

Coverage Control for Teams of Nonholonomic Agents

John Enright

Ketan Savla

Emilio Frazzoli

Abstract—Consider a coverage problem for a team of agents in the plane: target points appear sporadically over time in a bounded environment and must be visited by one of the agents. It is desired to minimize the expected elapsed time between the appearance of a target point, and the instant it is visited. For holonomic agents, this reduces to the continuous Weber problem, well studied in the locational optimization literature. In this paper, we consider a team of nonholonomic vehicles constrained to move with constant forward speed along paths of bounded curvature. We show that, in this case, the optimal policy depends on the *density* of vehicles in the environment. In low density scenarios, the optimal policy resembles that of holonomic agents: the environment is partitioned into subregions of dominance, and each agent is responsible for targets appearing in its own subregion (*territorial behavior*). As the density increases, the optimal policy exhibits a transition to a *gregarious behavior* in which the team loiters in a coordinated pattern, and each agent visits targets that appear immediately in front of it.

I. INTRODUCTION

Wide-area surveillance is one of the prototypical missions for Uninhabited Aerial Vehicles (UAVs), e.g., in environmental monitoring, security, or military setting. Low-altitude UAVs on such a mission must provide coverage of a region and investigate events of interest as they manifest themselves. In particular, we are interested in cases in which close-range information is required on targets detected by high-altitude aircraft, spacecraft, or ground spotters, and the UAVs must proceed to the locations to gather on-site information.

Variations of problems falling into this class have been studied in a number of papers in the recent past, e.g., [1], [2]. In these papers, as in ours, the UAV is modeled as a planar vehicle that can move on paths of bounded curvature. However, in some of these papers, the locations of targets are known a priori and a strategy is computed that attempts to optimize the cost of servicing the known targets. We address a scenario in which the target points are generated dynamically, with only prior statistics on their location, and a policy is designed to minimize the expected time a target waits to be visited. Furthermore, we assume that the targets appear sporadically throughout time, so the focus is not on planning efficient paths to visit many targets consecutively. Instead, the challenge is to design loitering patterns or policies in such a way that, when a target does appear, the expected wait for the “closest” UAV to arrive is minimized.

This paper is along the vehicle routing themes of [3], [4], and many of the new results presented here are shadowed in [5]. However, by focusing on the case in which targets appear rarely, this work is applicable to coverage problems [6], [7], [8], in which the agents spread out with some sense of

balance, or comb the environment efficiently. On the other hand, mathematically its structure has a strong resemblance with facility location problems [9], [10], [11], [12], [13], with the main difference that the facilities are vehicles constantly in motion with nontrivial dynamics.

The main contributions of this paper are the following. 1) We employ a novel approach to establish a lower bound on the achievable performance for the problem, related to the reachable set of the vehicle. The approach taken may be of independent interest as it can be applied to other non-holonomic vehicles performing similar tasks. 2) We design two strategies. One strategy assigns regions of dominance to the vehicles and requires the vehicles to *guard* their own territory, thus requiring them to spread out. This strategy is optimal when vehicle density is low. The other strategy requires the agents to move as a group in a coordinated pattern, balancing the amount of space directly in front of each agent. This strategy is within a constant factor of the optimal for high vehicle densities. We observe that the optimal system time shrinks with the cube root of the number of vehicles (as opposed to the square root in the holonomic case). 3) As a consequence of the aforementioned algorithm design and analysis, we recognize a phase transition in the optimal behavior mode, depending on the density of vehicles in the environment. We introduce a non-dimensional parameter which characterizes the phase of the system, and use both numerical and analytical techniques to determine the value which demarcates the transition itself.

Other researchers, including [14] and, more recently, [15], have investigated phase transitions in the behavior of large groups of agents. In these works, the individuals follow local interaction laws, sometimes inspired by physical laws, and it is shown to give rise to group behavior of different modes, depending on design parameters and possible noise input. In some cases, these types of approaches are useful for accomplishing large-scale multi-agent coverage and similar tasks, in a distributed, robust fashion. In contrast, we take a top-down approach. We begin with simple motion constraints (found in nature and technology) for the individual agents, and an explicit task and performance measurement, we aim to design control strategies achieving near optimal performance. As a result of our analysis, we have noticed that the optimal policy exhibits more than one distinct mode of behavior. We are not attempting to imitate nature. Rather, we are attempting to uncover a possible cause (efficiency) behind natural multi-agent systems which exhibit phase transitions between territorial and gregarious behavior [16].

The rest of the paper is organized as follows. In Sections II and III, we set up the problem and review relevant results

about the Dubins vehicle. In Section IV, we establish two lower bounds on the optimal coverage cost. In Section V, we propose two policies and prove their optimality for the limiting cases of low and high vehicle densities. In Section VI we study the phase transition between the two policies. Finally, we conclude with comments about future work in Section VII.

II. NOTATION AND PROBLEM FORMULATION

In this section, we present some preliminary notations and formulate the problem. Let $\mathcal{Q} \subset \mathbb{R}^2$ be a convex, compact domain on the plane, with non-empty interior; we will refer to \mathcal{Q} as the *environment*. Let A be the area of \mathcal{Q} .

