



Brief paper

A sub-modular receding horizon solution for mobile multi-agent persistent monitoring[☆]

Navid Rezazadeh*, Solmaz S. Kia

University of California, Irvine, United States of America



ARTICLE INFO

Article history:

Received 12 August 2019

Received in revised form 11 October 2020

Accepted 27 November 2020

Available online xxxx

Keywords:

Persistent monitoring

Optimal policy design

Submodular optimization

ABSTRACT

We consider persistent monitoring of a finite number of inter-connected geographical nodes by a group of heterogeneous mobile agents. We assign to each geographical node a concave and increasing reward function that resets to zero after an agent's visit. Then, we design the optimal dispatch policy of which nodes to visit at what time and by what agent by finding a policy set that maximizes a utility that is defined as the total reward collected at visit times. We show that this optimization problem is NP-hard and its computational complexity increases exponentially with the number of the agents and the length of the mission horizon. By showing that the utility function is a monotone increasing and submodular set function of agents' policy, we propose a suboptimal dispatch policy design with a known optimality gap. To reduce the time complexity of constructing the feasible search set and also to induce robustness to changes in the operational factors, we perform our suboptimal policy design in a receding horizon fashion. Then, to compensate for the shortsightedness of the receding horizon approach we add a new term to our utility, which provides a measure of nodal importance beyond the receding horizon. This term gives the policy design an intuition to steer the agents towards the nodes with higher rewards on the patrolling graph. Finally, we discuss how our proposed algorithm can be implemented in a decentralized manner. A simulation study demonstrates our results.

© 2020 Elsevier Ltd. All rights reserved.

1. Introduction

In recent years, coordinating the movement of mobile sensors to cover areas that have not been adequately sampled/observed has been explored in controls, wireless sensors, and robotic communities with problems related to coverage, exploration, and deployment. Many of the proposed algorithms strive to spread sensors to desired positions to obtain a stationary configuration such that the coverage is optimized, see e.g., Bullo, Carli, and Frasca (2012), Carron, Todescato, Carli, Schenato, and Pillonetto (2015), Chung and Kia (2020), Cortes, Martinez, Karatas, and Bullo (2004), Krause and Guestrin (2007), Krause, Singh, and Guestrin (2008), Schwager, Rus, and Slotine (2009) and Todescato, Carron, Carli, Pillonetto, and Schenato (2017). Some sensor placement problems such Carron et al. (2015), Chung and Kia (2020), Schwager et al. (2009) and Todescato et al. (2017) are context-aware

and include also a period of exploration and observation to increase the knowledge used to find the optimal residing position of the sensors. In this paper, instead of aiming to achieve an improved stationary network configuration as the end result of the sensors' movement, our objective is to explore context-aware mobility strategies that dynamically re-position the mobile sensors to maximize their utilization and contribution over a mission horizon. Motivating applications include persistent monitoring to discover forest fires (Yuan, Zhang, & Liu, 2015) or oil spillage in its early stages (Henry & Henry, 2015), locating endangered animals in a large habitat (Engler, Guisan, & Rechsteiner, 2004) and event detection in urban environments (Thomas & van Berkum, 2009). Specifically, we consider persistent monitoring of a set of finite \mathcal{V} inter-connected geographical nodes via a set of finite \mathcal{A} mobile sensors/agents, where $|\mathcal{V}| > |\mathcal{A}|$. The mobile agents are confined to a set of pre-specified edges $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$, e.g., aerial or ground corridors, to traverse from one node to another, see Fig. 1. Depending on their vehicle type, agents may have to take different edges to go from one node to another. Also, they may have different travel times along the same edge. We study dispatch policy that orchestrates the topological distribution of the mobile agents such that an optimized service for a global monitoring task is provided with a reasonable computational cost. To quantify the service objective we assign to each node $v \in \mathcal{V}$ the reward

☆ This work is supported by NSF, United States of America award IIS-SAS-1724331. The material in this paper was partially presented at the 8th IFAC Workshop on Distributed Estimation and Control in Networked Systems, September 16–17, 2019, Chicago, Illinois, USA. This paper was recommended for publication in revised form by Associate Editor Luca Schenato under the direction of Editor Christos G. Cassandras.

* Corresponding author.

E-mail addresses: nrezasad@uci.edu (N. Rezazadeh), solmaz@uci.edu (S.S. Kia).

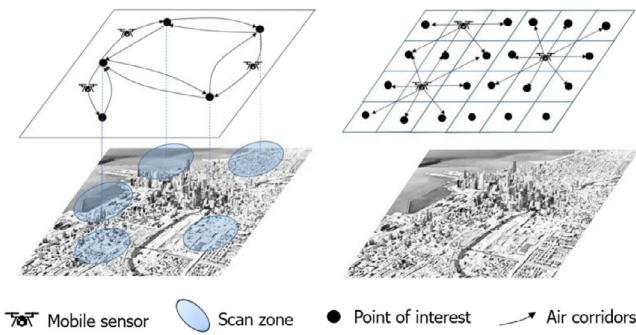


Fig. 1. Examples of a set of geographical nodes of interest and the edges between them. Finite number of nodes to monitor in a city can be restricted to some particular scanning zones (the picture on the left) or the cell partitioned map of the city (the picture on the right).

function,

$$R_v(t) = \begin{cases} 0, & t = \bar{t}_v, \\ \psi_v(t - \bar{t}_v), & t > \bar{t}_v, \end{cases} \quad (1)$$

where $\psi_v(t)$ is a nonnegative concave and increasing function of time and \bar{t}_v is the latest time node v is visited by an agent. For example, in the data harvesting or health monitoring, $\psi_v(\cdot)$ can be the weighted idle time of the node v , or in event detection, it can be the probability of at least one event taking place at inter-visit times. Optimal patrolling designs a dispatch policy (what sequence of nodes to visit at what times by which agents) to score the maximum collective reward for the team over the mission horizon. However, as we explain below, this problem is NP-hard. Our aim then is to design a suboptimal solution that has polynomial time complexity.

Related work: Dispatch policy design for patrolling/monitoring of geographical nodes can be divided into two categories: the edges to travel between the nodes are not specified (design in continuous edge space) or otherwise (design in discrete edge space). When there are no prespecified inter-node edges, the optimal patrolling policy design includes also finding the optimal inter-node trajectories that the agents should follow without violating their mobility limits. In some applications, however, the mobile agents are confined to travel through pre-specified known edges between the nodes. For example, in a smart city setting, regulations can restrict the admissible routes between the geographical nodes. In the dispatch policy design in discrete edge space, the complexity of finding the optimal policy for a single patrolling agent is the same as the complexity of solving the Traveling Salesman problem, where the computational complexity grows exponentially with the number of the nodes (Karp, 1972). In the case of multiple patrolling agents, the problem is even more complex, since each agent's policy design depends on the other agents' policy. This problem is formalized in earlier studies such as Almeida et al. (2004) and Machado, Ramalho, Zucker, and Drogoul (2002). Generally, when there are multiple edges to travel between every two nodes or when each node is connected to multiple other nodes, finding an optimal long term patrolling scheme is not tractable. Constraining the agents to travel through specific edges to traverse among the geographical nodes allows seeking optimal solutions for the problem. For example, when the connection topology between the geographical nodes is a path or a cyclic graph, optimal solutions for the problem are proposed in Chevaleyre (2004), Donahue, Rosman, Kotowick, Rus, and Baykal (2016), Pasqualetti, Franchi, and Bullo (2012) and Yu, Karaman, and Rus (2015). To overcome the complexity issue on generic graphs, Asghar, Smith, and Sundaram (2019) explore

