

Multi-Robot Monitoring in Dynamic Environments with Guaranteed Currency of Observations

Stephen L. Smith

Daniela Rus

Abstract—In this paper we consider the problem of monitoring a known set of stationary features (or locations of interest) in an environment. To observe a feature, a robot must visit its location. Each feature changes over time, and we assume that the currency, or accuracy of an observation decays linearly with time. Thus, robots must repeatedly visit the features to update their observations. Each feature has a known rate of change, and so the frequency of visits to a feature should be proportional to its rate. The goal is to route the robots so as to minimize the maximum change of a feature between observations. We focus on the asymptotic regime of a large number of features distributed according to a probability density function. In this regime we determine a lower bound on the maximum change of a feature between visits, and develop a robot control policy that, with probability one, performs within a factor of two of the optimal. We also provide a single robot lower bound which holds outside of the asymptotic regime, and present a heuristic algorithm motivated by our asymptotic analysis.

I. INTRODUCTION

Consider the following problem. An environment (such as a city, or a building) contains known static features of interest (such as intersections in the city, or rooms in the building). A group of robots is tasked with monitoring the features by visiting their locations. The environment is dynamic, and thus the properties of each feature change over time (i.e., the amount of traffic in each intersection, or the layout and number of people in each room). Features may change on different time scales. Thus, the robots must repeatedly visit the features to update their observations. The frequency of visits to each feature should be proportional to that feature's rate of change. The problem is to determine routes for the robots that allow them to guarantee the currency (or accuracy) of their most recent observations of each feature. That is, to determine routes that minimize the maximum change of a feature between visits (observations). We call this problem *persistent monitoring*.

In this paper we assume that we are given n robots and m features. Each feature consists of a known location and a constant rate of change (which may be different among features). We show that the problem of computing routes to minimize the maximum change of a feature between visits is NP-hard. However, in the regime of many randomly distributed features, we show that one can obtain strong performance guarantees. In particular, for this regime we present a lower bound on the optimal performance, and an algorithm that




Fig. 1. A scenario of three robots deployed in an environment. The yellow disks represent features in the map that must be repeatedly visited by the robots. The area of each disk represents its rate of change.

computes robot routes within a factor of two of this optimal. The lower bound and algorithm utilize a discretization of the rates of change of each feature from a number in $(0, 1]$ to a value in the set $\{1, 1/2, 1/4, \dots\}$. We show that this discretization results in at most a factor of two loss in performance. Thus, by developing an optimal routing algorithm for discretized problems, we create a two-approximation for non-discretized problems. The algorithm relies on computing shortest paths through subsets of features. We then use the algorithm to motivate a computationally efficient heuristic that operates in any regime. We show through simulation that the heuristic achieves very good performance.

Persistent monitoring is related to research in optimal search, coverage, patrolling, and dynamic vehicle routing. In optimal search [1], robots must move through an environment to find stationary targets. Given imprecise sensors, the search process may require repeatedly visiting areas that have high probability of containing a target. Recent search approaches include a decision-making framework [2], and a robust method for imprecise probabilistic descriptions [3].

In sweep coverage [4], a robot must move through the environment so as to cover the entire region with its sensor. Variants of this problem include on-line coverage [5], where the robot has no *a priori* information about the environment, and dynamic coverage [6], where each point in the environment requires a pre-specified “amount” of coverage. In [7], a dynamic coverage problem is considered where sensor continually surveys regions of interest by moving according to a Markov Chain. The problem considered in this paper differs from sweep coverage in that the features must be repeatedly visited, and differs from dynamic coverage in that we seek to maximize the currency of observations.

In patrolling [8], [9], a region must be continually surveyed by a group of robots. The goal is to visit each point in space with equal frequency. A variant of patrolling is

This material is based upon work supported in part by ONR-MURI Award N00014-09-1-1051.

S. L. Smith and D. Rus are with the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 (slsmith@mit.edu; rus@csail.mit.edu)

considered in [10] for continual target surveillance. Persistent monitoring generalizes patrolling by requiring that the frequency of visits to a feature depends on its rate of change.

Finally, persistent monitoring is related to vehicle routing and dynamic vehicle routing (DVR) problems. One example is the period routing problem [11], where each customer must be visited a specified number of times per week. A solution consists of an assignment of customers to days of the week, and a set of routes for the vehicles on each day. In DVR, new customers arrive sequentially over time, and the goal is to minimize the expected amount of time a customer remains waiting for service. Recently, DVR has been addressed for customers with different priority levels [12]. In order to balance the load of customers among vehicles, equitable partitioning policies have been developed [13].

The organization of this paper is as follows. In Section II we review shortest paths and tours. In Section III we formalize persistent monitoring, and introduce a method for discretizing each feature's rate of change. In Section IV we consider present a lower bound and a constant factor algorithm for the asymptotic regime of many randomly distributed features. Finally, in Section V we provide a general lower bound and a computationally efficient algorithm.

II. BACKGROUND MATERIAL

Consider a complete undirected, and weighted graph $G = (V, E, w)$ with vertex set V , edge set $E \subseteq V \times V$, and non-negative edge weights $w : E \rightarrow \mathbb{R}_{\geq 0}$. An induced subgraph of G is a graph $H = (V_H, E_H)$ and $E_H := \{\{\mathbf{q}_i, \mathbf{q}_k\} \in E : \mathbf{q}_i, \mathbf{q}_k \in V_H\}$.

The traveling salesman problem (TSP) on $G = (V, E, w)$ is to find a cycle that visits each vertex in V exactly once, and has minimum length, denoted $\text{TSP}(G)$ (the length of a cycle is the sum of its edge weights). This cycle is also known as the *minimum Hamiltonian cycle*. A similar problem is to find the minimum length *path* (rather than closed cycle). This is known as the *minimum Hamiltonian path*, and we denote its length by $\text{MHP}(G)$.

A. The Euclidean Traveling Salesperson Problem

A graph is *Euclidean* if the vertices are embedded in \mathbb{R}^d and the edge weights are given by Euclidean distances. The following result characterizes the asymptotic length of the TSP tour on a random Euclidean graph. For simplicity of notation, we write the tour length on a Euclidean graph as $\text{TSP}(V)$, rather than $\text{TSP}(G)$.

Theorem II.1 (TSP Tour Length, [14]) *Consider a Euclidean graph $G = (V, E, w)$, where the m vertices in V are independently and identically distributed in a compact set \mathcal{E} according to the probability distribution $f : \mathbb{R}^2 \rightarrow \mathbb{R}_{\geq 0}$. Then, there exists a constant β_{TSP} such that with probability one*

$$\lim_{m \rightarrow +\infty} \frac{\text{TSP}(V)}{\sqrt{m}} = \beta_{\text{TSP}} \int_{\mathcal{E}} \sqrt{\bar{f}(\mathbf{q})} d\mathbf{q}, \quad (1)$$

where \bar{f} is the absolutely continuous part of f .