Service request points are generated in the environment sporadically throughout time, with uniform spatial density. In order to service the request, one of the m UAVs has to move to the target point associated with it. The UAVs are modeled as nonholonomic vehicles constrained to move at unit speed along a path with bounded curvature; $1/\rho$ being the maximum curvature. In other words, let the configuration $g_i \in SE(2)$ of the i -th vehicle ($\forall i \in I_m$ where we denote $I_m = \{1, \dots, m\}$) be given in coordinates by $g_i = (x_i, y_i, \theta_i)$, where x_i, y_i are the projections of the vehicle's position along inertially fixed orthogonal axes, and θ_i is the orientation of the vehicle's longitudinal axis with respect to the x -axis; then the kinematics of the vehicle is described by the Dubins model:

$$\begin{aligned} \dot{x}_i &= \cos(\theta_i), \\ \dot{y}_i &= \sin(\theta_i), \\ \dot{\theta}_i &= \omega_i, \quad \omega_i \in [-1/\rho, 1/\rho]. \end{aligned} \quad (1)$$

The UAVs are identical, and have unlimited sensing range. We note that the above kinematic model of an airplane is very common in the literature on UAV motion planning; the model is very similar to the one studied in [17], with the difference that the vehicle we consider is constrained to move at constant speed. Results in terms of minimum-length paths for Dubins vehicle hold for our model, where they assume an additional connotation of being minimum-time paths as well.

In the course of this paper, we shall use the term *Dubins frame* to refer to a coordinate frame with origin fixed to the Dubins vehicle and x -axis along the longitudinal axis. Let $L_\rho(g_i, q) : SE(2) \times \mathbb{R}^2 \rightarrow \mathbb{R}$ be the minimum length of a path satisfying (1), steering a Dubins vehicle from initial configuration g in $SE(2)$ to a point q in the plane, without any constraints on the final heading, i.e., to any configuration in the set $\{q\} \times S^1 \subset SE(2)$. For a given configuration of the vehicles $g(t) = \{g_1(t), \dots, g_m(t)\} \in SE(2)^m$ at time t , we define the coverage cost function at time t as $\mathcal{W}(g(t)) := \int_{q \in \mathcal{Q}} \frac{1}{A} \min_{i \in I_m} L_\rho(g_i(t), q) dq$. The cost function quantifies the expected time for the closest UAV to reach a target point uniformly distributed in the environment at time t . If the UAVs were capable of stopping, this problem could be reduced to finding the optimal configuration of the m -vehicle system. Instead, we must design control policies

that keep the configuration near optimal in a time-averaged sense.

A control policy $\pi = \{\pi_1, \dots, \pi_m\}$ for the coverage task is specified by the initial configurations of the vehicles and the subsequent control inputs, i.e., $\pi_i = (g_i(0), (\omega_i(t) \in [-1/\rho, 1/\rho], t \geq 0)) \forall i \in I_m$. We define the coverage cost associated with policy π as $\mathcal{W}_\pi = E_{t \geq 0}[\mathcal{W}(g(t))]$, where, from Eq. (1), $g(t)$ is completely determined by the policy π for all $t \geq 0$.

We are now ready to state our problem. Given a region \mathcal{Q} and m vehicles modeled by Eq. (1), design control policies π such that the coverage cost associated with that policy achieves or approximates the theoretical optimal coverage cost given by $\mathcal{W}_{\text{opt}} = \inf_\pi \mathcal{W}_\pi$. In the following, we are interested in designing computationally efficient control policies that are within a constant factor of the optimal, i.e., policies π such that $\mathcal{W}_\pi \leq \kappa \mathcal{W}_{\text{opt}}$ for some constant κ . We conclude this section with some notation that is the standard concise way to state asymptotic properties. For $f, g : \mathbb{N} \rightarrow \mathbb{R}$, we say that $f \in O(g)$ (respectively, $f \in \Omega(g)$) if there exist $N_0 \in \mathbb{N}$ and $k \in \mathbb{R}_+$ such that $|f(N)| \leq k|g(N)|$ for all $N \geq N_0$ (respectively, $|f(N)| \geq k|g(N)|$ for all $N \geq N_0$). If $f \in O(g)$ and $f \in \Omega(g)$, then we use the notation $f \in \Theta(g)$.

III. OPTIMAL PATHS AND THE REACHABLE SET FOR THE DUBINS VEHICLE

Minimum-length Dubins paths between two configurations have been extensively studied, due to their importance in mobile robotics. A full characterization of optimal paths is given in [17]. For our purpose, we need optimal paths for a different type of boundary condition; the difference is that the final heading at the target point is not constrained a priori. The characteristics of paths of minimal length with such boundary conditions were studied in [18], where it is proved that all such paths are a concatenation of an arc of a minimum-radius circle (either in the positive or negative direction), with either an arc of a minimum-radius circle (in the opposite direction), or with a straight segment. Closed-form expressions for the lengths for such paths were derived in [3], [5].

We now present an upper bound on $L_\rho(g_i, q)$, originally from [3], which will be useful in analyzing one of the policies.

Lemma 3.1: The function $L_\rho(g_i, q)$ satisfies the following inequality, for all $q \in \mathbb{R}^2$:

$$L_\rho(g_i, q) \leq c_1 \rho + \|q - q_{\text{circ}}\|, \quad (2)$$

where, expressed in the Dubins frame, $q_{\text{circ}} = (0, \pm c_2 \rho)$. The constants are $c_1 \approx 3.75595$ and $c_2 \approx 2.91801$.

In other words, if the vehicle is circling point q_{circ} at radius $c_2 \rho$, it can reach any point $q \in \mathbb{R}^2$ within time $c_1 \rho + \|q - q_{\text{circ}}\|$.

We now introduce the *reachable set* for a Dubins vehicle and state a property which will be useful later on in the paper. Given a configuration $g_i \in SE(2)$, the reachable set $R_t(g_i)$ for a Dubins vehicle is a subset of $SE(2)$ such that

for every configuration $g_f \in R_t(g_i)$, there exist admissible controls to drive the vehicle from the initial configuration g_i to the final configuration $g_f \in R_t(g_i)$ within time t .

We are interested in the reachable region on the Euclidean plane, i.e., the projection of the reachable set onto \mathbb{R}^2 . For simplicity of notation we shall henceforth denote $R_t(g_i)$ as the reachable region on the plane.