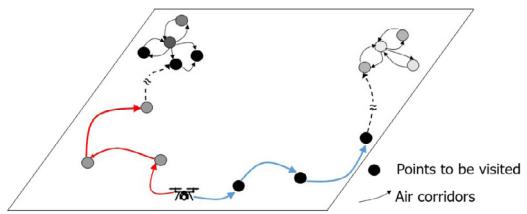


Fig. 2. An agent has two possible routes to take over the designated receding horizon. The nodes' color intensity shows their reward value. The blue route offers a higher reward over the receding horizon but it puts the agent close to an area with a lower amount of reward, while the red route results in lower total reward over the receding horizon but puts the agent near an area with higher amount of reward. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

forming different cycles in the graph and assigning agents to these cycles to patrol the nodes periodically and seeks to minimize the time that a node stays un-visited. Alternatively, Farinelli, Iocchi, and Nardi (2017) propose agents to move to the most rewarding neighboring node based on their current location.

Statement of contribution: In this paper, we propose a robust and suboptimal solution to the long term patrolling problem that we stated earlier. Instead of using the customary idle time, $\psi_v(t) = t$, as a reward function, which reduces the optimal dispatch policy design to the minimum latency problem (Blum et al., 1994), we consider reward functions described by an increasing concave function. This allows modeling a wider class of patrolling problems such as patrolling for event detection. We let the utility function to be the sum of the rewards collected over the mission horizon by the mobile agents. We discuss that the design of an optimal patrolling policy to maximize this utility over the mission horizon is an NP-hard problem. Specifically, we show that the complexity of finding the optimal policy increases exponentially with the mission horizon and number of agents. Next, we show that the utility function is a monotone increasing and submodular set function. To establish this result, we develop a set of auxiliary lemmas, presented in the appendix, based on Karamata's inequality (Kadelburg, Dukic, Lukic, & Matic, 2005). Given the submodularity of the utility function, we propose a receding horizon sequential greedy algorithm to compute a suboptimal dispatch policy with a polynomial computation cost and guaranteed bound on optimality. The receding horizon nature of our solution induces robustness to uncertainties of the environment. Our next contribution is to add a new term to our utility function to compensate for the shortsightedness of the receding horizon approach, see Fig. 2. When agents patrol a large set of interconnected nodes, this added term becomes useful by giving them an intuition of the existing reward in the farther nodes. In recent years, submodular optimization has been widely used in sensor and actuator placement problems (Clark, Bushnell, & Poovendran, 2014; Clark, Lee, Alomair, Bushnell, & Poovendran, 2018; Krause & Guestrin, 2007; Krause et al., 2008; Liu et al., 2018; Summers, Cortesi, & Lygeros, 2016). In comparison to the sensor/actuator placement problems, the challenge in our work is that the assigned policy per each mobile agent over the receding horizon is a dynamic scheduling problem rather than a static sensor placement. To deal with this challenge, we use the matroid constraint (Fisher, Nemhauser, & Wolsey, 1978) approach to design our suboptimal submodular-based policy. Finally, we discuss how our algorithm can be implemented in a decentralized manner. A simulation study demonstrates our results. Our notation is standard, though to avoid confusion, certain concepts and notation are defined as the need arises. This paper extends our preliminary work (Rezazadeh & Kia, 2019a) in detailed technical

treatment including all the proofs, introducing the notion of local importance to compensate for the shortsightedness of the receding horizon approach, decentralized implementation of our algorithm, and a new simulation study. Also, we consider a more generalized case of reward functions.

2. Problem formulation

To formalize our objective, we first introduce our notations and state our standing assumptions. For any node $v \in \mathcal{V}$, \mathcal{N}_v is a set consisting node v and all the neighboring nodes that are connected to node v via an edge in \mathcal{E} . If there exists a path connecting node $v \in \mathcal{V}$ to node $w \in \mathcal{V}$, we let $\tau_{v,w}^i \in \mathbb{R}_{>0}$ be the shortest travel time of agent $i \in \mathcal{A}$ from node v to w .

Assumption 1. Upon arrival of any agent $i \in \mathcal{A}$ at any time $\bar{t} \in \mathbb{R}_{>0}$ at node $v \in \mathcal{V}$, the agent immediately scans the node and the reward $R_v(\bar{t})$ is scored for the patrolling team \mathcal{A} and \bar{t}_v of node v in (1) is set to \bar{t} . If more than one agent arrives at node $v \in \mathcal{V}$ and scans it at the same time \bar{t} , the reward collected for the team is still $R_v(\bar{t})$. If an agent $i \in \mathcal{A}$ needs to linger over each node for $\delta^i \in \mathbb{R}_{\geq 0}$ amount of time to complete its scan, during this time the agent cannot scan the node again to score a reward for the team.

Let the tuple $\mathbf{p} = (\mathbf{V}_p, \mathbf{T}_p, a_p)$ be a dispatch policy of agent $a_p \in \mathcal{A}$ over the given mission time horizon, where \mathbf{V}_p and \mathbf{T}_p are the vectors that specify the nodes and the corresponding visit times assigned to agent a_p . Moreover, we let n_p be the total number of nodes visited by agent a_p , i.e., $n_p = \text{dim}(\mathbf{V}_p)$. We refer to n_p as the *length* of the policy \mathbf{p} . We refer to $(\mathbf{V}_p(l), \mathbf{T}_p(l))$, $l \in \{1, 2, \dots, n_p\}$, as the l th step of policy \mathbf{p} . Furthermore, for any agent $i \in \mathcal{A}$, we let \mathcal{P}^i be the set of all the admissible policies \mathbf{p} over the mission horizon such that $a_p = i$.

Assumption 2. For any policy \mathbf{p} , we have $\mathbf{V}_p(l+1) \in \mathcal{N}_{\mathbf{V}_p(l)}$, for all $l \in \{1, 2, \dots, n_p - 1\}$.

We let $\mathcal{P} = \bigcup_{i \in \mathcal{A}} \mathcal{P}^i$. Then, given any $\bar{\mathcal{P}} \subset \mathcal{P}$, the utility function $R : 2^{\mathcal{P}} \rightarrow \mathbb{R}_{>0}$ is

$$R(\bar{\mathcal{P}}) = \sum_{\forall p \in \bar{\mathcal{P}}} \sum_{l=1}^{n_p} R_{\mathbf{V}_p(l)}(\mathbf{T}_p(l)). \quad (2)$$

Given (2), the optimal policy to maximize the utility over a given mission horizon is given by

$$\mathcal{P}^* = \underset{\bar{\mathcal{P}} \subset \mathcal{P}}{\text{argmax}} R(\bar{\mathcal{P}}), \quad \text{s.t.} \quad (3a)$$

$$|\bar{\mathcal{P}} \cap \mathcal{P}^i| \leq 1 \quad i \in \mathcal{A}, \quad (3b)$$

where $|\cdot|$ returns the cardinality of a set. The constraint condition (3b) is in the so-called *partition matroid* form (Fisher et al., 1978) and restricts the choice of the optimal solution to be a set that contains of at most one member from each disjoint sets \mathcal{P}^i , $i \in \mathcal{A}$. A set value optimization problem of the form (3) is known to be NP-hard (Lovász, 1983). Lemma 3, whose proof is given in the appendix, gives the cost of constructing the feasible set \mathcal{P} and time complexity of solving optimization problem (3).

