Distributed Knowledge Leader Selection for Multi-Robot Environmental Sampling Under Bandwidth Constraints

Wenhao Luo*, Shehzaman S. Khatib*, Sasanka Nagavalli*, *Student Member, IEEE*,
Nilanjan Chakraborty†, *Member, IEEE*, Katia Sycara*, *Fellow, IEEE*

***Abstract*— In many multi-robot applications such as target search, environmental monitoring and reconnaissance, the multi-robot system operates semi-autonomously, but under the supervision of a remote human who monitors task progress. In these applications, each robot collects a large amount of task-specific data that must be sent to the human periodically to keep the human aware of task progress. It is often the case that the human-robot communication links are extremely bandwidth constrained and/or have significantly higher latency than inter-robot communication links, so it is impossible for all robots to send their task-specific data together. Thus, only a subset of robots, which we call the knowledge leaders, can send their data at a time. In this paper, we study the knowledge leader selection problem, where the goal is to select a subset of robots with a given cardinality that transmits the most informative task-specific data for the human. We prove that the knowledge leader selection is a submodular function maximization problem under explicit conditions and present a novel distributed submodular optimization algorithm that has the same approximation guarantees as the centralized greedy algorithm. The effectiveness of our approach is demonstrated using numerical simulations.**

I. Introduction

For human-supervised multi-robot applications where the human operator monitors the robotic group remotely, bandwidth constraints between the human and the robot team often make it impossible for all the robots to simultaneously communicate all their data to the human [1]. The multi-robot system often obtains a significant amount of redundant data such as video streams and images, which motivates the need to dynamically select a subset of the robots whose gathered information is most informative (not overly redundant) for transmission to the human. Instead of directly fusing the task-specific data between robots, which still requires a great amount of data transmission, it is beneficial to use an abstract model of that data for evaluating each robot's collected information and improving communication efficiency. Moreover, since the communication link to the human may also have significant latency (e.g. communication delay between a robot team on Mars and human on Earth), this further motivates the need to perform the computation in a distributed manner on the multi-robot team.

This work was supported in part by AFOSR award FA9550-15-1-0442 and an NSERC PGS D scholarship.

*The authors are with the Robotics Institute, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA. Email: {wenhaol, skhatib, snagaval}@andrew.cmu.edu, katia@cs.cmu.edu. †The authors are with the Department of Mechanical Engineering, Stony Brook University, Stony Brook, New York, USA. Email: nilanjan.chakraborty@stonybrook.edu.

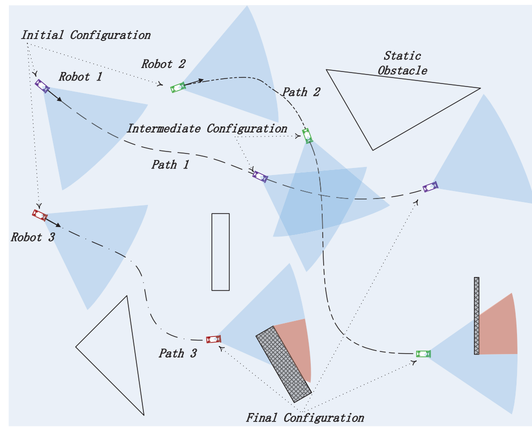

Fig. 1: Three robots explore the environment along their paths with limited field-of-view (FOV) sensing (dark blue sector regions). Dark red regions are examples of robots' blind zones due to occlusion from obstacles.

Our work considers a scenario where a group of robots under the supervision of a human operator is collecting data and have to coordinate among themselves via local communication using an abstract model of the data to periodically find a subset of k robots that have the most informative information, and then let those robots transmit their sensed data to the human so that the total environmental information available to the human is maximized. Note that we are not seeking to compute configurations or a control policy for the robots that will maximize the information collected as is traditionally done in either static or dynamic multi-robot dynamic coverage problems [2], [3], but rather we are selecting the best subset of robots to communicate the collected information at each configuration generated by some external controller [4] or planner [5], [6]. Due to possible latency issues in the human-robot communication channel [7], it is also desirable for the robots to autonomously and adaptively recompute the value of sensed information based on the accumulated data already reported to the human (e.g. redundant data from the same area sent to the human becomes less valuable), such that the robot team requires no human feedback after data transmission.

The leader selection problem under such considerations is NP-hard, so solutions with bounded suboptimality are acceptable. In this paper, we propose a distributed algorithm relying on local inter-robot communication that enables the robotic team to elect effective knowledge leaders with guaranteed bounds on the suboptimality of performance that also account for the human operator's dynamically accumulated knowledge. The contributions of this paper are (1) formal definition of the aforementioned coverage problem with consideration for a time-varying information potential function and visibility limitations due to obstacles, including a proof of submodularity under explicit conditions, and (2) a distributed greedy algorithm to solve the formulated submodular optimization problem with proof of convergence that relies only on the local interaction among robots but ensures the same suboptimality bound as the centralized algorithms.

