

Asynchronous and Dynamic Coverage Control Scheme for Persistent Surveillance Missions

Jeffrey R. Peters, Sean Wang, Amit Surana, and Francesco Bullo

Abstract—A *decomposition-based* coverage control scheme is proposed for multi-agent, persistent surveillance missions operating in a communication-constrained, dynamic environment. The proposed approach decouples high-level task assignment from low-level motion planning in a modular framework. Coverage assignments and surveillance parameters are managed by a central base station, and transmitted to mobile agents via unplanned and asynchronous exchanges. Coverage updates promote load balancing, while maintaining geometric and temporal characteristics that allow effective pairing with generic path planners. Namely, the proposed scheme guarantees that (i) coverage regions are connected and collectively cover the environment, (ii) subregions may only go uncovered for bounded periods of time, (iii) collisions (or sensing overlaps) are inherently avoided, and (iv) under static event likelihoods, the collective coverage regions converge to a Pareto-optimal configuration. This management scheme is then paired with a generic path planner satisfying loose assumptions. The scheme is illustrated through simulated surveillance missions.

I. INTRODUCTION

A. Decomposition-Based Multi-Agent Surveillance

Modern exploratory and surveillance missions often utilize autonomous vehicles or sensors to observe and monitor large geographic areas. Example application domains include search and rescue [14], environmental monitoring [32], warehouse logistics [37], and military reconnaissance [7]. Such scenarios require robust and flexible tools for autonomous coordination.

A number of planning strategies have been studied for single agents, ranging from simple *a priori* tour construction [23] to more complex methods involving Markov chains [34], optimization [17], or Fourier analysis [16]. Unfortunately, it is not straightforward to generalize single-agent strategies for use in multi-agent missions: Naive approaches where each agent follows an independent policy can result in poor performance and introduce collision risks, while sophisticated generalizations often require joint optimizations that are intractable for even modestly sized problems. Scaling issues are sometimes

This work has been sponsored by the U.S. Army Research Office and the Regents of the University of California, through Contract Number W911NF-09-D-0001 for the Institute for Collaborative Biotechnologies, and that the content of the information does not necessarily reflect the position or the policy of the Government or the Regents of the University of California, and no official endorsement should be inferred.

Jeffrey R. Peters and Sean Wang are with the Mechanical Engineering Department and the Center for Control, Dynamical Systems and Computation, University of California, Santa Barbara jrpeters@engr.ucsb.edu, seanwang@umail.ucsb.edu

Amit Surana is with the Systems Department at United Technologies Research Center, East Hartford, CT, SuranaA@utrc.utc.com

Francesco Bullo is with the Mechanical Engineering Department and the Center for Control, Dynamical Systems and Computation, University of California, Santa Barbara bullo@engr.ucsb.edu

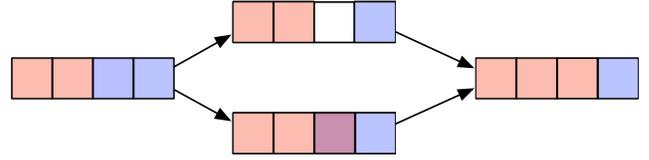


Fig. 1. Difficulties arise when complete or pairwise updates are impossible. With only single-region updates, two steps are required to move from the left-most to the right-most configuration. In the top path, blue is updated first, leaving an uncovered block at the intermediate step. In the bottom path, red is updated first, leaving a block that belongs to both regions simultaneously.

alleviated through distributed control; however, such setups are application specific and may require extensive efforts to pose a mathematical problem that is suitable for use with formal techniques [19]. In contrast, *decomposition-based* approaches, which decouple the assignment and routing problem by first dividing the workspace among agents, offer a straightforward, modular framework to reasonably accomplish the desired goals, despite sacrificing optimality in general.

Communication constraints can, however, make it difficult to effectively divide a dynamic workspace in real time, while still ensuring that the result is amenable for pairing with single-agent motion planners. For example, many unmanned missions require agents to transfer sensor data to a central base station for analysis. With difficult terrain or hardware limitations, these sporadic exchanges with the base station may provide the only means of sharing real-time information across agents. Applications that operate under this constraint include underwater gliders that must surface to communicate with a tower [28], data mules that periodically visit ground robots [31], and supervisory missions where human operators analyze data collected by unmanned vehicles [29]. Here, updated mission information can only be relayed to one agent at a time, rendering traditional dynamic partitioning schemes, which rely on complete or pairwise coverage updates, impossible. As such, designers are forced to either allow for overlap in coverage assignments, or sacrifice complete coverage (Fig. 1).

This work presents a decomposition-based surveillance framework that operates under asynchronous *one-to-base station* communication, a protocol in which data is transferred solely via sporadic exchanges between a central base station and autonomous agents [26]. The complete scheme consists of two components: a dynamic partitioning component and a single-agent routing component. The work herein focuses primarily on developing the partitioning component, which, in addition to dividing the workload among the agents, ma-

nipulates local variables in order to allow for effective pairing with single-agent route planners. This development naturally leads to a complete multi-agent framework in which high-level coverage is coordinated by the base station, while individual trajectories are generated independently via on-board planners.

B. Related Literature

Research relating to multi-agent coverage control is vast. Typical strategies involve optimization [27], auctions [10], biological meta-heuristics [4], potential fields [30], or space decomposition [20]. Of particular relevance are multi-agent *persistent surveillance* (*persistent monitoring*) problems, in which a mobile sensor team is tasked with continual surveillance of a region of interest, requiring each subregion to be visited multiple (or infinitely many) times with the goal of minimizing a cost, e.g., the time between visits or the likelihood of detecting stochastic events [19]. Persistent surveillance is a generalization of *patrolling*, where closed tours are sought for the purpose of protecting or supervising an environment. Patrolling has been studied extensively, with most current solutions being based on operations research algorithms, non-learning multi-agent systems, and multi-agent learning [1]. Patrolling formulations are often one-dimensional and solutions usually reduce to “back and forth” motions that do not readily extend to general scenarios (see, e.g. [24]).

The framework proposed herein is based on *workspace decomposition*, and uses an allocation strategy to reduce the multi-agent problem into a set of single-agent problems. This approach is common in multi-agent systems due to its simplicity and scalability [19]. For persistent surveillance of planar regions, decomposition-based approaches consist of two primary components: partitioning and single-agent routing. The most common approaches to optimal partitioning are based on Voronoi partitions [21]. Effective schemes exist for constructing centroidal Voronoi, equitable, or other types of optimal partitions under various communication, sensing, and workload constraints [5], [6], [25]. These strategies usually require convex environments. More general workspaces are usually addressed by discretizing the environment and considering the resulting graph, on which any number of graph partitioning schemes can be used [9]. In robotic contexts, the discrete partitioning problem is often considered under communication constraints, e.g., pairwise gossip [8] or asynchronous one-to-base station communication [26]. The proposed partitioning scheme most closely mirrors [26]; however, our approach employs additional logic to ensure the resultant coverage regions are amenable for use in decomposition-based surveillance.

Single-agent path planners for persistent surveillance most commonly operate over discrete environments, i.e., graphs [13], [35], and solutions to classical problems, e.g., the Traveling Salesperson Problem [11], may suffice. Stochasticity can be introduced using tools such as Markov chains [34]. Persistent surveillance strategies for non-discrete regions (in particular, open subsets of Euclidean space), are less common. Here, strategies include the *a priori* construction of motion routines [22], the adaptation of static coverage strategies [33], the use of random fields [12], and spectral decomposition [16].

The modular framework herein does not require any particular single-agent route planner, but rather can incorporate any planner satisfying a mild assumption set (see Section V).

Despite the vast research on both topics individually, remarkably few papers explicitly consider the implications of combining geometric partitioning with continuous routing in the context of multi-agent persistent surveillance. Research that does exist is mostly preliminary, considering ideal communication and employing simplistic methods. For example, the authors of [2] employ a sweeping algorithm for partitioning and guide vehicle motion via lawn-mower patterns. The authors of [20] use rectangular partitions, while employing a reactive routing policy which, in ideal cases, reduces to spiral search patterns. The work in [18] uses slightly more sophisticated partitioning in tandem with lawn-mower motion trajectories. Also relevant is [36], where partitions are based on the statistical expectation of target presence; however, ideal communication is assumed. Other works, e.g. [3], employ decomposition-based structures, but focus on task-assignment with no detailed treatment of the combined assignment/routing protocol.

C. Contributions

This work presents a decomposition-based, multi-agent coverage control framework for persistent surveillance, which requires only sporadic, unscheduled exchanges between agents and a central base station. In particular, we focus on developing a sophisticated partitioning and coordination scheme that is designed for pairing with generic single-agent trajectory planners within a modular framework. Our setup encompasses realistic constraints including restrictive communication, dynamic environments, and non-parametric event likelihoods.

Specifically, we develop a dynamic partitioning scheme that assumes only asynchronous, *one-to-base station* communication. This algorithm governs region assignments, while also introducing timing variables and manipulating the agents’ high-level surveillance parameters. We prove that our partitioning strategy has properties that make it amenable for use in a decomposition-based surveillance scheme: the produced coverage regions collectively form a connected m -covering, local likelihood functions have disjoint support, no subregion remains uncovered indefinitely, coverage regions evolve at a time-scale that allows for appropriate agent reactions, among others. For static likelihoods under mild assumptions, we show that the set of coverage regions and associated generators converges to a Pareto optimal partition in finite time. We combine our partitioning scheme with a generic single-agent trajectory, and show that this combination guarantees collision avoidance (no sensing overlap) when the trajectory planner obeys natural restrictions. We illustrate our framework through numerical examples.

For clarity and readability in what follows, we postpone all Theorem proofs until the Appendix.

