Distributed Knowledge Leader Selection for Multi-Robot Environmental Sampling

Wenhao Luo

CMU-RI-TR-16-05

May 2016

Submitted in partial fulfillment of the requirements for the degree of Master of Science in Robotics

Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

Thesis Committee:
Katia Sycara, Chair
Maxim Likhachev
Sasanka Nagavalli

Copyright © 2016 Wenhao Luo

For my parents

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 problem is a submodular optimization problem under some explicit conditions. We also present a novel distributed submodular optimization algorithm that has the same approximation guarantees as the centralized greedy algorithm for submodular function maximization. The effectiveness of our approach is demonstrated using numerical simulations.

Acknowledgements

First and foremost, I would like to express my sincere gratitude to my advisors Prof. Katia Sycara and Prof. Nilanjan Chakraborty (Stony Brook University), who has greatly helped in shaping my interests in the exciting field of multi-robot systems and always kept me inspired from interesting research problems. I am also extremely grateful to my thesis committee member Prof. Maxim Likhachev, who has provided immense help and feedback all along my Masters research. Their vision, guidance and knowledge has been invaluable in my work and this thesis. It has been a great honor to work with you all.

I would also like to thank Sasanka Nagavalli for his continuous support in overseeing my research. His dedication and experience has greatly helped me walking through numerous challenges and stressful times. Many of our interesting works grew out from countless discussions between us.

Further, I would like to thank all my fellow labmates, Yulong Pei, Shehzaman Salim Khatib, Anqi Li and Navyata Sanghvi, for many inspiring discussions and encouragements in my study and research, as well as balancing my life with loads of fun. It has been a great pleasure to be part of this amzing group.

Last but not least, I would express a deep sense of gratitude to my beloved parents and grandparents for their endless support and love throughout my life.

Contents

1 Introduction

Networked distributed multi-robot systems employ local communication and collaborative decision making to carry out a wide variety of large-scale applications such as search and rescue [23], exploration and dynamic coverage of unknown environments [12][4] and distributed environmental monitoring (e.g. water monitoring, flood response) [22]. However, for human supervised tasks 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 [13]. The multi-robot system often obtains a significant amount of redundant data such as a 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. Since the communication link to the human may also have significant latency [19] , this further motivates the need to perform the computation in a distributed manner on the multi-robot team.

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 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 of the robots that will maximize the information collected by the robots (as is traditionally done in either static or dynamic multi-robot dynamic coverage problems [7][12]), but rather we are selecting the best subset of robots to communicate the collected information at each configuration generated by some external controller [3] or planner [4], such that the overall information available to the human is maximized. Due to possible latency issues on human-robot communication channel [1], it is also desirous for the robotic group to be able to autonomously and adaptively recalibrate the value of sensed information based on accumulative data already reported to human (e.g. redundant data from the same area sent to human before becomes less valuable), such that no human feedback is needed for robot teams after data transmission.

In evaluating the quality of sensed data, we should account for (1) diminishing incremental gain due to overlapped sensing areas, (2) anisotropic sensor (e.g. camera) performance variation over the limited sensed region, and (3) sensing limitations due to occlusions from obstacles. The leader selection problem under such considerations is NP-hard that requires suboptimal solutions. In this paper we propose a fully distributed algorithm relying on local inter-robot communications that enables the robotic team to elect effective knowledge leaders with guaranteed suboptimal performance accounting for human operators' dynamically accumulative knowledge. Note that different from canonical understanding of robotic leaders such as formation leaders who control the entire swarm formation, knowledge leader here only refers to the robots with most knowledge who send their information back to human operator. The **contributions** of this paper are (1) formal definition of the aforementioned coverage problem with consideration for a diminishing information potential function (density 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

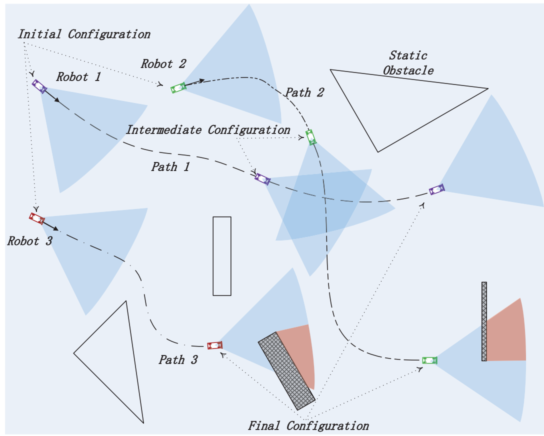

Figure 1: Three robots explore the environment along their respective paths with limited field-of-view (FOV) sensing (dark blue regions show FOV). In their final poses, robots' blind zones due to occlusion from obstacles are marked by dark red.

problem with proof of convergence that relies only on the local interaction among robots but is still able to achieve the same performance as the centralized algorithms.