The constant β_{TSP} has been estimated numerically as $\beta_{\text{TSP}} \approx 0.7120 \pm 0.0002$, [15]. In [16], the authors note that if \mathcal{E} is ‘‘fairly compact [square] and fairly convex’’, then equation (1) provides an adequate estimate of the optimal TSP tour length for values of m as low as 15. Defining

$$\Psi_f(\mathcal{E}) := \beta_{\text{TSP}} \int_{\mathcal{E}} \sqrt{\bar{f}(\mathbf{q})} d\mathbf{q}, \quad (2)$$

we can state the following result, whose proof is contained in the Appendix.

Lemma II.2 (TSP Length in Cell) *Consider a Euclidean graph where the m vertices in V are independently and identically distributed in a compact set $\mathcal{E} \subset \mathbb{R}^2$ according to a density f . For any set $C \subset \mathcal{E}$, with probability one, $\lim_{m \rightarrow +\infty} \text{TSP}(C \cap V) / \sqrt{m} = \Phi_f(C)$.*

B. The Minimum Hamiltonian Path

It is straightforward to use Theorem II.1 to verify that the asymptotic length of the Minimum Hamiltonian Path (MHP) is equal to that of the TSP tour, and thus with probability one, $\lim_{m \rightarrow +\infty} \text{MHP}(V) / \sqrt{m} = \Psi_f(\mathcal{E})$.

A variation on the MHP is the min-max n -MHP problem, where the goal is to find a partition of V into n sets V_1, \dots, V_n such that $\text{MHP}_n(V_1, \dots, V_n) := \max_i \text{MHP}(V_i)$ is minimized. Given m i.i.d. vertices with probability distribution f , Theorem II.1 can be used to show that

$$\lim_{m \rightarrow +\infty} \frac{\text{MHP}_n(V_1, \dots, V_n)}{\sqrt{m}} = \frac{\Psi_f(\mathcal{E})}{n}. \quad (3)$$

C. Methods for Solving the TSP

The TSP and MHP problems are NP-hard [17]. However, for metric problem instances (i.e., when the edge weights satisfy the triangle inequality) there exist many good approximation algorithms. One example is Christofides' Algorithm which finds a tour no longer than $3/2$ times the optimal in $O(m^3)$ computation time [17]. In addition, there are very good heuristic solvers, such as the Lin-Kernighan heuristic, which typically finds tours within 5% of the optimal in $O(m^{2.2})$ computation time [17].

III. PROBLEM STATEMENT AND APPROACH

Consider m features (i.e., points of interest) in a compact environment $\mathcal{E} \subset \mathbb{R}^2$. The environment, and its features can be represented as a complete, undirected, edge and vertex weighted graph $G = (V, E, w, \phi)$. The vertices $V = \{\mathbf{q}_1, \dots, \mathbf{q}_m\} \subset \mathcal{E}$ are the m distinct feature locations. The edges $E = \{\{\mathbf{q}_i, \mathbf{q}_j\} : \mathbf{q}_i, \mathbf{q}_j \in V\}$ connect all vertex pairs. The edge weights $w : E \rightarrow \mathbb{R}_{>0}$ give the travel time of each edge. We assume that w satisfies the triangle inequality. The vertex weights $\phi : V \rightarrow (0, 1]$ give the rate of change of each feature. For simplicity of notation we write $\phi_i := \phi(\mathbf{q}_i)$.

We consider n robots on the graph G . The infinite trajectories of the robots, denoted $t \mapsto (\mathbf{p}_1(t), \dots, \mathbf{p}_n(t)) =: \mathbf{p}(t)$, are paths on the graph. If a robot is located at vertex \mathbf{q}_i at time t_1 , then it can reach vertex \mathbf{q}_j at time $t_1 + w(\{\mathbf{q}_i, \mathbf{q}_j\})$. We write the initial robot positions as $\mathbf{p}(0) \subseteq V$. With this notation we can define a problem instance.




Fig. 2. A possible time evolution of the urgency of vertex \mathbf{q}_i .

Definition III.1 (Problem Instance) A problem instance is a tuple $I := (G, \mathbf{p}(0))$, where $G = (V, E, w, \phi)$ is a graph representing the features in the environment, and $\mathbf{p}(0) \subset V$ gives the initial robot positions.

A. The Maximum Urgency Performance Metric

A robot control policy is a map P that takes as an input a problem instance $I = (G, \mathbf{p}(0))$ and returns an n -tuple of robot trajectories $P : I \mapsto \mathbf{p}$. Given robot trajectories, we let $r_i(t)$ denote the most recent time $\tau \leq t$ at which vertex \mathbf{q}_i was visited (observed) by a robot. That is,

$$r_i(t) = \max \{ \tau \in [0, t] : \exists j \text{ with } \mathbf{p}_j(\tau) = \mathbf{q}_i \}, \quad (4)$$

where we define the maximum of the empty set to be zero. The *urgency* of a vertex \mathbf{q}_i at time $t \geq 0$ represents how out of date its most recent observation is (see Figure 2), and is defined as $U_i(t) := \phi_i(t - r_i(t))$.

We seek to minimize the maximum urgency of each vertex over time. For a policy P and problem instance I we define the infinite-horizon *maximum urgency* of vertex \mathbf{q}_i to be

$$M_i(P(I), I) := \sup_{t \in \mathbb{R}_{\geq 0}} \phi_i(t - r_i(t)),$$

where $P(I)$ are the robot trajectories generated by policy P . The corresponding cost function for the n -robots is

$$C(P(I), I) := \max_{i \in \{1, \dots, m\}} M_i(P(I), I). \quad (5)$$

We are now ready to state the problem.

Problem Statement: Determine the policy P^* such that for every problem instance I and policy P , we have $C(P^*(I), I) \leq C(P(I), I)$.

At times we will write $C^*(I)$ in place of $C(P^*(I), I)$. Note that the problem of determining P^* is NP-hard. This can be seen by noting that in simplest case where $n = 1$ and $\phi_i = 1$ for all $i \in \{1, \dots, m\}$, the problem reduces to the TSP. Because of this, we seek policies that perform within a constant factor of C^* .