II. MISSION OBJECTIVES AND SOLUTION APPROACH

A team of m mobile agents, each equipped with an on-board sensor, is tasked with endlessly monitoring a non-trivial, planar, large region of interest. The primary goal of

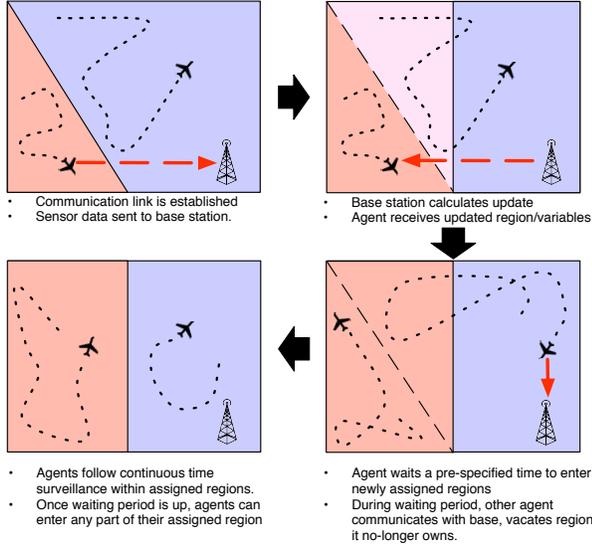


Fig. 2. An example illustrating the proposed strategy. The dynamic partitioning component (executed by the base station) manages coverage regions and introduces logic to prevent collisions, while the routing component (executed onboard each agent) governs agent motion.

the mission is to collect sensor data about some key dynamic event or characteristic, e.g., an intruder. Data collected by the agents is periodically relayed to a central base station for analysis. Agents must move within the region to obtain complete sensory information. Ideally, the agent motion/data collection strategy should be coordinated so that:

- 1) the workload is balanced across agents,
- 2) no subregion goes unobserved indefinitely,
- 3) agents never collide (have sensor overlap), and
- 4) the search is biased toward regions of greater interest.

To achieve these goals, we adopt a decomposition-based approach in which each agent’s motion is restricted to lie within a dynamically assigned *coverage region*.

The analysis herein focuses primarily on developing an algorithmic framework for partitioning that defines, maintains, and updates the agents’ coverage assignments in a manner that can be effectively paired with a typical single-agent route planner. More specifically, we seek a scheme to assign coverage responsibilities, as well as provide high-level restrictions on agent motion, so that lower-level, on-board trajectory planners can operate independently within the proposed boundaries to achieve global surveillance goals. In addition, we seek algorithms that do not require peer-to-peer communication or require the agents to be in constant contact with the base station. As such, we assume only asynchronous, *one-to-base station* communication (data transfer) [26], i.e., agents are only assumed to sporadically communicate with a central base station, subject to an upper bound on inter-communication times, which are not specified *a priori*.

Broadly, our proposed solution operates as follows (see Fig. 2). During each exchange with an agent, the base-station calculates the new coverage region solely for the communicating agent, and transmits the update. Updates are performed so as to guarantee connected coverage regions,

and, in the limit for static likelihoods, distribute the coverage load among the agents. During the exchange, the base station also updates timing variables and governs the local likelihood functions maintained by each agent so that, if no agent were ever located outside of the support of its local likelihood at any time, then the agents (sensors) would never be colocated. Timing considerations also guarantee that 1) no region remains uncovered, i.e., outside of the allowable area to which the agents can travel, for an unbounded amount of time, and 2) ample time is allotted for agents to vacate regions that are re-assigned before the newly assigned agent takes over. Once the update is complete, the link with the base is terminated and the agent follows the trajectory found via its onboard planner.

III. PROBLEM SETUP

The base station, as well as each agent, has its own local processor. “Global” information is stored by the base station, while each agent only stores information pertinent to itself.

Convention 1. *The subscripts i, j , or ℓ denote an entity relevant to agent i, j , or ℓ , resp. When attached to a set, the subscripts i, j, ℓ indicate the component relevant to agent i, j or ℓ , resp. The superscript ‘A’ indicates an entity that is stored on the agent’s local processor.*

A summary of the variables stored by relevant entities is shown in Table I. We expand on these, and define other relevant mathematical constructs here.

A. Agent Dynamics

Each agent (sensor) i is modeled as a point mass that moves with maximum possible speed $s_i > 0$. Define $s := \{s_i\}_{i=1}^m$.

B. Communication Protocol

There is a central base station with which the agents communicate. We assume the following:

- 1) each agent can identify itself to the base station and transmit sensor data,
- 2) there is a lower bound $\underline{\Delta} > 0$ on the time between any two successive exchanges involving the base station, and
- 3) there is an upper bound $\overline{\Delta} > 0$ on the time between any *single agent’s* successive exchanges with the base.

Note that 2) implies that no two agents can communicate with the base station simultaneously¹. These communication specifications collectively define the *one-to-base station* protocol.

C. Surveillance Region

Consider the surveillance region as a finite grid of disjoint, non-empty, simply-connected subregions. We represent the environment as a weighted graph $G(Q) := (Q, \mathcal{E})$, where Q is the set of vertices (each representative of a unique grid element), and \mathcal{E} is the edge set comprised of undirected, weighted edges $\{k_1, k_2\}$ spanning vertices representing *adjacent regions*². The weight assigned to $\{k_1, k_2\}$ is the distance

¹This constraint also prevents *zeno behavior* in the communication times

²Travel between the regions without entering other subregions is possible

between the appropriate regions. We assume without loss of generality that $Q \subset \mathbb{N}$. Environmental information, e.g., terrain, is largely unknown at the mission onset; however, the location of obstacles and prohibited areas, is known *a priori* and is not included in the graphical representation $G(Q)$.

Consider $Q' \subseteq Q$. A vertex $k_1 \in Q$ is *adjacent* to Q' if $k_1 \notin Q'$ and there exists $\{k_1, k_2\} \in \mathcal{E}$ with $k_2 \in Q'$. Define $G(Q') := (Q', \mathcal{E}')$, where $\mathcal{E}' := \{\{k_1, k_2\} \in \mathcal{E} \mid k_1, k_2 \in Q'\}$. A *path* on $G(Q')$ between $k_1, k_n \in Q'$ is a sequence (k_1, k_2, \dots, k_n) , where $k_1, k_2, \dots, k_n \in Q'$ and $\{k_r, k_{r+1}\} \in \mathcal{E}'$ for $r \in \{1, \dots, n-1\}$. We say Q' is *connected* if a path exists in $G(Q')$ between any $k_1, k_2 \in Q'$. Let $d_{Q'} : Q' \times Q' \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be the standard distance on $G(Q')$, i.e., the length of a shortest weighted path in $G(Q')$ (if none exists, $d_{Q'}$ takes value ∞). Notice $d_{Q'}(k_1, k_2) \leq d_Q(k_1, k_2)$ for any $k_1, k_2 \in Q'$. Also let $d_{Q'}$ denote the map $d_{Q'} : Q' \times 2^{Q'} \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$, where $d_{Q'}(k, Q'')$ is the length of a shortest weighted path in $G(Q')$ between k and any vertex in Q'' .

D. Coverage Regions

An *m-covering* of Q is a family $P = \{P_i \subseteq Q\}_{i=1}^m$ satisfying 1) $\bigcup_{i=1}^m P_i = Q$, and 2) $P_i \neq \emptyset$ for all i . Define $\text{Cov}_m(Q)$ as the set of all possible *m-coverings* of Q . An *m-partition* of Q is an *m-covering* that also satisfies 3) $P_i \cap P_j = \emptyset$, if $i \neq j$. An *m-covering* or *m-partition* P is *connected* if each P_i is connected. In what follows, the base station maintains an *m-covering* P of Q , and surveillance responsibilities are assigned by pairing each agent i with $P_i \in P$ (called agent i 's *coverage region*). Each agent maintains a copy P_i^A of P_i . The base station also stores a set $c := \{c_i \in Q\}_{i=1}^m$ (c_i is the *generator* of P_i), and each agent i maintains a copy c_i^A of c_i .

E. Identifiers, Timers, and Auxiliary Variables

The proposed algorithms introduce logic and timing considerations to ensure an effective overall framework. To each $k \in Q$, assign an *identifier* $\text{ID}_k \in \{1, \dots, m\}$. Define $\mathcal{ID} := \{\text{ID}_k\}_{k=1}^{|Q|}$, and let $P^{\mathcal{ID}} := \{P_i^{\mathcal{ID}}\}_{i=1}^m$, where $P_i^{\mathcal{ID}} := \{k \in Q \mid \text{ID}_k = i\}$. Notice $P^{\mathcal{ID}}$ is an *m-partition* of Q . For each agent i , define a timer T_i having dynamics $\dot{T}_i = -1$ if $T_i \neq 0$, and $\dot{T}_i = 0$ otherwise. The set $T := \{T_i\}_{i=1}^m$ is maintained by the base station. Similarly, each agent i maintains a local timing variable τ_i^A . Even though τ_i^A plays a similar role to the timer T_i , note that τ_i^A is constant unless explicitly changed by the algorithm, while T_i has autonomous dynamics. Next, the base station maintains a set $\omega := \{\omega_i\}_{i=1}^m$, where ω_i is the time of agent i 's last exchange with the base station. Each agent maintains a copy ω_i^A of ω_i . Finally, each agent stores locally a subset $P_i^{A,\text{pd}} \subseteq P_i^A$ which, loosely, collects vertices that have recently been added to P_i^A .

F. Likelihood Functions

A real-time estimate of the likelihood of events of interest occurring within any subset of the surveillance region is maintained by the base station in the form of a time-varying probability mass function³ $\Phi : Q \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$. Since it is

³For any t , we have $\sum_{k \in Q} \Phi(k, t) = 1$

TABLE I
STORAGE SUMMARY

Stored by Base Station	
Variable	Description
$G(Q)$	Graphical representation of the environment
P	<i>m-covering</i> of Q ($P \in \text{Cov}_m(Q)$)
c	Set of generators ($c \in Q^m$)
\mathcal{ID}	Set of identifiers ($\mathcal{ID} \in \{1, \dots, m\}^{ Q }$)
T	Set of Timers ($T(t) \in \mathbb{R}_{\geq 0}^m$)
ω	Set of most recent communication times ($\omega \in \mathbb{R}_{\geq 0}^m$)
Φ	Global likelihood function ($\Phi : Q \times \mathbb{R}_{> 0} \rightarrow \mathbb{R}_{> 0}$)
Stored by Agent i	
Variable	Description
$G(Q)$	Graphical representation of the environment
P_i^A	Coverage region ($P_i^A \subseteq Q$)
c_i^A	Coverage region generator ($c_i^A \in Q$)
$P_i^{A,\text{pd}}$	Set of "recently added" vertices ($P_i^{A,\text{pd}} \subseteq P_i^A$)
τ_i^A	Local timing parameter ($\tau_i^A \in \mathbb{R}$)
ω_i^A	Agent i 's most recent communication time ($\omega_i^A \in \mathbb{R}_{\geq 0}$)
Φ_i^A	Local likelihood function ($\Phi_i^A : Q \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$)

difficult to precisely define load-balancing requirements over arbitrarily time-varying environments, our analysis focuses on achieving optimal performance in the case of a static environment, i.e., when Φ is time-invariant. We then show through simulation that our proposed strategy still can be effectively employed in the case of a *quasi-static* environment, i.e., when $\Phi(k, \cdot)$ is piecewise constant for any k . This situation arises in many common scenarios, e.g., when estimates of Φ are only updated during exchanges between the base station and the mobile agents, which occur at discrete time-points.