2 Related Work

The multi-robot environmental monitoring or sampling tasks are often modeled as the multi-robot coverage problem that has received extensive attentions. Most of the existing related works addressed either static coverage problem or the dynamic coverage problem. The static coverage problem is related to the locational optimization where the main objective is to redeploy robots such that the given mission domain is covered optimally and the agents therein end up at a final configuration to accomplish the coverage objective [20], 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 swarm [12]. The extensive versions of the related problem have been studied based on different assumptions of the considered environment model, sensor models, density function and coverage metric, etc, to make the solution more practical in real-world application. For example, for the robots' sensing capability, isotropic disc-like sensing model has been widely used since [12], while recently the consideration of more realistic anisotropic sensors such as on-board cameras has been taken into both static coverage problem [11] and dynamic coverage problem [18]. In this paper, however, we employ the modified incremental dynamic coverage metric similar to that of [3] defined with anisotropic sensors as well as diminishing information potential function to establish a different problem, the knowledge leader selection that aims to maximizing the accumulated sensory information over time from a subset of robots, which is proved in this paper to be a submodular maximization problem.

Although maximizing submodular functions is NP-hard, [17] 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 functions such as information acquisition [2], sensor placement in networks [21], and leader-follower convergence in multi-agent systems [6]. However, since the standard greedy algorithm is centralized and requires global information of all the nodes (robots), decentralized optimization design is needed for solving our problem. Existing distributed submodular function maximization algorithms assume one of the following: (a) the data can be partitioned among the multiple computational nodes [15], (b) the communication graph is complete or star-shaped [10], and (c) global knowledge of current solution [5] is available at all the nodes. Other approaches such as [24] introduces local search in optimizing shared sensor networks merely based on local node-to-node communication. However, approximation bounds here are not as good as the centralized greedy algorithm, and there is a gap between the performance of the distributed algorithm and the centralized algorithm. In this paper we propose a distributed submodular optimization approach to address this problem and prove it to share the same optimality bound as the centralized greedy algorithm.

3 Problem Statement

Consider n mobile robots moving in a planar bounded space $\mathcal{A} \subset \mathbb{R}^2$, with the pose of each robot $i \in \{1, 2, \ldots, n\}$ 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 timepoint is $t = t_{aft}$, and hence the history of robot poses can be assumed to be recorded as $P(t) = \{p_1(t), \ldots, p_n(t)\}$ 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.

3.1 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 [18] 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.

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

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 with in the environment along its path (captured by $m_i(t,q)$). As before, $P(t) = \{p_1(t), \ldots, p_n(t)\}$ and $M(t,q) = \{m_1(t,q), \ldots, m_n(t,q)\}$.

3.2 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 [3]) 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 initial information potential at each point. We assume $\phi_0(\cdot, q)$ is given for all $q \in \mathcal{A}$.

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

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 $t_{j,pre}$ denotes the last time robot j was selected as a knowledge leader. Since the computation of $\phi_t\,(J, q)$ only depends on points sensed after time $t_{j,pre}$, each robot j only needs to store data collected since the last time it was selected as a knowledge leader and can discard all data collected prior to $t_{j,pre}$. The specification of k^* will be discussed in Section IV.

3.3 Objective Functions

As the robots move through the environment, our objective at each selection time point is to select a subset J of the robots that transmit their task-specific sensed data to the human to maximize the incremental gain in information from the perspective of the human. Our selection of the subset is subject to a predetermined cardinality constraint (i.e. we can only pick at most k knowledge leaders, so $|J| \leq k$). Then we have the incremental information gain for a subset of robots J from an arbitrary point $q \in \mathcal{A}$ over time interval $t \in [t_{pre}, t_{aft}]$ defined as follows.