II. Related Work

Most existing work models environmental monitoring or sampling tasks as either the multi-robot static or dynamic coverage problem. The static coverage problem is related to locational optimization where the main objective is to redeploy robots such that the mission domain is covered optimally and the agents move to a final configuration to accomplish the coverage objective [2], [8], while the goal of the dynamic coverage problem is to cover all the points in the mission domain to some predefined coverage level over time based on the mobility of the robot team [3], [4], [9]. Many versions of these problems have been studied based on different assumptions on the environmental model, sensor models, density function and coverage metric to make the solution more practical in real-world applications. For example, for the robots' sensing capability, an isotropic disc-like sensing model has been widely used since [3], while recently more realistic anisotropic sensors such as on-board cameras have been considered in [9]. In this paper, we employ the modified dynamic coverage metric similar to that of [4] defined with anisotropic sensors and a time-varying density function to establish a different problem: the knowledge leader selection. In this problem, we aim to maximize the accumulated sensory information over time from a subset of robots with given cardinality, which is proved in this paper to be a submodular maximization problem.

Although maximizing submodular functions is NP-hard, [10] has proved that a greedy algorithm could provide a solution with approximation ratio $1-1/e$. The greedy algorithm and the submodularity property have been studied recently in the context of leader selection problems with different submodular metrics for information acquisition [11], sensor placement in networks [12] and leader-follower convergence in multi-agent systems [13]. However, much of this work relies on a centralized greedy algorithm that requires global information from all the nodes (robots), while a decentralized optimization design is needed for solving our problem. Existing distributed submodular function maximization algorithms assume one or more of the following: (a) the data can be partitioned among the multiple computational nodes [14], (b) the communication graph is complete or star-shaped [15], or (c) global knowledge of current solution [16] is available at all the nodes. There is often a gap between the performance of the distributed algorithm and the centralized algorithm as in [14]. In this paper, we propose a distributed submodular optimization approach and prove it to share the same suboptimality bound as the centralized greedy algorithm.

III. Problem Statement

Consider n mobile robots moving in a planar bounded space $\mathcal{A} \subset \mathbb{R}^2$, with the pose of each robot $i \in \{1, 2, \dots, n\}$ at time t denoted by $p_i(t) = [x_i(t), y_i(t), \theta_i(t)]^T$ where $[x_i(t), y_i(t)]^T \in \mathbb{R}^2$ represents the position of each robot and $\theta_i(t) \in [-\pi, \pi)$ represents the orientation. Areas occupied by obstacles are defined by closed set $\mathcal{B} \subset \mathcal{A}$, so the traversable and observable space for robots is $\mathcal{T} = \mathcal{A} \setminus \mathcal{B}$. At regularly spaced time intervals, we select a subset of the robots to transmit information back to the human. For simplicity of exposition, assume the current selection time-point is $t = t_{aft}$, and hence the history of robot poses can be assumed to be recorded as $P(t) = \{p_1(t), \dots, p_n(t)\}$ for $t \in [t_0, t_{aft}]$. Each robot in the scenario shown in Figure 1 can evaluate the value of the collected data by considering the product of sensing strength (sensing model) and importance of the sensed area (information potential) that dynamically changes as the area is explored over time.

A. Sensing Model

For our work, the robots are equipped with limited field-of-view anisotropic sensors (e.g. cameras) that are used to gather task-specific environmental data. In particular, we adopt the limited field-of-view model utilized in [9] that incorporates degradation of effective sensing close to the boundaries of the sensing footprint, which is realistic for most sensors. Assume a homogeneous robotic team, where the sensing footprint of robot i is defined by a circular sector $\mathcal{S}_i$ with uniform radius $r \in \mathbb{R}^+$ and subtended by angle 2α, where $\alpha \in (0, \pi)$. If $p_i(t) = [x_i(t), y_i(t), \theta_i(t)]^T$ is the pose of robot i in world frame and $q \in \mathcal{A} : q = [\bar{x}, \bar{y}]^T$ is an arbitrary to-be-sampled point of interest, let $\psi_i(t, q)$ represent the bearing to point q in body frame of robot i.

$$\psi_i(t, q) = \text{atan2}(\bar{y} - y_i(t), \bar{x} - x_i(t)) - \theta_i(t) \tag{1}$$

We define the following functions for convenience.

$$\begin{aligned} c_{1i}(t, q) &= r^2 - (\bar{x} - x_i(t))^2 - (\bar{y} - y_i(t))^2 \\ c_{2i}(t, q) &= \alpha - \psi_i(t, q), \quad c_{3i}(t, q) = \alpha + \psi_i(t, q) \end{aligned} \tag{2}$$

All functions $c_{ji}(t, q)$ monotonically decrease as the point of interest q approaches the sensor footprint boundaries. We define the sensing performance function as follows.

$$f_i\left(p_i(t), m_i(t, q), q\right) = \frac{m_i(t, q) \prod_{j=1}^{3} \max\left(0, c_{ji}(t, q)\right)^2}{r^4 \alpha^4} \tag{3}$$

This function has range $[0, 1]$ and monotonically increases as robot i approaches point q. It is minimized (evaluates to 0) when q is on the boundary or outside the sensor footprint of robot i. The binary function $m_i(t, q)$ captures whether point q can be sensed by robot i at time t. Specifically, if $[p_i(t), a]$ is the line segment connecting $p_i(t)$ and a, we can define $m_i(t, q)$ as follows.

$$m_i(t, q) = \begin{cases} 1 & \text{if} \quad q \in \{a \in \mathcal{A} \mid [p_i(t), a] \in \mathcal{T}\} \\ 0 & \text{else} \end{cases} \tag{4}$$

This function captures the idea that environmental points that are occluded from a robots' view due to obstacles cannot be sensed. Initially, robots are unaware of obstacle locations, so each robot i cannot know $m_i(t,q)$ a priori.