Define each agent's *local likelihood* $\Phi_i^A : Q \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ as the function that, loosely, represents the agent's local belief regarding events of interest occurring in Q . The partitioning algorithm will manipulate the support of each Φ_i^A to ensure effective surveillance while also inherently preventing agent collisions (sensor redundancy). Specifically, define

$$\Phi_i^A(k, t) = \begin{cases} \Phi(k, t), & \text{if } k \in P_i^A \text{ and} \\ & (t - \omega_i^A \geq \tau_i^A \text{ or } k \notin P_i^{A,\text{pd}}), \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

In general, each function Φ_i^A will be different⁴ from Φ .

Remark 1 (Global Data). *If global knowledge of Φ is not available instantaneously to agent i , Φ_i^A can alternatively be defined by replacing $\Phi(k, t)$ in (1) by $\Phi(k, \omega_i^A)$. All subsequent theorems hold under this alternative definition.*

Remark 2 (Data Storage). *The cost of storing a graph as an adjacency list is $O(|Q| + |\mathcal{E}|)$. The generator set c , each element of P , and the identifier set \mathcal{ID} are stored as a integral*

⁴Note that Φ_i^A need not be normalized and thus may not be a time-varying probability mass function in a strict sense

vectors. The timer set T and the set ω are stored as real vectors. Typically, only the values of Φ at vertices in Q are needed, so Φ is stored as a time-varying real vector. Thus, the cost of storage at the base station is $O(m|Q| + |\mathcal{E}|)$. Similarly, each agent's local storage cost is $O(|Q| + |\mathcal{E}|)$.

IV. DYNAMIC COVERAGE UPDATE SCHEME

We adopt following convention for the remaining analysis.

Convention 2. Suppose that:

- 1) $\min \emptyset := \max \emptyset := 0$, and
- 2) when referring to a specific time instant, e.g., the instant when an update is executed, the notation e^+ and e^- refers to the value of the entity e immediately before and after the instant in question, resp.

A. Additive Subset

We start with a definition.

Definition 1 (Additive Subset). Given $k \in P_i^{\mathcal{I}\mathcal{D}}$, the additive subset $P_i^{\text{add}}(k) \subseteq Q$ is the largest connected subset satisfying:

- 1) $P_i^{\mathcal{I}\mathcal{D}} \subseteq P_i^{\text{add}}(k)$, and
- 2) for any $h \in P_i^{\text{add}}(k) \cap P_j$, where $j \neq i$:
 - a) $T_j = 0$, and
 - b) $\frac{1}{s_i} d_{P_i^{\text{add}}(k)}(h, k) < \frac{1}{s_j} d_{P_j}(h, c_j)$.

The following characterizes well-posedness of Definition 1.

Proposition 1 (Well-Posedness). If $P_i^{\mathcal{I}\mathcal{D}}$ is connected and disjoint from $\bigcup_{j \neq i} P_j$, then $P_i^{\text{add}}(k)$ exists and is unique for any $k \in P_i^{\mathcal{I}\mathcal{D}}$.

Proof. With the specified conditions, $P_i^{\mathcal{I}\mathcal{D}}$ is connected and satisfies conditions 1-3 in Definition 1. Thus, $P_i^{\text{add}}(k)$ is the unique, maximally connected superset of $P_i^{\mathcal{I}\mathcal{D}}$ satisfying the same conditions. \square

Under the conditions of Proposition 1, if $h \in P_i^{\text{add}}(k)$, then $\max\{T_j | j \neq i, h \in P_j\} = 0$ and there is a path from k to h in $G(P_i^{\text{add}}(k))$ that is shorter than the optimal path spanning c_j and h within $G(P_j)$, for any $j \neq i$ with $h \in P_j$. Also note that, by definition, additive subsets are connected.

B. Primary Update Algorithm

Define $\mathcal{H} : Q^m \times \text{Cov}_m(Q) \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ by

$$\mathcal{H}(c, P, t) = \sum_{k \in Q} \min \left\{ \frac{1}{s_i} d_{P_i}(k, c_i) \mid k \in P_i \right\} \Phi(k, t),$$

where Φ is the global likelihood. If (i) each agent is solely responsible for events within its own coverage region and (ii) events occur proportionally to Φ , then $\mathcal{H}(c, P, t)$ is understood as the expected time required for an appropriate agent to reach a randomly occurring event from its region generator at time t . Algorithm 1 defines the operations performed by the base station when agent i communicates at time t_0 . Here, the input $\Delta_H > 0$ is a constant mission-specific parameter.

Consider the following initialization assumptions.

Algorithm 1: Timed One-To-Base Station Update

Data: $t_0, P, c, \Phi, \omega, \bar{\Delta}, \Delta_H, T, \mathcal{I}\mathcal{D}, s$
Result: $P, c, P_i^A, c_i^A, P_i^{A,\text{pd}}, T, \tau_i^A, \Phi_i^A, \omega, \omega_i^A, \mathcal{I}\mathcal{D}$
begin

- 1 Initialize $P^* = P^{\text{test}} = P, c^* = c^{\text{test}} = c$
- 2 Set $P_i^* = P_i^{\text{test}} = P_i^{\mathcal{I}\mathcal{D}}$
- if** $T_i > 0$ **and** $P_i^* = P_i$ **then**
- 3 Set $\tau_i^A = \tau_i^A - t_0 + \omega_i$ **and** $\omega_i^A = \omega_i = t_0$
- else**
- for** $k \in P_i^{\mathcal{I}\mathcal{D}}$ **do**
- 4 Set $P_i^{\text{test}} = P_i^{\text{add}}(k)$ **and** $c_i^{\text{test}} = k$
- if** $\mathcal{H}(c_i^{\text{test}}, P_i^{\text{test}}, t) < \mathcal{H}(c^*, P^*, t)$ **then**
- 5 Set $P^* = P_i^{\text{test}}, c^* = c_i^{\text{test}}$
- 6 Set $P_i^{A,\text{pd}} = P_i^* \setminus P_i^{\mathcal{I}\mathcal{D}}$
- 7 Call Alg. 2 and obtain output $\Phi_i^A, \omega, T, \tau_i^A$
- 8 Set $P_i = P_i^A = P_i^*, c_i = c_i^A = c_i^*, \omega_i^A = \omega_i$
- 9 **for** $k \in P_i$ **do** Set $\text{ID}_k = i$
- 10 **return** $P, c, P_i^A, c_i^A, P_i^{A,\text{pd}}, T, \tau_i^A, \Phi_i^A, \omega, \omega_i^A, \mathcal{I}\mathcal{D}$

Algorithm 2: Timer Update

Data: $t_0, P, P^*, c^*, P_i^{A,\text{pd}}, \Phi, \omega, \bar{\Delta}, \Delta_H, T, s$
Result: $\Phi_i^A, \omega, T, \tau_i^A$
begin

- 1 $\Delta_i^{\text{Bf}} := \max \left\{ \frac{1}{s_i} d_{P_i}(k, P_i^* \setminus P_i^{A,\text{pd}}) \mid k \in P_i \setminus P_i^* \right\}$
- for** Each $j \neq i$ satisfying $P_j \cap P_i^* \neq \emptyset$ **do**
- 2 $\Delta_j^{\text{Bf}} := \max \left\{ \frac{1}{s_j} d_{P_j}(k, P_j \setminus P_i^*) \mid k \in P_j \cap P_i^* \right\}$
- 3 Set $T_j = \omega_j + \bar{\Delta} - t_0$
- 4 Find $\Delta_{\max}^{\text{Bf}} = \max_{j \neq i, P_j \cap P_i^* \neq \emptyset} \{\omega_j + \bar{\Delta} + \Delta_j^{\text{Bf}} - t_0\}$
- 5 Redefine $\Delta_{\max}^{\text{Bf}} = \max\{\Delta_{\max}^{\text{Bf}}, \Delta_i^{\text{Bf}}\}$
- 6 Set $T_i = \Delta_{\max}^{\text{Bf}} + \Delta_H, \tau_i^A = \Delta_{\max}^{\text{Bf}}, \omega_i = t_0$
- 7 Construct Φ_i^A , according to (1) with updated variables
- 8 **return** $\Phi_i^A, \omega, T, \tau_i^A$

Assumption 1 (Initialization). The following properties are satisfied when $t = 0$:

- 1) P is a connected m -partition of Q ,
- 2) $P = P^A = P^{\mathcal{I}\mathcal{D}}$, and
- 3) for all $i \in \{1, \dots, m\}$,
 - a) $c_i = c_i^A \in P_i^A$,
 - b) $P_i^{A,\text{pd}} = \emptyset$,
 - c) $T_i = \omega_i = \omega_i^A = 0$,
 - d) $\tau_i^A = -\Delta_H$, and
 - e) $\Phi_i^A(\cdot, 0)$ is defined according to Eq. 1.

Notice 1) and 3a) together imply that $c_i \neq c_j$ for any $j \neq i$. Our first result guarantees that the use of Algorithm 1 for coverage region updates is well-posed.

Theorem 1 (Well-Posedness). Under Assumption 1, an update scheme in which, upon each exchange, the base station and the communicating agent update their respective variables via Algorithm 1 is well-posed. That is, the operations required by

Algorithm 1 are well-posed at the time of execution.

Performing of updates with Algorithm 1 does not guarantee that individual coverage regions (elements of P), are disjoint from one another. It does, however, guarantee that the m -covering P , the local coverage region set $P^A := \{P_i^A\}_{i=1}^m$, and the local likelihoods $\{\Phi_i^A\}_{i=1}^m$ retain properties that are consistent with the goals of the decomposition-based coverage framework. Namely, the coverings P and P^A maintain connectivity, and each function Φ_i^A has support that is disjoint from that of all other local likelihoods, yet still evolves to provide reasonable global coverage. The manipulations in Algorithm 2 also ensure that agents are able to “safely” vacate areas that are re-assigned before newly assigned agents enter. We detail these and other properties for the remainder of this section.

C. Set Properties

The next result formalizes key set properties.