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

The integrand in Equation (6) gives the value of information obtained by a single robot j with pose $p_j(t)$ sensing point q at time $t = \tau$. This value depends on the sensing performance denoted by the function $f_i(\cdot)$ from Equation (3) and the information potential $\phi_t(\cdot)$ in Equation (5). Thus Equation (6) represents the total incremental information accumulated from point q by the robots J until time t_{aft} without "double-counting" the information gained prior to the last time those robots were selected as knowledge leaders. 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 may be written formally as follows.

$$\begin{aligned} \arg\max_{J} \quad & F(J) \\ \text{subject to} \quad & |J| \leq k \end{aligned} \tag{8}$$

4 Leader Selection using Distributed Submodular Optimization

Considering that the problem in (8) is an NP-hard combinatorial optimization problem whose exact optimal solution cannot be found in polynomial time, in this section we show the submodularity of the function $F(J)$ problem and propose a distributed submodular optimization approach to solve this problem with solutions as good as the standard centralized submodular approach.

4.1 Submodularity Analysis

4.1.1 Definition of Submodularity

Definition 1 (Submodularity [9]): Let V be a finite set. A function $f : 2^V \to \mathbb{R}$ is *submodular* if for every sets S and T with $S \subseteq T \subseteq V$ and every $v \notin T$,

$$f(S \cup \{v\}) - f(S) \geq f(T \cup \{v\}) - f(T) \tag{9}$$

4.1.2 Submodularity of Incremental Information Gain

In [9], 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 $F(J)$ (incremental information gain) in (8). The condition for making $F(J)$ a submodular function is given in the following Theorem.

Theorem 1: At each selection time $t = t_{aft}$, 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).

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

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

$$\begin{aligned}
&Q(I \cup \{i\}, q) - Q(I, q) \\
=& \sum_{j \in I \cup \{i\}} \int_{t_{j,pre}}^{t_{aft}} f_j(p_j, m_j, q)\phi_\tau(I \cup \{i\}, q)d\tau \\
&- \sum_{j \in I} \int_{t_{j,pre}}^{t_{aft}} f_j(p_j, m_j, q)\phi_\tau(I, q)d\tau \\
&Q(I' \cup \{i\}, q) - Q(I', q) \\
=& \sum_{j \in I' \cup \{i\}} \int_{t_{j,pre}}^{t_{aft}} f_j(p_j, m_j, q)\phi_\tau(I' \cup \{i\}, q)d\tau \\
&- \sum_{j \in I'} \int_{t_{j,pre}}^{t_{aft}} f_j(p_j, m_j, q)\phi_\tau(I', q)d\tau
\end{aligned} \tag{11}$$

It follows from (5) that the relation between $\phi_t(I \cup \{i\}, q)$ and $\phi_t(I, q)$ is

$$\phi_t(I \cup \{i\}, q) = \phi_t(I, q)e^{-k^* \int_{t_{i,pre}}^{t} f_i(p_i, m_i, q)d\tau} \tag{12}$$

The same relations hold for $\phi_t(I' \cup \{p\}, q)$ and $\phi_t(I', q)$. Hence we have

$$\begin{aligned}
\Delta =& (Q(I \cup \{i\}, q) - Q(I, q)) \\
&- (Q(I' \cup \{i\}, q) - Q(I', q)) \\
=& D_1 \cdot D_2 + D_3
\end{aligned} \tag{13}$$

where

$$
\begin{aligned}
D_1 &= Q(I,q) - Q(I',q) \\
&= \sum_{j \in I} \int_{t_{j,pre}}^{t_{aft}} f_j(p_j, m_j, q) \phi_\tau(I,q) d\tau \\
&\quad - \sum_{i \in I'} \int_{t_{j,pre}}^{t_{aft}} f_j(p_j, m_j, q) \phi_\tau(I',q) d\tau \\
D_2 &= e^{-k^* \int_{t_{i,pre}}^{t} f_i(p_i, m_j, q) d\tau} - 1 \\
D_3 &= f_i(p_i, m_i, q) e^{-k^* \int_{t_{i,pre}}^{t} f_i(p_i, m_i, q) d\tau} (\phi_t(I,q) - \phi_t(I',q))
\end{aligned}
\tag{14}
$$