B. Discretized Rates of Change

Given a problem instance I with non-discretized rates ϕ , we can create a problem instance \bar{I} with discretized rates $\bar{\phi}$. For each vertex i , the discretized rate $\bar{\phi}_i$ is defined as $\bar{\phi}_i := \frac{1}{2^{k-1}}$, where k is the largest integer such that $\bar{\phi}_i \geq \phi_i$. A discretization has ℓ levels, $1, 1/2, 1/4, \dots, (1/2)^{\ell-1}$, where level k has rate $(1/2)^{k-1}$. We let V_k denote the set of vertices in level k and $m_k = |V_k|$ denote the number of vertices in the level. By construction, we have $\sum_{k=1}^{\ell} m_k = m$, and

$$V = \cup_{k=1}^{\ell} V_k \quad \text{with } V_k \cap V_j = \emptyset \text{ for all } j \neq k.$$

The discretization is useful because each vertex in V_k must be visited twice as often as each vertex in V_{k+1} . Thus, we will seek policies P which minimize $C(P(\bar{I}), \bar{I})$ for discretized problem instances \bar{I} . But, how well do the trajectories $P(\bar{I})$, perform on the non-discretized problem I ? The next lemma shows that $C(P(\bar{I}), I)$ is within a factor of two of $C(P(\bar{I}), \bar{I})$, and the performance of an optimal policy on I is upper bounded by its performance on \bar{I} .

Lemma III.2 (Discretized Approximation) Consider a policy P , a problem instance I , and the corresponding discretized problem instance \bar{I} . Then,

$$\frac{1}{2}C(P(\bar{I}), \bar{I}) \leq C(P(\bar{I}), I) \leq C(P(\bar{I}), \bar{I}).$$

Moreover, for the optimal policy P^* we have $C^*(\bar{I})/2 \leq C^*(I) \leq C^*(\bar{I})$.

Proof: We begin by proving the first expression of the lemma. For the discretized problem \bar{I} , the policy P generates trajectories $P(\bar{I})$, and the maximum urgency on problem \bar{I} is $C(P(\bar{I}), \bar{I})$. Therefore, for every vertex i , we have $M_i(P(\bar{I}), \bar{I}) \leq C(P(\bar{I}), \bar{I})$, with equality for some vertex j . If we use the trajectories $P(\bar{I})$ on problem instance I , then for each vertex i $M_i(P(\bar{I}), I) = (\phi_i/\bar{\phi}_i)M_i(P(\bar{I}), \bar{I}) \leq C(P(\bar{I}), \bar{I})$, where we applied the fact that $\bar{\phi}_i \geq \phi_i$. Thus, $C(P(\bar{I}), I) \leq C(P(\bar{I}), \bar{I})$, and we have the RHS of the expression. In addition, for vertex j , we have $\bar{\phi}_j \leq 2\phi_j$, and thus $C(P(\bar{I}), \bar{I}) \geq C(P(\bar{I}), I)/2$, giving us the LHS.

To prove the ‘‘Moreover’’ part of the lemma, consider the trajectories $P(I)$ on problem instance I , and the maximum urgency $C(P(I), I)$. Using the trajectories $P(I)$ on the discretized instance \bar{I} , we obtain that for every vertex i , $M_i(P(I), \bar{I}) \leq (\bar{\phi}_i/\phi_i)C(P(I), I) \leq 2C(P(I), I)$, where we have used $\bar{\phi}_i \leq 2\phi_i$. Therefore,

$$C(P(I), \bar{I}) \leq 2C(P(I), I). \quad (6)$$

Finally, consider the optimal policy P^* . Since the policy is optimal it follows directly that

$$C(P^*(I), I) \leq C(P^*(\bar{I}), I), \quad \text{and} \quad (7)$$

$$C(P^*(\bar{I}), \bar{I}) \leq C(P^*(I), \bar{I}). \quad (8)$$

Combining (6), (7), and (8) with the RHS of the first expression in the lemma, we obtain the desired result. ■

C. Discretized Optimization Formulation

In this section we illustrate the advantages of dealing with discretized problem instances. Consider a policy P , a discretized problem instance \bar{I} , and suppose that $C(P(\bar{I}), \bar{I}) < B$ for some $B \in \mathbb{R}_{>0}$. Let the initial time be $t = 0$. Since the rate of change of all vertices in V_k is $(1/2)^{k-1}$, every vertex in V_k must be visited in the time interval $[0, 2^{k-1}B)$. If this does not hold, then there exists a vertex in V_k such that the time between successive visits is no smaller than $2^{k-1}B$, implying that the urgency is no smaller than B . Consider the time intervals $[(s-1)B, sB)$ for $s \in \{1, \dots, 2^{\ell-1}\}$. (Recall that ℓ is the number of levels in the discretization.)




Fig. 3. Discretization of the rates of change into $\ell = 3$ levels, 1, 1/2, 1/4. The subsets of vertices visited in each time interval are shown. The constraints (9) results in 4 constraints, listed on the right.

In each of the $2^{\ell-1}$ time intervals, some (possibly empty) subset of the vertices in V_k will be visited. Let us denote the vertices in V_k visited in time slot $[(s-1)B, sB)$ by $V_{k,s}$. Since $C(P(\bar{I}), \bar{I}) < B$, the sets $V_{k,s}$ must satisfy

$$\bigcup_{s=(j-1)2^{k-1}+1}^{j2^{k-1}} V_{k,s} = V_k \quad \text{for each } k \in \{1, \dots, \ell\} \quad (9)$$

$$j \in \{1, \dots, 2^{\ell-k}\}.$$

The constraints in (9) capture the fact that all vertices in V_k must be visited in each of the time intervals $[0, 2^{k-1}B)$, $[2^{k-1}B, 2^k B)$, and so on, to $[(2^{\ell-1} - 2^{k-1})B, 2^{\ell-1}B)$. An example is shown in Figure 3. Thus, the problem consists of

- (i) an assignment, satisfying (9), of the vertices in V_k to the sets $V_{k,1}, \dots, V_{k,2^{\ell-1}}$, and
- (ii) robot trajectories to visit the vertices $\bigcup_{k=1}^{\ell} V_{k,s}$, for each time slot $s \in \{1, \dots, 2^{\ell-1}\}$.

This problem can be written as an integer linear program (ILP) and is related to the *period routing problem*. However, the ILP can be solved only for small problem instances, and thus most research focuses on heuristics [11].

IV. ASYMPTOTIC LOWER BOUND AND POLICY

In this section we characterize scaling of cost function (5) with the number of features. To do this we consider a subset of the discretized problem instances \bar{I} , called *geometric discretized problem instances*. In these problem instances, the m vertices in V are embedded in the environment \mathcal{E} , and the weight on each edge $\{\mathbf{q}_i, \mathbf{q}_j\} \in E$ is equal to the length of the shortest path in \mathcal{E} from \mathbf{q}_i to \mathbf{q}_j . Note that if \mathcal{E} is convex, then $w(\{\mathbf{q}_i, \mathbf{q}_j\}) = \|\mathbf{q}_i - \mathbf{q}_j\|_2$.