It is noteworthy that with (1)-(4) any robot's sensing performance over time can be obtained by others merely based on its path through the environment (captured by $p_i(t)$) and the points it senses within the environment along its path (captured by $m_i(t,q)$). As before, $P(t)=\{p_1(t),\dots,p_n(t)\}$ and $M(t,q)=\{m_1(t,q),\dots,m_n(t,q)\}$.

B. Information Potential of Points in the Environment

Each point $q\in\mathcal{A}$ in the environment potentially holds new information. As the robots explore the environment, points that have not been previously sensed are expected have more information potential than points that have already been sensed (i.e. redundant sensing of a point results in less information gained each time it is sensed). To capture this idea, we define $\phi_t(J,q)$ (a modification of the density function in [4]) which measures the information potential at point q at time t for a subset of the robots $J\subseteq\{1,2,...,n\}$. Given $J_{pre}\subseteq\{1,2,...,n\}$ as the previous set of selected knowledge leaders and and t_{pre} as the time at which they were selected, we can compute $\phi_t(J,q)$ recursively. The base case of the recursive computation is $\phi_0(\cdot,q)$ which represents the predefined initial information potential at each point. We assume $\phi_0(\cdot,q)$ is given for all $q\in\mathcal{A}$.

$$\begin{aligned}\phi_t(J,q)&=\phi_{t_{pre}}(J_{pre},q)\,e^{-k^*A_t(J,q)}\\ A_t(J,q)&=\sum_{j\in J}\int_{t_{j,pre}}^{t}f_j(p_j(\tau),m_j(\tau,q),q)\,d\tau\end{aligned}\tag{5}$$

Here, $k^*\in\mathbb{R}^+$ is a design variable and the function $A_t(J,q)$ quantifies how well a certain point q has been geometrically explored by the subset of robots J since the last time they were selected as knowledge leaders. Note that $t_{j,pre}\le t_{pre}$ denotes the last time-point robot j was selected as a knowledge leader and sent collected data to human. As the computation of $\phi_t(J,q)$ only depends on points sensed after time $t_{j,pre}$, each robot j can then simply store data collected starting from $t_{j,pre}$ and discard all data collected prior to $t_{j,pre}$. The specification of k^* will be discussed in Section IV.

C. Objective Functions

Our selection of the subset is subject to a predetermined cardinality constraint (i.e. we can only pick at most k knowledge leaders, so $|J|\le k$). We define the incremental information gain, an abstract model of the collected task-specific data, for a subset of robots J from an arbitrary point $q\in\mathcal{A}$ over time interval $t\in[t_{pre},t_{aft}]$ as follows.

$$Q(J,q)=\sum_{j\in J}\int_{t_{j,pre}}^{t_{aft}}f_j(p_j(\tau),m_j(\tau,q),q)\phi_\tau(J,q)\,d\tau\tag{6}$$

Then the incremental information gain over the entire environment $\mathcal{A}$ for robots J is given by

$$F(J)=\int_{\mathcal{A}}Q(J,q)\,dq\tag{7}$$

and our objective at each selection time point can be written formally as follows.

$$\begin{aligned}J^*&=\arg\max_J F(J)\\ \text{subject to}\quad&|J|\le k\end{aligned}\tag{8}$$

IV. Leader Selection using Distributed Submodular Optimization

Considering that the problem in (8) is an NP-hard combinatorial optimization problem, in this section we show the submodularity of the function $F(J)$ and propose a distributed submodular optimization approach with solutions as good as the standard centralized greedy approach.

A. Submodularity Analysis

Definition 1 (Submodularity [17]): Let V be a finite set. A function $f:2^V\to\mathbb{R}$ is *submodular* if for all sets S and T with $S\subseteq T\subseteq V$, the following is satisfied.

$$\forall v\notin T: f(S\cup\{v\})-f(S)\ge f(T\cup\{v\})-f(T)\tag{9}$$

In [17], it is shown that a nonnegative weighted sum of submodular functions is submodular, which serves as a preliminary lemma used in the following analysis of submodularity of the objective function $F(J)$ in (8). The condition for making $F(J)$ a submodular function is given in the following Theorem.

Theorem 1: At each leader selection time-point $t=t_{aft}$, the function $F(J)$ is a monotone submodular function of robot set J if the following condition holds (k is the maximum number of knowledge leaders).

$$k^*\in\left[0,\frac{1}{k\left(t_{aft}-t_0\right)}\right]\tag{10}$$

Proof: First we consider the submodularity of the integrand of $F(J)$, namely the incremental information gain function $Q(J,q)$ on any point of interest q defined in (6). Let $I\subseteq I'\subseteq\{1,2,\dots,n\}$ and $i\in\{1,2,\dots,n\}\setminus I'$, which implies $i\in\{1,2,\dots,n\}\setminus I$. Then we have

$$\begin{aligned}\Delta=&\left(Q(I\cup\{i\},q)-Q(I,q)\right)-\left(Q(I'\cup\{i\},q)-Q(I',q)\right)\\ =&D_1\cdot D_2+D_3\end{aligned}\tag{11}$$

where

$$\begin{aligned}D_1&=Q(I,q)-Q(I',q)\\ &=\sum_{j\in I}\int_{t_{j,pre}}^{t_{aft}}f_j(p_j,m_j,q)\phi_\tau(I,q)d\tau\\ &\quad-\sum_{i\in I'}\int_{t_{j,pre}}^{t_{aft}}f_j(p_j,m_j,q)\phi_\tau(I',q)d\tau\\ D_2&=e^{-k^*\int_{t_{i,pre}}^{t}f_i(p_i,m_j,q)d\tau}-1\\ D_3&=f_i(p_i,m_i,q)e^{-k^*\int_{t_{i,pre}}^{t}f_i(p_i,m_i,q)d\tau}\left(\phi_t(I,q)-\phi_t(I',q)\right)\end{aligned}\tag{12}$$

Considering the non-negativity of $f_i(\cdot)$ and the non-increasing function $\phi_t(J,q)$, it is straightforward that $D_2\le 0$ and $D_3\ge 0$.