Theorem 2 (Set Properties). *Suppose Assumption 1 holds, and that, upon each exchange, the base station and the communicating agent update their respective variables via Algorithm 1. Then, the following⁵ hold at any time $t \geq 0$:*

- 1) $P^{\mathcal{I}\mathcal{D}}$ is a connected m -partition of Q ,
- 2) P is a connected m -covering of Q ,
- 3) $c_i \in P_i$ and $c_i \neq c_j$ for any $i \neq j$,
- 4) $\text{supp}(\Phi_i^A(\cdot, t)) \subseteq P_i$ for any i , and
- 5) $\bigcap_{i=1}^m \text{supp}(\Phi_i^A(\cdot, t)) = \emptyset$

Whenever additions are made to an agent’s coverage region during a call to Algorithm 1, the newly added vertices are not immediately included in the instantaneous support⁶ of the agent’s local likelihood. As such, if each agent’s movement is restricted to lie within the aforementioned instantaneous support, then exploration of newly added regions is temporarily prohibited following an update. This delay allows other agents to vacate before the newly assigned agent enters. Conversely, when regions are removed from an agent’s coverage region, Algorithm 1 guarantees a “safe” path, i.e., a path with no collision risk, exists and persists long enough for the agent to vacate. Let $\bar{d} := \max_i \frac{1}{s_i} \sum_{\{k_1, k_2\} \in \mathcal{E}} d_Q(k_1, k_2)$, and define the *prohibited region* of agent i at time t , $\text{Proh}_i(t)$, as the subset consisting of any newly added vertices that do not yet belong to $\text{supp}(\Phi_i^A(\cdot, t))$, i.e., $\text{Proh}_i(t) := \{k \in P_i^A, t - \omega_i^A < \tau_i^A \text{ and } k \in P_i^{\text{A, pd}}\}$, we formalize this discussion as follows.

Theorem 3 (Coverage Quality). *Suppose Assumption 1 holds, and that, upon each exchange, the base station and the communicating agent updates their respective variables via Algorithm 1. Then, for any $k \in Q$ and any $t \geq 0$:*

- 1) k belongs to at least one agent’s coverage region P_i ,
- 2) if $k \in \text{Proh}_i(t)$ for some i , then there exists t_0 satisfying $t < t_0 < t + \Delta + \bar{d}$ such that, for all $\bar{t} \in [t_0, t_0 + \Delta_H]$, the vertex k belongs to the set $P_i \setminus \text{Proh}_i(\bar{t})$, and
- 3) if k is removed from P_i at time t , then, for all

$$\bar{t} \in \left(t, t + \frac{1}{s_i} d_{P_i^-} \left(k, P_i^{\mathcal{I}\mathcal{D}, -} \right) \right],$$

⁵ $\text{supp}(f)$ denotes the support of a function f .

⁶The *instantaneous support* of Φ_i^A at time t is defined as $\text{supp}(\Phi_i^A(\cdot, t))$.

we have

- a) $P_i^{\mathcal{I}\mathcal{D}, -} \subseteq P_i$, and
- b) there exists a length-minimizing path on $G(P_i^-)$ from k into $P_i^{\mathcal{I}\mathcal{D}, -}$, and all of the vertices along any such path (except the terminal vertex) belong to the set $\text{Proh}_{ID_k^+}(\bar{t}) \setminus \bigcup_{j \neq ID_k^+} P_j$.

Theorems 2 and 3 allow Algorithm 1 to operate within a decomposition-based framework to provide reasonable coverage with inherent collision avoidance. Indeed, when each agent moves within its coverage region and avoids its prohibited region, the theorems imply that each agent 1) can visit its entire coverage region (connectedness), 2) allows adequate time for other agents to vacate newly assigned regions before entering, and 3) has a “safe” route into the remaining coverage region if its current location is removed during an update.

Remark 3 (Local Variables). *Theorems 2 and 3 also hold if we replace P with P^A and c with c^A in the theorem statement.*

Remark 4 (Bounds). *Theorem 3 also holds when \bar{d} is re-defined as any other upper bound on the subgraph distance between two arbitrarily chosen vertices of an arbitrarily chosen connected subgraph of G .*

D. Convergence Properties

Our proposed strategy differs from that of [26] mainly due to the presence of logic (e.g., timer manipulations) to ensure effective pairing with single-agent trajectory planners. Note also that \mathcal{H} differs from typical partitioning costs, since it uses subgraph, rather than global graph, distances. As such, convergence properties of the algorithms herein do not follow readily from existing results. Consider the following definition.

Definition 2 (Pareto Optimality). *The pair (c, P) is Pareto optimal at time t if (i) $\mathcal{H}(c, P, t) \leq \mathcal{H}(\bar{c}, P, t)$ for any $\bar{c} \in Q^m$, and (ii) $\mathcal{H}(c, P, t) \leq \mathcal{H}(c, \bar{P}, t)$ for any $\bar{P} \in \text{Cov}_m(Q)$.*

The following result characterizes the dynamic evolution of coverage regions with respect to Pareto optimality.

Theorem 4 (Convergence). *Suppose Assumption 1 holds and that, upon each exchange, the base station and the communicating agent updates their respective variables via Algorithm 1. If Φ is static, i.e., $\Phi(\cdot, t_1) = \Phi(\cdot, t_2)$ for all t_1, t_2 , then the m -covering P and the generators c converge in finite time to an m -partition P^* of Q and a set c^* , resp. The pair (c^*, P^*) is Pareto optimal at any time following convergence.*

Whenever the event likelihood is static (and Assumption 1 holds), Algorithm 1 causes coverage regions and generators to collectively converge in finite time to a Pareto optimal configuration. That is, the agent’s limiting coverage assignments are “optimal” in that they balance the coverage load in a manner that directly considers the likelihood Φ . Further, the entire operation only relies on sporadic and unplanned information exchange between agents and the base station.

Remark 5 (Voronoi Partitions). *It can be shown that Pareto optimality of (c^*, P^*) in Theorem 4 implies that, following*

convergence, P^* is a multiplicatively weighted Voronoi partition (generated by c^* , weighted by s , subject to density $\Phi(\cdot, t)$) by standard definitions (e.g., [26]). If the centroid set of each P_i is defined as $\arg \min_{h \in P_i} \sum_{k \in P_i} d_{P_i}(k, h) \Phi(k, t)$, then P^* is also centroidal.

V. DECOMPOSITION-BASED SURVEILLANCE.

This section pairs the proposed partitioning framework with a generic, single-vehicle trajectory planner, forming the complete, decomposition-based coverage control framework.

A. Complete Routing Algorithm

Theorems 2 and 3 provide a number of useful insights when using the dynamic partitioning updates of the previous section. Namely, Theorem 2 ensures that (i) the instantaneous support of each Φ_i^A lies entirely within the coverage region P_i^A , and (ii) the support of two distinct agents' local likelihoods do not intersect. Theorem 3 states that (i) length of any interval on which a given vertex is uncovered, i.e., belongs solely to agent prohibited regions, cannot exceed a finite upper bound, and (ii) the parameter Δ_H is a lower bound on the length of time that a recently uncovered vertex must remain covered before it can become uncovered again. Since $G(Q)$ is a discrete representation of the surveillance region, these facts together suggest that an intelligent routing scheme that carefully restricts agent motion according to the instantaneous support of the local likelihood functions could achieve adequate coverage while also accomplishing the ancillary goal of collision avoidance. This motivates the following assumption.

Assumption 2 (Agent Motion). *Each agent i has knowledge of its position at any time t , and its on-board trajectory planner operates under the following guidelines:*

- 1) *generated trajectories obey agent motion constraints,*
- 2) *trajectories are constructed incrementally and can be altered in real-time, and*
- 3) *the agent is never directed to leave regions associated with P_i or enter regions associated with $Proh_i(t)$.*

Each agent precisely traverses generated trajectories.

Algorithm 3 presents the local protocol for Agent i .

B. Collision Avoidance

Although Assumption 2 locally prevents agents from leaving assigned coverage regions or entering prohibited regions, dynamic coverage updates can still result in undesirable agent configurations. Indeed, an agent can be placed outside of its own coverage region if the vertex corresponding to its location is abruptly removed during an update. If this happens, Algorithm 3 constructs a route from the agent's location back into a region where there is no collision risk. With mild assumptions, Theorem 3 guarantees that this construction 1) is well-defined, and 2) does not present the agent with a transient collision risk. We formalize this result here.

Theorem 5 (Collision Avoidance). *Suppose Assumptions 1 and 2 hold, and that each agent's initial position lies within its initial coverage region P_i . Suppose further that each agent's*

Algorithm 3: Motion Protocol for Agent i

Data: $G(Q)$, Φ_i^A , P_i^A , c_i^A , $P_i^{A, pd}$, τ_i^A , ω_i^A
begin
 while True do
1 Increment trajectory via on-board planner
2 Follow trajectory
 if Communication with the base station then
3 Set $P_i^{test} = P_i^A$
4 Obtain updated variables from base station
 if Location lies within the set $P_i^A \setminus P_i^{test}$ then
5 Find a minimum-length path in $G(P_i^{test})$
 from the currently occupied node into P_i^A
 while Agent i is outside P_i^A do
6 Follow the aforementioned route

motion is locally governed according to Algorithm 3, where the update in line 4 is calculated by the base station via Algorithm 1. If the weight assigned to each edge in $G(Q)$ is an upper bound on the distance between the associated regions, then no two agents will ever collide.

By running the schemes of Section IV in conjunction with a motion planning scheme obeying Assumption 2, we obtain a complete strategy that 1) only requires weak communication assumptions, 2) provides dynamic load-balancing, and 3) has inherent collision avoidance/efficiency properties.

VI. NUMERICAL EXAMPLES

This section presents numerical examples to illustrate the functionality of the decomposition-based routing scheme. In all examples, high-level coverage assignment updates are performed by the base station via Algorithm 1 during each exchange with an agent, while each agent's local processor runs the motion protocol in Algorithm 3. For incremental trajectory construction (Algorithm 3, line 1), we implement a modified version of the *Spectral Multiscale Coverage* (SMC) scheme in [15], which creates agent trajectories that mimic ergodic dynamics while also locally constraining agent motion to lie within the appropriate sets. This planner satisfies Assumption 2. Initial region generators were selected randomly (enforcing the non-coincidence constraint), and each agent was initially placed at its region generator. The initial covering P was created from these generators by calculating a weighted Voronoi partition. The remaining parameters were chosen according to Assumption 1. During the simulation, randomly chosen agents (chosen via random number generator) sporadically communicated with the base station to receive coverage assignment updates. Communication times were randomly chosen, subject to a maximum inter-communication time $\bar{\Delta}$.

A. Static Likelihood

The first example is a 4 agent mission, which is executed over a 100 x 100 surveillance region subject to a static, Gaussian likelihood centered at the bottom left corner. For

discretization, the region is divided into 400, 5×5 subregions, and regions are considered adjacent if they share a horizontal or vertical edge. Each agent has a maximum speed of 1 unit distance per unit time, and the maximum inter-communication time is $\bar{\Delta} = 10$ time units. Figure 3 shows the evolution of the coverage regions at various time points for an example simulation run. Note that Figure 3 only shows each agent i 's *active* coverage region, i.e., the subset of P_i that does not intersect its prohibited region $\text{Proh}_i(t)$. The family of active coverage regions does not generally form an m -covering of Q ; however, elements of this family are connected and never intersect as a result of inherent collision avoidance properties.