The minimum-length path to any point within the interior of one of the minimum turning radius circles has a length of at least $\pi\rho$, and so any point reachable within time $t \leq \pi\rho/2$ is of type CL. In Fig. 1, we depict $R_t(g_i)$ for some time $t \leq \pi\rho/2$, along with a generic minimum-length path of length t . Let us denote the length of the C segment with s and the angle it sweeps out with θ . Since $s = \rho\theta$, the length of the L segment is $t - s = t - \rho\theta$. Also, denote the switching point between the C and L segments with (x_1, y_1) , the point at the end of the L segment with (x_2, y_2) . Using planar geometry, we can write $x_1 = \rho \sin \theta$, $y_1 = \rho(1 - \cos \theta)$, $x_2 = x_1 + (t - \rho\theta) \cos \theta$, and $y_2 = y_1 + (t - \rho\theta) \sin \theta$. For fixed ρ and t , $\text{Area}(R_t(g_i)) = 2 \left(\int_0^{t/\rho} y_1 \frac{dx_1}{d\theta} d\theta + \int_{t/\rho}^0 y_2 \frac{dx_2}{d\theta} d\theta \right)$, hence

$$\text{Area}(R_t(g_i)) = \frac{t^3}{3\rho}, \quad \text{for } t \leq \frac{\pi\rho}{2}. \quad (3)$$

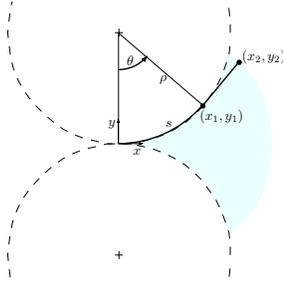


Fig. 1. The derivation of $\text{Area}(R_t(g_i))$ for $t \leq \pi\rho/2$.

Remark 3.2: Eq. (3) implies that area of the Dubins reachable set decreases faster than the area of the reachable set for a Euclidean vehicle as $t/\rho \rightarrow 0$.

IV. LOWER BOUNDS

We first show that a trivial lower bound is obtained by adopting the corresponding lower bound for a *Euclidean vehicle*, i.e., a vehicle moving with unit speed and having no motion constraints. In order to do that, we give a brief overview of a related problem from geometric optimization. Given a set $\mathcal{Q} \subset \mathbb{R}^2$ and a set of points $p = \{p_1, p_2, \dots, p_m\} \in \mathcal{Q}^m$, the *Voronoi partition* of \mathcal{Q} generated by p is $\mathcal{V}(p) = (\mathcal{V}_1(p), \mathcal{V}_2(p), \dots, \mathcal{V}_m(p))$, where $\mathcal{V}_i(p) = \{q \in \mathbb{R}^2 : \|q - p_i\| \leq \|q - p_j\|, \forall j, j \in I_m\}$. The expected distance between a random point q , sampled from a uniform distribution over \mathcal{Q} , and the closest point in p is given by $\mathcal{H}_m(p, \mathcal{Q}) := \sum_{i=1}^m \int_{\mathcal{V}_i(p)} \|p_i - q\| dq$. The problem of choosing p to minimize \mathcal{H}_m is known in the geometric optimization literature as the *continuous supply / continuous demand k -median problem* [12], [13].

Related problems include k -medians and centers for discrete supply and/or demand sets; see [9], [10], [11], [19], [20] and references therein. The m -median of the set \mathcal{Q} is the global minimizer $p_m^*(\mathcal{Q}) = \text{argmin}_{p \in \mathcal{Q}^m} \mathcal{H}_m(p, \mathcal{Q})$. We let $\mathcal{H}_m^*(\mathcal{Q}) = \mathcal{H}_m(p_m^*(\mathcal{Q}), \mathcal{Q})$ be the global minimum of \mathcal{H}_m . The map $p \mapsto \mathcal{H}_m(p, \mathcal{Q})$ with $m > 1$ is differentiable (whenever (p_1, \dots, p_m) are distinct) but not convex, thus making the solution of the continuous m -median problem hard in the general case. In fact, it is known [10] that the discrete version of the m median problem is NP-hard for $d \geq 2$. However, numerical techniques for finding local minima of continuous m -median problems can be designed using i) the "honeycomb heuristic" [12], and ii) the fact that p_i is the median of $\mathcal{V}_i(p)$ for each p_i in p^* . We will not pursue the issue of computation of the m -median and of the corresponding $\mathcal{H}_m^*(\mathcal{Q})$ further, but will assume that these values are available.

In the following result we characterize how the optimal value of the m -median function depends on the number of generators, i.e., how $\mathcal{H}_m^*(\mathcal{Q})$ depends on the integer m , for a fixed environment \mathcal{Q} . This is a variation of a more general result given in [12] for d dimensions, whose lower bound proof has much of the logical structure to be used in the Theorem 4.4, without the additional geometrical complexity due to the motion constraints associated with the Dubins vehicle.

Lemma 4.1: The function $\mathcal{H}_m^*(\mathcal{Q})$ is $\Theta(1/\sqrt{m})$ for a given \mathcal{Q} . In particular,

$$\frac{2}{3} \sqrt{\frac{A}{\pi m}} \leq \mathcal{H}_m^*(\mathcal{Q}) \leq \frac{c_{\text{sq}} \sqrt{3} L(\mathcal{Q})}{\sqrt{m}},$$

where $c_{\text{sq}} \approx 0.38$ and $L(\mathcal{Q})$ is the side-length of the smallest square enclosing \mathcal{Q} .

We now state the following trivial lower bound.

Theorem 4.2: The coverage cost satisfies the following lower bound.

$$\mathcal{W}_{\text{opt}} \geq \mathcal{H}_m^*(\mathcal{Q}).$$

Remark 4.3: Theorem 4.2 and Lemma 4.1 imply that $\mathcal{W}_{\text{opt}} \geq \frac{2}{3} \sqrt{\frac{A}{\pi m}}$. Note that this lower bound holds for any region and any number of vehicles.