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

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

$$
\frac{dg(x)}{dx} = (1 - k^* x) e^{-k^* x} \tag{15}
$$

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 $I \subseteq I'$.

$$
0 \leq k^* \leq \frac{1}{x} \tag{16}
$$

Recalling the definition of x, we have $\max\{x\} = k\mathcal{M}(t_{aft} - t_0)$ due to the cardinality constraint $|J| \leq k$ and $f_i(\cdot) \in [0, m_i]$, where $\mathcal{M} = \max\{\int_{\mathcal{A}} m_i dq\}$ for $i = 1, \ldots, n$. Hence, the analytic upper bound for k^* can be found as

$$
k^* \in \left[0, \frac{1}{k\mathcal{M}(t_{aft} - t_0)}\right] \tag{17}
$$

Then it follows from (13) that $\Delta \geq 0$ under the condition (17) and hence $Q(J,q)$ is a non-decreasing submodular function of J. Since the gain function $F(\cdot)$ can be regarded as the sum of $Q(J,q)$ over all $q \in \mathcal{A}$ if $\mathcal{A}$ 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* also ensures the other intuitive constraint that the decay rate of the information potential function $\phi_t(\cdot)$, determined by k^*, from adding new robot's information is constrained so that doing so will never decrease the existing information gain $F(\cdot)$ from the current robot set. This corresponds to the fact of taking information from as many robots as possible is always beneficial for increasing human operators' knowledge over the map despite of diminishing returns.

4.2 Distributed Knowledge Leader Selection

Consider the 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, then for any pair of robot nodes $v_i, v_j \in \mathcal{V}$, the edge $(v_i, v_j) \in \mathcal{E}$ if they are located in each other's commuincation range, leading to undirected communication graph (i.e. $(v_i, v_j) \in \mathcal{E} \Rightarrow (v_j, v_i) \in \mathcal{E}$). We also assume the system's communicating network is connected as one cluster without isolated robot. So the whole data set is naturally partitioned and in a distributed scheme, and the robots have to perform computations over partitioned data sets (since no robot will know the data collected by all other robots).

Due to the proved submodularity of our objective function in (8), it is convenient to use standard greedy algorithm as 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 [15], here we propose a novel distributed algorithm relying on local inter-robot communications that 1) implicitly constructs a hop-optimal spanning tree [16], and 2) uses standard greedy algorithm as 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 as done similarly in [14], with the exploitation of the spanning tree structure. Such algorithm provides a solution that is as good as the centralized greedy algorithm on the whole dataset.

4.2.1 Standard Greedy Algorithm

First we recall the subroutine centralized greedy (see Algorithm 1) as applied to our problem. The input is the candidate robots' index set J with cardinality of n, their path set $\{p_j\}_{j \in J}$, the corresponding value of point sets $\{m_j\}_{j \in J}$ over 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 value of point sets $\{m_j\}_{j \in J'}$. By iteratively considering all the robots for evaluating the objective function $F(\cdot)$ in (8) with corresponding sensing model $f(\cdot)$ and information potential function $\phi(\cdot)$ 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($J, \{p_j\}_{j \in J}, \{m_j\}_{j \in J}, k$)
  $J' \leftarrow \text{NIL}, i \leftarrow 0$
  while $(i < k) \wedge (i < n)$ **do**
    $j \leftarrow j \in (J \setminus J') : \arg\max_{j \in J \setminus J'} F(J' \cup \{j\})$
    $J' \leftarrow J' \cup \{j\}$
    $i \leftarrow i + 1$
  end while
end procedure

