

Stochastic and Dynamic Routing Problems for multiple UAVs

John J. Enright*

Kiva Systems, Woburn, MA, 01801, USA.

Ketan Savla[†] and Emilio Frazzoli[‡]

Massachusetts Institute of Technology, Cambridge, MA, 02139, USA.

Francesco Bullo[§]

University of California, Santa Barbara, CA, 93106, USA.

Consider a routing problem for a team of vehicles in the plane: target points appear randomly over time in a bounded environment and must be visited by one of the vehicles. It is desired to minimize the expected system time for the targets, i.e., the expected time elapsed between the appearance of a target point, and the instant it is visited. In this paper, such a routing problem is considered for a team of Uninhabited Aerial Vehicles (UAVs), modeled as vehicles moving with constant forward speed along paths of bounded curvature. Three algorithms are presented, each designed for a distinct set of operating conditions. Each is proven to provide a system time within a constant factor of the optimal when operating under the appropriate conditions. It is shown that the optimal routing policy depends on problem parameters such as the *workload* per vehicle and the vehicle *density* in the environment. Finally, there is discussion of a *phase transition* between two of the policies as the problem parameters are varied. In particular, for the case in which targets appear sporadically, a dimensionless parameter is identified which completely captures this phase transition and an estimate of the critical value of the parameter is provided.

*Algorithm Designer, Kiva Research Team, jenright@ucla.edu. AIAA Member.

[†]Postdoctoral Associate, Laboratory for Information and Decision Systems, ksavla@mit.edu. AIAA Member.

[‡]Associate Professor, Laboratory for Information and Decision Systems, Aeronautics and Astronautics Department, frazzoli@mit.edu. AIAA Senior Member.

[§]Professor, Mechanical Engineering Department and the Center for Control, Dynamical Systems and Computation, bullo@engineering.ucsb.edu. AIAA Member.

Nomenclature

A	area of the bounded domain, m^2
a	area per vehicle in the unbounded domain, m^2
d_ρ	nonholonomic vehicle density
$D_\rho(g_1, g_2)$	configuration-to-configuration Dubins distance function, m
$D(t)$	set of targets outstanding at time t
g	configuration (position and heading) of vehicle
$H_m(p, \mathcal{Q})$	multi-median function, m
$H_m^*(\mathcal{Q})$	minimum value of the multi-median function, m
I_m	the set of integers between 1 and m , included
$L_\rho(g, q)$	configuration-to-position Dubins distance function, m
m	number of UAVs
\bar{n}	steady-state number of outstanding targets
p	position of the Voronoi generators, m
q	point in the domain, m
\mathcal{Q}	a bounded planar domain
$R_t(g)$	the set of positions reachable from configuration g within time t
$SE(2)$	Special Euclidean group in two dimensions (i.e., the group of rigid planar motions)
t	time, s
\bar{T}	the steady-state system time, s
v	vehicle speed, ms^{-1}
W, H	width and height of a box bounding the domain \mathcal{Q} , m
λ	target generation rate, s^{-1}
π	routing policy
ρ	minimum turning radius, m
ω	angular velocity, rad s^{-1}
<i>Subscripts</i>	
i	vehicle index
j	target index

I. Introduction

Wide-area surveillance is one of the prototypical missions for Uninhabited Aerial Vehicles (UAVs), e.g., in environmental monitoring, security, or military settings. Low-altitude UAVs on such a mission must provide coverage of a region and investigate events of interest as they manifest themselves. In particular, it is of interest to consider 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 their locations to gather on-site information.

In many studies of problems falling into this class, the locations of targets are known a priori, and a strategy is computed that attempts to optimize the cost of servicing the known targets. In this paper, a scenario is considered 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. This formulation is a variation of the Dynamic Traveling Repairman Problem (DTRP)—introduced by Psaraftis¹ and thoroughly developed by Bertsimas and van Ryzin^{2–4}—with the addition of differential constraints on the vehicle’s motion. In particular, this paper concentrates on vehicles that are constrained to move along paths of bounded curvature, without reversing direction. Such vehicles, often referred to as Dubins vehicles, have been extensively studied in the robotics and control literature.^{5–7} Moreover, the Dubins vehicle model is widely accepted as reasonably accurate to represent aircraft kinematics, e.g., for air traffic control^{8,9} and UAV mission planning purposes.^{10–12}

The DTRP analysis in the literature^{2–4} is broken into two limiting cases: i) targets are generated sporadically (*light load*), and ii) targets are generated rapidly (*heavy load*). In heavy load, the optimal policy for holonomic vehicles relies upon Euclidean Traveling Salesman tours through large sets of target points. Policies based on Euclidean distance are known to provide very poor performance for UAVs modeled as nonholonomic vehicles for high target densities;¹³ in this case, we use recently-developed algorithms that aim at minimizing the travel distance, taking into account the turning cost between consecutive targets.¹⁴ In light load, rather than planning tours through targets, the challenge is to design loitering patterns or policies with the property that, when a target does appear, the expected wait for the “closest” UAV to arrive is minimized. This problem is inherently more complex for nonholonomic than for holonomic vehicles: for holonomic vehicles the light load DTRP reduces to a choice of waiting locations and solutions are known from the well-developed locational optimization literature.¹⁵

The recent literature concerning DTRP and aerial surveillance problems is vast. This paper is along the vehicle routing themes of some of our previous work.^{14,16,17} Many of the new results for the light load were introduced in Enright *et al.*¹⁸ and are applicable to coverage problems,^{19–22} in which the vehicles spread out uniformly, or comb the environment efficiently. The light load case discussed in this paper has connections to the Persistent Area Denial (PAD) and area coverage problems.^{12,19,23,24} Moreover, the mathematical structure of the light load case has a strong resemblance with continuous facility location problems,^{15,25,26} with the main difference in our case being that the facilities are vehicles constantly in motion with nontrivial dynamics. The heavy load case is strongly related to works concerned with the generation of efficient cooperative strategies for several vehicles moving through given

target points and possibly avoiding obstacles or threats.^{11,27-32} Trajectory efficiency in these cases is understood in terms of cost for the vehicles: in other words, efficient trajectories minimize the total path length, the time needed to complete the task, or the fuel/energy expenditure. A related problem has been investigated as the Weapon-Target Assignment (WTA) problem, in which vehicles are allowed to team up in order to enhance the probability of a favorable outcome in a target engagement.^{33,34} Other related works address UAV task-assignment:³⁵⁻³⁸ in this setup, targets locations are known and an assignment strategy is sought that maximizes a global performance metric such as success rate. Some other works, like this paper, address coverage and sensing tasks with only prior statistics on the operating environment.³⁹

The main contributions of the paper are the following.