To discuss the sign of D_1, we consider the continuous function $g(x)=xe^{-k^*x}$ whose monotonicity is identical to that of $Q(J,q)$ in Equation (6) (where $x=$

$\sum_{j=1}^{n} \int_{t_{j,pre}}^{t_{aft}} f_j(p_j, m_j, q)dt$ and $e^{-k^* x} \sim \phi_t(J, q)$). By taking the first derivative of $g(x)$ w.r.t. x, we have

$$\frac{dg(x)}{dx} = (1 - k^* x)e^{-k^* x} \tag{13}$$

To that end, we have the following condition for $dg(x)/dx \geq 0$, namely $Q(J, q)$ is non-decreasing, which thus renders $D_1 \leq 0$ since $I \subseteq I'$.

$$0 \leq k^* \leq \frac{1}{x} \tag{14}$$

Recalling the definition of x, we have $\max\{x\} \leq k(t_{aft} - t_0)$ due to the cardinality constraint $|J| \leq k$ and $f_i(\cdot) \in [0, 1]$. Hence, the explicit condition for k^* is as follows.

$$k^* \in \left[0, \frac{1}{k(t_{aft} - t_0)}\right] \tag{15}$$

Then it follows from (11) that $\Delta \geq 0$ under condition (15) and hence $Q(J, q)$ is a non-decreasing submodular function of J. Since the objective function $F(\cdot)$ can be regarded as the sum of $Q(J, q)$ over all $q \in \mathcal{A}$ given $\mathcal{A}$ is discretized, then by the aforementioned lemma from *Definition 1*, $F(\cdot)$ is a monotonially non-decreasing submodular function of set J, which concludes the proof. ■

Remark 1: The restriction on k^* in *Theorem 1* corresponds to the fact that taking information from as many robots as possible is always beneficial for increasing human operators' knowledge over the map despite diminishing returns.

B. Distributed Greedy Algorithm

Consider the connected communication graph of the robot team given as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with each node $v \in \mathcal{V}$ representing a robot in the graph. Assume each robot has the same limited communication range. There is an edge $(v_i, v_j) \in \mathcal{E}$ for any pair of robots $v_i, v_j \in \mathcal{V}$ within communication range of each other. Note that the communication graph is undirected (i.e. $(v_i, v_j) \in \mathcal{E} \Rightarrow (v_j, v_i) \in \mathcal{E}$).

Due to the proven submodularity of our objective function in (8), it is convenient to use the standard greedy algorithm as a subroutine for the computation described above. In order to avoid the pitfalls of the distributed algorithms that work on partitioned data sets, as suffered in [14], here we propose a novel distributed algorithm relying on local inter-robot communications that (1) implicitly constructs a hop-optimal spanning tree [18], and (2) uses the standard greedy algorithm as a subroutine to perform local leader computation as well as repeated backtracking verification to retrieve the solutions that are omitted during the greedy optimization over partitioned data until convergence (similarly to [19]) while exploiting the spanning tree structure. Such an algorithm provides a solution that is as good as the centralized greedy algorithm on the whole dataset.

First, consider the centralized greedy subroutine (see Algorithm 1) as applied to our problem. The input is the robots' index set J, their path set $\{p_j\}_{j \in J}$, the corresponding visibility $\{m_j\}_{j \in J}$ of points in the map for the considered time span, and the maximum number of knowledge leaders k. The output is the selected knowledge leader index set J^*, the corresponding path sets $\{p_j\}_{j \in J^*}$ and the corresponding visibility $\{m_j\}_{j \in J^*}$ of points in the map. By iteratively considering all the robots for evaluating the objective function $F(\cdot)$ in (8) with corresponding sensing model $f(\cdot)$ and information potential $\phi(\cdot)$ derived from $\{p_j\}_{j \in J^*}$ and $\{m_j\}_{j \in J^*}$, the near-optimal knowledge leader set J^* will be constructed after at most k iterations.

