eliminated, assuming that base stations are able to synchronize their data after each search round. Each base station can perform the same (deterministic) weighted Voronoi partitioning algorithm in parallel such that all base stations achieve the same partitioning.

3 Distributed Information Gathering

In this section, we propose three methodologies to address Problem 1. The first methodology consists of a decoupled two-step scheme: i) entropy-weighted Voronoi partitioning as a centralized method to decompose the environment in S_i spaces; and ii) mutual information-based planner from each base station d_i that maximizes local information gathering. In the second method, we assume no partitioning of the environment, thus information-based planning is executed in the search space $\mathcal{S}$. However, we deploy a team of robots simultaneously from various base stations d_i. Lastly, we discuss a global planning method that calculates expected belief maps for each agent i before deployment from the same base station d and then plans an information-driven path for all agents.

3.1 Distributed Entropy Voronoi Partition and Planner (DEVPP)

To address the obscure information obtained from a path-based sensor, [1] proposed integration of all potential location for a path-based sensor triggering. However, this increases the computational complexity of optimal path solutions beyond short path lengths. Therefore, we propose a distributed strategy with multiple agents that partitions the environment in a way that maximizes the expected information gain of the entire search space $\mathcal{S}$ without task duplicity. The structure of the proposed methodology is presented in Fig. 2.

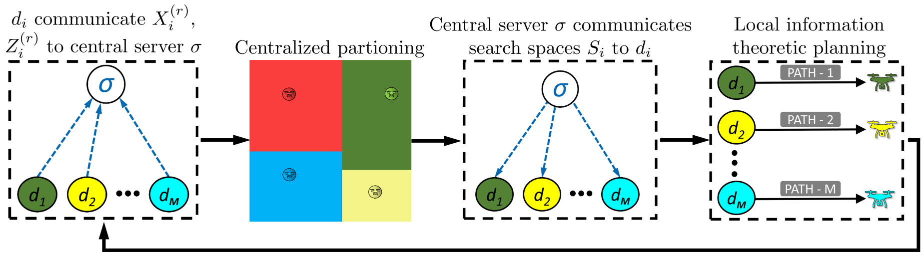

Fig. 2. The block diagram illustrates a decoupled two-step scheme of distributed local planning with centralized partitioning. In the leftmost block diagram base stations d_i communicate their belief maps $X_i^{(r)}$, $Z_i^{(r)}$ to the central server σ. Next, the central server partitions the environment $\mathcal{S}$ into sub-search spaces S_i and then transmits this information to each base station d_i. The sub-search spaces S_i are used locally for information theoretic planning.

Centralized Entropy Decomposition For centralized decomposition we use weighted Voronoi partitioning to decompose the search space $\mathcal{S}$ into M smaller search spaces $\{S_1, S_2, \ldots, S_M\}$, where each robot i is assigned S_i. This partitioning uses base stations d_i as *generators*, assigning each agent i to a specific search space S_i. Let $\boldsymbol{p}_i = [p_{i,x}, p_{i,y}]^\top \in \mathcal{S}$ denote the location of a cell in the search space and $\boldsymbol{p}_{d_j} \in \mathcal{S}$ denote the location of base station d_j, where $j = 2, 3, \ldots, \mathrm{card}(\mathcal{S})$ with $\mathrm{card}(\mathcal{S}) = ab$. The geodesic Voronoi partitioning of the search space is computed by,

$$S_i = \{\ \boldsymbol{p}_i \in \mathcal{S} \mid g(\boldsymbol{p}_i, \boldsymbol{p}_{d_j}) \leq g(\boldsymbol{p}_i, \boldsymbol{p}_{d_k}),\ \forall\ j \neq k\}, \tag{1}$$

where $g(\boldsymbol{p}_i, \boldsymbol{p}_d) = \|\boldsymbol{p}_i - \boldsymbol{p}_d\|$ is the distance between $\boldsymbol{p}_i$ and $\boldsymbol{p}_d$. Next, we redefine $g(\boldsymbol{p}_i, \boldsymbol{p}_d)$ to incorporate the Shannon entropy of the partition,