1. A novel approach to establish a lower bound on the system time for Dubins DTRP is formulated. The approach explores properties of the reachable set of the Dubins vehicle and may be of independent interest as it can be applied to other vehicles performing similar tasks.
2. Three algorithms have been proposed and it is proven that, for certain range of the problem parameters, they perform within a constant factor of the optimum, i.e., the upper bounds on their performances have the *same form* (in terms of team size m , speed v and target-generation rate λ) as the lower bounds. Specifically, two strategies for the light load are designed. 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 vehicles to move as a group in a coordinated pattern, balancing the amount of space directly in front of each vehicle. This strategy is within a constant factor of the optimal for high vehicle densities. In the heavy load, a previously proposed single-vehicle policy¹⁴ is extended to the multiple-vehicle setting.
3. In light load, the optimal system time is shown to be of order $1/(v\sqrt[3]{m})$, contrasting the holonomic case which yields $1/(v\sqrt{m})$. The added performance per vehicle is penalized by the nonholonomic motion constraints. In heavy load, the system time is shown to be of order $\lambda^2/(mv)^3$, i.e., the system time is more sensitive to the workload per vehicle for nonholonomic vehicles than for holonomic vehicles (where the system time is of order $\lambda/(mv)^2$).
4. A consequence of the aforementioned analytical results is recognized: an optimal routing policy for the UAVs exhibits *phase transitions* between these three algorithms as one varies the problem parameters. For the light load case, a dimensionless parameter

that completely captures the phase transition is identified. An estimate of the critical value of the parameter that separates the phases is also provided.

Other researchers, including Vicsek *et al.*⁴⁰ and, more recently, Spears *et al.*,⁴¹ have investigated phase transitions in the behavior of large groups of vehicles. In these works, the individuals follow local interaction laws, sometimes inspired by physical laws, which are shown to give rise to different modes of group behavior, depending on design parameters and possible input noise. In some cases, these types of approach are useful for accomplishing large-scale multi-vehicle coverage and similar tasks, in a distributed, robust fashion. In contrast, a top-down approach is taken in this paper. Beginning with simple motion constraints (found in nature and technology) for the individual vehicles, and an explicit task and performance measurement, the aim is to design control strategies achieving near optimal performance. As a result of the analysis, the optimal policy is noticed to exhibit more than one distinct mode of behavior. Rather than attempting to imitate nature, an attempt is made in this paper to divulge a possible cause (efficiency) behind natural multi-agent systems which exhibit phase transitions between territorial and gregarious behavior.⁴²

The rest of the paper is organized as follows. In Section II, the problem is set up and some preliminary concepts and notations are introduced. In Section III, the case in which targets appear sporadically is addressed, and in Section IV, the case in which targets are generated rapidly is investigated. A phase transition in the optimal routing policy is studied in Section V. Finally, some concluding remarks are made and directions for future research are discussed in Section VI. A detailed exposition of some relevant results about minimum-length Dubins paths is given in Appendix A, and some technical lemmas along with their proofs are given in Appendix B.

II. Problem Formulation and Preliminary Concepts

II.A. Problem Formulation

In this section, the *Dubins DTRP* is formulated and some preliminary concepts for minimum-time paths and reachable sets for the Dubins vehicle are presented. Let $\mathcal{Q} \subset \mathbb{R}^2$ be a convex, compact domain on the plane, with non-empty interior; \mathcal{Q} is assumed to be at the ground level, and will be referred to as the *environment*. Let A be the area of \mathcal{Q} . A spatio-temporal Poisson process, with finite time intensity $\lambda > 0$ and uniform spatial density in \mathcal{Q} , generates *targets* inside the environment. This Poisson process can be thought of as a collection of functions $\{\mathcal{P} : \overline{\mathbb{R}}_+ \rightarrow 2^{\mathcal{Q}}\}$ such that, for any $t \geq 0$, $\mathcal{P}(t)$ is a random collection of points in \mathcal{Q} , representing the targets generated in the time interval $[0, t)$, and such that

- the number of targets generated in two disjoint time-space regions are *independent*

random variables;

- the *expected* number of targets generated in a region $S \subseteq \mathcal{Q}$ during a time interval of length Δt is

$$\mathbb{E}[\text{card}((\mathcal{P}(t + \Delta t) \setminus \mathcal{P}(t)) \cap \mathcal{S})] = \lambda \Delta t \cdot \frac{\text{Area}(\mathcal{S})}{A}.$$

The information about the locations of these dynamically generated targets is provided by surveillance assets (e.g., high-altitude aircraft or spacecraft, ground spotters, sensor networks, etc.), and is assumed to be broadcast to the UAVs instantaneously. Each target corresponds to a service request. These service requests are identical and are to be fulfilled by a team of m UAVs. A service request is fulfilled when one of m UAVs flies over the associated target point (assumed to be on the ground). The UAVs operate on horizontal planes, at different altitudes to avoid collisions. Also, we assume that all these horizontal planes are high enough to clear all terrain and obstructions on the ground. We do not model no-fly zones, and leave this to future work. Since the flying altitude does not play a significant role in the course of the paper, the motions of all the UAVs are projected onto the plane containing \mathcal{Q} and restrict our attention to this plane. Henceforth, the location of the UAV shall be identified with its projection onto the ground plane. Finally, we assume that the UAVs maintain a constant speed; this assumption is consistent with standard operational practice, as aircraft typically fly most of their mission at a constant cruising speed, calculated in such a way to maximize range or endurance.⁴³

The motion of the UAVs will be modeled by a kinematic model in \mathbb{R}^2 . In particular, the UAVs are modeled as nonholonomic vehicles constrained to move on the plane at constant speed, $v > 0$, along paths with a minimum radius of curvature $\rho > 0$. In other words, if the configuration $g_i \in SE(2)$ of the i -th vehicle ($\forall i \in I_m$ where $I_m = \{1, \dots, m\}$) is 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 heading with respect to the x -axis—then the model of the vehicle is given by:

$$\begin{aligned} \dot{x}_i &= v \cos(\theta_i), \\ \dot{y}_i &= v \sin(\theta_i), \\ \dot{\theta}_i &= \omega_i, \quad |\omega_i| \leq v/\rho. \end{aligned}$$

This model is often referred to as the Dubins vehicle, in recognition of Dubins’ work in computing minimum-length paths for such model.⁵ A *feasible path for the Dubins vehicle* is defined as a path that is twice differentiable almost everywhere, and such that its radius of curvature is bounded below by ρ . The UAVs are assumed to be identical, and have unlimited range. The above kinematic model is very common in the literature on UAV

motion planning.^{10–12} Results in terms of minimum-length paths for Dubins vehicle⁵ hold for our model, where they assume an additional connotation of being minimum-time paths for vehicles with constant speeds. In the course of this paper, the term *Dubins frame* shall be used to refer to a coordinate frame with the origin attached to the Dubins vehicle and x -axis along the vehicle's velocity vector.