The lower bound in Theorem 4.2 was obtained trivially by adopting the corresponding lower bound for the Euclidean vehicle. It is reasonable to expect that one can obtain a tighter lower bound by explicitly considering the motion constraints of the Dubins vehicle. Before deriving a new lower bound, we introduce a few concepts. To that effect, we first define $t_\rho(A) = \{t | \text{Area}(R_t(g_i)) = A\}$ to be the time for which the Dubins reachable set has area A . With a slight abuse of notation, in the following, we shall use an alternate notation to denote the Dubins reachable set parametrized by its area: $R_A(g_i) = R_{t_\rho(A)}(g_i)$. The approach of the lower bound heavily exploits the following fact: for a given configuration, g_i , the region of area A with least expected Dubins distance is the reachable region of area A . Given a set $\mathcal{Q} \subset \mathbb{R}^2$ and a set of Dubins configurations $g = \{g_1, g_2, \dots, g_m\} \in SE(2)^m$, let us define

$\mathcal{DV}(g) = \{\mathcal{DV}_1(g), \mathcal{DV}_2(g), \dots, \mathcal{DV}_m(g)\}$ as the *Dubins Voronoi* partition of the set \mathcal{Q} generated by the configurations g . In other words, $q \in \mathcal{DV}_i(g)$ if $L_\rho(g_i, q) \leq L_\rho(g_k, q)$, for all $k \in I_m$. Note that a vehicle's Dubins Voronoi region is *exactly* its instantaneous region of dominance. We are now ready to state a new lower bound on \mathcal{W}_{opt} , based on the idea of finding the optimal configuration for the m -vehicle system, as if it were capable of stopping motion.

Theorem 4.4: For any m , the coverage cost satisfies

$$\mathcal{W}_{\text{opt}} \geq \frac{m}{A} \int_{R_{A/m}(g_i)} L_\rho(g_i, q) dq.$$

Proof: In the following, we use the notation $A_i = \text{Area}(\mathcal{DV}_i(g))$ and $\mathcal{A} = \{A_1, \dots, A_m\}$. We begin with

$$\begin{aligned} \mathcal{W}_{\text{opt}} &\geq \min_{g \in SE(2)^m} \int_{\mathcal{Q}} \frac{1}{A} \min_{i \in I_m} L_\rho(g_i, q) dq \\ &= \min_{g \in SE(2)^m} \sum_i^m \int_{\mathcal{DV}_i(g)} \frac{1}{A} L_\rho(g_i, q) dq \\ &\geq \min_{g \in SE(2)^m} \sum_i^m \int_{R_{A_i}(g_i)} \frac{1}{A} L_\rho(g_i, q) dq \\ &\geq \min_{\mathcal{A} \in \mathbb{R}_+^m} \sum_i^m \int_{R_{A_i}(g_i)} \frac{1}{A} L_\rho(g_i, q) dq \text{ s.t. } \sum_i A_i = A. \end{aligned}$$

Let us define $f: \mathbb{R}^+ \rightarrow \mathbb{R}^+$ as $f(A) = \int_{R_A(g_i)} L_\rho(g_i, q) dq$. It is easy to verify that the function f is a continuous, strictly increasing function of A . One can also show that f is convex. Thus by using the Karush-Kuhn-Tucker conditions [21], one can show that the quantity $\min_{\mathcal{A} \in \mathbb{R}_+^m} \frac{1}{A} \sum_i^m f(A_i)$ subject to $\sum_i^m A_i = A$ is minimized with an equitable partition, i.e., A_i is A/m , $\forall i \in I_m$. This yields the stated result. ■

Note that the above lower bound holds for any region and any number of vehicles. Although, in this integral form, the lower bound's dependency on parameters such as m and ρ is not transparent, we have shown that a lower bound can be obtained by dividing the environment into equally sized domains of responsibility, i.e., balance the workload in terms of coverage.

We now introduce a non-dimensional parameter, which we call the *nonholonomic vehicle density* as $d_\rho = \rho^2 m / A$. The motivation of this parameter is shown in the following theorem.

Theorem 4.5: The coverage cost for any \mathcal{Q} of area A , $\rho \in \mathbb{R}_+$ and $m \in \mathbb{N}$ such that $d_\rho \geq 24/\pi^3$ satisfies

$$\mathcal{W}_{\text{opt}} \geq \frac{3}{4} \left(\frac{3\rho A}{m} \right)^{1/3}.$$

Proof: Rearranging Eq. (3), we see that $t_\rho(A/m) = (3\rho A/m)^{1/3}$ for $t_\rho(A/m) \leq \pi\rho/2$, i.e., for $d_\rho \geq 24/\pi^3$. Also under this condition, $L_\rho(g_i, q) \geq x$, and thus,

$$\begin{aligned} \int_{R_{A/m}(g_i)} L_\rho(g_i, q) dq &\geq 2 \int_0^{y_2} \int_{x_1(y)}^{t_\rho(A/m)} x dx dy \\ &= \frac{t_\rho^4(A/m)}{4\rho} + o(t_\rho^4(A/m)) \approx \frac{1}{4\rho} \left(\frac{3\rho A}{m} \right)^{4/3}, \end{aligned}$$

where $x_1(y) = \sqrt{\rho^2 - (y - \rho)^2}$ and $y_2 = \rho(1 - \cos \frac{t}{\rho})$. Combining with Theorem 4.4 yields the result. ■

Remark 4.6: Theorem 4.5 shows that $\mathcal{W}_{\text{opt}} = \Omega(1/\sqrt[3]{m})$.