$$g_w(\boldsymbol{p}_i, \boldsymbol{p}_d; X_{\boldsymbol{p}_i}, X_{S_{\boldsymbol{p}_d}}) = w(X_{\boldsymbol{p}_i}, X_{S_{\boldsymbol{p}_d}}) g(\boldsymbol{p}_i, \boldsymbol{p}_d), \tag{2}$$

where $w(X_{\boldsymbol{p}_i}, X_{S_{\boldsymbol{p}_d}}) = (H(X_{\boldsymbol{p}_i}) + H(X_{S_{\boldsymbol{p}_d}}))/(\mathrm{card}(X_{S_{\boldsymbol{p}_d}}) + 1)$ is a weight representing the average entropy of the expected partition with entropy given as $H(X_{\boldsymbol{p}_i}) = -\int_{\mathcal{X}} \mathbb{P}(X_{\boldsymbol{p}_i}) \log \mathbb{P}(X_{\boldsymbol{p}_i})\, \mathrm{d}x$. Thus, by substituting (2) to (1) the entropy-weighted Voronoi partitioning yields,

$$S_i = \{\ \boldsymbol{p}_i \in \mathcal{S} \mid w(X_{\boldsymbol{p}_i}, X_{S_j}) g(\boldsymbol{p}_i, \boldsymbol{p}_{d_j}) \leq w(X_{\boldsymbol{p}_i}, X_{S_k}) g(\boldsymbol{p}_i, \boldsymbol{p}_{d_k}),\ \forall\ j \neq k\}. \tag{3}$$

Note that for the initial decomposition of the environment with no prior information, the entropy-weighted Voronoi partitioning (3) is identical to the geodesic Voronoi partitioning (1), i.e., $w(X_{\boldsymbol{p}_i}, X_{S_j}) = w(X_{\boldsymbol{p}_i}, X_{S_k})$ for all $j \neq k$.

Local Information-theoretic Planning Centralized entropy-weighted Voronoi partitioning decomposes the problem into single-agent sub-problems of information gathering. Here, each agent i stationed at its base station d_i is now assigned

Algorithm 1 Distributed Entropy Voronoi Partitioning and Planner (DEVPP)

Inputs: $\{X_i^{(0)}\}_{i=1}^M$, $\{Z_i^{(0)}\}_{i=1}^M$, $\mathcal{S}$, M, D, $r_{\max}$
Output: $\{\zeta_i^{(r)*}\}_{i=1}^M$

for $i = 1, \dots, M$ **do** ▷ Initial decomposition and planning
    $S_i^{(0)} \leftarrow$ `geodesicPartition`$(\boldsymbol{p}, \boldsymbol{p}_{d_i})$ (1)
    $\zeta_i^{(0)*} \leftarrow$ `calculatePath`$(X_i^{(0)}, Z_i^{(0)}, S_i^{(0)})$
    $X_i^{(1)}, Z_i^{(1)} \leftarrow$ `beliefUpdate`$(X_i^{(0)}, \boldsymbol{Y}_i^{(0)}, Z_i^{(0)}, \zeta_i^{(0)*})$
end for
for $r = 2$ to $r_{\max}$ **do**
    for $i = 1, \dots, M$ **do** ▷ Communication to server
        communicate $X_i^{(r-1)}$, $Z_i^{(r-1)}$ from base stations to central server
    end for
    $S_i^{(r)} \leftarrow \{\emptyset\}$
    for $\boldsymbol{p} \in \mathcal{S}$ **do** ▷ Centralized Entropy Decomposition
        for $i = 1, \dots, M$ **do**
            $g_w \leftarrow$ `EntropyWeight`$(\boldsymbol{p}, \boldsymbol{p}_{d_i}; X_{\boldsymbol{p}}, X_{S_{\boldsymbol{p}_{d_i}}})$ (2)
        end for
        $\boldsymbol{p}_{d_i} = \arg\min_{\boldsymbol{p}_{d_i} \in D}\{g_w\}$
        $i \leftarrow$ `index`$(\boldsymbol{p}_{d_i})$; $S_i^{(r)} = S_i^{(r)} \cup \boldsymbol{p}$
    end for
    broadcast $S_i^{(r)}$ from central server to base stations ▷ Communication from server
    for $i = 1, \dots, M$ **do** ▷ Local Information Theoretic Planning
        $\zeta_i^{(r)*} \leftarrow$ `calculatePath`$(X_i^{(r-1)}, Z_i^{(r-1)}, S_i^{(r)})$
        $\Theta_{\zeta_i}^{(r)}, \boldsymbol{Y}_i^{(r)} \leftarrow$ `traversePath`$(\zeta_i^{(r)*})$
        $X_i^{(r)}, Z_i^{(r)} \leftarrow$ `beliefUpdate`$(X_i^{(r-1)}, \boldsymbol{Y}_i^{(r)}, Z_i^{(r-1)}, \Theta_{\zeta_i}^{(r)}, \zeta_i^{(r)*})$
    end for
end for