Let $\mathcal{D} : t \rightarrow 2^{\mathcal{Q}}$ indicate a realization of the stochastic process obtained by combining the target-generation process \mathcal{P} and the removal process caused by the vehicles visiting the outstanding requests. The random set $\mathcal{D}(t) \subset \mathcal{Q}$ represents the *demand*, i.e., the service requests outstanding at time t ; $n(t)$ is defined to be the cardinality, i.e., the number of elements, of the set $\mathcal{D}(t)$.

A motion coordination policy is a function that determines the actions of each vehicle over time. For the time being, denote these functions will be denoted as $\pi = (\pi_1, \pi_2, \dots, \pi_m)$, without explicitly stating their domain; the output of these functions is a steering command for each vehicle. The objective is the design of motion coordination strategies that allow the vehicles to fulfill service requests efficiently.

A policy $\pi = (\pi_1, \pi_2, \dots, \pi_m)$ is said to be *stabilizing* if, under its effect, the expected number of outstanding targets does not diverge over time, i.e., if

$$\bar{n}_\pi = \lim_{t \rightarrow +\infty} \mathbb{E}[n(t) : \text{vehicle } i \text{ executes policy } \pi_i, i \in I_m] < +\infty.$$

Intuitively, a policy is stabilizing if the vehicles are able to visit targets at a rate that is, on average, at least as fast as the rate at which new service requests are generated. Let T_j be the time elapsed between the generation of the j -th service request, and the time it is fulfilled. If the system is stable, then the following balance equation (also known as Little's formula⁴⁴) holds:

$$\bar{n}_\pi = \lambda \bar{T}_\pi, \tag{1}$$

where $\bar{T}_\pi := \lim_{j \rightarrow +\infty} \mathbb{E}[T_j]$ is the system time under policy π , i.e., the expected time a service request must wait before being fulfilled, given that the vehicles follow the strategy defined by π . The system time \bar{T}_π can be thought of as a measure of the quality of service collectively provided by the vehicles.

At this point the problem can be finally stated as: we wish to devise policies that are (i) stabilizing, and (ii) yield a quality of service (i.e., system time) achieving, or approximating, the theoretical optimal performance given by

$$\bar{T}_{\text{opt}} = \inf_{\pi \text{ stabilizing}} \bar{T}_\pi$$

It is of interest to design computationally efficient control policies that are within a constant factor of the optimal, i.e., policies π such that $\bar{T}_\pi \leq \kappa \bar{T}_{\text{opt}}$ for some constant κ . In the rest of the paper, the terms policy and algorithm are used interchangeably.

Since it is difficult to characterize the optimal achievable performance for general values of the problem parameters, the focus is on particular range of problem parameters. To that purpose, $\frac{\lambda}{m}$ is defined as the *vehicle workload*. This parameter characterizes the average rate at which each vehicle must service requests in order for the number of outstanding targets to remain bounded. It is found that the asymptotic cases of $\lambda/m \rightarrow 0^+$ (light load) and $\lambda/m \rightarrow +\infty$ (heavy load) lend themselves to succinct and meaningful analysis.

This section is concluded by describing some notation providing a concise way of expressing the asymptotic dependence of the system time \bar{T} on the problem parameters. For $f, g : \mathbb{N} \rightarrow \mathbb{R}$, $f \in O(g)$ (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 the notation $f \in \Theta(g)$ is used.

II.B. Minimum-length 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 Dubins;⁵ a further classification is given in Shkel and Lumelsky.⁴⁵ Our purpose in this section is to study the length of optimal paths given different boundary conditions. Namely, the problem of steering a Dubins vehicle from a given initial configuration to a point in the plane will be considered; the difference with the original problem posed by Dubins⁵ is that the final heading at the target point is not constrained a priori. Let $L_\rho(g, q) : \text{SE}(2) \times \mathbb{R}^2 \rightarrow \mathbb{R}^+$ be the length of shortest feasible path for a Dubins vehicle to steer from initial configuration g in $\text{SE}(2)$ to a point q in the plane, without any constraints on the final heading, i.e., to any configuration in the set $\{q\} \times S^1 \subset \text{SE}(2)$. The characteristics of paths of minimal length with such boundary conditions were studied in Thomaschewski,⁴⁶ 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. The length of the circular arcs is at most $2\pi\rho$, and all subpaths are allowed to have zero length. In the Dubins formalism, such paths are called either of CC (circle, circle) or of CL (circle, line) type, and include the subtypes C or L , arising when the length of either one of the two subpaths is zero. Clearly, both subpaths have zero length in the trivial case in which q is the initial position of the vehicle. Furthermore, the optimal paths are of type CC when q lies within the interior of one of the two circles of radius ρ that are tangent to the vehicle's direction of motion at the initial configuration, and are of type CL otherwise. Closed-form expressions for the lengths of such paths, and a detailed

exposition on the function $L_\rho(g_i, q)$, are given in Appendix A. For now, an upper bound on $L_\rho(g_i, q)$ which will be useful in analyzing one of the presented policies is presented.

Proposition 1 *For any $q \in \mathbb{R}^2$ and $g_0 = (0, 0, 0)$, there exist constants $c_1 > 0, c_2 > 0$ such that the function $L_\rho(g_0, q)$ satisfies the following inequality.*

$$L_\rho(g_0, q) \leq c_1\rho + \|q - (0, \pm c_2\rho)\|. \quad (2)$$

In particular, the assertion is true for $c_1 \approx 9.39$ and $c_2 = 1$.