A. Asymptotic Lower Bound

Recall that for a discretized problem instance, the m vertices in V are discretized into the sets V_1, \dots, V_{ℓ} , which contain m_1, \dots, m_{ℓ} vertices, with $m = \sum_{k=1}^{\ell} m_k$. Lemma III.2 showed that $C^*(I) \geq C^*(\bar{I})/2$, and thus the following theorem also provides a lower bound on $C^*(I)$.

Theorem IV.1 (Asymptotic Lower Bound) *Consider a geometric discretized problem instance \bar{I} , where the m vertices in V are independently and identically distributed in \mathcal{E} according to the density f . Then,*

$$C^*(\bar{I}) \geq \frac{\Psi_f(\mathcal{E})}{n} \sum_{k=1}^{\ell} \max\{2^{-k}, 2^{1-\ell}\} \left(\sum_{j=1}^k m_j \right)^{1/2},$$

with probability one as $\min_k m_k \rightarrow +\infty$, where $\Psi_f(\mathcal{E})$ is defined in equation (2).

Proof: Consider the problem instance \bar{I} and let us lower bound the shortest path distance w by the Euclidean distance w_E . Now, suppose that for a policy P , the maximum urgency $C(P(\bar{I}), \bar{I}) < B$. Then, using the notation introduced in Section III-C, we see that for each $s \in \{1, \dots, 2^{\ell-1}\}$, all vertices in $\bigcup_{k=1}^{\ell} V_{k,s}$ must be visited in the time interval $[(s-1)B, sB)$. From Section II-B, the time to visit all vertices in $\bigcup_{k=1}^{\ell} V_{k,s}$ is given by the min-max n -MHP, and thus

$$\text{MHP}_n \left(\bigcup_{k=1}^{\ell} V_{k,s} \right) < B \quad \text{for each } s \in \{1, \dots, 2^{\ell-1}\}.$$

We can restate the above conditions as $\max_{s \in \{1, \dots, 2^{\ell-1}\}} \text{MHP}_n \left(\bigcup_{k=1}^{\ell} V_{k,s} \right) < B$. Let $m_{k,s}$ denote the number of vertices in $V_{k,s}$. Then, letting $\min_k m_k$ become very large, we see that $\sum_{k=1}^{\ell} m_{k,s}$ is very large for each s , and we can apply the result on the asymptotic length of MHP_n in equation (3) to obtain

$$\max_{s \in \{1, \dots, 2^{\ell-1}\}} \frac{\Psi_f(\mathcal{E})}{n} \sqrt{\sum_{k=1}^{\ell} m_{k,s}} < B. \quad (10)$$

Now, given a set of positive numbers, the average is necessarily a lower bound on the maximum. Thus, a necessary condition for equation (10) to be satisfied is that

$$\frac{\Psi_f(\mathcal{E})}{n 2^{\ell-1}} \sum_{s=1}^{2^{\ell-1}} \sqrt{\sum_{k=1}^{\ell} m_{k,s}} < B. \quad (11)$$

The minimizer of the following optimization problem provides a lower bound on the maximum urgency $C(P(\bar{I}), \bar{I})$ for every policy P :

$$\begin{aligned} & \text{minimize} && \sum_{s=1}^{2^{\ell-1}} \sqrt{\sum_{k=1}^{\ell} m_{k,s}}, \\ & \text{subject to} && \sum_{s=(j-1)2^{k-1}+1}^{j2^{k-1}} m_{k,s} = m_k \\ & && \text{for each } k \in \{1, \dots, \ell\}, j \in \{1, \dots, 2^{\ell-k}\}. \end{aligned}$$

The cost function is monotonically increasing, concave, and sub-additive on the non-negative real numbers. The minimization is subject to a set of linear equality constraints. To determine a minimizer, begin by looking at the problem when $\ell = 2$: minimize $\sqrt{m_1 + m_{2,1}} + \sqrt{m_1 + m_{2,2}}$, subject to $m_{2,1} + m_{2,2} = m_2$. Since the square root is sub-additive, the minimum value is given by $\sqrt{m_1 + m_2} + \sqrt{m_1}$ and the corresponding solutions are $m_{2,1} = 0$ and $m_{2,2} = m_1$ or $m_{2,1} = m_1$ and $m_{2,2} = 0$. This argument can easily be extended to k variables, from which we see that a minimizer of the original optimization problem is to set $m_{k,s} = m_k$ if $s \pmod k \equiv 1$ and $m_{k,s} = 0$ otherwise (for each s and k), where the modulo notation denotes arithmetic modulo k , i.e. $k+1 \pmod k \equiv 1$. Upon simplifying, the corresponding value of the cost function is given by

$$\sum_{k=1}^{\ell-1} 2^{\ell-1-k} \sqrt{m_1 + \dots + m_k} + \sqrt{m_1 + \dots + m_{\ell}}.$$

Combining the cost function value with the lower bound from equation (11), we obtain the desired result. ■

Remark IV.2 (Comparison to a simple policy) To get a better feel for the lower bound in Theorem IV.1, consider a discretized problem \bar{I} , with one robot, and with $\bar{\phi}$ taking values in $1, 1/2, 1/4$. The asymptotic lower bound yields

$$C^*(\bar{I}) \geq \frac{\Psi_f(\mathcal{E})}{4} (\sqrt{m_1 + m_2 + m_3} + \sqrt{m_1 + m_2} + 2\sqrt{m_1}).$$

If one were to use a simple policy (SP) of performing cycles of a single TSP tour computed through all vertices, we would have $C(\text{SP}(\bar{I}), \bar{I}) = \Psi_f(\mathcal{E})\sqrt{m_1 + m_2 + m_3}$. Comparing these bounds we get

$$\frac{C^*(\bar{I})}{C(\text{SP}(\bar{I}), \bar{I})} \geq \frac{1}{4} \left(1 + \frac{\sqrt{m_1 + m_2} + 2\sqrt{m_1}}{\sqrt{m_1 + m_2 + m_3}} \right).$$

Thus, if $m_1 \gg \max\{m_2, m_3\}$, the simple policy performs nearly optimally. However, when $m_3 \gg \max\{m_1, m_2\}$, the lower bound is approximately 1/4 of the simple policy. •

B. An Asymptotically Optimal Policy

We now introduce the PARTITION-TOUR policy which, in the limit as the number of features becomes very large, attains the lower bound in Theorem IV.1. The policy operates as follows. The region is partitioned into n regions, one for each robot. Within each region, $2^{\ell-1}$ tours are computed through its vertices. Each tour visits all vertices in V_1 , 1/2 of the vertices in V_2 , 1/4 of the vertices in V_3 , and so on, visiting $(1/2)^{\ell-1}$ of the vertices in V_ℓ . Each tour uses the same “large-scale” structure, defined by a macro-TSP, but utilizes local re-optimizations to shorten tour lengths. An example of the policy is illustrated in Figure 4.