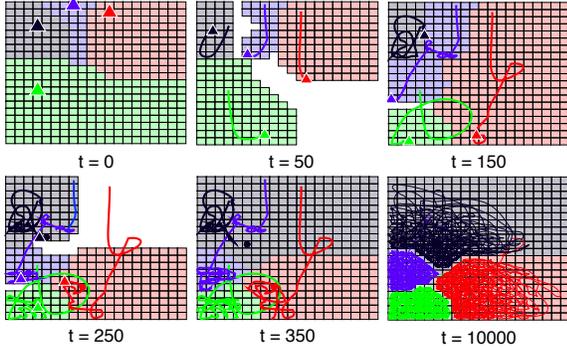


Fig. 3. Snapshots of a 4 agent surveillance mission, assuming a static Gaussian likelihood. Each agent's position, past trajectory, and *active coverage region* are indicated by the colored triangle, line, and squares, resp..

The left plot in Figure 4 depicts the maximum amount of time that any individual subregion went uncovered, i.e. the subregion did not belong to any agent's active covering $P_i \setminus \text{Proh}_i(t)$, during each of 50 simulation runs. Here, the

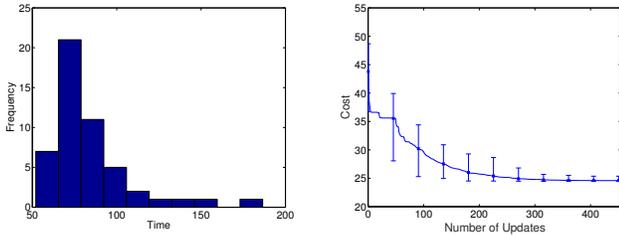


Fig. 4. The maximum amount of time that any subregion went uncovered in each of 50 simulation runs (left), and the value of the cost H as a function of time, averaged over the same 50 runs (right).

maximum amount of time that any region went uncovered was 186 units, though most trials had maximums of less than 75 units. This is well-below the loose bound $\bar{\Delta} + \bar{d} = 770$ predicted by Theorem 3 (see Remark 4). The right plot in Figure 4 shows the mean values of the cost function \mathcal{H} as a function of time, calculated over the same 50 simulations runs. Here, error bars represent the range of cost values achieved at select time points. The variance between runs is due to the stochastic nature of the data-exchange patterns between the agents and the base station. Notice that the cost is a non-increasing function of time, as predicted in the proof of Theorem 4, eventually settling as the coverage

regions/generators reach their limiting configuration, e.g., see Figure 3. These configurations are guaranteed to be Pareto optimal and, by Remark 5, to form a multiplicatively weighted Voronoi partition. Since vehicles are assumed to have identical maximum speeds, we see from the limiting configuration in Figure 3 that the resultant coverage assignments provide load-balancing that takes into account the event likelihood. If the low-level trajectory planner biases trajectories according to the event likelihood, this results in desirable coverage properties. Under the modified SMC planner used here, the temporal distribution of agent locations closely resembles the spatial likelihood distribution in the limit, as shown in Figure 5.

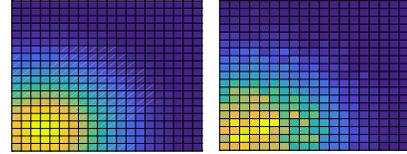


Fig. 5. Comparison between the (static) event likelihood Φ (left), and the proportion of time that some agent occupied each subregion after significant time has passed (10000 units) (right).

Further, during the simulation, no two agents ever occupied the same space due to the careful parameter manipulations employed by Algorithm 1. Figure 6 illustrates the logic governing these manipulations through a simplistic example. During the

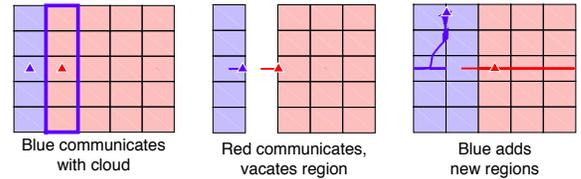


Fig. 6. Simplified example illustrating how Algorithm 1 manipulates timing parameters to prevent agent collisions.

first update, the blue agent acquires some of the red agent's coverage region. Rather than immediately adding these regions to its active covering, the blue region waits until sufficient time has passed to guarantee that the red agent has updated and moved out of the reassigned regions. Under Algorithm 3, once the red agent communicates with the base, it immediately vacates the re-assigned regions, after which the blue agent can add the region to its active covering. This procedure guarantees that no two agents will never have overlapping active coverings and thus never collide (Theorem 5). This same logic is results in inherent collision prevention over more complex scenarios.

B. Quasi-static Likelihood

We now illustrate how the proposed coverage framework reacts to abrupt changes in the underlying likelihood, i.e., when the likelihood is quasi-static. This type of scenario is common in realistic missions, e.g., when the base-station's estimate of the underlying likelihood is only re-formulated if some agent's sensor data indicates a drastic change in the underlying landscape. For this purpose, we adopt identical parameters as in the first example, with the exception of the likelihood Φ , whose spatial distribution abruptly switches at

select time-points. If the switches are sufficiently spaced in comparison to the rate of convergence, then the coverage regions dynamically adjust to an optimal configuration that is reflective of the current state. For example, Figure 8 shows the coverage region evolution after the underlying likelihood undergoes a single switch from the initial to the final density shown in Figure 7 at time $t = 2000$. In contrast, when

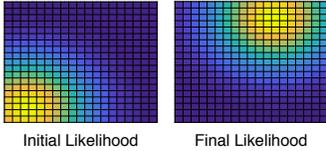


Fig. 7. The initial and final likelihood $\Phi(\cdot, t)$.

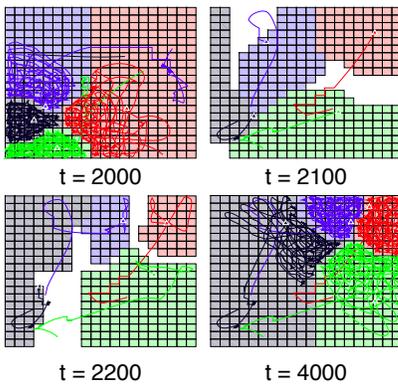


Fig. 8. Coverage regions after the likelihood switches (see Fig. 7)

the underlying likelihood changes faster than the rate of convergence, coverage regions are constantly in a transient state. Despite this, the proposed framework still results in some degree of load-balancing. To illustrate, Figure 9 shows the value of the cost \mathcal{H} during a simulation in which the underlying likelihood switches at 12 randomly chosen time-points over a 1000 unit horizon. Each switch re-defined the spatial likelihood as a Gaussian distribution centered at a randomly selected location. Notice that the value of the cost function monotonically decreases between the abrupt spikes caused by changes in the underlying likelihood. Despite the fact that convergence is not reached, coverage regions quickly shift away from high-cost configurations, as indicated by the sharp decreases in the cost shortly following each switch.

VII. CONCLUSION

This work presents a modular, decomposition-based, coverage control framework for communication-constrained multi-agent surveillance missions. In particular, our approach uses a dynamic partitioning strategy to balance the surveillance load across available agents, requiring only sporadic and unplanned exchanges between individual agents and a base station. The partitioning update algorithm also manages high-level timing and logic parameters to guarantee that the resulting coverage assignments have geometric and temporal properties that are amenable for combination with generic single vehicle trajectory planners. Under appropriate assumptions, the proposed

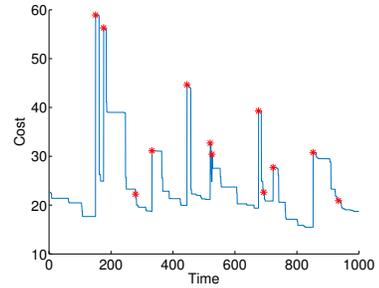


Fig. 9. Evolution of the cost \mathcal{H} using a quasi-static likelihood with 12 random switches (switch times are indicated by the stars).

algorithms will produce a set of coverage regions forming a Pareto optimal partition, while also ensuring collision avoidance and quality of coverage guarantees.

Future work should further relax communication assumptions to reflect additional realistic hardware limitations, e.g., use of directional antennae for wireless transmission. Other areas of future research include the addition of peer-to-peer, in addition to central, communication, performance comparisons between specific trajectory planners when used within our framework, e.g., those involving ergodic Markov chains, and further theoretical characterizations of performance.

APPENDIX: PROOFS

Proposition 2 (Set Membership). *Suppose Assumption 1 holds, and that, at the time of each exchange occurring prior to the fixed time $\bar{t} \geq 0$, required algorithmic constructions are well-posed so that the base station and the communicating agent are able to update their respective variables via Algorithm 1. Then, for any $k \in Q$ at any time $t \leq \bar{t}$:*

- 1) $k \in P_{ID_k}$,
- 2) k belongs to at most 2 elements of P ,
- 3) if $T_{ID_k} = 0$, then $k \notin P_j$ for any $j \neq ID_k$, and
- 4) if $k \in P_j$, $j \neq ID_k$, then $P_j \cap P_\ell^{TD} = \emptyset$ for $\ell \notin \{j, ID_k\}$

Proof. Fix $\bar{t} \geq 0$, $k \in Q$. When $t = 0$, $P = P^{TD}$ is an m -partition of Q , implying the proposition. Since k is not removed from P_{ID_k} or added to any P_i with $i \neq ID_k$ until its first reassignment, i.e., when ID_k is changed. Thus, the proposition is true for all t prior to the first reassignment. Suppose the proposition holds for all t prior to the p th reassignment, which occurs at $t = t_0$. Suppose that $ID_k^- = j$, $ID_k^+ = i \neq j$. Algorithm 1 defines $P_i^{TD,+} = P_i^+ = P_{ID_k^+}^+$. Thus, $k \in P_{ID_k^+}^+ = P_i^+$ and remains in these sets until another reassignment. Statement 1 holds, therefore, for all t prior to the $p + 1$ st reassignment. Now note that, by Alg. 2, reassignment cannot occur at t_0 unless $T_j^- = 0$. By inductive assumption, statement 3 of the proposition holds when $t = t_0^-$, implying that $k \notin P_\ell^-$ for any $\ell \neq j$. Upon reassignment, the timers T_j, T_i are modified such that $T_j^+, T_i^+ > \omega_j^+ + \bar{\Delta} - t_0$. Since 1) ID_k is unchanged as long as $T_j > 0$, and 2) agent j communicates with the base prior to time $\omega_j^+ + \bar{\Delta}$ (when k is removed from P_j), we deduce that k solely belongs to P_j and P_i until the $p + 1$ st reassignment. By the same logic, at

any time $t \geq t_0^+$ at which $T_i = 0$ and the $p+1^{\text{st}}$ reassignment has not yet occurred, $k \in P_i$ exclusively (addition to other sets in P without reassignment is impossible). We deduce that statements 2 and 3 hold for any t prior to the $p+1^{\text{st}}$ reassignment. Now note that, 1) statement 3 holds prior at time t_0^+ , implying $P_j^- = P_i^{\text{TD},-}$, 2) no agent claims additional points vertices from P_j^+ unless $T_j = 0$, and 3) k cannot be added to a coverage region without reassignment. As such, $P_j \cap P_\ell^{\text{TD}} = \emptyset$ for any $\ell \notin \{j, \text{ID}_k = i\}$ prior to the $p+1^{\text{st}}$ reassignment. From persistent communication and the bound $\underline{\Delta}$, the proposition follows by induction. \square