The above result is simply the integral in Theorem 4.4 carried out for a specific range of the problem parameters. Thus, it is reasonable to assume that a policy approximating this lower bound would somehow allow the regions of dominance to remain balanced throughout time. In a low density scenario, this requires the vehicles to simply spread out. However, as the density increases, a balanced coverage requires that a vehicle's region of dominance be very small, i.e., the area immediately in front of it. Thus the optimal policy in this phase involves dynamic regions of dominance.

V. ALGORITHMS AND UPPER BOUNDS

In this section, we propose control policies for the coverage task and then analyze their performances. The first policy – the Median Circling Policy essentially attempts to imitate the optimal policy for Euclidean vehicles, assigning static regions of responsibility. The algorithm is formally described as follows.

The Median Circling (MC) Policy

Each agent is associated with a generator p_i , $i \in I_m$. Let p^* be the m -median of \mathcal{Q} , and define the *loitering station* for the i -th agent as a circular trajectory of radius $c_2\rho$ centered at p_i^* . Each agent visits all targets in the Voronoi region of its own generator $\mathcal{V}_i(p^*)$ in the order in which they arrive. When no targets are available, the vehicle returns to its loitering station; the direction in which the orbit is followed is not important, and can be chosen in such a way that the station is reached in minimum time. A depiction of the MC policy is shown in Fig. 2.

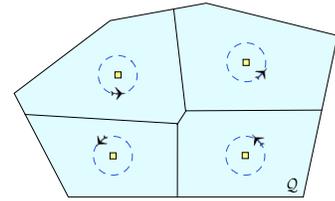


Fig. 2. Depiction of the Median Circling policy. The yellow squares represent p^* , the m -median of \mathcal{Q} . Each UAV loiters about its respective generator at a radius $c_2\rho$. The regions of dominance are demarcated according to the Voronoi partition generated by p^* .

Note that there may be vehicles closer to a given target in terms of Euclidean distance or Dubins minimum-length path. However, we find the target-assignment strategy described above lends itself to tractable analysis.

Theorem 5.1: For any convex, compact $\mathcal{Q} \subset \mathbb{R}^2$ of area A and $\rho \in \mathbb{R}^+$, the coverage cost of the Median Circling policy satisfies

$$\mathcal{W}_{\text{MC}} \leq H_m^*(\mathcal{Q}) + c_1\rho. \quad (4)$$

Proof:

$$\begin{aligned}
\mathcal{W}_{\text{MC}} &\leq \min_{g \in SE(2)^m} \int_{\mathcal{Q}} \frac{1}{A} \min_{i \in I_m} L_\rho(g_i, q) dq \\
&\leq \min_{p \in \mathcal{Q}^m} \int_{\mathcal{Q}} \frac{1}{A} \min_{i \in I_m} (\|p_i - q\| + c_1 \rho) dq \\
&= \sum_{i=1}^m \int_{\mathcal{V}_i(p^*)} \frac{1}{A} \|p_i^* - q\| dq + c_1 \rho \\
&= H_m^*(\mathcal{Q}) + c_1 \rho.
\end{aligned}$$

Remark 5.2: In other words, we have shown that the system time achieved by the MC policy is within a constant additive factor $c_1 \rho \approx 3.76\rho$ from the optimal. The additive factor, which can be considered a penalty due to the non-holonomic constraints imposed on the vehicle's dynamics, depends linearly on the minimum turn radius ρ .

Note that the upper bound stated above does not tend to zero as the size of the team grows large. We now compare the performance of the Median Circling Policy with the optimal limit shown in Theorem 4.2.

Theorem 5.3: The system time of the Median Circling policy in light load satisfies $\mathcal{W}_{\text{MC}}/\mathcal{W}_{\text{opt}} \leq 1 + 10\sqrt{d_\rho}$. In particular, $\lim_{d_\rho \rightarrow 0^+} (\mathcal{W}_{\text{MC}}/\mathcal{W}_{\text{opt}}) = 1$.

Proof: Since $\frac{2}{3}\sqrt{\frac{A}{\pi m}} \leq \mathcal{H}_m^*(\mathcal{Q})$ from Proposition 4.1, we have that

$$c_1 \rho \leq c_1 \rho \frac{\mathcal{H}_m^*(\mathcal{Q})}{\frac{2}{3}\sqrt{\frac{A}{\pi m}}}.$$

The first statement in the theorem then follows by substituting this into Equation (4) and applying Theorem 4.2. The second result follows by taking the limit as $d_\rho \rightarrow 0^+$ in the first statement. ■

Remark 5.4: Theorem 5.3 implies that the Median Circling Policy is an efficient policy for scenarios where we have low nonholonomic vehicle density. In low density scenarios, Euclidean distance dominates the cost, and thus, a policy imitating that of Euclidean vehicles is near optimal.

The next strategy aims to i) maintain balanced coverage despite the requirement of dynamic regions of dominance due to high vehicle density, and ii) achieve a system time that provably tends to zero as the size of the team grows to $+\infty$.

The Strip Loitering (SL) Policy

The Strip Loitering policy is based on the following idea. Bound the environment \mathcal{Q} with a rectangle of minimum height, where we use height to denote the smaller of the two side lengths of a rectangle. Let W and H be the width and height of this bounding rectangle respectively. Divide \mathcal{Q} into strips of width w where

$$w = \min \left\{ \left(\frac{4}{3\sqrt{\rho}} \frac{WH + 10.38\rho H}{m} \right)^{2/3}, 2\rho \right\}. \quad (5)$$

Orient the strips along the side of length W . Construct a closed Dubins-feasible path which runs along the longitudinal bisector of each strip, visiting all strips from top-to-bottom, making U-turns between strips at the edges of \mathcal{Q} ,

and finally returning to the initial configuration. The m UAVs loiter on this path, equally spaced, in terms of path length.