The PARTITION-TOUR (PT) policy

Input: A discretized problem instance \bar{I} with m features distributed in the environment \mathcal{E} according to density f and a positive integer α .


Output: n robot trajectories.

- 1: Partition \mathcal{E} into n regions R_1, \dots, R_n such that $\Phi_f(R_p) = \Phi_f(\mathcal{E})/n$ for each $p \in \{1, \dots, n\}$.
 - 2: Assign one robot to each region.
 - 3: **for** $p = 1$ to n **do**
 - 4: Partition R_p into $M = \alpha 2^{\ell-1}$ cells C_1, \dots, C_M , such that $\Phi_f(C_i) = \Phi_f(R_p)/M$, for each $i \in \{1, \dots, M\}$.
 - 5: Compute a macro-TSP tour to visit each cell C_i . Relabel cells according to their order on this tour.
 - 6: **for** $T = 1$ to $2^{\ell-1}$ **do** /* creating tours */
 - 7: Include the vertices $V_j \cap C_i$ in tour T if and only if $i + T - 1 \pmod{2^{j-1}} \equiv 1$.
 - 8: In each cell C_i , compute a TSP tour through all included vertices.
 - 9: Stitch the cell tours together using the macro-TSP to create tour T .
 - 10: Robot p 's trajectory is a shortest path from $\mathbf{p}_p(0)$ to a vertex on tour 1 followed by the sequence tour 1, tour 2, ..., tour $2^{\ell-1}$, tour 1, ...
-


Note that by partitioning \mathcal{E} such that each $\Phi_f(R_p)$ is equal (and partitioning R_p such that each $\Phi_f(C_i)$ is equal),



(a) The left figure shows a partitioning into 16 cells. The right figure shows the macro-TSP which orders the cells.



(b) Construction of Tour 1. On the left are the tours in each cell, and the right shows the complete tour. The tour includes all vertices in V_1 and the vertices of V_2 in every other cell.



(c) Construction of Tour 2. On the left are the tours in each cell, and the right shows the complete tour. The tour includes all vertices in V_1 and the vertices of V_2 which were not included in Tour 1.

Fig. 4. The PARTITION-TOUR policy for $\ell = 2$. Yellow disks have rate of change 1 and blue disks have a rate of change 1/2.

we are creating an equitable partition [13]. The partitions are equitable in the sense that the tour lengths through the vertices in each region are equal, see Lemma II.2.

We now characterize the asymptotic performance of the PT policy. The proof is contained in the Appendix.

Theorem IV.3 (PARTITION-TOUR Upper Bound)

Consider a discretized problem instance \bar{I} , where the m vertices in V are independently and identically distributed in a convex region \mathcal{E} according to the density f . Then, with probability one,

$$\frac{C(\text{PT}(\bar{I}), \bar{I})}{C^*(\bar{I})} \rightarrow 1^+, \quad \text{as } \min_k m_k \rightarrow +\infty. \quad (12)$$

We can extend Theorem IV.3 to non-convex regions that can be written as the union of a finite number of convex sets.

Corollary IV.4 (Non-convex region) Consider a problem instance I , where the m vertices in V are independently and identically distributed according to f , in a region \mathcal{E} that is the union of a finite number of convex sets. Then, limit (12) still holds with probability one.

Finally, combining Theorem IV.3 with Lemma IV.3, we can extend our result to non-discretized problem instances.

Corollary IV.5 (Non-Discretized Performance) Consider a problem instance I , where the m vertices in V are independently and identically distributed in \mathcal{E} according to the density f . Then, with probability one,

$$\frac{C(\text{PT}(\bar{I}), I)}{C^*(I)} \leq 2, \quad \text{as } \min_k m_k \rightarrow +\infty.$$

V. GENERAL LOWER BOUND AND POLICY

In this section we provide a general lower bound, a computationally efficient policy, and simulation results.

A. Lower Bound for a Single Robot

Here we provide a single robot lower bound that holds outside of the asymptotic regime. In a multiset S , the second smallest value is defined as the minimum in $S \setminus \{\min S\}$.

Proposition V.1 (Single robot lower bound) Consider a single robot problem instance I . Then, for every $\bar{V} \subseteq V$,

$$C^* \geq \Phi_2(\bar{V})\text{TSP}(\bar{V}),$$

where $\Phi_2(\bar{V})$ is the second smallest value in the multi-set $\{\phi(\mathbf{q}) : \mathbf{q} \in \bar{V}\}$.

Proof: Let C^* be the optimal value, and consider the corresponding policy P^* . Take any vertex set $\bar{V} \subseteq V$, and consider the induced subgraph of G , denoted $H = (\bar{V}, \bar{E})$. Select a vertex $\mathbf{q} \in \bar{V}$ that has minimum weight, i.e., $\phi(\mathbf{q}) \leq \phi(\mathbf{v})$ for all $\mathbf{v} \in \bar{V}$. Now, consider a time \bar{t} during the evolution of policy P^* when the robot is located at \mathbf{q} . Since the maximum urgency is C^* , the robot must return to \mathbf{q} by $\mathcal{T} := \bar{t} + C^*/\phi(\mathbf{q})$.

For each vertex $\mathbf{v} \in \bar{V} \setminus \mathbf{q}$, the time of the most recent visit prior to \bar{t} is $t_{\text{last}}(\mathbf{v}) \leq \bar{t} - w(\{\mathbf{v}, \mathbf{q}\})$. The time between visits for vertex \mathbf{v} can be no more than $C^*/\phi(\mathbf{v})$. Therefore, every vertex in \bar{V} must be visited by time \mathcal{T} .

Let us choose the last vertex in $\bar{V} \setminus \mathbf{q}$ to be visited in $[\bar{t}, \bar{t} + \mathcal{T})$, call it vertex \mathbf{v} , and denote the time of its first visit after time \bar{t} as $t_{\text{next}}(\mathbf{v})$. Every vertex in $\bar{V} \setminus \{\mathbf{q}, \mathbf{v}\}$ is visited between time \bar{t} and the visit to \mathbf{v} . Thus, the earliest time at which \mathbf{v} could be visited is $t_{\text{next}}(\mathbf{v}) \geq \text{MHP}(\mathbf{q}, \bar{V} \setminus \{\mathbf{q}, \mathbf{v}\}, \mathbf{v}) + \bar{t}$, where $\text{MHP}(\mathbf{q}, \bar{V} \setminus \{\mathbf{q}, \mathbf{v}\}, \mathbf{v})$ is the length of the shortest path that starts at \mathbf{q} , passes through all vertices in $\bar{V} \setminus \{\mathbf{q}, \mathbf{v}\}$, and terminates at \mathbf{v} . For vertex \mathbf{v} , the time between successive visits is

$$t_{\text{next}}(\mathbf{v}) - t_{\text{prev}}(\mathbf{v}) \geq \text{MHP}(\mathbf{q}, \bar{V} \setminus \{\mathbf{q}, \mathbf{v}\}, \mathbf{v}) + w(\{\mathbf{v}, \mathbf{q}\}) \geq \text{TSP}(\bar{V}).$$