Algorithm 1 Greedy Algorithm

procedure GREEDY($J, \{p_j\}_{j \in J}, \{m_j\}_{j \in J}, k$)
  $J^* \leftarrow \varnothing$
  while $(|J^*| < k) \wedge (|J^*| < |J|)$ **do**
    $j \leftarrow \arg\max_{j \in J \setminus J^*} F(J^* \cup \{j\})$
    $J^* \leftarrow J^* \cup \{j\}$
  end while
end procedure

Algorithm 2 Distributed Greedy Algorithm

1: **procedure** DISTRIBUTEDGREEDY(u, $\mathcal{C}_u$, $\mathcal{N}_u$)
2: $\mathcal{C}_{\mathcal{J}^*} \leftarrow \mathcal{C}_u$, $l \leftarrow u$, $h \leftarrow 0$, $m \leftarrow$ NIL
3: **for all** $j \in \mathcal{N}_u$ **do**
4: SENDMSG($j, u, h, l, \mathcal{C}_{\mathcal{J}^*}$)
5: **end for**
6: **while** $\{j', h', l', \mathcal{C}_{\mathcal{J}^{*\prime}}\} \leftarrow$ RECVMSG() **do**
7: **if** $(l > l') \vee ((l = l') \wedge (h > h' + 1))$ **then**
8: $l \leftarrow l'$, $h \leftarrow h' + 1$, $m \leftarrow j'$, $\mathcal{C}_{\mathcal{J}^*} \leftarrow \mathcal{C}_u$
9: **for all** $j \in \mathcal{N}_u$ **do**
10: SENDMSG($j, u, h, l, \mathcal{C}_{\mathcal{J}^*}$)
11: **end for**
12: **else if** $(l = l') \wedge (h < h')$ **then**
13: $\mathcal{C}_{\mathcal{J}^*} \leftarrow$ GREEDY($\mathcal{C}_{\mathcal{J}^*} \cup \mathcal{C}_{\mathcal{J}^{*\prime}} \cup \mathcal{C}_u$)
14: SENDMSG($j', u, h, l, \mathcal{C}_{\mathcal{J}^*}$)
15: **if** $m \neq$ NIL **then**
16: SENDMSG($m, u, h, l, \mathcal{C}_{\mathcal{J}^*}$)
17: **end if**
18: **else if** $(l = l') \wedge (m = j') \wedge (\mathcal{C}_{\mathcal{J}^*} \neq \mathcal{C}_{\mathcal{J}^{*\prime}})$ **then**
19: $\mathcal{C}_{\mathcal{J}^*} \leftarrow$ GREEDY($\mathcal{C}_u \cup \mathcal{C}_{\mathcal{J}^{*\prime}}$)
20: **for all** $j \in \mathcal{N}_u$ **do**
21: SENDMSG($j, u, h, l, \mathcal{C}_{\mathcal{J}^*}$)
22: **end for**
23: **end if**
24: **end while**
25: **end procedure**

By using the standard greedy algorithm as our subroutine, the distributed greedy algorithm for selecting k knowledge leaders is proposed in Algorithm 2, which contains three interleaved stages of data processing for each robot node: implicit spanning tree construction (line 7–11), information propagation (line 12–17) and backtracking (line 18–23). Assume each robot v_i in the communication graph has a unique identifier (UID) and assume UID(v_i)=i for simplicity of exposition. The UIDs of the communication graph neighbors of robot v_i are denoted by $\mathcal{N}_i = \{j \mid v_j \in \mathcal{V} : (v_i, v_j) \in \mathcal{E}\}$. The algorithm takes as inputs the robot's own UID u, its own UID-stamped information $\mathcal{C}_u = \{u, p_u, m_u\}$ and the set of its direct neighbour UIDs $\mathcal{N}_u$ within its communication range. On line 2, it initializes its current leader information set $\mathcal{C}_{\mathcal{J}^*}$, root UID l, number of hops h from root and master UID m, and then sends them to its direct neighbors on lines 3–5. The robots will not start to perform the subroutine greedy algorithm until reaching a consensus on the lowest UID l as the root of the spanning tree and each robot has been assigned the lowest possible number of hops (line 7).

Although the root is identified by every robot, we do not assume it will collect every robot's information and then compute the leader set in a centralized manner, since that method has no bound on message size and is not applicable in a bandwidth constrained environment. Instead, exploiting the structure of the constructed spanning tree in which each robot has a unique master (spanning tree parent), we utilize each robot's dual roles on processing the incoming messages as either the non-child node (line 12) or the child node (line 18) to switch between the information propagation process and the verification process, which repeatedly applies the standard greedy algorithm (Algorithm 1) and sends the updated information to different nodes as necessary to collaboratively and efficiently obtain the final solution in a decentralized manner. It is also noted that the output of each robot following this protocol is always its current estimate of leader set with cardinality limited to k, which ensures that the outgoing message size will never exceed one containing k robots, their path sets and visibility of points in the map.

Theorem 2: Algorithm 2 always converges to the identical leader set obtained from the standard centralized greedy algorithm (Algorithm 1).

Proof: First, let $J^* = \{j_1, \ldots, j_k\}$ be the leader set obtained from Algorithm 1 over all robots $\{1, \ldots, n\}$, with robot j_1 computed in the 1st iteration and robot j_k in the kth iteration. Then consider the distributed leader set computation in Algorithm 2 (see lines 13 and 19). Let J^{j_1} denote the merged candidate set robot j_1 provides as input to the GREEDY() subroutine. Since J^{j_1} always contains j_1, j_1 will also be included in its output leader set ${J^{j_1}}^*$ because $j_1 = \underset{j \in \{1,\ldots,n\}}{\arg\max}\ F(\{j\})$ and $J^{j_1} \subset \{1, \ldots, n\}$. As this output set propagates through the robot network, j_1 will always be preserved in the output set from other robots as well. In particular, after robot j_2 receives the candidate set containing j_1 and merges it into set input J^{j_2}, j_2 will be included into its output set ${J^{j_2}}^*$ and always preserved because $j_2 = \underset{j \in \{1,\ldots,n\} \setminus \{j_1\}}{\arg\max}\ F(\{j_1\} \cup \{j\})$ and $J^{j_2} \subset \{1, \ldots, n\} \setminus \{j_1\}$. When this output set containing j_1 and j_2 reaches robot j_3, then j_3 will be added and preserved in further information propagation, and so forth. In this way, the output leader set $\{j_1, \ldots, j_k\}$ will be recursively obtained and agreed upon by each robot, after which Algorithm 2 will eventually terminate as no more messages are sent (see line 18). ■