A depiction of the Strip Loitering policy can be viewed in Figure 3. Moreover, in Figure 4 we define two distances that are important in the analysis of this policy. Variable d_2 is the length of the shortest path departing from the loitering path and ending at the target (a circular arc of radius ρ). The UAV responsible for visiting the target is the one closest in terms of loitering path length (variable d_1) to the point of departure, at the time of target-arrival. Note that there may be UAVs closer to the target in terms of Euclidean distance or Dubins minimum-length path. However, we find the target-assignment strategy described above lends itself to tractable analysis.

After a UAV has serviced a target, it must return to its place in the loitering pattern. We now describe a method to accomplish this task through the example shown in Fig. 4. After making a left turn of length d_2 to service the target, the UAV makes a right turn of length $2d_2$ followed by another left turn of length d_2 , returning it to the loitering path. However, the UAV has fallen behind in the loitering pattern. To rectify this, as it nears the end of the current strip, it takes its U-turn early.

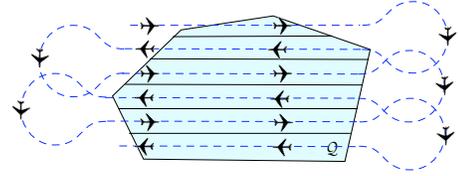


Fig. 3. Depiction of the Strip Loitering policy. The segment providing closure of the loitering path (returning the UAVs from the end of the last strip to the beginning of the first strip) is not shown here for clarity of the drawing.

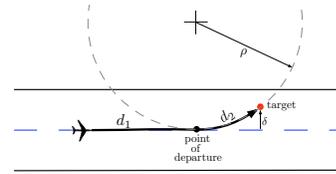


Fig. 4. Close-up of the Strip Loitering policy with construction of the point of departure and the distances δ , d_1 , and d_2 for a given target, at the instant of appearance.

Theorem 5.5: For any $\rho \in \mathbb{R}_+$ and a convex, compact $\mathcal{Q} \subset \mathbb{R}^2$ with a bounding rectangle of width W and height H , the system time of the Strip Loitering Policy satisfies

$$\mathcal{W}_{\text{SL}} \leq \begin{cases} 1.24 \left(\frac{\rho WH + 10.4\rho^2 H}{m} \right)^{1/3} + \frac{W+H+6.2\rho}{m} & \text{if } m \geq m_c, \\ \frac{WH + 10.4\rho H}{4\rho m} + \frac{W+H+6.2\rho}{m} + 1.06\rho & \text{otherwise,} \end{cases} \quad (6)$$

where $m_c = 0.47 \left(\frac{WH}{\rho^2} + \frac{10.4H}{\rho} \right)$. In particular, for any $\rho \in \mathbb{R}_+$ and \mathcal{Q} ,

$$\limsup_{m \rightarrow +\infty} \mathcal{W}_{\text{SL}} m^{1/3} \leq 1.24(\rho WH + 10.4\rho^2 H)^{1/3}. \quad (7)$$

Also,

$$\limsup_{\frac{\rho}{W} \rightarrow +\infty} \mathcal{W}_{\text{SL}} \frac{m}{\rho} \leq 6.2. \quad (8)$$

Proof: It is easy to verify that $N_{\text{strips}} = \lceil \frac{H}{w} \rceil \leq \frac{H}{w} + 1$ and $L_{\text{strip}} \leq W + 2\rho$. Bounds from [22] give us that $L_{\text{u-turn}} \leq w + \kappa\pi\rho$ and $L_{\text{closure}} \leq W + H + 2\rho + \kappa\pi\rho$, where $\kappa \approx 8/3$. The length of the closed path, L_1 , satisfies

$$L_1 \leq N_{\text{strips}}L_{\text{strip}} + (N_{\text{strips}} - 1)L_{\text{u-turn}} + L_{\text{closure}}.$$

Due to equal spacing of the UAVs along the loitering path, $E[d_1] = L_1/2m$. Combining these bounds we get that

$$E[d_1] \leq \frac{WH + 10.38\rho H}{2mw} + \frac{W + H + 6.19\rho}{m}.$$

To calculate $E[d_2]$ we define δ as the smallest distance from the target to any point on the loitering path (see Fig. 4). Since $d_2(s) = 2\rho \sin^{-1}(\sqrt{\frac{s}{2\rho}})$ for $s \leq \rho$ and δ is uniformly distributed between 0 and $w/2$, $E[d_2] = \frac{4\rho}{w} \int_0^{w/2} \sin^{-1}\left(\sqrt{\frac{s}{2\rho}}\right) ds$. From which, one can show that $E[d_2] \leq (3/4)\sqrt{\rho w}$. The coverage cost satisfies $\mathcal{W}_{\text{SL}} \leq E[d_1] + E[d_2]$. Combining this with the bounds on $E[d_1]$ and $E[d_2]$ we get that for $w \leq 2\rho$

$$\mathcal{W}_{\text{SL}} \leq \frac{WH + 10.38\rho H}{2mw} + \frac{W + H + 6.19\rho}{m} + \frac{3}{4}\sqrt{\rho w}. \quad (9)$$

Minimizing the right hand side of Equation (9) with respect to w subject to the constraint that $w \leq 2\rho$, we arrive at Eq. (5) and Eq. (6). The results in Eqs. (7) and (6) follow by taking the associated limits on Eq. (6). ■

Remark 5.6: Theorem 5.5 and Theorem 4.5 imply that the coverage cost belongs to $\Theta(1/m^{1/3})$.

Simulations

Simulations of the MC policy, in Fig. 5 (left), confirm theoretical predictions that the system time shrinks with $1/\sqrt{m}$: the log-log plots of system time versus number of vehicles have slopes ranging from -0.48 to -0.52 , implying a power law of approximately $-1/2$,

As the minimum turning radius becomes very small, the performance of the MC policy approximates the lower bound valid for a vehicle without kinematics constraints, i.e., as $\rho \rightarrow 0$, $\mathcal{W}_{\text{MC}} \rightarrow H^*(\mathcal{Q})$.