Algorithm 2 Distributed Greedy Algorithm

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

4.2.2 Distributed Greedy Algorithm

By using the standard greedy algorithm as our subroutine, the distributed greedy algorithm of selecting k knowledge leaders is proposed in the Algorithm 2, which contains three asynchronous stages of data processing for each robot node such as spanning tree construction (line 7-11), information propagation (line 12-17) and backtracking process (line 18-22). Assume each robot v_i in communication graph has unique identifiers (UIDs) with UID(v_i)= i. The UIDs of 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 2 takes as inputs the robot's own UID u, its own UID-stamped information $\mathcal{C}_u = \{u, p_u, m_u\}$ and set of its direct neighbour UIDs $\mathcal{N}_u$ within its communication range. On line 2, it initializes its current leader information set $\mathcal{C}_\mathcal{S}$, leader UID l, number of hops h from leader and master UID m, and then send them to its direct neighbors on line 3-5. The robots won't start to perform the subroutine greedy algorithm until reaching the consensus of lowest UID l as the root/leader of the spanning tree and has been assigned the lowest possible number of hops (line 7). Although a root node is identified by every robot, we do not assume it to collect every robot's information and then compute the leader set in a centralized manner, which has no bound on message size and not applicable in bandwidth constrained environment. Instead, with the structure of the constructed spanning tree in which each robot has a unique master, 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 conduct the standard greedy algorithm (Algorithm 1) and send the updated information to different kinds of 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 estimates of leader set with cardinality of k or less (when input set cardinality is less than k), which ensures that the message size will never exceed the one containing k robots and their essential information of path sets and value of point sets. The algorithm 2 will always converge to the identical leader set obtained from the standard greedy algorithm 1 over the entire robotic swarm whenever all the robots in the network stop sending messages.

Proof: Assuming the leader set J' is obtained by running standard greedy Algorithm 1 over the entire robot set $[n]$ in a centralized manner and denoting its members by $j_1, j_2, \ldots, j_k$ robots, where the order of robots preserve as j_1 is the first selected leader and j_k is the last. Then we have $j_1 = \arg\max\{F(j_i)\}$ for $i = 1, \ldots, n$, $j_2 = \arg\max\{F(J_i \cup j_1)\}$ for $i = 2, \ldots, n$, etc. Since algorithm 1 functions as the subroutine of Algorithm 2 whenever the information of current robot and its direct neighbors are merged by running line 13 or line 19 in Algorithm 2, the gain function $F(\cdot)$ of robot j_1 will always be the maximum among any local cluster in the first round computation of Algorithm 1 as well, making j_1 remain in the leader set flowing through the networks. Recursively, all the members in $j_1, \ldots, j_k$ will be included in each child node when they runs line 19, and after each robot agrees to this identical leader set, the network stops communication since the child node will directly discards the information containing the same leader set.

Bound on Optimality: In [17] it has been proved that the standard greedy algorithm (Algorithm 1) is able to obtain a solution at least a constant fraction of $(1 - 1/e)$ of the optimal value for solving the maximization of submodular function with cardinality constraint, and in [8] such bound was proved to be tight and improving this bound of approximation is NP-hard. As a byproduct of the above proof, the solutions obtained from the proposed distributed greedy algorithm shares the same optimality bound.

5 Results

5.1 Example

In the first simulation shown in Figure 2, we consider a homogeneous swarm consisting of 5 mobile robots equipped with cameras moving in a rectangular region $[-2,\ 35] \times [-8,\ 30]$ with four static obstacles over A maneuver time span $t \in [0, 40s]$. The task is to select 2 robots every $20s$ as knowledge leaders whose accumulated sensed information is currently maximum and send it to the operator. The leader selection time-points are then specified by $t = 20s$ and $t = 40s$ respectively. Each robot is able to recognize the selection time-points, memorize its traveled path and exchange this information with the corresponding information gain along its path with its direct neighbors within the limited communication range through their communication graph (grey) given in Figure 2b. At $t = 0$ each robot can get access to the initial value of the information potential (density function) 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 will communicate their accumulated sensed information gain, 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 will be able to know the leaders' information after convergence, their information potential (density function) 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 map is updated again as shown in Figure 3c. Each robot's respective information gain 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 of robot 2 and 5 have 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. In Table I the resulting performance comparison of different combinations among robots with highest map coverage area are given at each selection time-point, which validates the selected leaders by Algorithm 2 are the optimal solution.

5.2 Numerical Study

To further analyze the performance and computation cost of the proposed distributed greedy algorithm with other existing work, we conduct 50 simulation trials with each trial consisting of a randomly distributed robotic swarm containing 40 mobile robots. For each trial we conduct our proposed distributed greedy algorithm, the distributed algorithm GREEDI in [15], the standard greedy algorithm, random selection and 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 can converge to the solution from 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 straightforward that the computational cost for optimal selection algorithm grows exponential as number of leaders increases, which makes it impractical in large scale swarm application. For the GREEDI algorithm, since it always performs two-stage standard greedy algorithms on each subset of dataset and then the union of the solution set, in dealing with small scale problem it may not be more efficient than the standard greedy algorithm. For our distributed greedy algorithm, however, 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 swarm scale, the computation cost is hence ensured to significantly decrease compared to the standard greedy algorithm. The communication-related metrics results are reported in Fig. 6, in which the centralized algorithms represent any algorithms that requires global information of all the nodes sent to the root robot of the tree. It is noted that although the distributed greedy algorithm consumes more number of messages, the total amount of data floating within the networks is much less due to the bounded message size.