Lemma 3 (Time complexity of Problem (3a)). *The cost of constructing the feasible set \mathcal{P} of optimization problem (3a) is of order $O(\sum_{i \in \mathcal{A}} D^{\bar{n}^i})$, where $D = \max_{v \in \mathcal{V}} (|\mathcal{N}_v|)$ and $\bar{n}^i = \max\{n_p\}_{p \in \mathcal{P}^i}$. Furthermore, the time complexity of solving optimization problem (3a) is $O(\prod_{i \in \mathcal{A}} D^{\bar{n}^i})$.*

If the system parameters, such as the number of the mobile agents or the nodes, or the parameters of $\psi_v(\cdot)$ of the reward function at any node v , change after the optimal policy design, the optimization problem (3) should be solved again over the remainder of the mission horizon under the new conditions. Our objective in this paper is to construct a suboptimal solution to solve the persistent monitoring problem given by (3) with polynomial time complexity. Moreover, we seek a solution that has intrinsic robustness to changes that can happen during the mission horizon.

We close this section by introducing some definitions and notations used subsequently. For any set function $g : 2^{\mathcal{Q}} \rightarrow \mathbb{R}$, we let $\Delta_g(q|\bar{Q}) = g(\bar{Q} \cup q) - g(\bar{Q})$, for $\forall \bar{Q} \in 2^{\mathcal{Q}}$ and $\forall q \in \mathcal{Q}$, where Δ_g shows the increase in value of the set function g going from set \bar{Q} to $\bar{Q} \cup q$. Recall that $g : 2^{\mathcal{Q}} \rightarrow \mathbb{R}$ is *submodular* if and only if for two sets \mathcal{Q}_1 and \mathcal{Q}_2 satisfying $\mathcal{Q}_1 \subset \mathcal{Q}_2 \subset \mathcal{Q}$, and for $q \notin \mathcal{Q}_2$ we have $\Delta_g(q|\bar{Q}_1) \geq \Delta_g(q|\bar{Q}_2)$ (Fisher et al., 1978). Thus submodularity is a property of set functions that shows diminishing reward as new members are being introduced to the system. We say $g : 2^{\mathcal{Q}} \rightarrow \mathbb{R}$ is *monotone increasing* if for all $\mathcal{Q}_1, \mathcal{Q}_2 \subset \mathcal{Q}$ we have $\mathcal{Q}_1 \subset \mathcal{Q}_2$ if and only if $g(\mathcal{Q}_1) \leq g(\mathcal{Q}_2)$ (Fisher et al., 1978). We denote a sequence of m real numbers (t_1, \dots, t_m) by $(t)_1^m$. Given two increasing (resp. decreasing) sequences $(t)_1^n$ and $(v)_1^m$, $(t)_1^n \oplus (v)_1^m$ is their concatenated increasing (resp. decreasing) sequence, i.e., for $(u)_1^{n+m} = (t)_1^n \oplus (v)_1^m$, any u_k , $k \in \{1, \dots, n+m\}$ is either in $(t)_1^n$ or $(v)_1^m$ or is in both. We assume that $(u)_1^{n+m}$ preserves the relative labeling of $(t)_1^n$ or $(v)_1^m$, i.e., if t_k and t_{k+1} , $k \in \{1, \dots, n-1\}$ (resp. v_k and v_{k+1} , $k \in \{1, \dots, m-1\}$) correspond to u_p and u_q in $(u)_1^{n+m}$, then $p < q$.

3. Suboptimal policy design

According to Lemma 3, the time complexity of finding an optimal patrolling policy in (3) increases exponentially by the maximum length, \bar{n}^i , of the admissible policies of any agent $i \in \mathcal{A}$ and also by the number of the exploring agents M . In light of this observation, to reduce the computational cost, we propose the following suboptimal policy design. Since the maximum policy length \bar{n}^i is proportional to the length of the mission horizon, we first propose to trade-off optimality and divide the planning horizon into multiple shorter horizons so that the policy design can be carried out in a consecutive manner over these shorter horizons. Then, to reduce the optimality gap and also to induce robustness to the online changes that can occur during the mission time, we propose to implement this approach in a receding horizon fashion where we calculate the policy over a specified shorter horizon but execute only some of the initial steps of the policy, and then we repeat the process. However, a receding horizon approach suffers from what we refer to as *shortsightedness*. That is, over large inter-connected geographical node sets, a receding horizon design is oblivious to the reward distribution of the nodes that are not in the feasible policy set in the planning horizon. Then, the optimal policy over the planning horizon can inadvertently steer the agents away from the distant nodes with a higher reward, see Fig. 2. To compensate for this shortcoming, we introduce the notion of *nodal importance* and augment the reward function (2) over the design horizon with an additional term that given an admissible policy, provides a measure of how close an agent at the final step of the policy is to a cluster of geographical nodes with a high concentration of reward.

Let the augmented reward, whose exact form will be introduced below, over the planning horizon be \bar{R} . Then, the optimal policy design over each receding horizon is

$$\mathcal{P}^* = \underset{\bar{\mathcal{P}} \subset \mathcal{P}}{\text{argmax}} \bar{R}(\bar{\mathcal{P}}), \quad \text{s.t.} \quad |\bar{\mathcal{P}} \cap \mathcal{P}^i| \leq 1 \quad i \in \mathcal{A}, \quad (4)$$

Algorithm 1 Sequential Greedy Algorithm

```

1: procedure  $\text{SGOpt}(\mathcal{P}^i, i \in \mathcal{A})$ 
2:   Init:  $\bar{\mathcal{P}} \leftarrow \emptyset$ ,  $i \leftarrow 0$ ,  $\{\bar{t}_v^0\}_{v \in \mathcal{V}}$ 
3:   for  $i \in \mathcal{A}$  do
4:      $p^{i*} = \underset{p \in \mathcal{P}^i}{\text{argmax}} \Delta_{\bar{\mathcal{P}}}(p|\bar{\mathcal{P}})$ .
5:      $\bar{\mathcal{P}} \leftarrow \bar{\mathcal{P}} \cup p^{i*}$ .
6:   end for
7:   Return  $\bar{\mathcal{P}}$ .
8: end procedure

```

where hereafter $\mathcal{P} = \bigcup_{i \in \mathcal{A}} \mathcal{P}^i$ is the set of the union of the admissible policies of the agents $\mathcal{P}^i, i \in \mathcal{A}$, over the planning horizon. Hereafter, we let \bar{t}_v^0 be the last time node $v \in \mathcal{V}$ was visited before a planning horizon starts.

Next, to reduce the computational burden further, we propose to use Algorithm 1, which is a sequential greedy algorithm with a polynomial cost in terms of the number of the agents to obtain a suboptimal solution for (4). In what follows, we show that since the objective function (4) is a submodular set function, Algorithm 1 comes with a known optimality gap. We also show that with proper inter-agent communication coordination, Algorithm 1 can be implemented in a decentralized manner.