Proof: Let $D_\rho(g_1, g_2)$ be the length of shortest feasible path for a Dubins vehicle to steer from initial configuration g in $SE(2)$ to a final configuration $g_2 \in SE(2)$. Since, $L_\rho(g_0, q)$ is the length of the *shortest* path from g_0 to q , one can verify that it satisfies the following relation for any $q_c \in \mathbb{R}^2$.

$$L_\rho(g_0, q) \leq \max_{\theta \in [0, 2\pi)} D_\rho(g_0, (q_c, \theta)) + \|q - q_c\|. \quad (3)$$

The following bound is proven in Savla *et al.*:¹⁴

$$D_\rho(g_0, (q_c, \theta)) \leq \|q_c\| + 8.39\rho \quad \forall \theta \in [0, 2\pi). \quad (4)$$

The result follows by substituting Equation (4) into Equation (3) for $q_c = (0, \pm\rho)$. ■

Remark 2 *An implication of Proposition 1 is that, if a vehicle is moving along a circle of radius $c_2\rho$ about a point q_{circ} , then it can reach any point $q \in \mathbb{R}^2$ with a path of length no greater than $c_1\rho + \|q - q_{\text{circ}}\|$ at any time during its circular motion.*

Remark 3 *Using numerical optimization, it is found that, by selecting $c_2 \approx 2.91$, one could reduce c_1 to approximately 3.76.*

Remark 4 *An upper bound on $L_\rho(g_0, q)$ under the simplifying assumption of $\|q\| \geq 2\rho$ is also obtained in Rathinam *et al.*⁴⁷*

The *reachable set* for a Dubins vehicle and a property that will be useful later in the paper are defined next. Given an initial condition g_i , the reachable set $R_{\leq t}(g_i) \subset SE(2)$ for a Dubins vehicle is the set of final configurations g_f such that there exist admissible controls to drive the vehicle from the initial configuration g_i to the final configuration g_f within time t . It is well known that the Dubins vehicle is not *small-time locally controllable*, i.e., there exist $t > 0$ for which the reachable set $R_{\leq t}(g_i)$ does not contain an ϵ -neighborhood of g_i . It is of interest in this paper to focus on the reachable region on the Euclidean plane. For any

$g = (p, \theta) \in \text{SE}(2)$, define the projection operator $P : \text{SE}(2) \rightarrow \mathbb{R}^2$ as $P(g) = p$. Furthermore, for any $R \subseteq \text{SE}(2)$ let $P(R) = \cup_{g \in R} P(g)$. For simplicity of notation the reachable region on the plane, $P(R_{\leq t}(g_i))$, shall be henceforth indicated with $R_t(g_i)$.

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 the shortest path to any point reachable within time $t \leq \pi\rho/(2v)$ is of type CL. In Fig. 1, $R_t(g_i)$ for some time $t \leq \pi\rho/(2v)$ is depicted, along with a generic minimum-length path of length vt . Let the length of the C segment be denoted by s and the angle it sweeps out with θ . Since $s = \rho\theta$, the length of the L segment is $vt - s = vt - \rho\theta$. Also, the switching point between the C and L segments is denoted with (x_1, y_1) , the point at the end of the L segment with (x_2, y_2) . Elementary




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

planar geometry computations yield

$$\begin{aligned} x_1 &= \rho \sin \theta, \\ y_1 &= \rho(1 - \cos \theta), \\ x_2 &= x_1 + (vt - \rho\theta) \cos \theta, \\ y_2 &= y_1 + (vt - \rho\theta) \sin \theta. \end{aligned}$$

Moreover, for ρ and t such that $t \leq \pi\rho/(2v)$,

$$\begin{aligned} \text{Area}(R_t(g_i)) &= 2 \left(\int_0^{vt/\rho} y_1 \frac{dx_1}{d\theta} d\theta + \int_{vt/\rho}^0 y_2 \frac{dx_2}{d\theta} d\theta \right) \\ &= 2\rho^2 \int_0^{vt/\rho} (\sin^2 \theta (vt/\rho - \theta)^2 + \sin \theta (1 - \cos \theta) (vt/\rho - \theta) + \cos \theta (1 - \cos \theta)) d\theta. \end{aligned}$$

The above integral can be computed as

$$\text{Area}(R_t(g_i)) = \frac{v^3 t^3}{3\rho} \quad \text{for } t \leq \frac{\pi\rho}{2v}. \quad (5)$$

Remark 5 Equation (5) shows that $\text{Area}(R_t(g_i))$ approaches zero as $(vt)^3/\rho$, i.e., the area of this reachable set decreases faster than the area of a circle of radius vt as $vt/\rho \rightarrow 0^+$.

Equation (5) also implies the following result.

Proposition 6 For a given $a, \rho > 0$ and $g_i \in \text{SE}(2)$, let τ be such that $\text{Area}(R_\tau(g_i)) = a$. Then,

$$\tau = \frac{1}{v} (3\rho a)^{1/3} \quad \text{for } a/\rho^2 \leq \pi^3/24.$$

III. Light Load

In this section the case in which new targets appear very rarely is addressed, and lower bounds and constructive upper bounds on the system time are provided.

III.A. Lower Bounds

A review of the known results for the Euclidean case is provided first, which are then used to derive lower bounds.

Euclidean Lower Bound

A trivial lower bound on the system time \bar{T} can be 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, a brief overview of a related problem from geometric optimization is necessary. Given a convex, compact 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 expected distance between a random point q , sampled from a uniform distribution over \mathcal{Q} , and the closest point in p is given by

$$\mathcal{H}_m(p, \mathcal{Q}) := \frac{1}{A} \int_{\mathcal{Q}} \min_{i \in I_m} \|p_i - q\| \, dq = \frac{1}{A} \sum_{i=1}^m \int_{\mathcal{V}_i(p)} \|p_i - q\| \, dq,$$

where $\mathcal{V}(p) = (\mathcal{V}_1(p), \mathcal{V}_2(p), \dots, \mathcal{V}_m(p))$ is the *Voronoi (Dirichlet) partition*⁴⁸ of \mathcal{Q} generated by the points in p , i.e.,

$$\mathcal{V}_i(p) = \{q \in \mathcal{Q} : \|q - p_i\| \leq \|q - p_j\|, \forall i, j \in I_m\}.$$