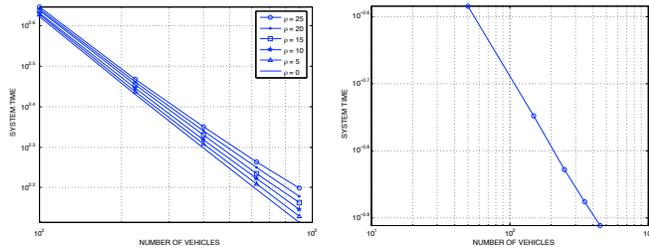


Fig. 5. Performance of the MC policy as a function of the number of vehicles (left) with $A = 1 \times 10^8$ and $d_\rho \leq 0.0056$. Simulation performance results for the SL policy (right) with $A = 1$ and $\rho = 0.2$

Simulations of the SL policy, in Fig. 5 (right), confirm that the system time shrinks with $1/m^{1/3}$: the log-log plot has

a slope of -0.34 , implying a power law of approximately $-1/3$. Interestingly, this data matches the theoretical asymptotic results over a very broad range of vehicle densities, including values as low as $d_\rho \geq 2$.

VI. THE PHASE TRANSITION

In Section V, we proposed two policies. These policies exhibit different modes of behavior. We have shown that the territorial MC policy is optimal as $d_\rho \rightarrow 0^+$ and the gregarious SL policy is constant-factor optimal as $m \rightarrow +\infty$. This suggests the existence of a phase transition in the optimal policy as one increases the number of vehicles for a fixed ρ and A .

We use simulation to compare the performances of the two policies under varying conditions. Fig. 6 shows which policy performs better under varying conditions. We vary the number of vehicles m between 1 and 100 and the ratio A/ρ^2 between 1 and about 2000. In these simulations, the environment is a square of unit area. The transition line has a slope of approximately 10.8, implying that the optimal policy switches behavior when $d_\rho^{\text{sim}} \approx 0.0925$.

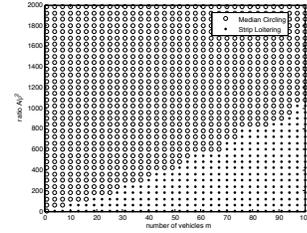


Fig. 6. Simulation results aimed at demarcating the phase transition. The environment is a square of unit area.

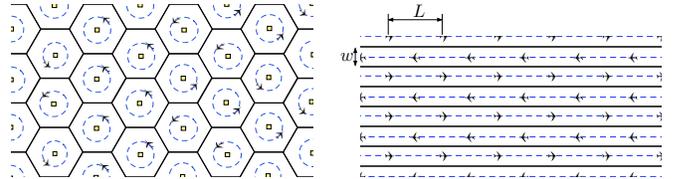


Fig. 7. The MC (left) and SL (right) policies with an infinite number of vehicles on an unbounded domain.

We are interested in fundamental factors driving the transition, ignoring its dependence on the shape of \mathcal{Q} . Towards this end, we envision an infinite number of vehicles operating on the unbounded plane, where the system is characterized by the vehicle density. Let us denote the inverse of the density with a , the area per vehicle. Depictions of the MC and SL policies operating on an unbounded domain are shown in Fig. 7. In this case, the configuration p^* yielding the minimum of the function $H_1(p^*)$ is that in which the Voronoi partition induced by p^* is a network of regular hexagons [12], each of area a . It is known that if \mathcal{Q} is a regular hexagon of area a then $H_1^*(\mathcal{Q}) \approx 0.377\sqrt{a}$. The system time of the policy satisfies $\mathcal{W}_{\text{MC}} \leq 0.377\sqrt{a} + 3.76\rho$. In this scenario, the SL policy reduces to vehicles moving straight on infinite strips, where the design criteria are the

width of the strips $w \leq 2\rho$, and the distance between consecutive vehicles in the same strip L . The system time of the policy satisfies $\mathcal{W}_{\text{SL}} \leq L/2 + 3\sqrt{\rho w}/4$. The area per vehicle is related to the design parameters by $Lw = a$. We substitute to get $\mathcal{W}_{\text{SL}} \leq a/2w + 3\sqrt{\rho w}/4$. Minimizing with respect to w , with the constraint that $w \leq 2\rho$, we choose $w = \min\{(4a/3\sqrt{\rho})^{2/3}, 2\rho\}$, yielding

$$\mathcal{W}_{\text{SL}} \leq \begin{cases} 1.238(\rho a)^{1/3} & \text{for } d_\rho \geq 0.471, \\ a/4\rho + 1.06\rho & \text{otherwise,} \end{cases} \quad (10)$$

where we have substituted the nonholonomic vehicle density with $d_\rho = \rho^2/a$. Equating the upper bound on the MC policy with the upper bound on the SL policy for $d_\rho < 0.471$, we get $0.377\sqrt{a} + 3.76\rho = a/4\rho + 1.06\rho$. Dividing both sides by ρ and substituting $x = \sqrt{a}/\rho$, the quadratic formula gives $x \approx 4.127$ and hence the critical nonholonomic vehicle density is given by $d_\rho^{\text{unbd}} = 1/x^2 \approx 0.0587$. This is a lower critical density than implied by the simulation results, favoring the SL policy. It would seem that the relaxation of the environment boundary has a greater impact on the performance of the SL policy, no longer requiring U-turns.

To gain intuition on the nature of this critical condition, consider the area per vehicle $a^{\text{crit}} = \rho^2/d_\rho^{\text{crit}}$, which yields $a^{\text{sim}} \approx 3.44\pi\rho^2$ and $a^{\text{unbd}} \approx 5.42\pi\rho^2$. In other words, the transition occurs when each vehicle is responsible for a region of area 3.44 or 5.42 times that of a minimum turning-radius disk.