Table 1: Comparison of accumulated information gain and selecting 2 knowledge leaders

	Robot Index #					Leaders, Gain
	1	2	3	4	5	
Info. Gain at $20s$	1.03	1.07	1.11	1.12	1.26	$[2, 5]$, 2.34
$(\times 10^4)$	1.06	1.17	1.11	1.01	1.26	$[4, 5]$, 2.27
Info. Gain at $40s$	2.13	0.99	2.11	2.07	1.19	$[3, 4]$, 4.18
$(\times 10^4)$	2.06	1.01	2.21	2.07	1.18	$[1, 4]$, 4.13

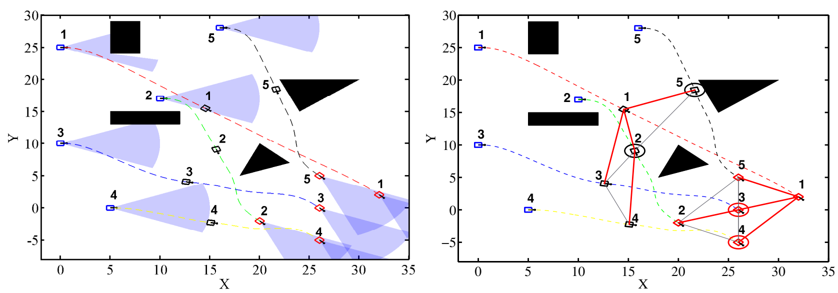

(a) (b)

Figure 2: Simulation of 5 mobile robots exploring the environment with 4 static obstacles (black) over time $[0, 40s]$. The time-points for leader selection is set to be 20s and 40s respectively. (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.

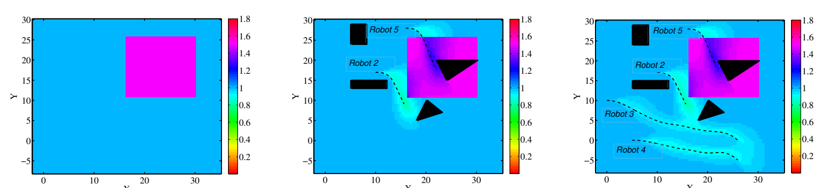

(a) (b) (c)

Figure 3: Heat map of diminishing information potential (density function) changes over time. (a) Initial interest distribution at $t = 0s$. Information on the magenta area gives robots higher value 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 again at $t = 40s$ after robot 3 and 4 are selected as new leaders. Polygons with black edges marks the positions of static obstacles in the heat maps.

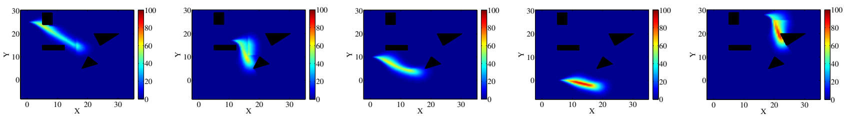

(a) Robot 1 (b) Robot 2 (c) Robot 3 (d) Robot 4 (e) Robot 5

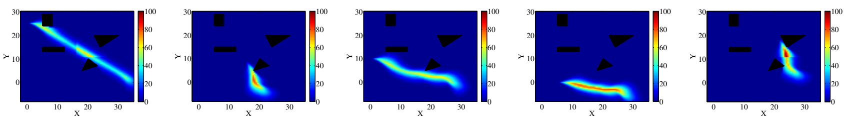

(f) Robot 1 (g) Robot 2 (h) Robot 3 (i) Robot 4 (j) Robot 5

Figure 4: Snapshot of accumulated information gain of robots 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 (density function) 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.