to find a path ζ_i^* in its corresponding search space S_i that maximizes the expected information gained about targets X and hazards Z, given the observed target sensor measurement Y and path-based sensor measurements Θ_{ζ_i} along a path ζ_i at the exploration round r,

$$\zeta_i^{(r)*} = \arg\max_{\zeta_i \in \Omega_i} \left\{ I(X_i^{(r-1)};\ \boldsymbol{Y}_i^{(r)} \mid \Theta_{\zeta_i}^{(r)}, S_i^{(r)}) + I(Z_i^{(r-1)}; \Theta_{\zeta_i}^{(r)} \mid S_i^{(r)}) \right\}, \quad (4)$$

where Ω_i is the space all possible paths of agent i in its search space $S_i^{(r)}$, $\boldsymbol{Y}_i^{(r)} = [Y_{i,1}^{(r)}, \dots, Y_{i,l}^{(r)}]^\intercal \in \mathbb{R}^l$ denotes all l observations collected at search round r by agent i, and Θ_{ζ_i} is the space of observation of agent i about the path-based sensor triggering. This local path planning (4) is addressed using the methodology in [1].

We illustrate the proposed methodology for partitioning the environment with M agents and M base stations in Fig. 2. DEVPP (Algorithm 1) describes the centralized partitioning and local planning methodology, where brown color

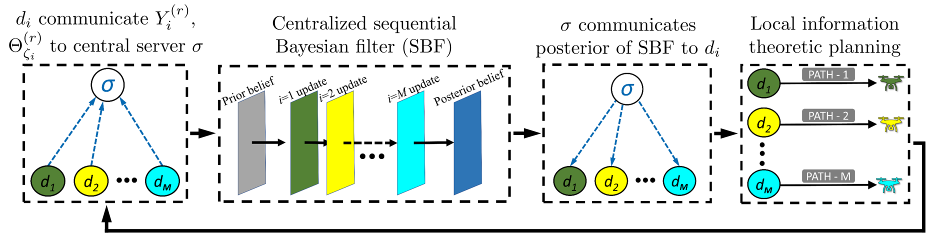

Fig. 3. The block diagram illustrates a distributed local planning with centralized sequential Bayesian filter (SBF). In the leftmost block all base stations d_i communicate their local observations $\boldsymbol{Y}_i^{(r)}$, $\Theta_{\zeta_i}^{(r)}$ to the central server σ. Next, the central server updates the belief map sequentially, where the posterior belief map of agent i is used as prior for the i+1 update. The posterior belief map is broadcasted to the base stations which use them for local information theoretic planning.

indicates action from central server σ, and blue color local action from a base station d_i. After the environment has been partitioned successfully at the central entity, each agent i plans a path using `calculatePath` algorithm introduced in [1]. DEVPP terminates after a number of rounds $r_{\max}$.