VII. CONCLUSIONS

In this paper, we considered a coverage problem for a team of nonholonomic agents. We prove that the system time belongs to $\Theta(1/m^{1/3})$. It is interesting to compare this with the Euclidean case, where the system time belongs to $\Theta(1/m^{1/2})$. This shows that, in dynamic vehicle routing scenarios with low target-generation rate, the nonholonomic constraints make the system time less sensitive to the team size. This is in stark contrast to a result in [3], [5]: nonholonomic dynamics cause greater sensitivity to the team size when the target generation rate is high.

We showed that the optimal policy exhibits a phase transition: in low density scenarios, the optimal policy is one of territorial behavior, but as the density increases, a gregarious behavior is required to reap the benefits of a large team. We identified a parameter we call the nonholonomic vehicle density as a key in characterizing the mode of the system, and used numerical experiment and analytical techniques to study the value of the parameter which demarcates the phase transition. Furthermore, we aim to design decentralized strategies for the individual agents to obey, maintaining near optimal performance for the system as a whole. A game-theoretic approach might lend itself naturally to a completely distributed control policy.

We also would like to compare our results with those found empirically by observing biological systems [16].

ACKNOWLEDGMENTS

The authors thank Marco Pavone for insightful feedback during the development of this work. This research was done

with support from the Michigan/AFRL Collaborative Center on Control Sciences. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the supporting organizations.

REFERENCES

- [1] W. Li and C. Cassandras, "A cooperative receding horizon controller for multivehicle uncertain environments," *IEEE Trans. on Automatic Control*, vol. 51, no. 2, pp. 242–257, 2006.
- [2] X. Ma and D. A. Castanon, "Receding horizon planning for dubins traveling salesman problems," in *IEEE Conference on Decision and Control*, (San Diego, CA), Dec 2006.
- [3] J. J. Enright, E. Frazzoli, K. Savla, and F. Bullo, "On multiple UAV routing with stochastic targets: Performance bounds and algorithms," in *Proc. of the AIAA Conf. on Guidance, Navigation, and Control*, (San Francisco, CA), August 2005.
- [4] K. Savla, F. Bullo, and E. Frazzoli, "The coverage problem for loitering Dubins vehicles," (New Orleans, LA), pp. 1398–1403, Dec. 2007.
- [5] J. J. Enright, K. Savla, E. Frazzoli, and F. Bullo, "Dynamic routing of UAV teams," *AIAA Journal of Guidance, Control, and Dynamics*, 2008 (submitted).
- [6] J. Cortés, S. Martínez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," vol. 20, no. 2, pp. 243–255, 2004.
- [7] H. Choset, "Coverage of known spaces: the boustrophedon cellular decomposition," *Autonomous Robots*, vol. 9, no. 3, pp. 247–253, 2000.
- [8] Y. Guo and M. Balakrishnan, "Complete coverage control for non-holonomic mobile robots in dynamic environments," (Orlando, FL), pp. 1704–1709, 2006.
- [9] S. P. Fekete, J. S. B. Mitchell, and K. Weinbrecht, "Continuous weber and k-median problems," in *16th Annual ACM Symposium on Computational Geometry*, pp. 70–79, 2000.
- [10] P. K. Agarwal and M. Sharir, "Efficient algorithms for geometric optimization," *ACM Computing Surveys*, vol. 30, no. 4, pp. 412–458, 1998.
- [11] Z. Drezner, ed., *Facility Location: A Survey of Applications and Methods*. Springer Series in Operations Research, New York: Springer Verlag, 1995.
- [12] E. Zemel, "Probabilistic analysis of geometric location problems," *Annals of Operations Research*, vol. 1, October 1984.
- [13] C. H. Papadimitriou, "Worst-case and probabilistic analysis of a geometric location problem," *SIAM Journal on Computing*, vol. 10, August 1981.
- [14] T. Vicsek, A. Czirok, E. Ben-Jacob, I. Cohen, and O. Shochet, "Novel type of phase transition in a system of self-driven particles," *Physical Review Letters*, vol. 75, pp. 1226–1229, Aug 1995.
- [15] W. Spears, J. Spears, J. Hamann, and R. Heil, "Distributed, physics-based control of swarms of vehicles," *Autonomous Robots*, vol. 17, no. 2-3, pp. 137–162, 2004.
- [16] J. Buhl, D. J. Sumpter, I. D. Couzin, J. Hale, E. Despland, E. Miller, and S. J. Simpson, "From disorder to order in marching locusts," *Science*, vol. 312, pp. 1402–1406, 2006.
- [17] L. Dubins, "On curves of minimal length with a constraint on average curvature and with prescribed initial and terminal positions and tangents," *American Journal of Mathematics*, vol. 79, pp. 497–516, 1957.
- [18] B. Thomaschewski, "Dubins' problem for free terminal direction," Preprint M09/2001, Technical University of Ilmenau, Ilmenau, Germany, 2001.
- [19] N. Megiddo and K. J. Supowit, "On the complexity of some common geometric location problems," *SIAM Journal on Computing*, vol. 13, no. 1, pp. 182–196, 1984.
- [20] S. Masuyama, T. Ibaraki, and T. Hasegawa, "The computational complexity of the m-center problem on the plane," *Trans. IECE of Japan*, vol. 64, no. 2, pp. 57–64, 1981.
- [21] S. Boyd and L. Vandenberghe, *Convex Optimization*. 2004.
- [22] K. Savla, E. Frazzoli, and F. Bullo, "Traveling salesperson problems for the dubins vehicle," *IEEE Trans. on Automatic Control*, vol. 53, no. 6, pp. 1378–1391, 2008.