The problem of choosing p to minimize $\mathcal{H}_m(p, \mathcal{Q})$ is known in the geometric optimization literature as the *continuous supply / continuous demand multi-median problem*.^{49,50} Related problems include computing m -medians and centers for discrete supply and/or demand sets and references therein.^{15,25,26,51,52} The m -median of the set \mathcal{Q} is the global minimizer $p^*(\mathcal{Q}) = \operatorname{argmin}_{p \in \mathcal{Q}^m} \mathcal{H}_m(p, \mathcal{Q})$. Let $\mathcal{H}_m^*(\mathcal{Q}) = \mathcal{H}_m(p^*(\mathcal{Q}), \mathcal{Q})$ be the global minimum of $\mathcal{H}_m(p, \mathcal{Q})$ for fixed \mathcal{Q} and m . It is straightforward to show that the map $p \mapsto \mathcal{H}_1(p, \mathcal{Q})$ is differentiable and strictly convex on \mathcal{Q} . Therefore, it is a simple computational task to compute $p_1^*(\mathcal{Q})$. However, 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 known²⁶ 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’’⁴⁹ discussed below, and ii) the fact that p_i is the median of $\mathcal{V}_i(p)$ for each p_i in $p^*(\mathcal{Q})$. An adaptive sampling-based algorithm to compute p , converging to local minima or saddle points of $\mathcal{H}_m(p, \mathcal{Q})$ is given in.⁵³ The issue of computation of the m -median and of the corresponding $\mathcal{H}_m^*(\mathcal{Q})$ will not be pursued further, but it will be assumed that these values are available.

In course of the paper, the dependence of $\mathcal{H}_m^*(\mathcal{Q})$ on m will play a crucial role in the design and analysis of algorithms. However, finding the exact relationship for the general case is difficult. Hence, now the problem of providing some tight bounds on $\mathcal{H}_m^*(\mathcal{Q})$ is considered. This problem is studied thoroughly in Papadimitriou⁵⁰ for square regions and in Zemel⁴⁹ for more general compact regions. It is known that, in the asymptotic case ($m \rightarrow +\infty$), $\mathcal{H}_m^*(\mathcal{Q}) = c_{\text{hex}} \sqrt{A/m}$ almost surely,⁴⁹ where $c_{\text{hex}} \approx 0.377$ is the first moment of a hexagon of unit area about its center. This optimal asymptotic value is achieved by placing the m points at the centers of the hexagons in a regular hexagonal lattice within \mathcal{Q} (the honeycomb heuristic). Working towards the above result, it is also known⁴⁹ that for any $m \in \mathbb{N}$:

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

where $c(\mathcal{Q})$ is a constant depending on the shape of \mathcal{Q} . The above result is mirrored in the following Proposition for several reasons. First, the proof technique that will be used for the lower bound foreshadows much of the logical structure used in the Dubins lower bound of Theorem 10, without the additional complexity due to the motion constraints of the Dubins vehicle. Second, the proof of the upper bound is omitted from the published work.⁴⁹ The proof for a specific example of such a constant $c(\mathcal{Q})$ is provided next.

Proposition 7 *For any convex, compact $\mathcal{Q} \subset \mathbb{R}^2$ of area A , the function $\mathcal{H}_m^*(\mathcal{Q})$ belongs to*

$\Theta(1/\sqrt{m})$. 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 containing \mathcal{Q} .

Proof: The proof of the upper bound is given first. If the environment \mathcal{Q} is a square of side-length L , it is known⁵⁴ that $H_1^*(\mathcal{Q}) = c_{\text{sq}}L$, with $c_{\text{sq}} \approx 0.38$. If the environment is a square and m is a perfect square (i.e., \sqrt{m} is integer), then the environment can be divided into m equally sized squares, each with one of the m points at its center, and so $\mathcal{H}_m^*(\mathcal{Q}) \leq c_{\text{sq}}L/\sqrt{m}$. For an environment of general shape and a perfect square m ,

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

where $L(\mathcal{Q})$ is the side-length of the smallest square containing \mathcal{Q} . This holds because the optimal value for a square containing \mathcal{Q} is strictly greater than the optimal value for \mathcal{Q} itself. For general \mathcal{Q} and m only $\lfloor\sqrt{m}\rfloor^2$ of the m points are utilized and the same technique as earlier is used to get:

$$\mathcal{H}_m^*(\mathcal{Q}) \leq \frac{c_{\text{sq}}L(\mathcal{Q})}{\lfloor\sqrt{m}\rfloor}.$$

Finally, noting that $\lfloor\sqrt{m}\rfloor \geq \sqrt{m}/\sqrt{3}$ for any $m \in \mathbb{N}$, the stated upper bound is obtained.

Now the lower bound is proved. Let $A_i = \text{Area}(\mathcal{V}_i(p))$ be the area of the Voronoi region belonging to vehicle i and let $C_{A_i}(p_i)$ be a disk of area A_i centered at p_i . Then

$$\begin{aligned} \mathcal{H}_m^*(\mathcal{Q}) &= \min_{p \in \mathcal{Q}^m} \int_{\mathcal{Q}} \frac{1}{A} \min_{i \in I_m} \|p_i - q\| dq \\ &= \min_{p \in \mathcal{Q}^m} \sum_{i=1}^m \int_{\mathcal{V}_i(p)} \frac{1}{A} \|p_i - q\| dq \\ &\geq \min_{p \in \mathcal{Q}^m} \sum_{i=1}^m \int_{C_{A_i}(p_i)} \frac{1}{A} \|p_i - q\| dq \\ &\geq \min_{\{A_1, A_2, \dots, A_m\} \in \mathbb{R}^m} \sum_{i=1}^m \int_{C_{A_i}(p_i)} \frac{1}{A} \|p_i - q\| dq, \quad \text{subject to } \sum_i^m A_i = A, \end{aligned}$$

where the fact that for a given p_i , the region of area A_i with the least first moment about p_i is a disk of that area is used. Note that the radius of such a disk is $r(A_i) = \sqrt{A_i/\pi}$. It is known that the expected distance between a uniformly distributed point within a disk of

radius r and its center is $2r/3$. Thus,

$$\begin{aligned} \mathcal{H}_m^*(\mathcal{Q}) &\geq \min_{\{A_1, A_2, \dots, A_m\} \in \mathbb{R}^m} \sum_{i=1}^m \frac{A_i}{A} \frac{2}{3} r(A_i), \quad \text{subject to} \quad \sum_{i=1}^m A_i = A, \\ &= \frac{2}{3A\sqrt{\pi}} \min_{\{A_1, A_2, \dots, A_m\} \in \mathbb{R}^m} \sum_{i=1}^m A_i^{3/2}, \quad \text{subject to} \quad \sum_{i=1}^m A_i = A. \end{aligned}$$

The above quantity is minimized when $A_i = A/m$, for all i , yielding the stated result. Note that the above lower bound holds for any region and any number of vehicles. ■

The following result was first proved in Bertsimas and Van Ryzin² for holonomic vehicles, and as such is trivially valid for nonholonomic vehicles. A short proof is reported for completeness.