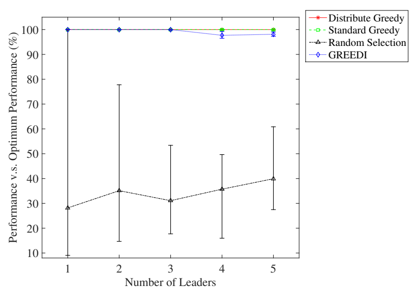

(a)

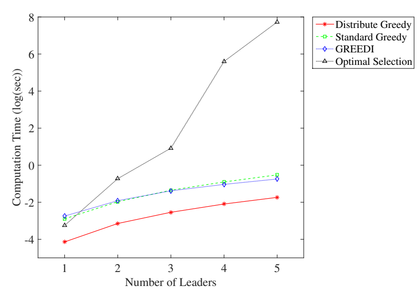

(b)

Figure 5: Performance comparison of the proposed distributed greedy algorithm and other submodular optimization algorithms from 50 independent trials on randomly generated swarms 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 [15] vs. the benchmark performance of centralized combinatorial optimization algorithm (NP hard). (b) The average computation time (log(sec)) among the four algorithms. Note that for GREEDI algorithm the cardinality of each subset is chosen to be 8.

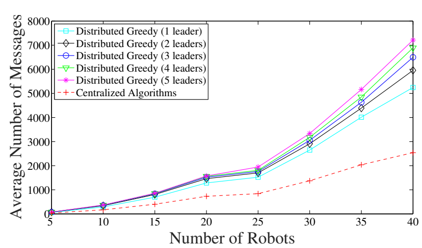

(a)

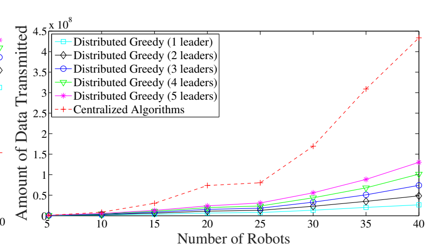

(b)

Figure 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.

6 Conclusion

We formulated the knowledge leader selection problem as an optimization problem of a submodular function and provided a distributed greedy algorithm to find the best knowledge leaders. Our proposed distributed submodular optimization approach is guaranteed to provide the same approximate solution as the centralized greedy algorithm. Numerical simulations is given to compare our distributed knowledge leader selection algorithm to GREEDI algorithm [15], standard greedy algorithm and random leader selection algorithm on computational time and performance with respect to an optimal selection strategy, which validates the effectiveness of the proposed algorithm.

References

[1] Ian F. Akyildiz, Dario Pompili, and Tommaso Melodia. "Underwater acoustic sensor networks: research challenges". In: *Ad Hoc Networks* 3.3 (2005), pp. 257–279. ISSN: 1570-8705. DOI: `http://dx.doi.org/10.1016/j.adhoc.2005.01.004`.

[2] N. Atanasov, J. Le Ny, K. Daniilidis, and G. J. Pappas. "Information acquisition with sensing robots: Algorithms and error bounds". In: *IEEE Int'l Conf. on Robotics and Automation*. May 2014, pp. 6447–6454. DOI: `10.1109/ICRA.2014.6907811`.

[3] Gokhan M Atinc, Dusan M Stipanovic, Petros G Voulgaris, and Mansour Karkoub. "Swarm-based dynamic coverage control". In: *IEEE 53rd Annual Conf. on Decision and Control*. 2014, pp. 6963–6968.

[4] Jonathan Butzke and Maxim Likhachev. "Planning for multi-robot exploration with multiple objective utility functions". In: *IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*. 2011, pp. 3254–3259.

[5] Andrew Clark, Basel Alomair, Linda Bushnell, and Radha 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*. 2014, pp. 397–404.

[6] Andrew Clark, Basel Alomair, Linda Bushnell, and Radha Poovendran. "Minimizing convergence error in multi-agent systems via leader selection: A supermodular optimization approach". In: *IEEE Transactions on Automatic Control* 59.6 (2014), pp. 1480–1494.

[7] Jorge Cortés, Sonia Martınez, Timur Karatas, and Francesco Bullo. "Coverage control for mobile sensing networks". In: *IEEE Int'l Conf. on Robotics and Automation* 20.2 (2004).

[8] Uriel Feige. "A threshold of ln n for approximating set cover". In: *Journal of the ACM* 45.4 (1998), pp. 634–652.