For $v \in \mathcal{V}$, let \mathcal{N}_v^r be the set consisted of node v itself and its r -hop neighbors. This set can be computed using the Breadth-first search in time $O(|\mathcal{E}| + |\mathcal{V}|)$ (Thomas, Leiserson, Rivest, & Stein, 2009). Then, for every node $v \in \mathcal{V}$, we define the nodal importance with radius r at time τ as $L(v, \tau, r) = \sum_{w \in \mathcal{N}_v^r} R_w(\tau)$. Next, given an agent $i \in \mathcal{A}$ that is at node $w \in \mathcal{V}$ at time $\hat{\tau} \in \mathbb{R}_{\geq 0}$, we define the *relative nodal importance* of a node $v \in \mathcal{V}$ with respect to agent i as

$$L(v, w, \hat{\tau}, i) = L(v, \hat{\tau} + \tau_{w,v}^i, r) / \tau_{w,v}^i.$$

Here, $\tau_{w,v}^i$ can be computed via A^* algorithm in time $O(|\mathcal{E}|)$ (Duchon et al., 2014). Then, $L(v, \mathbf{V}_p(n_p), \mathbf{T}_p(n_p), a_p)$ is a measure of the relative size of the awards concentration around any node $v \in \mathcal{V}$ that takes into account also the travel time of agent a_p from the final step of policy $p = (\mathbf{V}_p, \mathbf{T}_p, a_p) \in \mathcal{P}$ to v . Let $L(v, p)$ be the shorthand notation for $L(v, \mathbf{V}_p(n_p), \mathbf{T}_p(n_p), a_p)$. To compensate for the shortsightedness of the receding horizon design, then we revise the utility function to

$$\bar{R}(\bar{\mathcal{P}}) = R(\bar{\mathcal{P}}) + \alpha \sum_{p \in \bar{\mathcal{P}}} \max_{v \in \bar{\mathcal{V}}} L(v, p), \quad \alpha \in \mathbb{R}_{\geq 0}. \quad (5)$$

The weighting factor $\alpha \in \mathbb{R}_{\geq 0}$ defines how much significance we want to assign to the distribution of the reward beyond the receding horizon. We should note that using a large α can gravitate the agents to move towards the nodes close to the anchor nodes, and make them oblivious to the rest of the nodes. For computational efficiency, instead of incorporating the relative nodal importance of all the nodes, which can be achieved by setting $\bar{\mathcal{V}}$ equal to \mathcal{V} , we propose to use only $\bar{\mathcal{V}}$ subset of the nodes. We refer to nodes in $\bar{\mathcal{V}}$ as *anchor nodes*. The anchor nodes can be selected to be the nodes with higher reward return or to be a set of nodes that are scattered uniformly on the graph. It is interesting to note that the relative nodal importance term in (5) is reminiscent of terminal cost used in the model predictive control (MPC). In MPC, a terminal cost that is used to achieve an infinite horizon control with closed-loop stability guarantees (Garcia, Prett, & Morari, 1989) in some way also compensates for the shortsightedness of the design over a finite planning horizon. Next, we show that the reward function (5) is submodular over any given feasible policy set \mathcal{P} in every planning horizon.

Theorem 4 (Submodularity of the Reward Function (5)). For any weighting factor $\alpha \in \mathbb{R}_{\geq 0}$, the reward function $\bar{R} : 2^{\mathcal{P}} \rightarrow \mathbb{R}_{\geq 0}$ in (5) is a monotone increasing and submodular set function over \mathcal{P} .

Proof. Let $c(v, \mathcal{Q}) : \mathcal{V} \times 2^{\mathcal{Q}} \rightarrow \mathbb{Z}_{\geq 0}$ be the total number of visits to the geographical node v , and $\mathcal{I}_{\mathcal{Q}} \subset \mathcal{V}$ be the set of the nodes that are visited when a policy set $\mathcal{Q} \subset \mathcal{P}$ is implemented. Furthermore, let the increasing sequence $(t^v(\mathcal{Q}))_1^{c(v, \mathcal{Q})} = (t_1^v(\mathcal{Q}), t_2^v(\mathcal{Q}), \dots, t_{c(v, \mathcal{Q})}^v(\mathcal{Q}))$ be the sequence of time that node $v \in \mathcal{I}_{\mathcal{Q}}$ was visited when agents implement \mathcal{Q} . Now consider the reward function \bar{R} in (5). Then, the first summand of \bar{R} expands as $R(\bar{\mathcal{P}}) = \sum_{v \in \mathcal{I}_{\bar{\mathcal{P}}}} (\sum_{j=1}^{c(v, \bar{\mathcal{P}})} \psi_v(\Delta t_j^v(\bar{\mathcal{P}})))$, where $\Delta t_j^v(\bar{\mathcal{P}}) = t_j^v(\bar{\mathcal{P}}) - t_{j-1}^v(\bar{\mathcal{P}})$ is the time between two consecutive visits of node v , and $t_0^v(\bar{\mathcal{P}}) = \bar{t}_0^v$. Next, consider the monitoring policy sets $\mathcal{Q}_1, \mathcal{Q}_2$ and monitoring policy q with $\mathcal{Q}_1 \subset \mathcal{Q}_2 \subset \mathcal{P}$, $q \in \mathcal{P}$, $q \notin \mathcal{Q}_1$, and $q \notin \mathcal{Q}_2$. Because $(t^v(\mathcal{Q}_1))_1^{c(v, \mathcal{Q}_1)}$ is a sub-sequence of $(t^v(\mathcal{Q}_2))_1^{c(v, \mathcal{Q}_2)}$, using Lemma A.8 and the fact that $\psi_v(\cdot)_v$ is a normalized increasing concave function, we conclude that $\sum_{j=1}^{c(v, \mathcal{Q}_1 \cup q)} \psi_v(\Delta(t_j^v(\mathcal{Q}_2 \cup q))) - \sum_{j=1}^{c(v, \mathcal{Q}_1)} \psi_v(\Delta(t_j^v(\mathcal{Q}_2))) \geq 0$ for $\forall v \in \mathcal{I}_{\bar{\mathcal{P}}}$. Therefore, $\Delta_R(p|\mathcal{Q}_1) \geq 0$ which shows that $R(\bar{\mathcal{P}})$ is a monotone increasing set function. Furthermore, using Lemma A.9 we can write $\sum_{j=1}^{c(v, \mathcal{Q}_2 \cup q)} \psi_v(\Delta(t_j^v(\mathcal{Q}_2 \cup q))) - \sum_{j=1}^{c(v, \mathcal{Q}_2)} \psi_v(\Delta(t_j^v(\mathcal{Q}_2))) \leq \sum_{j=1}^{c(v, \mathcal{Q}_1 \cup q)} \psi_v(\Delta(t_j^v(\mathcal{Q}_1 \cup q))) - \sum_{j=1}^{c(v, \mathcal{Q}_1)} \psi_v(\Delta(t_j^v(\mathcal{Q}_1)))$. Hence, $\Delta_R(q|\mathcal{Q}_1) \geq \Delta_R(q|\mathcal{Q}_2)$ which shows that $R(\bar{\mathcal{P}})$ is a submodular set function. Then, since the second summand of \bar{R} , $\sum_{v \in \bar{\mathcal{V}}} \max_{p \in \bar{\mathcal{P}}} L(v, p)$, is trivially positive and modular, the proof is concluded. \square

Due to Theorem 4, the suboptimal dispatch policy of Algorithm 1, which has a polynomial computational complexity, has the following well-defined optimality gap.