Theorem 8 *For any convex, compact $\mathcal{Q} \subset \mathbb{R}^2$ of area A and $\lambda, \rho, v \in \mathbb{R}^+$ and $m \in \mathbb{N}$, the system time for the Dubins DTRP satisfies*

$$\bar{T}_{\text{opt}} \geq \frac{\mathcal{H}_m^*(\mathcal{Q})}{v}. \quad (6)$$

Proof: Consider vehicle i with configuration g_i and position p_i . Trivially, $L_\rho(g_i, q) \geq \|p_i - q\|$ for any q in \mathbb{R}^2 . Thus,

$$\begin{aligned} \bar{T}_{\text{opt}} &\geq \inf_{g \in \text{SE}(2)^m} \int_{\mathcal{Q}} \frac{1}{A} \min_{i \in I_m} \frac{1}{v} L_\rho(g_i, q) dq \geq \inf_{p \in (\mathbb{R}^2)^m} \int_{\mathcal{Q}} \frac{1}{A} \min_{i \in I_m} \frac{1}{v} \|p_i - q\| dq \\ &= \min_{p \in \mathcal{Q}^m} \int_{\mathcal{Q}} \frac{1}{A} \min_{i \in I_m} \frac{1}{v} \|p_i - q\| dq = \frac{\mathcal{H}_m^*(\mathcal{Q})}{v}. \end{aligned}$$

■

Remark 9 *Proposition 7 and Theorem 8 imply that, for a given λ and ρ , \bar{T}_{opt} belongs to $\Omega\left(\frac{1}{v}\sqrt{\frac{A}{m}}\right)$.*

Dubins Lower Bound

In the previous section, a lower bound is derived by approximating the Dubins distance with the Euclidean distance. It is natural to expect that such a lower bound is conservative. In this section, a new lower bound is derived by explicitly considering the turning cost associated with the motion of a Dubins vehicle. Consider the simple scenario of a single Dubins vehicle responsible for the environment \mathcal{Q} , in the case where targets are generated very infrequently. Since the vehicle cannot stop, a control policy must include a loitering path while waiting for targets to appear. This may be a predetermined pattern, or possibly a randomized or adaptive path. Towards a lower bound, consider a Dubins vehicle which is capable of waiting

in a static configuration (alternatively, consider that for any loitering path, there exists a moment in the pattern for which the expected Dubins distance to a target in \mathcal{Q} is minimal). If the UAVs were capable of stopping, this problem could be reduced to finding the optimal configuration of the m -vehicle system. Instead, it is required to design control policies that keep the configuration near optimal in a time-averaged sense.

Considering a single-vehicle system, towards a lower bound, one gets that

$$\bar{T}_{\text{opt}} \geq \inf_{g_i \in \text{SE}(2)} \int_{\mathcal{Q}} \frac{1}{Av} L_{\rho}(g_i, q) dq.$$

Let $t_{\rho}(A) = \{t \mid \text{Area}(R_t(g_i)) = A\}$. The function t_{ρ} is a single-valued function. This is because the area of a reachable set $R_t(g_i)$ is a strictly increasing function of t . For clarity, in the following theorems, let us specify a reachable region by its area, i.e., $R_A(g_i) = R_{t_{\rho}(A)}(g_i)$. Now, towards a lower bound, allow the environment to take on any shape, so long as its area remains unchanged. For a given configuration, g_i , the region S of area A with least expected Dubins distance from g_i , i.e., $\int_S \frac{1}{Av} L_{\rho}(g_i, q) dq$, is the reachable region of area A , i.e., $R_A(g_i)$.

Given a set $\mathcal{Q} \subset \mathbb{R}^2$ and a set of Dubins configurations $g = \{g_1, g_2, \dots, g_m\} \in \text{SE}(2)^m$, $\mathcal{DV}(g) = \{\mathcal{DV}_1(g), \dots, \mathcal{DV}_m(g)\}$ is defined 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 $L_{\rho}(g_i, q) \leq L_{\rho}(g_j, q)$, for all $i, j \in I_m$. Note that a vehicle's Dubins Voronoi region is *exactly* its instantaneous region of dominance, i.e., the set of points which can be reached by the vehicle faster than any other. A new lower bound on the Dubins DTRP can be obtained based on the idea of finding the optimal configuration for the m -vehicle system, as if they were capable of stopping:

Theorem 10 *For any convex, compact $\mathcal{Q} \subset \mathbb{R}^2$ of area A and $\lambda, \rho, v \in \mathbb{R}^+$ and $m \in \mathbb{N}$, the system time for the Dubins DTRP satisfies*

$$\bar{T}_{\text{opt}} \geq \frac{m}{Av} \int_{R_{\frac{A}{m}}(g_i)} L_{\rho}(g_i, q) dq. \quad (7)$$

Proof: In the following, the notation $A_i = \text{Area}(\mathcal{DV}_i(g))$ is used. Beginning with the

definition of \bar{T}_{opt} , one has that

$$\begin{aligned}
\bar{T}_{\text{opt}} &\geq \inf_{g \in \text{SE}(2)^m} \int_{\mathcal{Q}} \frac{1}{A} \min_{i \in I_m} \frac{1}{v} L_{\rho}(g_i, q) dq \\
&= \inf_{g \in \text{SE}(2)^m} \sum_i^m \int_{\mathcal{DV}_i(g)} \frac{1}{Av} L_{\rho}(g_i, q) dq \\
&\geq \inf_{g \in \text{SE}(2)^m} \sum_i^m \int_{R_{A_i}(g_i)} \frac{1}{Av} L_{\rho}(g_i, q) dq \\
&\geq \min_{\{A_1, A_2, \dots, A_m\} \in \mathbb{R}^m} \sum_i^m \frac{1}{Av} f(A_i), \quad \text{subject to} \quad \sum_i^m A_i = A \quad \text{and} \quad A_i \geq 0 \quad \forall i \in \{1, \dots, m\},
\end{aligned}$$

where $f : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is defined by

$$f(A) = \int_{R_A(g_i)} L_{\rho}(g_i, q) dq. \quad (8)$$

It is easy to verify that the function f is a continuous, strictly increasing function of A . Lemma 31 in Appendix B shows that f is convex. Thus, by using the Karush-Kuhn-Tucker conditions,⁵⁵ one can show that the quantity

$$\min_{\{A_1, A_2, \dots, A_m\} \in \mathbb{R}^m} \frac{1}{Av} \sum_i^m f(A_i) \quad \text{subject to} \quad \sum_i^m A_i = A$$

is minimized with an equitable partition, i.e., $A_i = A/m$, $\forall i$. This yields the stated result. ■