3.2 Multi-agent Distributed Information-theoretic Planner (MA-DITP)

Unlike DEVPP, this approach does not perform any partitioning of the search space. Thus, each agent i explores the entire search space $\mathcal{S}$ and plan its path based on the prior belief of hazards and targets in the search space $\mathcal{S}$ at every search round r. In other words, the planning is global in the entire search space $\mathcal{S}$ and not local as in Section 3.1. After a search round, each agent i transmits the target observations $\boldsymbol{Y}_i^{(r)}$ and the path-based sensor measurement $\Theta_{\zeta_i}^{(r)}$ to its base station d_i which subsequently uploads this information to the central server σ. The central server then receives observations from all base stations $\{\Theta_{\zeta_i}^{(r)}\}_{i=1}^{M}$, $\{\boldsymbol{Y}_i^{(r)}\}_{i=1}^{M}$ and updates the belief maps using sequential Bayesian filtering (SBF). Next, the central server broadcasts the posterior belief maps to the base stations. Lastly, each base station d_i computes an information theoretic path $\zeta_i^{(r+1)*}$ in the search space $\mathcal{S}$ and the robots are assigned to traverse the paths and collect measurements. The proposed MA-DITP method is presented in Algorithm 2, where brown color indicates action from the central entity, and blue color local action from a base station d_i. The block diagram is illustrated in Fig. 3.

3.3 Multi-agent Global Information-theoretic Planner (MA-GITP)

In this section, we discuss a global planning method where a team of robots is deployed simultaneously from the same base station d instead of multiple base stations as in DEVPP and MA-DITP. Using MA-DITP in a single station with

Algorithm 2 Multi-Agent Distributed Information Theoretic Planner (MA-DITP)

Input: $\{X_i^{(0)}\}_{i=1}^M$, $\{Z_i^{(0)}\}_{i=1}^M$, $\mathcal{S}$, M, $r_{\max}$
Output:$\{\zeta_i^{(r)*}\}_{i=1}^M$

```
for r = 1 to r_max do
    for i = 1, ..., M do                                   ▷ Local information-theoretic planning
        ζ_i^(r)* ← calculatePath(X_i^(r−1), Z_i^(r−1), S)
        Θ_{ζ_i}^(r), Y_i^(r) ← traversePath(ζ_i^(r)*)
        communicate Θ_{ζ_i}^(r), Y_i^(r) from base station to central server
    end for
    X_i^(r), Z_i^(r) ← beliefUpdate(X_i^(r−1), Y_i^(r), Z_i^(r−1), Θ_{ζ_i}^(r), ζ_i^(r)*)   ▷ Centralized SBF
    broadcast X_i^(r), Z_i^(r) from central server to base stations
end for
```

multiple-agent scenario would result in redundant and repeated exploration of search space $\mathcal{S}$, hampering the information gathering process. To address this problem, we introduce the idea of using an *expected belief map* for each agent. The expected belief map for the first agent is the true belief map computed by the sequential Bayesian filter (SBF) using observations from the previous search round $r-1$. Next, the base station makes a prediction for each possible path-based sensor observation of agent 1 ($\Theta_{\zeta_1}^{(r)} = 1$ for destruction or $\Theta_{\zeta_1}^{(r)} = 0$ for survival) and computes the weighted average belief map which generalizes to,

$$Z_{i+1}^{(r)} = p_1^{(r)} Z_i^{(r)}(\Theta_{\zeta_i}^{(r)} = 1) + p_0^{(r)} Z_i^{(r)}(\Theta_{\zeta_i}^{(r)} = 0), \ \forall i = 2, \ldots, M, \tag{5}$$

where $p_1^{(r)} = \mathbb{P}(\Theta_{\zeta_i}^{(r)} = 1)$ is the probability of destruction, $p_0^{(r)} = \mathbb{P}(\Theta_{\zeta_i}^{(r)} = 0)$ is the probability of survival, $Z_i^{(r)}(\Theta_{\zeta_i}^{(r)} = 1)$ is the belief map in case of destruction, and $Z_i^{(r)}(\Theta_{\zeta_i}^{(r)} = 0)$ is the belief map in case of survival of the previous agent i. Each expected belief map $Z_{i+1}^{(r)}$ is used to compute an information-theoretic path $\zeta_{i+1}^{(r)}$. The multi-agent global information-theoretic planning method (MA-GITP) from a single base station is presented in Algorithm 3. Note that Algorithm 3 includes only local updates, thus the all actions are illustrated in blue color.