Theorem 5 (Optimality gap of Algorithm 1). Let \mathcal{P}^* be an optimal solution of (4) and $\bar{\mathcal{P}}$ be the output of Algorithm 1. Then, $\bar{R}(\bar{\mathcal{P}}) \geq \frac{1}{2} \bar{R}(\mathcal{P}^*)$.

Proof. Since the objective function of (4) is monotone increasing and submodular over \mathcal{P} , the proof follows by invoking (Fisher et al., 1978, Theorem 5.1).

3.1. Comments on decentralized implementations of Algorithm 1

To implement Algorithm 1, given the current position of each agent and $\{\bar{t}_v^0\}_{v \in \mathcal{V}}$ at the beginning of each planning horizon, the admissible set of policies \mathcal{P}^i for each agent $i \in \mathcal{A}$ should be calculated. Let every agent know $\{\psi_v(t)\}_{v \in \mathcal{V}}$. A straightforward decentralized implement of Algorithm 1 then is a multi-centralized solution. In this solution, agents transmit the feasible policy sets across the entire network until each agent knows all the policy sets $\mathcal{P}^i, \forall i \in \mathcal{A}$ (flooding approach). Then, each agent acts as a central node and runs a copy of Algorithm 1 locally. Although reasonable for small-size networks, the communication and storage costs of this approach scale poorly with the network size. The sequential structure of Algorithm 1 however, offers an opportunity for a communicationally and computationally more efficient decentralized implementation, as described in steps 1 to 9 of Algorithm 2. Step 10 of Algorithm 2 is included for receding horizon implementation purpose, where the execution plan can be for example one or all of the agents visit at least one node. To implement Algorithm 2, we assume that the agents \mathcal{A} can form a bidirectional connected communication graph $\mathcal{G}^a = (\mathcal{A}, \mathcal{E}^a)$, i.e., there is a path from every agent to every other agent on \mathcal{G}^a . Then, there always exists a route $\text{SEQ} = s_1 \rightarrow \dots \rightarrow s_i \rightarrow \dots \rightarrow s_k, s_k \in \mathcal{A}, k \in \{1, \dots, K\}, K \geq |\mathcal{A}|$, that visits all the agents (not necessarily only one time), see Fig. 3(a). The

Algorithm 2 Decentralized Implementation of Sequential Greedy Algorithm

```

1: Init:  $\bar{\mathcal{P}} \leftarrow \emptyset$ ,  $i \leftarrow 1$ ,  $\{\bar{t}_v^0\}_{v \in \mathcal{V}}$ 
2: while  $i \leq K$  do
3:   if  $s_i$  is being called for the first time then
4:     agent  $s_i$  computes  $p^{s_i*} = \operatorname{argmax}_{p \in \mathcal{P}^{s_i}} \Delta_{\bar{\mathcal{P}}}(p|\bar{\mathcal{P}})$ .
5:    $\bar{\mathcal{P}} \leftarrow \bar{\mathcal{P}} \cup p^{s_i*}$ .
6:   end if
7:   agent  $s_i$  pass  $\bar{\mathcal{P}}$  to  $s_{i+1}$ .
8:    $i \leftarrow i + 1$ .
9: end while
10: agent  $s_K$  based on the execution plan of the receding horizon operation updates  $\{\bar{t}_v^0\}_{v \in \mathcal{V}}$  and communicates it to the team

```

agents follow SEQ to share their information while implementing Algorithm 2. The communication cost to execute Algorithm 2 can be optimized by picking SEQ to be the shortest path (Lawler, Lenstra, Kan, & Shmoys, 1985) that visits all the agents over graph \mathcal{G}^a . If \mathcal{G}^a has a Hamiltonian path, the optimal choice for SEQ is a Hamiltonian path. Recall that a Hamiltonian path is a path that visits every agent on \mathcal{G}^a only once (Thomas et al., 2009). When, there is a SEQ that visits every agent on \mathcal{G}^a , the directed information graph $\mathcal{G}^l = (\mathcal{A}, \mathcal{E}^l)$ of Algorithm 2, which shows the information access of each agent while implementing Algorithm 2, is full, see Fig. 3. That is, each agent in SEQ is aware of the previous agents' decision. Therefore, the solution obtained by Algorithm 2 is an exact sequential greedy algorithm and its optimality gap is 1/2. Recall that the labeling order of the mobile agents does not have an effect on the optimality gap guaranteed by Theorem 5 (Gharesifard & Smith, 2018). If an agent $i \in \mathcal{A}$ appears repeatedly in SEQ (e.g., the blue agent in Fig. 3), with a slight increase in computation cost, we can modify Algorithm 2 to allow agent i to redesign and improve its sub-optimal policy p^{i*} by re-executing step 4 of Algorithm 2.

Another form of decentralized implementation of Algorithm 1, which may be more relevant in urban environments, is through a client-server framework implemented over a cloud. In this framework, agents (clients) connect to shared memory on a cloud (server) to download or upload information or use the cloud's computing power asynchronously. Let $\{\mathcal{T}^i\}$, $i \in \mathcal{A}$, be the set of disjoint time slots that is allotted respectively to agents \mathcal{A} , see Fig. 4. To implement Algorithm 1, agent $i \in \mathcal{A}$ connects to the server at the beginning of \mathcal{T}^i to check out $\bar{\mathcal{P}}$ and $\{\bar{t}_v^0\}_{v \in \mathcal{V}}$. Then, it completes steps 4 and 5 of Algorithm 1, and checks in the updated $\bar{\mathcal{P}}$ to the server before \mathcal{T}^i elapses fully. The last agent based on the execution plan of the receding horizon operation updates $\{\bar{t}_v^0\}_{v \in \mathcal{V}}$ and checks it in the cloud memory for the next receding horizon planning. Since the time slots assigned to the agents do not overlap, agent i has access to policy p^{k*} of all agents k which has already communicated to the cloud. Thus, the information graph \mathcal{G}^l is full, and the optimality gap of 1/2 holds.

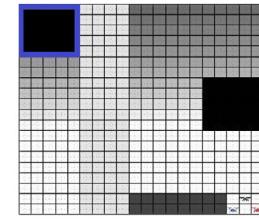
If there is a message dropout while executing Algorithm 2 or in the decentralized server-client based operation an agent j takes a longer time than \mathcal{T}^j to complete and check-in $\bar{\mathcal{P}}$ to the cloud, the information graph becomes incomplete, see for example Fig. 4. Then, the corresponding decentralized implementation deviates from the exact sequential greedy Algorithm 2. For such cases, Gharesifard and Smith (2018) show that the optimality gap instead of 1/2 becomes $\frac{1}{M - \omega(\mathcal{G}^l) + 2}$, where $\omega(\mathcal{G}^l)$ is the clique number of \mathcal{G}^l (Gharesifard & Smith, 2018). Recall that the clique number of a graph is equal to the number of the nodes in the largest sub-graph such that adding an edge will cause a cycle (Bondy, Murty, et al., 1976).