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

Let the *nonholonomic vehicle density* be a non-dimensional parameter defined as

$$d_{\rho} = \frac{\rho^2 m}{A}. \quad (9)$$

The motivation of this parameter is shown in the proof of the following lower bound, which holds for any policy, and is tightest for small vehicle workloads.

Theorem 11 *For any convex, compact $\mathcal{Q} \subset \mathbb{R}^2$ of area A , $\lambda, \rho, v \in \mathbb{R}^+$, and $m \in \mathbb{N}$, the*

system time for the Dubins DTRP satisfies

$$\liminf_{d_\rho \rightarrow +\infty} \bar{T}_{\text{opt}} \left(\frac{m}{\rho A} \right)^{1/3} \geq \frac{3\sqrt[3]{3}}{4v}.$$

Proof: Proposition 6 implies that, for $d_\rho \geq 24/\pi^3$,

$$t_\rho(A/m) = \frac{1}{v} \left(\frac{3\rho A}{m} \right)^{1/3}. \quad (10)$$

The Dubins distance to any point q in the reachable set $R_{\frac{A}{m}}(g_i)$, expressed as (x, y) in the Dubins frame, can be bounded by $L_\rho(g_i, q) \geq x$. Let $(x_1(\theta), y_1(\theta)) = (\rho \sin \theta, \rho(1 - \cos \theta))$ represent parametrized points on the minimum turning circle, as depicted in Figure 1. Let $\tilde{R}_{\frac{A}{m}}(g_i) := \left\{ q \in R_{\frac{A}{m}}(g_i) \mid x \leq x_1(vt_\rho(A/m)/\rho) \right\}$, i.e., $\tilde{R}_{\frac{A}{m}}(g_i)$ is the set of points in $R_{\frac{A}{m}}(g_i)$ which lie to the left of the vertical line $x = x_1(vt_\rho(A/m)/\rho)$. Then,

$$\int_{R_{\frac{A}{m}}(g_i)} L_\rho(g_i, q) dq \geq \int_{R_{\frac{A}{m}}(g_i)} x dq \geq \int_{\tilde{R}_{\frac{A}{m}}(g_i)} x dq.$$

The last integral in the equation above can be evaluated to get the following bound:

$$\begin{aligned} \int_{R_{\frac{A}{m}}(g_i)} L_\rho(g_i, q) dq &\geq 2 \int_0^{vt_\rho(A/m)/\rho} x_1(\theta) y_1(\theta) \frac{dx_1(\theta)}{d\theta} d\theta \\ &= 2\rho^3 \int_0^{vt_\rho(A/m)/\rho} \sin \theta \cos \theta (1 - \cos \theta) d\theta \\ &= \rho^3 \int_0^{vt_\rho(A/m)/\rho} (\theta^3 + o(\theta^3)) d\theta = \frac{v^4 t_\rho^4(A/m)}{4\rho} + \rho^3 o(v^4 t_\rho^4(A/m)/\rho^4). \end{aligned} \quad (11)$$

Note that, as $d_\rho \rightarrow +\infty$, $vt_\rho(A/m)/\rho$ goes to zero from Equation (5). Therefore, applying Equation (10) in Equation (11), one can deduce that, as $d_\rho \rightarrow +\infty$,

$$\int_{R_{\frac{A}{m}}(g_i)} L_\rho(g_i, q) dq \geq \frac{3A}{4m} \left(\frac{3\rho A}{m} \right)^{1/3}.$$

Substituting into Equation (7) yields the result. ■

Remark 12 *Theorem 11 shows that \bar{T}_{opt} belongs to $\Omega\left(\frac{1}{v} \left(\frac{\rho A}{m}\right)^{1/3}\right)$. In particular, for fixed ρ and A , \bar{T}_{opt} belongs to $\Omega\left(\frac{1}{v\sqrt[3]{m}}\right)$.*

The above result is simply the integral in Theorem 10 carried out for a particular asymp-

otic region of the problem parameter space. 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. This behavior is exhibited by the algorithms proposed in the next section.

III.B. Policies and Upper Bounds

In this section, control policies designed for the light load are proposed and then analyzed for their performance. The first policy—called 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 vehicle 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 vehicle as a circular trajectory of radius $c_2\rho$ centered at p_i^* . Each vehicle 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. An illustration of the MC policy is shown in Fig. 2.

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

Theorem 13 *For any convex, compact $\mathcal{Q} \subset \mathbb{R}^2$ of area A and $\rho, v \in \mathbb{R}^+$, the Median Circling policy is a stabilizing policy in the light load, i.e., as $\lambda/m \rightarrow 0^+$. Moreover, the system time of the Median Circling policy satisfies the following, as $\lambda/m \rightarrow 0^+$.*

$$\bar{T}_{\text{MC}} \leq \frac{\mathcal{H}_m^*(\mathcal{Q}) + c_1\rho}{v}. \quad (12)$$

Proof: For $\lambda/m \rightarrow 0^+$, after an initial transient, all targets will be generated with the vehicles in the assigned loitering station, and an empty target queue with probability one. Therefore, the waiting time for a new target after that is finite. Hence, the Median Circling




Figure 2. Illustration 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^* . In this figure, a target has appeared in the subregion roughly in the upper-right quarter of the domain. The UAV responsible for this subregion has left its loitering station and is en route to service the target.

policy is stabilizing. Moreover, the expected waiting time for a new target is given by

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

■

Remark 14 *In other words, it is shown that the system time achieved by the MC policy is within a constant additive factor $c_1\rho/v \approx 9.39\rho/v$ from the optimal. The additive factor, which can be considered a penalty due to the nonholonomic 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. A comparison is now made between the performance of the Median Circling Policy and the optimal limit shown in Theorem 8.