Since the maximum urgency is C^* , we must have that $\text{TSP}(\bar{V}) \leq C^*/\phi(\mathbf{v})$. However, the vertex \mathbf{q} has the smallest

value of $\phi(\mathbf{q})$, and thus $\phi(\mathbf{v}) \geq \Phi_2(\bar{V})$, where $\Phi_2(\bar{V})$ is the second smallest value in $\{\phi(\mathbf{q}) : \mathbf{q} \in \bar{V}\}$. Combining the previous two inequalities we obtain the desired result. ■

Proposition V.1 provides $2^m - m - 1$ lower bounds, one for every subset $\bar{V} \subseteq V$ that has at least two elements (for a set \bar{V} with one or zero elements, the bound trivially returns zero). However, in the following lemma we show that only a small subset of these constraints need actually be computed.

Lemma V.2 (Computing constraints) Assume, without loss of generality, that $\phi_1 \geq \phi_2 \geq \dots \geq \phi_m$. Then, the largest lower bound in Theorem V.1 can be computed by searching over the $m(m-1)/2$ sets of the form $\bar{V} = \{\mathbf{q}_1, \dots, \mathbf{q}_i, \mathbf{q}_j\}$ where $j > i$.

Proof: Assume that $\phi_1 \geq \phi_2 \geq \dots \geq \phi_m$ and consider an arbitrary set $\hat{V} \subseteq V$ whose vertices have indices $\hat{J} \subseteq \{1, \dots, m\}$. Let us construct a set $\bar{V} := \{\mathbf{q}_1, \dots, \mathbf{q}_i, \mathbf{q}_j\}$ ($j > i$) with index set $\bar{J} \supseteq \hat{J}$ that does not decrease the lower bound. To do this, let $j := \max\{\hat{J}\}$, let $i := \max\{\hat{J} \setminus \{j\}\}$, and let $\bar{V} = \{\mathbf{q}_1, \dots, \mathbf{q}_i, \mathbf{q}_j\}$. Notice that $\Phi_2(\bar{V}) = \Phi_2(\hat{V}) = \phi_i$. In addition, $\hat{V} \subseteq \bar{V}$ which implies that $\text{TSP}(\hat{V}) \leq \text{TSP}(\bar{V})$ and completes the proof. ■

Note that if we use a method for computing approximate TSP tours that runs in $O(m^p)$ time, then the lower bound can be computed in $O(m^{2p})$.

B. A Computationally Efficient Heuristic

The PARTITION-TOUR policy is asymptotically optimal for randomly distributed vertices, and it gives us insight into the structure of routing for persistent monitoring. However, it requires knowledge of the distribution f , and it requires a large number of features. It also requires that $2^{\ell-1}$ separate tours be computed, which is expensive when ℓ is large. Here we present a simple method for computing the $2^{\ell-1}$ tours that is computationally efficient and does not require information on the feature distribution. It uses a single tour, and then visits subsets of vertices on this tour. The method is inspired by insertion/deletion heuristics for the TSP [17]. Due to page constraints we provide only a brief outline.

Consider a discretized problem instance with vertices $V = \cup_k V_k$. Here, we describe the policy for one robot. For n robots we can use a clustering method (such as k -medians) to partition the vertices in V into n sets.

Computationally Efficient Tour Construction

Input: A single robot problem instance \bar{I} .

Output: $2^{\ell-1}$ robot tours.

- 1: Compute a single TSP tour through all points in V .
 - 2: For each k , move along $\text{TSP}(V)$, and alternately assign vertices in V_k to $V_{k,1}, \dots, V_{k,k}$.
 - 3: Tour T contains $V_1 \cup V_{2,2 \pmod T} \cup \dots \cup V_{\ell, \ell \pmod T}$, and visits each vertex using the order from $\text{TSP}(V)$.
-

As in the PARTITION-TOUR policy, each tour visits all vertices in V_1 , half of the vertices in V_2 , and so on, visiting $1/2^{\ell-1}$ of the vertices in V_ℓ . The computation for this policy is bounded by the computation of a single TSP tour, which from Section II-C, can be performed efficiently.

C. Simulations Results

We evaluate in simulation the performance of the heuristic algorithm relative to the lower bounds. Since our overall approach is to partition the environment, with each robot operating independently in its own partition, we show Monte Carlo simulations only for a single robot. This allows us to evaluate the lower bound in Prop. V.1. For the simulations we uniformly distributed m features in a convex environment. The rate of change of each feature is a uniform random variable in $(0, 1]$. For each value of m we performed 50 trials, and for each trial we made the following comparisons: 1) the heuristic policy (HP) to the general lower bound (Gen-LB) in Prop. V.1; and 2) the HP to the asymptotic lower bound (Asym-LB) in Theorem IV.1 (i.e., $C^*(\bar{I})/2$); 3) the HP on the discretized problem to the discretized asymptotic lower bound in Theorem IV.1; and 4) the PARTITION-TOUR (PT) policy on the discretized problem to the discretized asymptotic lower bound. The results are shown in Table I. Each entry records the mean over the 50 trials. The value in parentheses gives the standard deviation.

m	Non-discretized		Discretized	
	HP/Gen-LB	HP/Asym-LB	HP/Asym-LB	PT/Asym-LB
10	1.42 (0.22)	2.64 (0.34)	1.45 (0.18)	1.52 (0.23)
20	1.54 (0.17)	2.69 (0.24)	1.44 (0.13)	1.47 (0.14)
40	1.73 (0.17)	2.58 (0.19)	1.35 (0.09)	1.37 (0.10)
60	1.90 (0.14)	2.59 (0.15)	1.35 (0.07)	1.31 (0.08)
80	1.95 (0.16)	2.57 (0.09)	1.32 (0.05)	1.27 (0.06)

TABLE I

COMPARISON OF THE HEURISTIC POLICY (HP), TO THE TWO LOWER BOUNDS: GENERAL (GEN-LB) AND ASYMPTOTIC (ASYMP-LB).

The results reveal several properties. First, for the values considered, the performance of the heuristic policy is typically within a factor of two of optimal. For discretized problem instances its performance is within approximately 4/3 of optimal. Second, for lower values of m , the general lower bound is tighter than the asymptotic lower bound. Finally, for low values of m (i.e., $m < 40$), the heuristic policy actually outperforms the PT policy.