Proof of Theorem 1. It suffices to show that Def. 1 is well-posed (Prop. 1) whenever additive sets are required. We proceed by induction. When $t = 0$, $P^{\text{TD}} = P$ is a connected m -partition of Q and thus, for any i , P_i^{TD} is disjoint from $\bigcup_{j \neq i} P_j$. The same holds prior to the first exchange between the base station and some agent. Thus, by Prop. 1, the first call to Alg. 1 is well-posed. Now assume that, for all times t prior to the p^{th} call to Alg. 1, 1) P^{TD} is a connected m -partition of Q , and 2) if an exchange that requires construction of P_i^{add} occurs, then $P_i^{\text{TD}} \cap (\bigcup_{j \neq i} P_j) = \emptyset$. This implies that Prop. 2 holds at any time t prior to the $p+1^{\text{st}}$ exchange. Assume the p^{th} communication occurs when $t = t_0$ and involves agent i . By the uniqueness of identifiers, $P^{\text{TD},+}$ is an m -partition of Q . To show that $P^{\text{TD},+}$ is connected, first note that $P_i^{\text{TD},+} = P_i^+$. Since $P_i^+ = P_i^{\text{add}}(c_i^+)$ (connected by Def. 1), or $P_i^{\text{TD},+} = P_i^{\text{TD},-}$ (connected by inductive assumption) connectivity of $P_i^{\text{TD},+}$ follows. Now consider P_j^{TD} , $j \neq i$. If $T_j^- \neq 0$, then $P_j^{\text{TD},-} = P_j^{\text{TD},+}$ and connectivity of $P_j^{\text{TD},+}$ follows. Suppose $T_j^- = 0$ and $P_j^{\text{TD},+}$ is not connected. Then, P_j^+ is not connected: if it were, there would exist $k \in P_j^+$ with $\text{ID}_k^+ \notin \{i, j\}$, contradicting Prop. 2, statement. Thus, there exists $k_1 \in P_j^{\text{TD},+}$ such that 1) $k_1 \notin P_i^{\text{add}}(c_i^+)$, and 2) any length-minimizing path in $G(P_j^-)$ spanning k_1 and c_j^+ contains some $k_2 \in P_i^{\text{add}}(c_i^+) = P_i^+(c_i^+)$. Select one such path and vertex k_2 . Assume without loss of generality that there exists an edge $\{k_1, k_2\} \in \mathcal{E}$. Def. 1 implies $\frac{1}{s_i} d_{P_i^+}(k_2, c_i^+) < \min\{\frac{1}{s_\ell} d_{P_\ell^+}(k_2, c_\ell^+) | \ell \neq i, k_2 \in P_i^+\}$ and thus $\frac{1}{s_i} d_{P_i^+ \cup \{k_1\}}(k_1, c_i^+) < \frac{1}{s_j} d_{P_j^+}(k_1, c_j^+)$. Since $T_j^- = 0$ and $\text{ID}_{k_1}^- = j$, Prop. 2 implies $\frac{1}{s_i} d_{P_i^+ \cup \{k_1\}}(k_1, c_i^+) < \frac{1}{s_j} d_{P_j^+}(k_1, c_j^+) = \min\{\frac{1}{s_\ell} d_{P_\ell^+}(k_1, c_\ell^+) | \ell \neq i, k_1 \in P_\ell^+\}$, contradicting $k_1 \notin P_i^{\text{add}}(c_i^+)$. Thus, $P_j^{\text{TD},+}$ is connected. Invoking Prop. 2 statement 3, the inductive assumption therefore holds for all times prior to the $p+1^{\text{st}}$ exchange, thereby implying well-posedness of the first $p+1$ exchanges. \square

Proof of Theorem 2.

Statement 1: The proof of Thm. 1 implies the statement.

Statement 2: P is an m -covering of Q since P^{TD} is always an m -partition of Q (statement 1), and $P_i^{\text{TD}} \subseteq P_i$ for any i (Prop. 2, statement 1). P is connected, since $P_i = P_i^{\text{TD}}$ (connected by statement 1) immediately following any update, and P_i is unchanged in between updates.

Statement 3: It suffices to show that $\text{ID}_{c_i} = i$ for any t and any i : this would imply $c_i \neq c_j$ for any $i \neq j$, and $c_i \in P_i$ (Prop. 2). By assumption, $\text{ID}_{c_i} = i$ for all i at $t = 0$. The same

holds for any t prior to the first exchange between any agent and the base station. Suppose $\text{ID}_{c_i} = i$ for all i (thus $c_i \neq c_j$ for any $i \neq j$) prior to the p^{th} exchange. If agent i is the p^{th} communicating agent, lines 2 and 9 of Alg. 1 imply $\text{ID}_{c_i^+}^+ = i$. Since $d_{P_j^-}(c_j^-, c_j^-) = 0$ for any j , we have $c_j^+ \notin P_i^{\text{add}}(c_i^+)$. Thus, $\text{ID}_{c_j^+}^+ = j$, and induction proves the statement.

Statements 4 and 5: Statement 4 follows from (1), noting that $P_i^A = P_i$. Statement 5 holds by assumption when $t = 0$. Let $k \in Q$, and consider times when ID_k changes (k is re-assigned). Since $\text{supp}(\Phi_j^A(\cdot, t)) = P_j = P_j^A$ for any j at $t = 0$, statement 4 implies that, for any t prior to the first reassignment, k belongs exclusively to $\text{supp}(\Phi_{\text{ID}_k}^A(\cdot, t))$, implying statement 5. Suppose statement 5 holds for all t prior to the p^{th} reassignment (occurring at time t_0), and $\text{ID}_k^- = j$, $\text{ID}_k^+ = i \neq j$. Then, $T_j^- = 0$ and k belongs exclusively to P_j^- when $t = t_0^-$ (Prop. 2). By Alg. 1 and 2, $k \in P_i^{\text{A,pd},+}$ and $T_i^+ > \omega_j^+ + \bar{\Delta} - t_0 \geq \tau_i^{\text{A},+}$. Since $\text{supp}(\Phi_i^A(\cdot, t))$ is not redefined for a duration of at least $T_i^+ \geq \tau_i^{\text{A},+}$, (1) implies $k \notin \text{supp}(\Phi_i^A(\cdot, t))$ when $t \in [t_0^+, t_0^+ + \tau_i^{\text{A},+}]$. Since k is re-assigned when $t = t_0$, $k \in P_i^+ \setminus P_i^{\text{TD},-}$ and $T_j^+ = \omega_j^+ + \bar{\Delta} - t_0$. Agent j will communicate with the base at some time $t_1 < t_0 + T_j^+ = \omega_j^+ + \bar{\Delta} < t_0 + T_i^+$. Thus, $T_i > 0$ when $t = t_1$, and k is removed from both P_j and $\text{supp}(\Phi_j^A(\cdot, t))$. Thus, for all $t > t_0 + \tau_i^+$ and before the $p+1^{\text{st}}$ reassignment, k belongs exclusively to $\text{supp}(\phi_i(\cdot, t))$. \square

Proof of Theorem 3. Thm. 2 implies statement 1.

Statement 2: For any i , 1) $T_i = 0$ when $t = 0$, and 2) $\frac{1}{s_i} d_{Q'}(k_1, k_2) \leq \bar{d}$ for any $Q' \subseteq Q$, $k_1, k_2 \in Q'$. Thus, it is straightforward to show that, for any i and any T , we have the bound $T_i \leq \bar{\Delta} + \Delta_H + \bar{d}$. We show that, for any i , the bound $\tau_i^A - t + \omega_i^A \leq T_i - \Delta_H$ holds by induction: $T_i = 0$ and $\tau_i^A = -\Delta_H$ when $t = 0$, so $\tau_i^A - t + \omega_i^A = \tau_i^A \leq T_i - \Delta_H$. The bound similarly holds prior to the first exchange involving any agent, since $\tau_i^A - t + \omega_i^A \leq -\Delta_H \leq T_i - \Delta_H$ at any such time. Assume the bound holds prior to the p^{th} update (occurring at $t = t_0$). Consider 2 cases: if agent i is the communicating agent, then $\tau_i^{\text{A},+} - t + \omega_i^{\text{A},+} = \tau_i^+ := T_i^+ - \Delta_H$; if not, then $\tau_i^{\text{A},+} = \tau_i^{\text{A},-}$ and either 1) $T_i^- = T_i^+$ implying the desired bound, or 2) $T_i^- = 0$ and $\tau_i^{\text{A},+} - t_0 + \omega_i^{\text{A},+} = \tau_i^{\text{A},-} - t + \omega_i^{\text{A},-} \leq T_i^- - \Delta_H = -\Delta_H \leq (\omega_i^{\text{A},+} + \bar{\Delta} - t_0) - \Delta_H = T_i^+ - \Delta_H$. This logic extends to all times prior to the $p+1^{\text{st}}$ exchange and the desired bound follows by induction.

Using the bounds from the previous paragraph, we have $\tau_i^A + \omega_i^A \leq t + \bar{\Delta} + \bar{d}$. Fix t and $k \in \text{Proh}_i(t)$. Then, $k \in P_i^{\text{A},+} = P_i^+$, $k \in P_i^{\text{A,pd},+}$, and $t - \omega_i^{\text{A},+} < \tau_i^{\text{A},+}$ ('+' is with respect to the fixed time t). Further, over the interval $[t, \omega_i^{\text{A},+} + \tau_i^{\text{A},+}]$, the vertex k is not re-assigned, P_i is not augmented, and τ_i^A is unchanged. Therefore, $k \notin \text{Proh}_i(t)$ at time $\omega_i^{\text{A},+} + \tau_i^{\text{A},+}$. If $t_0 := \omega_i^{\text{A},+} + \tau_i^{\text{A},+}$, we have $t < t_0 \leq t + \bar{\Delta} + \bar{d}$. Since $T_i \geq \tau_i^{\text{A},+} + \Delta_H$ at time $\omega_i^{\text{A},+}$, k is not re-assigned during the interval $[\omega_i^{\text{A},+}, \omega_i^{\text{A},+} + T_i^+] \supseteq [\omega_i^{\text{A},+}, t_0 + \Delta_H] \supseteq [t_0, t_0 + \Delta_H]$. Thus $k \in P_i \setminus \text{Proh}_i(\cdot)$ over the same interval.