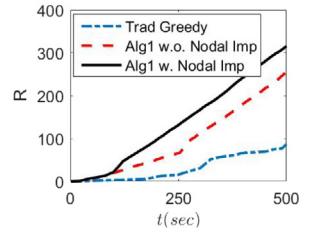
Fig. 3. The plot on the left shows the bi-directional communication graph \mathcal{G}^a in black along with an example SEQ path in red. The plot on the right shows the complete information sharing graph \mathcal{G}^l if agents follow SEQ while implementing Algorithm 2. Arrow going from agent i to agent j means that agent j receives agent i 's information. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)



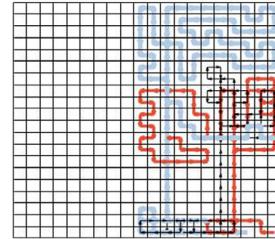
Fig. 4. $\{\mathcal{T}^i\}_{i \in \mathcal{A}}$, $\mathcal{A} = \{1, 2, 3, 4, 5\}$ are the time slots allotted to each agent to connect to the cloud. The arrows show the time each agent took to do their calculations for an example scenario. Here, the associated information graph \mathcal{G}^l is as the incomplete graph on the right with clique number of 3.



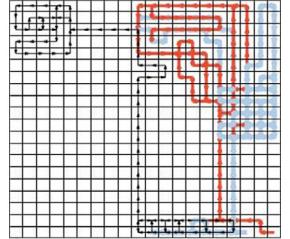
(a) Reward map



(b) The collected reward



(c) Agents' path when they follow Algorithm 1 and use $\alpha = 0$ over $[0, 150]$ seconds



(d) Agents' path when they follow Algorithm 1 and use $\alpha = 0.1$ over $[0, 150]$ seconds

Fig. 5. Three agents patrol a field, divided into 20 by 20 cells. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

4. Numerical example

We consider persistent monitoring using 3 agents for event detection over an area that is divided into 20 by 20 grid map as shown in Fig. 5(a). The geographical nodes of interest \mathcal{V} are the center of the cells in Fig. 5(a). The agents can travel from a cell to the neighboring cells in the right, left, bottom, and top. The agents are homogeneous and the travel time between any neighboring nodes for all the agents are identical and equal to 1 second. The agents start their patrolling task from the nodes where they are depicted in Fig. 5(a). We model the event occurrence in each geographical node as a Poisson process and define our reward function at each node $v \in \mathcal{V}$ as (1) with $\psi_v(t) = 1 - e^{-\lambda_v t}$ where $\lambda_v \in \mathbb{R}_{>0}$ is the arrival rate of the event; for more details see Rezazadeh and Kia (2019a). Fig. 5(a) shows the reward value of the nodes at $t = 120$ seconds when there is no monitoring. The color intensity of the cells in Fig. 5(a) is proportional to λ_v ; the higher λ_v , the darker the color v . The region enclosed by

the blue rectangle initially has a low reward but after 100 seconds its reward value is increased to a higher value by changing λ_v of the corresponding cells. An animated depiction of the change in the reward map because of the different dispatch policies we discuss below is available in Rezazadeh and Kia (2019b). We compare the performance of Algorithm 1, implemented in a receding horizon fashion, and a conventional greedy algorithm where each agent always moves to the neighboring node that has the instantaneous highest reward value. In implementing Algorithm 1 in a receding horizon fashion, we assume that the planning horizon is 4 seconds and the execution horizon is 1 second. We consider both the case of including ($\alpha = 0.1$) and excluding ($\alpha = 0$) the nodal importance measure in the reward function (5). Fig. 5(b) shows that the traditional greedy cell selection performs poorly compared to the other two planning algorithms. The reason is that the three agents' decision becomes the same after a while, i.e., they start choosing the same cell after a while and moving together, therefore all three agents act as if one agent is patrolling (recall Assumption 1). The performance of Algorithm 1 is better than a standard greedy cell selection because the effect of agent i 's patrolling policy is taken into account when agent $i + 1$ is designed. Therefore, the chances that all three agents go to the same cell together and move together is narrow. Furthermore, we can note that implementing Algorithm 1 by considering the effect of nodal importance delivers a better outcome. The reason is that in the case that there is no nodal importance, the agents are drawn to the region of high importance near them and stay there as Fig. 5(c) shows. However, there are other important regions with higher values that are farther away, especially the area on the left top corner which is separated by a low rate stripe from where agents start. Incorporating nodal importance, as Fig. 5(d) shows steers the agents to the regions with a higher rate of reward that are beyond the receding horizon's sight.

5. Conclusion

We presented a multi-agent dispatch policy design for persistent monitoring of a set of finite inter-connected geographical nodes. Our design relied on assigning an increasing and concave reward function of time to each node that reset to zero after a visit by an agent. We defined our design utility function as the sum of the rewards scored for the team when agents visit the geographical nodes. By showing that the utility function is a monotone increasing and submodular set function, we laid the ground to propose a suboptimal solution with a known optimality gap for our dispatch policy design, which was NP-hard. To induce robustness to the changes in the problem parameters, we proposed our suboptimal solution in a receding horizon setting. Next, to compensate for the shortsightedness of the receding horizon approach, we added a new term, called the relative nodal importance, to the utility function as a measure to incorporate a notion of the importance of the regions beyond the feasible solution set of the receding horizon optimization problem. Our numerical example demonstrated the benefit of introducing this term. Lastly, we discussed how our suboptimal solution can be implemented in a decentralized manner. Our future work is to investigate decentralized algorithms that allow agents to communicate synchronously with each other in order to have a consensus on a policy with a known optimality gap.

Appendix

Proof of Lemma 3. The time complexity of constructing the admissible policy set \mathcal{P}^i is of order of the number of possible paths that agent $i \in \mathcal{A}$ can traverse over the mission horizon

while respecting Assumption 2, which is of order D^{n^i} . Thus, the time complexity of constructing the feasible set $\mathcal{P} = \bigcup_{i \in \mathcal{A}} \mathcal{P}^i$ is $O(\sum_{i \in \mathcal{A}} D^{n^i})$. Next, let $\bar{\mathcal{P}}$ be any subset of \mathcal{P} that satisfies constraint (3b). Due to Assumption 1, the reward scored by implementing policy $p = (v_p, T_p, a_p) \in \bar{\mathcal{P}}$ cannot be calculated independent from the all the other policies in $\bar{\mathcal{P}} \setminus \{p\}$. Hence, to solve optimization problem (3a), we need to evaluate all the possible policy sets $\bar{\mathcal{P}}$ satisfying the constraint (3b). Since $\bar{\mathcal{P}}$ can have at most one policy from the policy set \mathcal{P}^i of $i \in \mathcal{A}$ and \mathcal{P}^i has $O(D^{n^i})$ members, then $O(\prod_{i \in \mathcal{A}} D^{n^i})$ different possibilities of $\bar{\mathcal{P}}$ exist which determines the time complexity of solving (3a). \square

We develop the auxiliary results below to use in the proof of Theorem 4. These results show some of the properties of the sum of evaluation of a concave and increasing function over increasing sequences and their concatenation. The decreasing sequence $(\delta t)_1^n$ majorizes the decreasing sequence $(\delta v)_1^n$, if $\delta t_1 \geq \delta t_2 \geq \dots \geq \delta t_n$, $\delta v_1 \geq \delta v_2 \geq \dots \geq \delta v_n$, $\delta t_1 + \dots + \delta t_i \geq \delta v_1 + \dots + \delta v_i$ for $i \in \{1, \dots, n-1\}$ and $\delta t_1 + \dots + \delta t_n = \delta v_1 + \dots + \delta v_n$ hold.

Lemma A.6. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a concave and increasing function with $f(0) = 0$. If sequences $(\delta t)_1^n$ and $(\delta v)_1^m$ with $n \leq m$ satisfy $\delta t_1 + \dots + \delta t_i \geq \delta v_1 + \dots + \delta v_i$, $\forall i \in \{1, \dots, n-1\}$ and $\delta t_1 + \dots + \delta t_n = \delta v_1 + \dots + \delta v_m$ then $f(\delta t_1) + \dots + f(\delta t_n) \leq f(\delta v_1) + \dots + f(\delta v_m)$ holds.