VI. CONCLUSIONS AND FUTURE DIRECTIONS

We considered a problem in routing robots to persistently monitor a set of dynamically changing features. We focused on the asymptotic regime of a large number of randomly distributed features. In this regime, we provided a lower bound and policy which performs within a factor two of the optimal. We also provided a general lower bound which holds outside of the asymptotic regime, and a heuristic algorithm motivated by our asymptotic analysis.

There are many areas for future research. We are working on extending the general lower bound to n robots, and on rigorously characterizing the performance of the heuristic algorithm. We would also like to consider a distributed version of the problem. Finally, we are interested in relating our model for the rates of change to a Kalman filter setup, where each feature evolves according to a linear system.

APPENDIX

Proof: [Proof of Lemma II.2] The probability that vertex i lies in C is given by $\int_C f(\mathbf{q})d\mathbf{q}$. Thus, the number of vertices in C is a binomially distributed random variable m_C , with mean $m \int_C f(\mathbf{q})d\mathbf{q}$. By the strong law of large numbers, with probability one,

$$\lim_{m \rightarrow +\infty} \frac{m_C}{m} = \int_C f(\mathbf{q})d\mathbf{q}. \quad (13)$$

Now, the conditional density of vertices in C is given by $f(\mathbf{q})/(\int_C f(\mathbf{q})d\mathbf{q})$. From equation (1) we know that

$$\text{TSP}(C \cap V) \sim \sqrt{m_i} \int_C \sqrt{\frac{\bar{f}(\mathbf{q})}{\int_C f(\mathbf{q})d\mathbf{q}}} d\mathbf{q}.$$

Dividing through by \sqrt{m} , and applying (13), we obtain the desired result. ■

We are now ready to prove Theorem IV.3.

Proof: [Proof of Theorem IV.3] Let us consider a vertex k in $V_j \cap C_i$ that is visited in tour T of region R_p . This vertex will be visited once in every 2^{j-1} tours. The time between successive visits to this vertex can be written as the sum of five components:

- (i) travel in cell C_i from vertex k to the last vertex in cell C_i using tour T ;
- (ii) travel through cells C_{i+1}, \dots, C_M using tour T ;
- (iii) travel for tours $T+1, \dots, T+2^{j-1}-1$;
- (iv) travel through cells C_1, \dots, C_{i-1} on tour $T+2^{j-1}$;
- (v) travel in cell C_i from vertex 1 to k using tour $T+2^{j-1}$.

We now bound the time for each of these five components, $\Delta t_{(i)}, \dots, \Delta t_{(v)}$. Components (i) and (v) will turn out to be negligible and we can use the conservative upper bound of

$$\Delta t_{(i)}, \Delta t_{(v)} \leq \text{TSP}(\cup_{j=1}^{\ell} V_j \cap C_i). \quad (14)$$

Now, recall that the partition in region R_p was constructed such that for each cell C_i we have $\Phi_f(C_i) = \Phi_f(R_p)/M$. Moreover, the partition of \mathcal{E} into regions satisfies $\Phi_f(R_p) = \Phi_f(\mathcal{E})/n$, and thus $\Phi_f(C_i) = \Phi_f(\mathcal{E})/(nM)$. As $\min_k m_k \rightarrow +\infty$, by applying Lemma II.2, we see that the tour through $\cup_{j=1}^k V_j$ in cell C_i is given by

$$\begin{aligned} \text{TSP}(\cup_{j=1}^k V_j \cap C_i) &= \Phi_f(C_i) \sqrt{m_1 + \dots + m_k} \\ &= \frac{\Phi_f(\mathcal{E})}{nM} \sqrt{m_1 + \dots + m_k}, \end{aligned} \quad (15)$$

with probability one. Thus, by the construction of the partition, the tour length is independent of the cell. From equation (15) we can write equation (14) as

$$\Delta t_{(i)}, \Delta t_{(v)} \leq \frac{\Phi_f(\mathcal{E})}{nM} \sqrt{m_1 + \dots + m_{\ell}}$$

For component (iii) we need to bound the length of each tour $T+1, \dots, T+2^{j-1}-1$. First, note that the worst-case length of a tour through M points in an environment \mathcal{E} is $\text{const} \sqrt{M|\mathcal{E}|}$, see [18]. Therefore, the length of each tour T is upper bounded by

$$\sum_{i=1}^M \text{TSP} \left(\bigcup_{\{j \mid i+T-1 \pmod{2^{j-1}} \equiv 1\}} V_j \cap C_i \right) + \text{const} \sqrt{M|\mathcal{E}|},$$

where the first term gives the sum of the length of the tours in each cell, and the second term gives the length of stitching together each cell. Recall that there are $M = \alpha 2^{\ell-1}$ cells, where $\alpha \in \mathbb{N}$ is a positive integer. Looking at the condition $i + T - 1 \pmod{2^{j-1}} \equiv 1$ we see that for each tour $T \in \{1, \dots, 2^{\ell-1}\}$, exactly $2^{\ell-2}\alpha$ cells include only the vertices from V_1 , exactly $2^{\ell-3}\alpha$ cells include only the vertices from $V_1 \cup V_2$, and so on, with exactly α cells including the vertices from $\cup_{j=1}^{\ell-1} V_j$, and exactly α cells including the vertices from $\cup_{j=1}^{\ell} V_j$. Combining this with equation (15), we see that the time for each tour in component (iii) is given by

$$\Delta t_T = \frac{\alpha \Phi_f(\mathcal{E})}{nM} \left(\sum_{k=1}^{\ell-1} 2^{\ell-k-1} \sqrt{m_1 + \dots + m_k} + \sqrt{m_1 + \dots + m_\ell} \right) + \text{const} \sqrt{M|\mathcal{E}|}$$

Since $\alpha = M/2^{\ell-1}$, the above equation simplifies to

$$\Delta t_T = \frac{\Phi_f(\mathcal{E})}{n} \sum_{k=1}^{\ell} \max\{2^{-k}, 2^{1-\ell}\} \left(\sum_{j=1}^k m_j \right)^{1/2} + \text{const} \sqrt{M|\mathcal{E}|}.$$