Statement 3: Fix t and suppose $k \in P_i^- \setminus P_i^+$ (in this proof, '+, -' are with respect to t). Then, 1) ID_k changed (k was re-assigned) at time $t_0 < t$, 2) agent i communicates with the base at time t , and 3) no exchanges involving agent i occurred

during the interval $[t_0, t)$. Upon reassignment at time t_0 , Alg. 2 specifies that 1) T_i is reset to value $\omega_i^{A,-} + \bar{\Delta} - t_0$, thus P_i^{ID} is unchanged over the interval $[t_0, t)$, 2) k is added to $P_{\text{ID}_k}^{A,\text{pd}}$, and 3) $\tau_{\text{ID}_k}^A, T_{\text{ID}_k}$ are given values of at least

$$\tilde{\omega} := \max_{\bar{k} \in P_i^- \setminus P_i^+} \left\{ \omega_i^{A,-} + \bar{\Delta} + \frac{1}{s_i} d_{P_i^-}(\bar{k}, P_i^{\text{ID},-}) - t_0 \right\},$$

implying that $P_{\text{ID}_k}, \text{Proh}_{\text{ID}_k}(\cdot)$ remain unchanged over the interval $(t_0, \tilde{\omega}] \supseteq (t_0, t + \frac{1}{s_i} d_{P_i^-}(k, P_i^{\text{ID},-})] \supseteq (t_0, t]$.

Since coverage regions are connected and non-empty (Thm. 2) and $P_i^- \cap P_\ell^{\text{ID}} = \emptyset$ for any $\ell \notin \{i, \text{ID}_k^+\}$ on the interval $(t_0, t]$ (Prop. 2), 1) there exists a path of length $d_{P_i^-}(k, P_i^{\text{ID},-})$ from k into $P_i^{\text{ID},-}$ and every vertex along any such path (except the terminal vertex) lies within $P_i^- \setminus P_i^+$, and 2) $P_i^- \setminus P_i^+ \subseteq \text{Proh}_{\text{ID}_k^+}$ over the interval $(t_0, \tilde{\omega}] \supseteq [t, t + \frac{1}{s_i} d_{P_i^-}(k, P_i^{\text{ID},-})]$. Since 1) each vertex belongs to no more than two sets simultaneously (Prop. 2), 2) $k \in P_i^- \setminus P_i^+$, and 3) no agent claims any vertex in the set $\text{Proh}_{\text{ID}_k^+}$ when $T_{\text{ID}_k^+} > 0$, vertices along the path (excluding the terminal vertex) do not belong to P_j with $j \neq \text{ID}_k$ over the interval $[t, t + \frac{1}{s_i} d_{P_i^-}(k, P_i^{\text{ID},-})]$. To complete the proof, note that Alg. 2 implies $T_i^+ > \frac{1}{s_i} d_{P_i^-}(k, P_i^{\text{ID},-})$, and thus $P_i^{\text{ID},-} \subseteq P_i^{\text{ID}}$ over $[t, t + \frac{1}{s_i} d_{P_i^-}(k, P_i^{\text{ID},-})]$. \square

Proposition 3 (Cost). *Suppose Assumption 1 holds and that, upon each exchange, the base station and the communicating agent update their respective variables via Alg. 1. If $\Phi(\cdot, t_1) = \Phi(\cdot, t_2)$ for all t_1, t_2 , then $\mathcal{H}(c, P^{\text{ID}}, \cdot) = \mathcal{H}(c, P, \cdot)$.*

Proof. Since Φ is static, $\mathcal{H}(\cdot, \cdot, t_1) = \mathcal{H}(\cdot, \cdot, t_2)$ for any t_1, t_2 . When $t = 0$, $P = P^{\text{ID}}$ and thus $\mathcal{H}(c, P^{\text{ID}}, 0) = \mathcal{H}(c, P, 0)$. The same is true prior to the first exchange between any agent and the base. Suppose that immediately prior to the p^{th} exchange (occurring at $t = t_0$, involving agent i), we have $\mathcal{H}(c^-, P^{\text{ID},-}, t_0) = \mathcal{H}(c^-, P^-, t_0)$. Recall that, for any agent j , P_j and P_j^{ID} coincide immediately following any exchange involving agent j and, if agent j claims vertices from P_i , then Alg. 2 guarantees that agent i will communicate with the base station before any additional vertices are claimed by other agents. At the time of the p^{th} update, this logic, along with Prop. 2, implies that $P_i^{\text{ID},-} \cap P_j^- = \emptyset$, for all $j \neq i$. Noting that $c_i^+ \in P_i^{\text{ID},-}$, we deduce that any single $k \in P_i^{\text{ID},-}$ contributes equally to $\mathcal{H}(c^+, P^{\text{ID},+}, t_0)$ and $\mathcal{H}(c^+, P^+, t_0)$. If $k \in P_i^{\text{add}}(c_i^+) \setminus P_i^{\text{ID},-}$, then for any $j \neq i$ such that $k \in P_j^+$, we have $\frac{1}{s_i} d_{P_i^+}(k, c_i^+) < \frac{1}{s_j} d_{P_j^+}(k, c_j^+)$ (Def. 1). Therefore, k contributes equivalently to the values of both $\mathcal{H}(c^+, P^{\text{ID},+}, t_0)$ and $\mathcal{H}(c^+, P^+, t_0)$. Now suppose $k \in P_j^+ \setminus P_i^+$, where $P_j^+ \cap P_i^+ \neq \emptyset$. We show that $d_{P_j^{\text{ID},+}}(c_j^+, k) = d_{P_j^+}(c_j^+, k)$: if a length-minimizing path in $G(P_j^+)$ between c_j^+ and k is also contained in $G(P_j^{\text{ID},+})$, then the result is trivial. Suppose that every such minimum length path does leave $G(P_j^{\text{ID},+})$. From Prop. 2 statement 4, every $\bar{k} \in P_j^+$ must satisfy either $\text{ID}_k^+ \in \{i, j\}$. Therefore, we assume without loss of generality that k is adjacent to P_i^+ . Let $\bar{k} \in P_i^+$ be a vertex that is adjacent to k and lies along a minimum-length path in $G(P_j^+)$

spanning c_j^+ and k . Since $\bar{k} \in P_i^+ \setminus P_i^{\text{ID},-}$, we must have $\bar{k} \in P_i^{\text{add}}(c_i^+)$ as constructed during the update, which implies $\frac{1}{s_i} d_{P_i^+}(\bar{k}, c_i^+) < \min\{\frac{1}{s_\ell} d_{P_\ell^+}(\bar{k}, c_\ell^+) | \ell \neq i, \bar{k} \in P_\ell^+\}$ and thus $\frac{1}{s_i} d_{P_i^+ \cup \{k\}}(k, c_i^+) < \frac{1}{s_j} d_{P_j^+}(k, c_j^+)$. Since $T_j^- = 0$ and $\text{ID}_k^- = j$, Prop. 2 implies $\frac{1}{s_i} d_{P_i^+ \cup \{k\}}(k, c_i^+) < \frac{1}{s_j} d_{P_j^+}(k, c_j^+) = \min\{\frac{1}{s_\ell} d_{P_\ell^+}(k, c_\ell^+) | \ell \neq i, P_\ell^+\}$, contradicting $k \notin P_i^{\text{add}}(c_i^+) \subset P_i^+$. Thus, $d_{P_j^{\text{ID},+}}(c_j^+, k) = d_{P_j^+}(c_j^+, k)$, which, by inductive assumption, implies that k contributes equally to the value of both $\mathcal{H}(c^+, P^{\text{ID},+}, t_0)$ and $\mathcal{H}(c^+, P^+, t_0)$. We conclude that $\mathcal{H}(c^+, P^{\text{ID},+}, t_0) = \mathcal{H}(c^+, P^+, t_0)$. Since P, P^{ID} , and c are static between updates, the statement follows by induction. \square

Proof of Theorem 4. The value of $\mathcal{H}(c, P, t)$ is static in between base station exchanges since P and c do not change in between updates. Consider an update occurring at $t = t_0$ involving agent i . Noting Prop. 3, we have $\mathcal{H}(c^+, P^+, t_0) \leq \mathcal{H}(c^-, P^{\text{ID},-}, t_0) = \mathcal{H}(c^-, P^-, t_0)$. Thus, the value $\mathcal{H}(c, P, t)$ is non-increasing as $t \rightarrow \infty$. Since $\text{Cov}_m(Q)$ is finite, there must exist some time t_0 after which the value of \mathcal{H} is static. Consider fixed $t > t_0$ at which some agent i communicates with the base. Since the value of \mathcal{H} does not change during the update, Alg. 1 implies that P^{ID} and c will be unchanged by the update. It follows that c and P^{ID} converge in finite time. Further, since $P_i^{\text{ID}} \subseteq P_i$ for any i (Prop. 2), we have $P_i^{\text{ID},-} = P_i^{\text{ID},+} = P_i^{\text{ID},-} \cup P_i^{\text{add}}(c_i^+) = P_i^+$. Noting the persistence of communication imposed by $\bar{\Delta}$, the same logic implies that after some finite time, P and P^{ID} are concurrent.

We show that the limiting configuration is Pareto optimal. Consider t_0 , such that for all $t > t_0$, c and P are static and P is an m -partition of Q . Since timers T_i are only reset when P is altered, we assume without loss of generality that $T_i = 0$ for all $i \in \{1, \dots, m\}$ at any $t > t_0$. Suppose that agent i communicates with the base station at time $t > t_0$. With these assumptions, Alg. 1 implies that there exists no $k \in P_i$ such that $\sum_{h \in P_i} d_{P_i}(h, k) \Phi(h, t) < \sum_{h \in P_i} d_{P_i}(h, c_i) \Phi(h, t)$ (otherwise the value of \mathcal{H} would be lowered by moving c_i).