Proof. We note that the sequence $(\delta u)_1^m$ defined as $\delta u_i = \delta u_i$ for $i \in \{1, \dots, n\}$ and $\delta u_i = 0$ for $i \in \{n+1, \dots, m\}$ majorizes any sequence $(\delta v)_1^m$ defined in the lemma statement. Then, since $f(0) = 0$, the proof follows from the Karamata's inequality (Kadelburg et al., 2005).

Corollary A.7. Let $f : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ be a monotone increasing and concave function. Then for any $a, b, c, d \in \mathbb{R}_{\geq 0}$ such that $0 \leq a \leq c$ and $0 \leq b \leq d$, then $f(c) + f(d) - f(c+d) \leq f(a) + f(b) - f(a+b)$ holds.

Proof. Note that since $a \leq c$ and $b \leq d$, we have $a+b \leq c+d$. Let $(\delta t)_1^3 = (c, d, a+b)$ and $(\delta v)_1^3 = (a, b, c+d)$. Then, the proof follows from Lemma A.6. \square

Lemma A.8. For any $(q)_1^l$, let $g((q)_1^l) = \sum_{i=1}^{l-1} f(\Delta q_i)$, where $\Delta q_i = q_{i+1} - q_i$ and f be a concave and increasing function with $f(0) = 0$. Now, consider two increasing sequences $(t)_1^n$ and $(u)_1^l$, and their concatenation $(a)_1^{n+l} = (t)_1^n \oplus (u)_1^l$. Then, $g((a)_1^{n+l}) - g((t)_1^n) \geq 0$ holds.

Proof. If $a_p = t_1$ and $a_q = t_n$, then since $(a)_1^{n+l}$ is an increasing sequence, $p < q$. Let the sub-sequence of $(a)_1^{n+l}$ ranging from index p to q be $(v)_1^m$ where $m \geq n$. Letting $\Delta v_i = v_{i+1} - v_i$ and $\Delta t_i = t_{i+1} - t_i$, we rearrange Δv_i 's and Δt_i 's in a descending order to form the sequences $(\delta v)_1^{l-1}$ and $(\delta t)_1^{n-1}$. Since $a_p = t_1$ and $a_q = t_n$, we have $\sum_{i=1}^{m-1} \Delta v_i = \sum_{i=1}^{m-1} \delta v_i = \sum_{i=1}^{n-1} \delta t_i = \sum_{i=1}^{n-1} \Delta t_i = t_n - t_1$. Because $(a)_1^{n+l} = (t)_1^n \oplus (u)_1^l$, then $\forall i \in \{1, \dots, n\}$ there exists $S_i \subset \{1, \dots, m\}$ such that $\sum_{j \in S_i} \delta v_j = \delta t_i$, where $S_i \cap S_k = \emptyset$, $i \neq k$. Consequently, for $r \in \{1, \dots, m\}$, we have $\sum_{i=1}^r \delta v_i = \sum_{j \in S} \delta t_j$ for $S \subset \{1, \dots, n\}$ and $|S| \leq r$. Since $(\delta t)_1^{n-1}$ is a decreasing sequence, we can write $\sum_{i=1}^r \delta v_i \leq \sum_{i=1}^r \delta t_i$. Thus, $f(\delta t_1) + \dots + f(\delta t_{n-1}) \leq f(\delta v_1) + \dots + f(\delta v_{m-1})$ holds as a result of Lemma A.6. Given that $f(\delta t_1) + \dots + f(\delta t_{n-1}) = \sum_{i=1}^{n-1} f(\Delta t_i)$ and $f(\delta v_1) + \dots + f(\delta v_{m-1}) = \sum_{i=1}^{m-1} f(\Delta v_i) \leq \sum_{i=1}^{n-1} f(\Delta a_i)$, then $\sum_{i=1}^{n-1} f(\Delta t_i) \leq \sum_{i=1}^{n-1} f(\Delta a_i)$, which concludes the proof. \square

Lemma A.9. For any $(q)_1^l$, let $g((q)_1^l) = \sum_{i=1}^l f(\Delta q_i)$ where $\Delta q_i = q_{i+1} - q_i$ and f is a concave and increasing function with

$f(0) = 0$. Now, consider three increasing sequences $(t)_1^n$ and $(v)_1^m$ and $(u)_1^l$ and concatenations $(a)_1^{n+l} = (t)_1^n \oplus (u)_1^l$ and $(b)_1^{m+l} = (v)_1^m \oplus (u)_1^l$ where $(v)_1^m$ is a sub-sequence of $(t)_1^n$, then $(g((b)_1^{m+l}) - g((v)_1^m)) - (g((a)_1^{n+l}) - g((t)_1^n)) \geq 0$.

Proof. Let the sequence $(u)_1^p$ be the first p elements of $(u)_1^l$. Then, we can form $\Delta S_p = (g((v)_1^m \oplus (u)_1^p) - g((v)_1^m \oplus (u)_1^{p-1})) - (g((t)_1^n \oplus (u)_1^p) - g((t)_1^n \oplus (u)_1^{p-1}))$, where $(u)_1^0$ to be an empty sequence with no members. Since $(v)_1^m$ is a sub-sequence of $(t)_1^n$ and $(u)_1^p$ having one member more over $(u)_1^{p-1}$, then we have $\Delta S_p = (f(\Delta S_1) + f(\Delta S_2) - f(\Delta S_1 + \Delta S_2)) - (f(\Delta S_3) + f(\Delta S_4) - f(\Delta S_3 + \Delta S_4))$ with $0 \leq \Delta S_3 \leq \Delta S_1$ and $0 \leq \Delta S_4 \leq \Delta S_2$. From Corollary A.7, we can conclude that $\Delta S_p \geq 0$. Then, given $\sum_{p=1}^l \Delta S_p = (g((b)_1^{m+l}) - g((v)_1^m)) - (g((a)_1^{n+l}) - g((t)_1^n))$, the proof is concluded. \square

References

Almeida, A., Ramalho, G., Santana, H., Tedesco, P., Menezes, T., Corruble, V., et al. (2004). Recent advances on multi-agent patrolling. In *Brazilian symposium on artificial intelligence* (pp. 474–483). Springer.

Asghar, A., Smith, S., & Sundaram, S. (2019). Multi-robot routing for persistent monitoring with latency constraints. arXiv preprint arXiv:1903.06105.

Blum, A., Chalasani, P., Coppersmith, D., Pulleyblank, B., Raghavan, P., & Sudan, M. (1994). The minimum latency problem. In *Proceedings of the twenty-sixth annual ACM symposium on theory of computing* (pp. 163–171).

Bondy, J., Murty, U., et al. (1976). *Graph theory with applications* (vol. 290). Macmillan London.

Bullo, F., Carli, R., & Frasca, P. (2012). Gossip coverage control for robotic networks: Dynamical systems on the space of partitions. *SIAM Journal on Control and Optimization*, 50(1), 419–447.

Carron, A., Todescato, M., Carli, R., Schenato, L., & Pillonetto, G. (2015). Multi-agents adaptive estimation and coverage control using Gaussian regression. In *2015 European control conference* (pp. 2490–2495). IEEE.