Bound on Suboptimality: In [10], it has been proven that the standard greedy algorithm will obtain a solution set that will cause the submodular objective function to evaluate to at least $(1 - 1/e)$ times the value produced by the optimal solution set of the same cardinality, and this bound was proven to be tight in [20]. As a byproduct of *Theorem 2*, the solutions obtained from our proposed distributed greedy algorithm share the same suboptimality bound.

V. RESULTS

In the first simulation shown in Figure 2, we consider a homogeneous team of 5 mobile robots equipped with cameras moving in trajectories generated by external planners [6] in a region with four static obstacles over time span $t \in [0, 40s]$. The task is to select 2 robots every $20s$ as knowledge leaders to maximize the new accumulated sensed information sent to the operator. The leader selection time-points are then specified by $t = 20s$ and $t = 40s$ respectively. At $t = 0$ each robot can access the initial value of the information potential over the entire environment shown in Figure 3a and use it to evaluate their coverage over time. At the first selection time-point $t = 20s$, robots communicate their information, construct the spanning tree rooted at robot 1, and then converge to the selected knowledge leader set (robot 2 and 5 in black circle) in Figure 2b by running Algorithm 2. Since each robot knows the leaders' information after convergence, their knowledge of the information potential will be updated by (5) in which the value of areas covered by selected leaders decreases as shown in Figure 3b. Following the same process, in the next selection round, at $t = 40s$, robot 3 and 4 are selected as new knowledge leaders and each robot's information potential value over the map is updated again as shown in Figure 3c. Each robot's respective information gain is evaluated by the updated information potential at $t = 20s$ and $t = 40s$ are shown in Figure 4a-4j. It is noted that at $t = 40s$ since the information from robot 2 and 5 has already been sent to human at a previous selection time-point, their information gain before $t = 20s$ is reset to zero, as shown in Figure 4g and Figure 4j.

To further compare with other existing work, we conduct 50 simulation trials with each trial consisting of a randomly distributed robotic group containing 40 mobile robots. For each trial, we execute our proposed distributed greedy algorithm, the distributed algorithm GREEDI in [14], the standard greedy algorithm, random selection and an optimum selection algorithm to pick up the knowledge leaders. The comparisons on performance and computation time are shown in Figure 5. It is noted in Figure 5a that the proposed distributed greedy algorithm will converge to the same solution as the standard centralized greedy algorithm, which is not ensured for the GREEDI algorithm and random selection algorithm, especially when the error accumulates as the required number of leaders increases. It should also be noted that although for the 50 trials our proposed distributed greedy algorithm can always reach the optimal solution, as the property of standard greedy algorithm, it can only ensure an approximation of $(1 - 1/e)$ to the optimal performance in general cases. The computation time comparison is given in Figure 5b, and it is clear that the computational cost for optimal selection algorithm grows exponentially as number of leaders increases, which makes it impractical in large scale multi-robot application. For the GREEDI algorithm, since it always performs two-stage standard greedy algorithms on each subset of the dataset and then the union of the solution set, in dealing with small scale problems, it may not be more efficient than the standard greedy algorithm. However, for our distributed greedy algorithm, since the cardinality of total inputs to the subroutine greedy algorithm at each iteration will never exceed $2k$, where k is the required number of knowledge leaders and independent of the robotic group

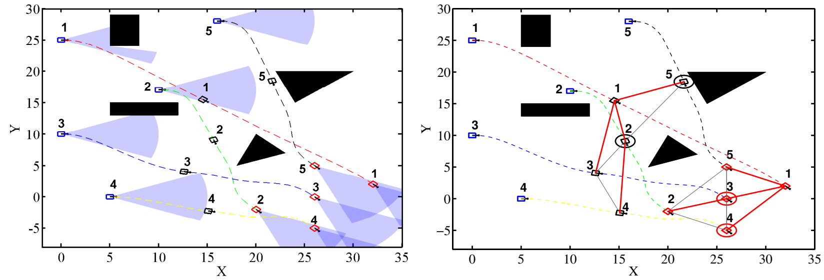

(a) (b)

Fig. 2: Simulation of 5 mobile robots exploring the environment with 4 static obstacles (black) over time $[0, 40s]$. The time-points for leader selection are 20s and 40s. (a) Robot's trajectories (dashed line), visible range examples (blue sectors) and snapshots of their positions and orientations at starting time $t = 0s$ (blue) and the two leader selection time-points $t = 20s$ (black) and $t = 40s$ (red). (b) Inter-robot communication graph (grey) and constructed spanning tree (red) at leader selection time-points $t = 20s$ and $t = 40s$. Selected leaders are marked by circles.