The multi-agent global planner is inherently resilient to communication channel attacks, because there is no communication and all the processing is performed at a base station for all agents. In addition, this approach is practical as in most search and rescue missions we typically have access to a single station for exploration.

Algorithm 3 Multi-Agent Global Information Theoretic Planner (MA-GITP)

Inputs: $X^{(0)}$, $\{Z_i^{(0)}\}_{i=1}^{M}$, $\mathcal{S}$, M, $r_{\max}$
Output: $\{\zeta_i^{(r)*}\}_{i=1}^{M}$
for $r = 1$ to $r_{\max}$ **do**
  for $i = 1, \ldots, M$ **do** ▷ Weighted Average Belief Map
    $\zeta_i^{(r)*} \leftarrow$ `calculatePath`$(X^{(r-1)}, Z_i^{(r-1)})$
    $Z_{i+1}^{(r-1)} \leftarrow$ `weightedAvgBelief`$(Z_i^{(r-1)})$ (5)
  end for
  for $i = 1, \ldots, M$ **do** ▷ Sequential Bayesian Filter
    $\Theta_{\zeta_i}^{(r)}, \boldsymbol{Y}_i^{(r)} \leftarrow$ `traversePath`$(\zeta_i^{(r)*})$
    $X^{(r)}$, $Z_i^{(r)} \leftarrow$ `beliefUpdate`$(X^{(r-1)}, \boldsymbol{Y}_i^{(r)}, Z_i^{(r-1)}, \Theta_{\zeta_i}^{(r)}, \zeta_i^{(r)*})$
  end for
end for

4 Experiments and Results

Experimental Setup: To evaluate the performance of the proposed methodologies we implement 30 Monte Carlo (MC) replications for all experiments, where each MC replication has $r = 150$ search rounds. The spatial environment is of dimensions $a \times b = 15 \times 15$ cells. A cell in the environment can be either empty (contain no hazard and no target), or contain a hazard, or a target, or both. The movement of each agent in the environment is determined by a 9-grid connectivity, i.e., an agent can move anywhere to the 8-neighboring cells or decide to stay at its current cell for the next timestep.

Let each agent to have a malfunction probability of 5% per timestep and the target sensor to have a false positive and false negative rate of 10% per timestep. We consider environments with different hazard lethalities of 50%, 70%, and 90%. Lethality refers to the probability of an agent destruction when it visits a cell that contains a hazard. We evaluate the proposed methodologies for various fleet sizes, ranging from $M = 1$ to $M = 5$ agents. In all numerical experiments, we set the maximum number of timesteps for each path to be twice the Manhattan distance between its base station and the furthest corner of its search space S_i, i.e., $l_i = 2 \max_{\boldsymbol{p}_i \in S_i} \|\boldsymbol{p}_{d_i} - \boldsymbol{p}_i\|_1$.

Experiment 1: In this set of experiments, we compare the efficiency of DEVPP against [1] with different fleet sizes. Note that for the case of $M = 1$ agent, DEVPP is identical to [1]. We deploy up to $M = 5$ robots at $\text{card}(D) = 5$ distinct base stations. In Fig. 4, we demonstrate the progress of information entropy for the environment during 150 search rounds. As the number of agents increases, the information entropy reduces faster for all adversary lethalities cases. Contrary to the notion that adding more agents may increase the computational cost of the experiments, Fig. 5 (Left) presents that the time elapsed per agent to complete 150 search rounds decreases with an increasing number of agents. This significant reduction in computational time is attributed to the smaller search sub-spaces for planning with multiple robots. Although the time elapsed for