Chevaleyre, Y. (2004). Theoretical analysis of the multi-agent patrolling problem. In *Intelligent agent technology* (pp. 302–308). IEEE.

Chung, Y., & Kia, S. S. (2020). A distributed service-matching coverage via heterogeneous mobile agents. Available at <https://arxiv.org/pdf/2009.11943.pdf>.

Clark, A., Bushnell, L., & Poovendran, R. (2014). A supermodular optimization framework for leader selection under link noise in linear multi-agent systems. *IEEE Transactions on Automatic Control*, 59(2), 283–296.

Clark, A., Lee, P., Alomair, B., Bushnell, L., & Poovendran, R. (2018). Combinatorial algorithms for control of biological regulatory networks. *IEEE Transactions on Control of Network Systems*, 5(2), 748–759.

Cortes, J., Martinez, S., Karatas, T., & Bullo, F. (2004). Coverage control for mobile sensing networks. *IEEE Transactions on Automatic Control*, 20(2), 243–255.

Donahue, M., Rosman, G., Kotowick, K., Rus, D., & Baykal, C. (2016). Persistent surveillance of transient events with unknown statistics: Tech. rep., MIT Lincoln Laboratory Lexington United States.

Duchoň, F., Babinec, A., Kajan, M., Beňo, P., Florek, M., Fico, T., et al. (2014). Path planning with modified a star algorithm for a mobile robot. *Procedia Engineering*, 96, 59–69.

Engler, R., Guisan, A., & Rechsteiner, L. (2004). An improved approach for predicting the distribution of rare and endangered species from occurrence and pseudo-absence data. *Journal of Applied Ecology*, 41(2), 263–274.

Farinelli, A., Iocchi, L., & Nardi, D. (2017). Distributed on-line dynamic task assignment for multi-robot patrolling. *Autonomous Robots*, 41(6), 1321–1345.

Fisher, L., Nemhauser, G., & Wolsey, L. (1978). An analysis of approximations for maximizing submodular set functions—II. In *Polyhedral combinatorics* (pp. 73–87). Springer.

Garcia, E., Prett, D., & Morari, M. (1989). Model predictive control: Theory and practice—a survey. *Automatica*, 25(3), 335–348.

Gharesifard, B., & Smith, S. (2018). Distributed submodular maximization with limited information. *IEEE Transactions on Control of Network Systems*, 5(4), 1635–1645.

Henry, N., & Henry, O. (2015). Wireless sensor networks based pipeline vandalism and oil spillage monitoring and detection: Main benefits for nigerian oil and gas sectors. *The SIJ Transactions on Computer Science Engineering & its Applications (CSEA)*, 3(1), 1–6.

Kadelburg, Z., Dukic, D., Lukic, M., & Matic, I. (2005). Inequalities of Karamata, Schur and Muirhead, and some applications. *The Teaching of Mathematics*, 8(1), 31–45.

Karp, R. (1972). Reducibility among combinatorial problems. In *Complexity of computer computations* (pp. 85–103). Springer.

Krause, A., & Guestrin, C. (2007). Near-optimal observation selection using submodular functions. In *American association for artificial intelligence (vol. 7)* (pp. 1650–1654).

Krause, A., Singh, A., & Guestrin, C. (2008). Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9(Feb), 235–284.

Lawler, E., Lenstra, J., Kan, A. H. R., & Shmoys, D. (1985). *The traveling salesman problem; a guided tour of combinatorial optimization*. Chichester: Wiley.

Liu, Z., Clark, A., Lee, P., Bushnell, L., Kirschen, D., & Poovendran, R. (2018). Submodular optimization for voltage control. *IEEE Transactions on Power Systems*, 33(1), 502–513.

Lovász, L. (1983). Submodular functions and convexity. In *Mathematical programming the state of the art* (pp. 235–257). Springer.

Machado, A., Ramalho, G., Zucker, J., & Drogoul, A. (2002). Multi-agent patrolling: An empirical analysis of alternative architectures. In *International workshop on multi-agent systems and agent-based simulation* (pp. 155–170).

Pasqualetti, F., Franchi, A., & Bullo, F. (2012). On cooperative patrolling: Optimal trajectories, complexity analysis, and approximation algorithms. *IEEE Transactions on Robotics*, 28(3), 592–606.

Rezazadeh, N., & Kia, S. S. (2019a). A sub-modular receding horizon approach to persistent monitoring for a group of mobile agents over an urban area. *IFAC-PapersOnLine*, 52(20), 217–222.

Rezazadeh, N., & Kia, S. S. (2019b). A sub-modular receding horizon solution for mobile multi-agent persistent monitoring: Animated numerical example. <https://youtu.be/8NE28UjyLOQ>.

Schwager, M., Rus, D., & Slotine, J. (2009). Decentralized, adaptive coverage control for networked robots. *International Journal of Robotics Research*, 28(3), 357–375.

Summers, T., Cortesi, F., & Lygeros, J. (2016). On submodularity and controllability in complex dynamical networks. *IEEE Transactions on Control of Network Systems*, 3(1), 91–101.

Thomas, T., & van Berkum, E. (2009). Detection of incidents and events in urban networks. *IET Intelligent Transport Systems*, 3(2), 198–205.

Thomas, T., Leiserson, C., Rivest, R., & Stein, C. (2009). *Introduction to algorithms*. MIT Press.

Todescato, M., Carron, A., Carli, R., Pillonetto, G., & Schenato, L. (2017). Multi-robots Gaussian estimation and coverage control: From client-server to peer-to-peer architectures. *Automatica*, 80, 284–294.

Yu, J., Karaman, S., & Rus, D. (2015). Persistent monitoring of events with stochastic arrivals at multiple stations. *IEEE Transactions on Robotics*, 31(3), 521–535.

Yuan, C., Zhang, Y., & Liu, Z. (2015). A survey on technologies for automatic forest fire monitoring, detection, and fighting using unmanned aerial vehicles and remote sensing techniques. *Canadian Journal of Forest Research*, 45(7), 783–792.



Navid Rezazadeh is a Ph.D. candidate in the Department of Mechanical and Aerospace Engineering, University of California, Irvine (UCI). He obtained his M.Sc. in Mechanical and Aerospace Engineering in 2017 from UCI, and his B.Sc. in Mechanical engineering in 2014 from Sharif University of Technology. His research interests include distributed systems, distributed optimization, discrete optimization, probabilistic learning and privacy.



Solmaz S. Kia is an Associate Professor in the Department of Mechanical and Aerospace Engineering, University of California, Irvine (UCI). She has a joint appointment with the Computer Science Department of UCI. She obtained her Ph.D. degree in Mechanical and Aerospace Engineering from UCI, in 2009, and her M.Sc. and B.Sc. in Aerospace Engineering from the Sharif University of Technology, Iran, in 2004 and 2001, respectively. She was a senior research engineer at SySense Inc., El Segundo, CA from Jun. 2009 Sep. 2010. She held postdoctoral positions in the Department of Mechanical and Aerospace Engineering at the UC San Diego and UCI. She was the recipient of UC President's Postdoctoral Fellowship from 2012–2014. She is also a recipient of 2017 NSF CAREER award. Dr. Kia is a senior member of IEEE, and serves as associate editor of IEEE Sensors Letters. Dr. Kia's main research interests, in a broad sense, include nonlinear control theory, distributed optimization/coordination/estimation, and probabilistic robotics.