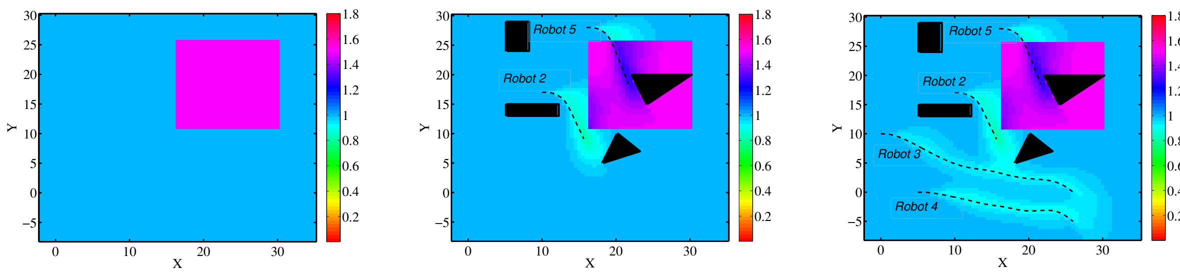

(a) (b) (c)

Fig. 3: Heat map of time-varying information potential. (a) Initial information potential at $t = 0s$. Viewing the magenta area gives robots more information than the surrounding light blue area. (b) Information potential updates at $t = 20s$ after robot 2 and 5 are selected as leaders. (c) Information potential updates at $t = 40s$ after robot 3 and 4 are selected as new leaders. Black polygons mark the positions of static obstacles in the heat maps.

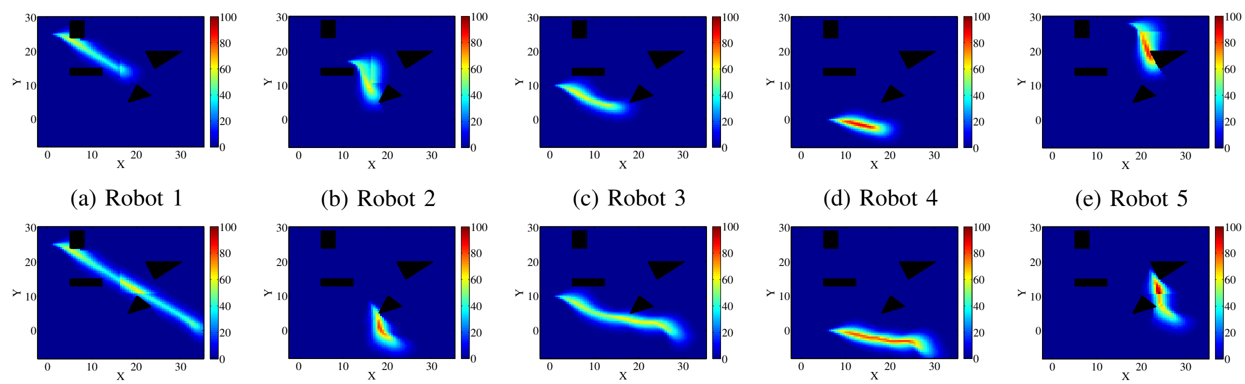

(a) Robot 1 (b) Robot 2 (c) Robot 3 (d) Robot 4 (e) Robot 5

(f) Robot 1 (g) Robot 2 (h) Robot 3 (i) Robot 4 (j) Robot 5

Fig. 4: Snapshot of each robot's accumulated information at sampling time-points $t = 20s$ and $t = 40s$. (a)-(e) Information gain of robot 1-5 at $t = 20s$ computed with updated information potential from robot 2 and 5. (f)-(j) Information gain of robot 1-5 at $t = 40s$ computed with updated information potential from robot 3 and 4.

scale, the computation cost is hence ensured to significantly decrease compared to the standard greedy algorithm. The communication-related results are reported in Fig. 6, in which the centralized algorithms represent any algorithms that requires global information from all the nodes to be sent to the root robot of the tree. It is noted that although the distributed greedy algorithm consumes a larger number of messages, the total amount of data moving through the networks is much less due to the bounded message size.

VI. Conclusion

We formulated the knowledge leader selection as a submodular maximization problem and proposed a distributed greedy algorithm with bounded suboptimality to find knowledge leaders. Such an approach is guaranteed to provide the same approximate solution as the centralized greedy algorithm. We account for diminishing sensory information gain due to overlapped sensing areas, constrained anisotropic sensing performance, and sensing limitations due to occlusions from obstacles. Numerical simulations were performed

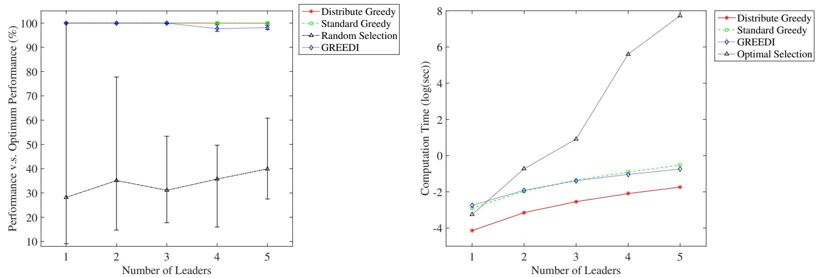

(a) (b)

Fig. 5: Performance comparison of the proposed distributed greedy algorithm and other submodular optimization algorithms from 50 independent trials on randomly generated groups of 40 robots with their paths. (a) The maximum, minimum and average ratio of performance for global objective functions in (8) of proposed distributed greedy, standard greedy, random leader selection and GREEDI, which is another distributed greedy algorithm proposed in [14] vs. the benchmark performance of centralized combinatorial optimization algorithm (NP-hard). (b) The average computation time (log(sec)) among the four algorithms. For the GREEDI algorithm, the cardinality of each subset was chosen to be 8.