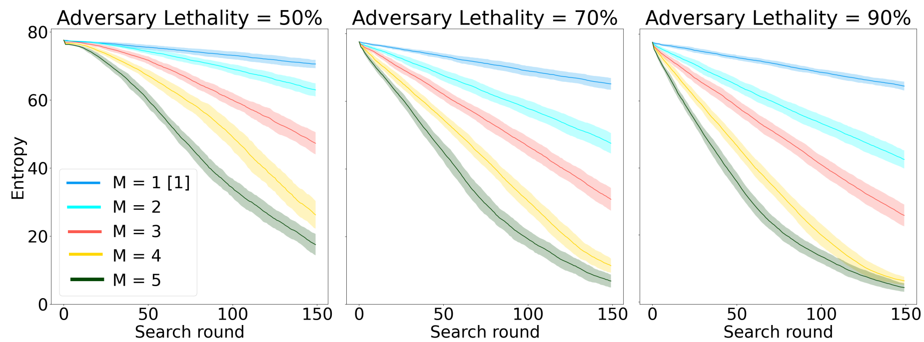

Fig. 4. Information entropy of the environment with hazard lethality of 50%, 70%, 90% using the DEVPP method. As the fleet size increases the information entropy of the environment decreases faster.

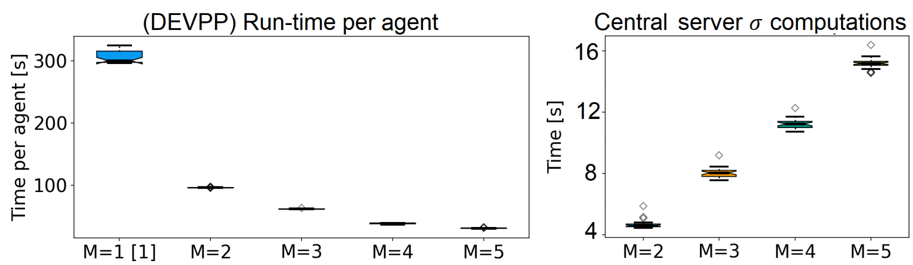

Fig. 5. (Left) Computational time required for each agent to perform local computations for $r = 150$ search rounds using DEVPP for $M \geq 2$ and 5 for $M = 1$. As the fleet size increases the computation is distributed among agents. (Right) Computational time required by the central server σ to partition the environment using (3). As the size of the fleet increases, the time required to compute the partitioned search spaces $S_i \in \mathcal{S}$ increases.

the entire planning and execution decreases with increasing number of robots, the time taken by the central server σ to compute the partitioning increases with addition of robots in the search space $\mathcal{S}$, as shown in Fig. 5 (Right). The computations of the central server σ are insignificant compared to the local computations of the agents for small fleet size ($M = 2, 3$), but increases at similar level for local computations with higher fleet size ($M = 4, 5$). In any case, even if we account the central computations for all fleet sizes, the proposed DEVPP method outperforms [1].

Experiment 2: Next we compare all methods with each other and with [1]. The lethality rate is fixed to 90% and we consider three fleet sizes $M = 2$, $M = 4$, and $M = 5$ robots. Note that when $M = 1$ agent all methods are identical to [1]. In Fig. 6, we present the progress of information entropy through $r = 150$ search rounds. The results show that DEVPP outperforms all methods; MA-DITP gathers competitive information to DEVPP; while MA-GITP performs