Similarly, for any $k \in P_j$ with $j \neq i$ that is adjacent to P_i , we must have $\frac{1}{s_i} d_{P_i \cup \{k\}}(c_i, k) \geq \frac{1}{s_j} d_{P_j}(c_j, k)$. Indeed, if this were not so, there would exist $k \in P_i^{\text{add}}(c_i^+) \setminus P_i^-$, contradicting the convergence assumption, since $P_i^{\text{add}}(c_i^+) = P_i^+$. As such, for any i , there is no $Q' \subset Q \setminus P_i$ such that $\sum_{k \in Q'} \frac{1}{s_i} d_{P_i \cup Q'}(c_i, k) < \sum_{k \in Q'} \min\{\frac{1}{s_j} d_{P_j}(c_j, k) | k \in P_j, j \neq i\}$, which implies statement (ii) of Def. 2. \square

Proof of Theorem 5. By Assumption 2, no agent ever leaves its assigned coverage region or enters its prohibited region, provided no abrupt changes to these regions occur during an update. Therefore, if no update ever occurs in which the vertex corresponding to the communicating agent's location is removed from the relevant agent's coverage region, then the statement is immediate. Suppose now that, at some time $t = t_0$, agent i , whose location is associated with some $k \in P_i^{A,-}$, communicates with the base station and k is removed, i.e., $k \notin P_i^{A,+}$. At time t^+ , agent i executes lines 5 and 6 of Alg. 3. Thm. 3, however, guarantees that 1) there will exist a path in $G(P_i^{A,-})$ between k and the set $P_i^{\text{ID},-}, 2)$

that all vertices along this path belong to $\text{Proh}_{\text{ID}_k^+} \setminus \bigcup_{j \neq \text{ID}_k^+} P_j^+$ during the time period $(t_0, t_0 + \frac{1}{s_i} d_{P_i^{A,-}}(k, P_i^{\text{TD},-}))$, and 3) $P_i^{\text{TD},-} \subseteq P_i^A := P_i^A$ over the same interval. Therefore, if agent i immediately starts moving along the path, its location will lie exclusively within $\text{Proh}_{\text{ID}_k^+}$ until it reaches $P_i^{A,+}$.

It remains to show that no agent i ever enters $\text{Proh}_i(t)$. Whenever agent i is executing lines 1–2 of Alg. 3, it follows readily that it will not enter $\text{Proh}_i(t)$. We show that the same holds when agent i is forced to execute line 5 and 6 of Alg. 3. Without loss of generality, consider the update at time t_0 previously described. Since k is re-assigned prior to the update at time t_0 , we have $\text{Proh}_i^- = \emptyset$ (since vertices in P_i cannot be claimed unless $T_i = 0$, implying $t_0 - \omega_i^{A,-} > \tau_i^{A,-}$). Using Prop. 2, we deduce that $T_{\text{ID}_k^+}^+ > 0$ and thus no vertices in $P_i^{A,-} \cap P_{\text{ID}_k^+}^+$ can belong to $P_i^{A,\text{pd},+}$, and no vertex on the constructed path between k and $P_i^{\text{TD},-}$ belongs to $\text{Proh}_i(t_0^+)$. Since $T_i^+ > \tau_i^{A,+} > \frac{1}{s_i} d_{P_i^{A,-}}(k, P_i^{\text{TD},-})$, $\text{Proh}_i(\cdot)$ remains unchanged over the interval $(t_0, t_0 + \frac{1}{s_i} d_{P_i^{A,-}}(k, P_i^{\text{TD},-}))$. \square

REFERENCES

- [1] A. Almeida, G. Ramalho, H. Santana, P. Tedesco, T. Menezes, V. Corruble, and Y. Chevalere. Recent advances on multi-agent patrolling. In *Advances in Artificial Intelligence*, volume 3171 of *Lecture Notes in Computer Science*, pages 474–483. Springer, 2004.
- [2] J. F. Araujo, P. B. Sujit, and J. B. Sousa. Multiple UAV area decomposition and coverage. In *IEEE Symposium on Computational Intelligence for Security and Defense Applications*, pages 30–37, Singapore, April 2013.
- [3] B. Bethke, M. Valenti, and J. P. How. UAV task assignment. *IEEE Robotics & Automation Magazine*, 15(1):39–44, 2008.
- [4] E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Number 1 in Santa Fe Institute Studies in the Sciences of Complexity. Oxford University Press, 1999.
- [5] J. Cortés. Coverage optimization and spatial load balancing by robotic sensor networks. *IEEE Transactions on Automatic Control*, 55(3):749–754, 2010.
- [6] J. Cortés, S. Martínez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, 20(2):243–255, 2004.
- [7] S. R. Dixon and C. D. Wickens. Automation reliability in unmanned aerial vehicle control: A reliance-compliance model of automation dependence in high workload. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 48(3):474–486, 2006.
- [8] J. W. Durham, R. Carli, P. Frasca, and F. Bullo. Discrete partitioning and coverage control for gossiping robots. *IEEE Transactions on Robotics*, 28(2):364–378, 2012.
- [9] P. O. Fjällström. Algorithms for graph partitioning: A survey. *Linköping Electronic Articles in Computer and Information Science*, 3(10), 1998.
- [10] B. P. Gerkey and M. J. Mataric. Sold!: Auction methods for multi-robot coordination. *IEEE Transactions on Robotics and Automation*, 18(5):758–768, 2002.
- [11] G. Gutin and A. P. Punnen. *The Traveling Salesman Problem and Its Variations*. Springer, 2007.
- [12] X. Lan and M. Schwager. Planning periodic persistent monitoring trajectories for sensing robots in Gaussian random fields. In *IEEE Int. Conf. on Robotics and Automation*, pages 2415–2420, Karlsruhe, Germany, May 2013.
- [13] G. Laporte. The vehicle routing problem: An overview of exact and approximate algorithms. *European Journal of Operational Research*, 59(3):345–358, 1992.
- [14] A. Macwan, G. Nejat, and B. Benhabib. Optimal deployment of robotic teams for autonomous wilderness search and rescue. In *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, pages 4544–4549, San Francisco, CA, USA, September 2011.
- [15] G. Mathew and I. Mezić. Spectral multiscale coverage: A uniform coverage algorithm for mobile sensor networks. In *IEEE Conf. on Decision and Control*, pages 7872–7877, Shanghai, China, December 2009.
- [16] George Mathew, Amit Surana, and Igor Mezić. Uniform coverage control of mobile sensor networks for dynamic target detection. In *Decision and Control (CDC), 2010 49th IEEE Conference on*, pages 7292–7299. IEEE, 2010.
- [17] N. Mathew, S. L. Smith, and S. L. Waslander. Multirobot rendezvous planning for recharging in persistent tasks. *IEEE Transactions on Robotics*, 31(1):128–142, 2015.
- [18] I. Maza and A. Ollero. Multiple UAV cooperative searching operation using polygon area decomposition and efficient coverage algorithms. In R. Alami, R. Chatila, and H. Asama, editors, *Distributed Autonomous Robotic Systems 6*, pages 221–230. Springer, 2007.
- [19] N. Nigam. The multiple unmanned air vehicle persistent surveillance problem: A review. *Machines*, 2(1):13–72, 2014.
- [20] N. Nigam and I. Kroo. Persistent surveillance using multiple unmanned air vehicles. In *IEEE Aerospace Conference*, pages 1–14, Big Sky, MT, USA, March 2008.
- [21] A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. Wiley Series in Probability and Statistics. John Wiley & Sons, 2 edition, 2000.
- [22] J. Ousingsawat and M. G. Earl. Modified lawn-mower search pattern for areas comprised of weighted regions. In *American Control Conference*, pages 918–923, New York, USA, July 2007.
- [23] F. Pasqualetti, J. W. Durham, and F. Bullo. Cooperative patrolling via weighted tours: Performance analysis and distributed algorithms. *IEEE Transactions on Robotics*, 28(5):1181–1188, 2012.
- [24] F. Pasqualetti, F. Zanella, J. R. Peters, M. Spindler, R. Carli, and F. Bullo. Camera network coordination for intruder detection. *IEEE Transactions on Control Systems Technology*, 22(5):1669–1683, 2014.
- [25] R. Patel, P. Frasca, and F. Bullo. Centroidal area-constrained partitioning for robotic networks. *ASME Journal of Dynamic Systems, Measurement, and Control*, 136(3):031024, 2014.
- [26] R. Patel, P. Frasca, J. W. Durham, R. Carli, and F. Bullo. Dynamic partitioning and coverage control with asynchronous one-to-base-station communication. *IEEE Transactions on Control of Network Systems*, 3(1):24–33, 2016.
- [27] J. Peng and S. Akella. Coordinating multiple robots with kinodynamic constraints along specified paths. *International Journal of Robotics Research*, 24(4):295–310, 2005.
- [28] A. Pereira, H. Heidarsson, C. Oberg, D. Caron, B. Jones, and G. Sukhatme. A communication framework for cost-effective operation of AUVs in coastal regions. In A. Howard, K. Iagnemma, and A. Kelly, editors, *Field and Service Robotics*, volume 62 of *Tracts in Advanced Robotics*, pages 433–442. Springer, 2010.
- [29] J. Peters, V. Srivastava, G. Taylor, A. Surana, M. P. Eckstein, and F. Bullo. Mixed human-robot team surveillance: Integrating cognitive modeling with engineering design. *IEEE Control Systems*, 35(6):57–80, 2015.
- [30] J. H. Reif and H. Wang. Social potential fields: A distributed behavioral control for autonomous robots. *Robotics & Autonomous Systems*, 27(3):171–194, 1999.
- [31] R. C. Shah, S. Roy, S. Jain, and W. Brunette. Data MULEs: modeling and analysis of a three-tier architecture for sparse sensor networks. *Ad Hoc Networks*, 1(2-3):215–233, 2003.
- [32] R. N. Smith, Y. Chao, P. P. Li, D. A. Caron, B. H. Jones, and G. S. Sukhatme. Planning and implementing trajectories for autonomous underwater vehicles to track evolving ocean processes based on predictions from a regional ocean model. *International Journal of Robotics Research*, 29(12):1475–1497, 2010.
- [33] D. E. Soltero, M. Schwager, and D. Rus. Generating informative paths for persistent sensing in unknown environments. In *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, pages 2172–2179, Vilamoura, Portugal, October 2012.
- [34] V. Srivastava, F. Pasqualetti, and F. Bullo. Stochastic surveillance strategies for spatial quickest detection. *International Journal of Robotics Research*, 32(12):1438–1458, 2013.
- [35] P. Toth and D. Vigo, editors. *The Vehicle Routing Problem*. Monographs on Discrete Mathematics and Applications. SIAM, 2001.
- [36] J. Wood and J. K. Hedrick. Space partitioning and classification for multi-target search and tracking by heterogeneous unmanned aerial system teams. *Infotech@ Aerospace*, March 2011.
- [37] P. R. Wurman, R. D’Andrea, and M. Mountz. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI Magazine*, 29(1):9–20, 2008.