The time for component (iii) is then given by $\Delta t_{(iii)} = (2^{j-1} - 1)\Delta t_T$. Finally, we need to bound the time for components (ii) and (iv). As shown above, we need only count the number of cells in tour T that include a particular combination of V_1, \dots, V_ℓ . First, consider cells C_1, \dots, C_{i-1} . For any tour T , the combination $\cup_{j=1}^{\ell-1} V_j$ can be visited in at most $\lceil i\alpha/M \rceil$ cells. Similarly, $\cup_{j=1}^{\ell-1-k} V_j$ can be visited in at most $\lceil 2^{k-1}i\alpha/M \rceil = \lceil 2^{k-\ell}i \rceil$ cells. Thus, applying equation (15), the asymptotic length of the path in step (ii) is

$$\Delta t_{(ii)} \leq \frac{\Phi_f(\mathcal{E})}{nM} \left(\sum_{k=1}^{\ell-1} \lceil 2^{-k}i \rceil \sqrt{m_1 + \dots + m_k} + \lceil 2^{1-\ell}i \rceil \sqrt{m_1 + \dots + m_\ell} \right) + \text{const} \sqrt{M|\mathcal{E}|}.$$

Similarly, the tour in step (iv) can be upper bounded by

$$\Delta t_{(iv)} \leq \frac{\Phi_f(\mathcal{E})}{nM} \left(\sum_{k=1}^{\ell-1} \lceil 2^{-k}(M-i) \rceil \sqrt{m_1 + \dots + m_k} + \lceil 2^{1-\ell}(M-i) \rceil \sqrt{m_1 + \dots + m_\ell} \right) + \text{const} \sqrt{M|\mathcal{E}|},$$

where we simply replaced i by $M - i$. Thus, we can upper bound the sum of steps (ii) and (iv) by $(M+2)\Delta t_T/M$.

Now, we have bounded all five components in the limit as $m_1, \dots, m_\ell \rightarrow +\infty$. The time between visits to the vertex in $V_j \cap C_i$ is upper bounded by $\Delta t_{(i)} + \Delta t_{(ii)} + \Delta t_{(iii)} + \Delta t_{(iv)} + \Delta t_{(v)}$. (The time for a robot to travel from $\mathbf{p}_p(0)$ onto the tour is negligible.) Let $M \rightarrow +\infty$ such that $\min\{m_1, \dots, m_\ell\}^2 M \rightarrow \text{const} \in \mathbb{R}_{>0}$. That is, $M \rightarrow +\infty$ (i.e., $\alpha \rightarrow +\infty$) but more slowly than each m_k . In this case the terms $\text{const} \sqrt{M|\mathcal{E}|}$ become negligible, as do the contributions of components (i) and (v). In addition,

$(M+2)/M \rightarrow 1^+$, and the time between visits to the vertex in V_j is upper bounded by

$$2^{j-1} \frac{\Phi_f(\mathcal{E})}{n} \sum_{k=1}^{\ell} \max\{2^{-k}, 2^{1-\ell}\} \left(\sum_{j=1}^k m_j \right)^{1/2}.$$

The rate of change of each vertex in V_j is $(1/2)^{j-1}$. Therefore, the maximum urgency of the vertex in $V_j \cap C_i$ is upper bounded by

$$\frac{\Phi_f(\mathcal{E})}{n} \sum_{k=1}^{\ell} \max\{2^{-k}, 2^{1-\ell}\} \left(\sum_{j=1}^k m_j \right)^{1/2}.$$

Since V_j , C_i , and R_p were arbitrary, the result holds for all vertices in V . Combining the above expression with Theorem IV.1 we obtain the desired result. ■

REFERENCES

- [1] L. D. Stone, *Theory of Optimal Search*. Operations Research Society of America, 1975.
- [2] T. H. Chung and J. W. Burdick, "A decision-making framework for control strategies in probabilistic search," in *Proc ICRA*, Roma, Italy, Apr. 2007, pp. 4386–4393.
- [3] L. F. Bertuccelli and J. P. How, "Robust UAV search for environments with imprecise probability maps," in *Proc CDC-ECC*, Seville, Spain, Dec. 2005, pp. 5680–5685.
- [4] H. Choset, "Coverage for robotics – A survey of recent results," *Annals of Mathematics and Artificial Intelligence*, vol. 31, no. 1-4, pp. 113–126, 2001.
- [5] Y. Gabriely and E. Rimon, "Competitive on-line coverage of grid environments by a mobile robot," *Computational Geom.: Theory and App.*, vol. 24, no. 3, pp. 197–224, 2003.
- [6] I. Hussein and D. M. Stipanovič, "Effective coverage control for mobile sensor networks with guaranteed collision avoidance," *IEEE Trans Control Systems Technology*, vol. 15, no. 4, pp. 642–657, 2007.
- [7] A. Tiwari, M. Jun, D. E. Jeffcoat, and R. M. Murray, "Analysis of dynamic sensor coverage problem using Kalman filters for estimation," in *Proc IFAC World C*, Prague, Czech Republic, Jul. 2005.
- [8] Y. Chevaleyre, "Theoretical analysis of the multi-agent patrolling problem," in *IEEE/WIC/ACM Int. Conf. Intelligent Agent Technology*, Beijing, China, Sep. 2004, pp. 302–308.
- [9] Y. Elmaliach, N. Agmon, and G. A. Kaminka, "Multi-robot area patrol under frequency constraints," in *Proc ICRA*, Roma, Italy, Apr. 2007, pp. 385–390.
- [10] L. Caffarelli, V. Crespi, G. Cybenko, I. Gamba, and D. Rus, "Stochastic distributed algorithms for target surveillance," in *Intelligent Systems and Design Applications*, Tulsa, OK, Aug. 2003, pp. 137–148.
- [11] N. Christofides and J. E. Beasley, "The period routing problem," *Networks*, vol. 14, no. 2, pp. 237–256, 1984.
- [12] S. L. Smith, M. Pavone, F. Bullo, and E. Frazzoli, "Dynamic vehicle routing with priority classes of stochastic demands," *SIAM JCO*, vol. 48, no. 5, pp. 3224–3245, 2010.
- [13] M. Pavone, A. Arsie, E. Frazzoli, and F. Bullo, "Equitable partitioning policies for mobile robotic networks," *IEEE Trans Automatic Ctrl*, 2010, (Submitted Dec 2008 and Aug 2009) to appear.
- [14] J. Beardwood, J. Halton, and J. Hammersley, "The shortest path through many points," in *Proceedings of the Cambridge Philosophy Society*, vol. 55, 1959, pp. 299–327.
- [15] A. G. Percus and O. C. Martin, "Finite size and dimensional dependence of the Euclidean traveling salesman problem," *Physical Review Letters*, vol. 76, no. 8, pp. 1188–1191, 1996.
- [16] R. C. Larson and A. R. Odoni, *Urban Operations Research*. Prentice Hall, 1981.
- [17] B. Korte and J. Vygen, *Combinatorial Optimization: Theory and Algorithms*, 4th ed., ser. Algorithmics and Combinatorics. Springer Verlag, 2007, vol. 21.
- [18] L. Few, "The shortest path and the shortest road through n points in a region," *Mathematika*, vol. 2, pp. 141–144, 1955.