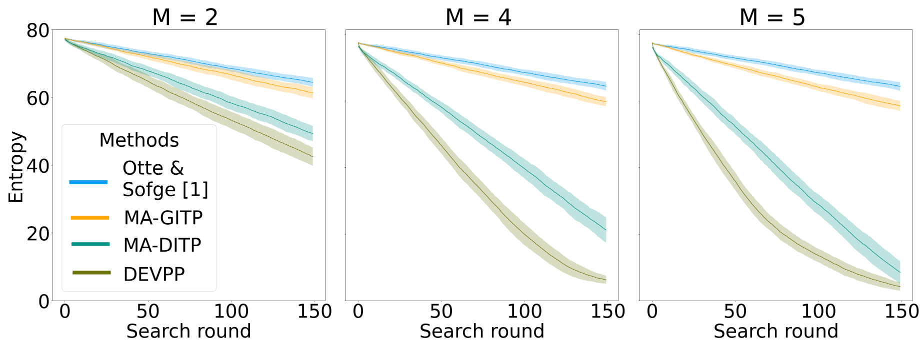

Fig. 6. Information entropy of the environment with hazard lethality of 90% for all proposed methods. We consider three fleet sizes $M = 2$, $M = 4$, and $M = 5$ agents, while all cases are identical to [1] when $M = 1$ agent.

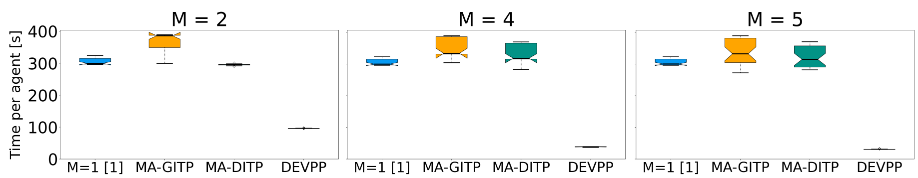

Fig. 7. The boxplots illustrate the execution time per agent in seconds with hazard lethality of 90% for fleet sizes $M = 2$, $M = 4$, and $M = 5$ agents. When of $M = 1$ agent all methods report identical execution time to [1].

slightly better to [1] for all fleet sizes. In Fig. 7, we demonstrate the execution time required per agent after $r = 150$ search rounds. The execution time of DEVPP is significantly lower from all other methods and with minimal variation. To this end, not only information gathered using DEVPP is more efficient, but also it is executed significantly faster. The other proposed methods (MA-DITP and MA-GITP) are executed in similar or slightly more time than [1].

Main result: The numerical experiments reveal that the proposed DEVPP method with multiple robots is more efficient in terms of information gathered compared to the case of a single robot [1]. This result is attributed to the shorter paths produced for each agent on the assigned subspaces, which subsequently reduces the likelihood of robot destruction and significantly reduces execution time. To this end, simultaneous deployment of multiple robots leads to efficient information gathering in all proposed methodologies.

5 Conclusion

This paper introduces DEVPP, a distributed information gathering approach in an environment with path-based sensors using a team of robots simultaneously.

We demonstrate the efficacy of this approach in various experiments with Monte Carlo replications. We also compare this approach against two information-gathering methods that employ multiple robots simultaneously (MA-DITP and MA-GITP) and [1] in the same environment. We experimentally show that the information gathered for hazards and targets using DEVPP outperforms the rest approaches in all cases. In addition, DEVPP reduces the computational cost, thereby accelerates the overall execution time of the information-theoretic planner. Although MA-GITP performs slower than the rest methods, no communication is required, and thus it is resilient to malicious attacks in the communication system. All proposed multi-robot methodologies, reduce the entropy of the environment faster than the case where only a single robot is deployed.

Acknowledgments

This study was conducted at the Motion and Teaming Laboratory, University of Maryland (UMD), and was funded by the Maryland Robotics Center (MRC) and the Office of Naval Research (ONR). Alkesh K. Srivastava and George P. Kontoudis were partially supported by the MRC to conduct this research with "Pathway to the Ph.D." and "MRC Postdoctoral Fellowship" programs, respectively. This work has been possible with the support of ONR under the grant "Experimentally Verified Autonomous Path Planning for Information Gathering in Lethally Hostile Environments with Severely Limited Communication" (N0001420WX01827). The views, positions, and conclusions contained in this document are solely those of the authors and do not explicitly represent those of ONR.