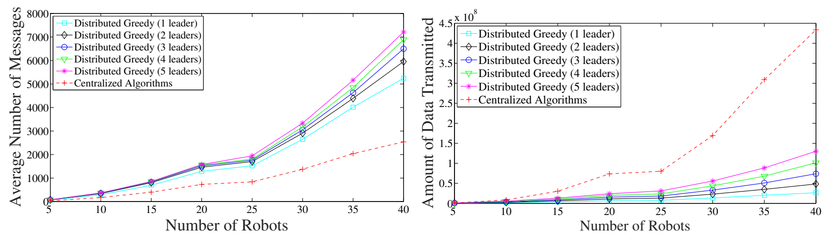

(a) (b)

Fig. 6: Communication-related simulation results from 50 independent trials on different number of robots/leaders using the proposed distributed greedy algorithm and centralized algorithms. (a) Average number of messages transmitted. (b) Average amount of data transmitted, as computed by the multiplication of number of messages and average message size.

to compare our proposed distributed algorithm to other methods on computational time and performance w.r.t. an optimal selection strategy and communication-related metrics. Results validated the effectiveness of the proposed algorithm.

References

[1] A. Kolling, P. Walker, N. Chakraborty, K. Sycara, and M. Lewis, "Human interaction with robot swarms: A survey," *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 1, pp. 9–26, 2016.

[2] J. Cortés, S. Martínez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, p. 243, 2004.

[3] I. Hussein and D. Stipanovic, "Effective coverage control for mobile sensor networks with guaranteed collision avoidance," *IEEE Transactions on Control Systems Technology*, vol. 15, no. 4, pp. 642–657, July 2007.

[4] G. M. Atinc, D. M. Stipanovic, P. G. Voulgaris, and M. Karkoub, "Swarm-based dynamic coverage control," in *IEEE 53rd Annual Conf. on Decision and Control*. IEEE, 2014, pp. 6963–6968.

[5] J. Butzke and M. Likhachev, "Planning for multi-robot exploration with multiple objective utility functions," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 3254–3259.

[6] W. Luo, N. Chakraborty, and K. Sycara, "Distributed dynamic priority assignment and motion planning for multiple mobile robots with kinodynamic constraints," in *American Control Conference*. IEEE, 2016, pp. 148–154.

[7] I. F. Akyildiz, D. Pompili, and T. Melodia, "Underwater acoustic sensor networks: research challenges," *Ad hoc networks*, vol. 3, no. 3, pp. 257–279, 2005.

[8] M. Schwager, D. Rus, and J.-J. Slotine, "Decentralized, adaptive coverage control for networked robots," *The International Journal of Robotics Research*, vol. 28, no. 3, pp. 357–375, 2009.

[9] D. Panagou, D. M. Stipanovic, and P. G. Voulgaris, "Vision-based dynamic coverage control for nonholonomic agents," in *IEEE 53rd Annual Conf. on Decision and Control*. IEEE, 2014, pp. 2198–2203.

[10] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, "An analysis of approximations for maximizing submodular set functions-1," *Mathematical Programming*, vol. 14, no. 1, pp. 265–294, 1978.

[11] N. Atanasov, J. Le Ny, K. Daniilidis, and G. J. Pappas, "Information acquisition with sensing robots: Algorithms and error bounds," in *IEEE International Conference on Robotics and Automation*. IEEE, 2014, pp. 6447–6454.

[12] T. H. Summers, F. L. Cortesi, and J. Lygeros, "On submodularity and controllability in complex dynamical networks," *IEEE Transactions on Control of Network Systems*, vol. 3, no. 1, pp. 91–101, 2016.

[13] A. Clark, B. Alomair, L. Bushnell, and R. Poovendran, "Minimizing convergence error in multi-agent systems via leader selection: A supermodular optimization approach," *IEEE Transactions on Automatic Control*, vol. 59, no. 6, pp. 1480–1494, 2014.

[14] B. Mirzasoleiman, A. Karbasi, R. Sarkar, and A. Krause, "Distributed submodular maximization: Identifying representative elements in massive data," in *Advances in Neural Information Processing Systems*, 2013, pp. 2049–2057.

[15] D. Golovin, M. Faulkner, and A. Krause, "Online distributed sensor selection," in *Proceedings of the 9th ACM/IEEE Int'l Conf. on Information Processing in Sensor Networks*. ACM, 2010, pp. 220–231.

[16] A. Clark, B. Alomair, L. Bushnell, and R. Poovendran, "Distributed online submodular maximization in resource-constrained networks," in *12th IEEE Int'l Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks*. IEEE, 2014, pp. 397–404.

[17] S. Fujishige, *Submodular functions and optimization*. Elsevier, 2005, vol. 58.

[18] S. Nagavalli, A. Lybarger, L. Luo, N. Chakraborty, and K. Sycara, "Aligning coordinate frames in multi-robot systems with relative sensing information," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2014, pp. 388–395.

[19] W. Luo, S. S. Khatib, S. Nagavalli, N. Chakraborty, and K. Sycara, "Asynchronous distributed information leader selection in robotic swarms," in *IEEE International Conference on Automation Science and Engineering*. IEEE, 2015, pp. 606–611.

[20] U. Feige, "A threshold of ln n for approximating set cover," *Journal of the ACM*, vol. 45, no. 4, pp. 634–652, 1998.