[9] Satoru Fujishige. *Submodular functions and optimization*. Vol. 58. Elsevier, 2005.

[10] Daniel Golovin, Matthew Faulkner, and Andreas Krause. "Online distributed sensor selection". In: *Proceedings of the 9th ACM/IEEE Int'l Conf. on Information Processing in Sensor Networks*. 2010, pp. 220–231.

[11] Bruno Hexsel, Nilanjan Chakraborty, and Katia Sycara. "Coverage control for mobile anisotropic sensor networks". In: *IEEE Int'l Conf. on Robotics and Automation*. 2011, pp. 2878–2885.

[12] I.I. Hussein and D.M. Stipanovic. "Effective Coverage Control for Mobile Sensor Networks With Guaranteed Collision Avoidance". In: *IEEE Transactions on Control Systems Technology* 15.4 (July 2007), pp. 642–657.

[13] A. Kolling, P. Walker, N. Chakraborty, K. Sycara, and M. Lewis. "Human Interaction With Robot Swarms: A Survey". In: *IEEE Transactions on Human-Machine Systems* 46.1 (Feb. 2016), pp. 9–26. ISSN: 2168-2291. DOI: `10.1109/THMS.2015.2480801`.

[14] Wenhao Luo, Shehzaman Salim Khatib, Sasanka Nagavalli, Nilanjan Chakraborty, and Katia Sycara. "Asynchronous Distributed Information Leader Selection in Robotic Swarms". In: *IEEE Int'l Conf. on Automation Science and Engineering*. Aug. 2015.

[15] Baharan Mirzasoleiman, Amin Karbasi, Rik Sarkar, and Andreas Krause. "Distributed submodular maximization: Identifying representative elements in massive data". In: *Advances in Neural Information Processing Systems*. 2013, pp. 2049–2057.

[16] Sasanka Nagavalli, Andrew Lybarger, Lingzhi Luo, Nilanjan Chakraborty, and Katia Sycara. "Aligning coordinate frames in multi-robot systems with relative sensing information". In: *IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*. 2014, pp. 388–395.

[17] George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. "An analysis of approximations for maximizing submodular set functions-1". In: *Mathematical Programming* 14.1 (1978), pp. 265–294.

[18] Dimitra Panagou, Dusan M Stipanovic, and Petros G Voulgaris. "Vision-based dynamic coverage control for nonholonomic agents". In: *IEEE 53rd Annual Conf. on Decision and Control*. 2014, pp. 2198–2203.

[19] Gregg Podnar, John Dolan, Alberto Elfes, and Marcel Bergerman. "Human Telesupervision of Very Heterogeneous Planetary Robot Teams". In: *Proceedings of the AIAA Space 2007 Conference, Long Beach, CA*. 2007, pp. 18–20.

[20] Mac Schwager, Daniela Rus, and Jean-Jacques Slotine. "Decentralized, adaptive coverage control for networked robots". In: *The Int'l Journal of Robotics Research* 28.3 (2009), pp. 357–375.

[21] T. Summers, F. Cortesi, and J. Lygeros. "On Submodularity and Controllability in Complex Dynamical Networks". In: *IEEE Transactions on Control of Network Systems* PP.99 (2015), pp. 1–1. ISSN: 2325-5870. DOI: 10.1109/TCNS.2015.2453711.

[22] Abhinav Valada, Christopher Tomaszewski, Balajee Kannan, Prasanna Velagapudi, George Kantor, and Paul Scerri. "An intelligent approach to hysteresis compensation while sampling using a fleet of autonomous watercraft". In: *Intelligent Robotics and Applications*. Springer, 2012, pp. 472–485.

[23] Huadong Wang, Andreas Kolling, Nathan Brooks, Sean Owens, Shafiq Abedin, Paul Scerri, Pei-ju Lee, Shih-Yi Chien, Michael Lewis, and Katia Sycara. "Scalable target detection for large robot teams". In: *ACM/IEEE Int'l Conf. on Human-Robot Interaction*. 2011, pp. 363–370.

[24] You Xu, Abusayeed Saifullah, Yixin Chen, Chenyang Lu, and Sangeeta Bhattacharya. "Near optimal multi-application allocation in shared sensor networks". In: *Proceedings of the 11th ACM Int'l symposium on Mobile ad hoc networking and computing*. 2010, pp. 181–190.