References

1. Otte, M.,Sofge, D.: Path-based sensors: Paths as Sensors, Bayesian updates, and Shannon information gathering. IEEE Transactions on Automation Science and Engineering **18**, 946-967 (2021).
2. Okabe,A., Boots, B., Sugihara, K.:Nearest neighbourhood operations with generalized Voronoi diagrams: A review. International Journal of Geographical Information Systems **8.1**, 43-71 (1994).
3. Choset, H.: Coverage for robotics – A survey of recent results. Annals of Mathematics and Artificial Intelligence **31**, 113–126 (2001).
4. Nair, V. G., Guruprasad, K. R.: GM-VPC: An algorithm for multi-robot coverage of known spaces using generalized Voronoi partition. Robotica **38(5)**, 845-860 (2020).
5. Jager, M. and Nebel, B.: Dynamic decentralized area partitioning for cooperating cleaning robots. IEEE International Conference on Robotics and Automation **4** 3577–3582 (2002).
6. Rankin, E. S., Rekleitis, I., New, A. P. and Choset, H.: Efficient Boustrophedon multi-robot coverage: An algorithmic approach. Annals of Mathematics and Artificial Intelligence **52**, 109–142 (2008).
7. Hazon, N. and Kaminka, G. A.: On redundancy, efficiency, and robustness in coverage for multiple robots. Robotics and Autonomous Systems **56**, 1102–1114 (2008).

8. Agmon, N., Hazon, N. and Kaminka, G. A.: The giving tree: Constructing trees for efficient offline and online multi-robot coverage. Annals of Mathematics and Artificial Intelligence **52(2–4)**, 43–168 (2009).
9. Zheng, X., Koenig, S., Kempe, D. and Jain, S.: Multirobot forest coverage for weighted and unweighted terrain. IEEE Transactions on Robotics **26(6)**, 1018–1031 (2010).
10. Wilson, Z., Whipple, T. and Dasgupta, P.: Multi-robot coverage with dynamic coverage information compression. International Conference on Informatics in Control, Automation and Robotics (2011), 236–241.
11. Michel, D. and McIsaac, K.: New path planning scheme for complete coverage of mapped areas by single and multiple robots. IEEE International Conference on Mechatronics and Automation (2012), 1233–1240.
12. Macharet, D., Azpurua, H., Freitas, G. and Campos, M.: Multi-robot coverage path planning using hexagonal segmentation for geophysical surveys. Robotica **36(8)**, 1144–1166 (2018).
13. Guruprasad, K. R., Wilson, Z. and Dasgupta, P.: Complete coverage of an initially unknown environment by multiple robots using Voronoi partition. International Conference on Advances in Control and Optimization in Dynamical Systems, Bangalore, India (2012).
14. Kurabayashi, D., Ota, J., Arai, T., Yoshida, E.: Cooperative sweeping by multiple mobile robots. IEEE International Conference on Robotics and Automation. **2** 1744-1749 (1996).
15. Bhattacharya, S., Michael, N., Kumar, V.: Distributed coverage and exploration in unknown non-convex environments. Distributed Autonomous Robotic Systems. Springer Tracts in Advanced Robotics **83** 61-75 (2013).
16. Otte, M.,Sofge, D.: Path planning for information gathering with lethal hazards and no communication. International Workshop on the Algorithmic Foundations of Robotics 389-405 (2018).
17. Schwager, M., Dames, P., Rus, D., Kumar, V.: A multi-robot control policy for information gathering in the presence of unknown hazards. Springer Robotics Research, pp. 455-472 (2017).
18. Julian, B. J., Angermann, M., Schwager, M., Rus, D.: Distributed robotic sensor networks: An information-theoretic approach. The International Journal of Robotics Research, **31(10)** 1134–1154 (2012).
19. Dames, P., Schwager, M., Kumar, V., Rus, D.: A decentralized control policy for adaptive information gathering in hazardous environments. IEEE Conference on Decision and Control, 2807–2813 (2012).
20. Dames, P., Schwager, M.,Rus, D., Kumar, V.: Active magnetic anomaly detection using multiple micro aerial vehicles. IEEE Robotics and Automation Letters, **1** 153–160 (2016).