Theorem 15 *The system time of the Median Circling policy in light load satisfies*

$$\limsup_{\frac{\lambda}{m} \rightarrow 0^+} \frac{\bar{T}_{\text{MC}}}{\bar{T}_{\text{opt}}} \leq 1 + 25\sqrt{d_\rho}.$$

In particular,

$$\lim_{d_\rho \rightarrow 0^+} \limsup_{\frac{\lambda}{m} \rightarrow 0^+} \frac{\bar{T}_{\text{MC}}}{\bar{T}_{\text{opt}}} = 1.$$

Proof: Since $\frac{2}{3}\sqrt{\frac{A}{\pi m}} \leq \mathcal{H}_m^*(\mathcal{Q})$ from Proposition 7,

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

The first statement in the theorem then follows by substituting this into Equation (12) and applying Theorem 8. The second result follows by taking the limit as $d_\rho \rightarrow 0^+$ in the first statement. ■

Remark 16 *Theorem 15 implies that the Median Circling Policy is an efficient policy for scenarios where the nonholonomic vehicle density is low. In low density scenarios, Euclidean distance dominates the cost, and thus, a policy imitating that of Euclidean vehicles is near optimal.*

Simulation Results for the MC policy

Simulations of the MC policy were run for various values of ρ and m . In the simulations, here and throughout the paper, speed and turning radius were set to $v = 50\text{ms}^{-1}$, and $\rho = 600$ m, respectively; these values are typical of current operational UAVs.⁵⁶

For each combination of conditions, the system time shown in Fig. 3 is the average of the waiting times of 10^5 target points uniformly generated in the environment. These numerical results confirm theoretical predictions that the system time decreases as $1/\sqrt{m}$ and grows with ρ : the log-log plots of system time versus number of vehicles (left) have slopes ranging from -0.48 to -0.52 , implying a power law of approximately $-1/2$, and the linear-linear plots of system time versus minimum turning radius (right) have slopes ranging from 0.96 to 1.04 , implying a linear dependence on ρ with a coefficient of approximately unity. 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 $\rho \rightarrow 0^+$, $\bar{T}_{\text{MC}} \rightarrow \mathcal{H}^*(\mathcal{Q})/v$.

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




Figure 3. Simulations of the MC policy were run with a square environment of area $A = 9 \times 10^{10} \text{ m}^2$ and vehicle speed $v = 50 \text{ ms}^{-1}$. The performance of the MC policy is plotted as a function of i) the number of vehicles (left), and ii) the minimum turning radius of the vehicles (right). The nonholonomic vehicle densities for the scenarios of all data points shown lie within the range $0 \leq d_\rho \leq 0.0056$.

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 height denotes 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

$$w = \min \left\{ \left(\frac{4}{3\sqrt{\rho}} \frac{WH + 10.38\rho H}{m} \right)^{2/3}, 2\rho \right\}. \quad (13)$$

Orient the strips along the side of length W . Construct a closed Dubins path which runs along the longitudinal bisector of each strip, visiting all strips in top-to-bottom sequence, making U-turns between strips at the edges of \mathcal{Q} , and finally returning to the initial configuration. Note that the path must cover \mathcal{Q} , but it need not cover the entire bounding box of \mathcal{Q} . The bounding box is merely a construction used to place an upper bound on the total path length. The m UAVs loiter on this path, equally spaced, in terms of path length. An illustration of the Strip Loitering policy is shown in Figure 4. Two distances that are important in the analysis of this policy are shown in Figure 5. At the instant a target arrives, a circle of radius ρ is constructed, tangent to the loitering path and intersects the target. The length of the arc departing from the loitering path and ending at the target is denoted by d_2 . The UAV responsible for visiting the target is the one closest in terms of loitering path length 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, 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. The following method to accomplish this task is described using the example shown in Fig. 5. 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 by a distance $4(d_2 - \rho \sin \frac{d_2}{\rho})$. To rectify this, as it nears the end of the current strip, it takes its U-turn a distance $2(d_2 - \rho \sin \frac{d_2}{\rho})$ early.




Figure 4. Illustration 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.




Figure 5. 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 17 For any convex, compact $Q \subset \mathbb{R}^2$ of area A and $\rho, v \in \mathbb{R}_+$, the Strip Loitering policy is a stabilizing policy in the light load case, i.e., when $\lambda/m \rightarrow 0^+$. Moreover, if Q has a bounding rectangle of width W and height H , the system time of the Strip Loitering Policy

satisfies the following, as $\lambda/m \rightarrow 0^+$.

$$\bar{T}_{\text{SL}} \leq \begin{cases} \frac{1.238}{v} \left(\frac{\rho W H + 10.38 \rho^2 H}{m} \right)^{1/3} + \frac{W + H + 6.19 \rho}{mv} & \text{for } m \geq 0.471 \left(\frac{W H}{\rho^2} + \frac{10.38 H}{\rho} \right), \\ \frac{W H + 10.38 \rho H}{4 \rho m v} + \frac{W + H + 6.19 \rho}{mv} + \frac{1.06 \rho}{v} & \text{otherwise.} \end{cases} \quad (14)$$

Furthermore, for any $\lambda, \rho, v \in \mathbb{R}_+$,

$$\limsup_{m \rightarrow +\infty} \bar{T}_{\text{SL}} m^{1/3} \leq \frac{1.238}{v} (\rho W H + 10.38 \rho^2 H)^{1/3}. \quad (15)$$

Proof: For $\lambda/m \rightarrow 0^+$, after an initial transient, all targets are generated with i) the vehicles at their designated position on the loitering pattern, and ii) an empty target queue, with probability one. Therefore, the waiting time for a new target after that is finite. Hence, the Strip Loitering policy is stabilizing in light load. In the following, an expression for the system time is derived. Denote the length of the closed path as L_1 . Due to the equal spacing of the UAVs along the loitering path,

$$\mathbb{E}[d_1] = \frac{L_1}{2m}. \quad (16)$$

Let N_{strips} be the number of strips, L_{strip} be the length traveled along a single strip, $L_{\text{u-turn}}$ be the length of a u-turn and L_{closure} be the length of the closure path segment. With these notations, L_1 can be bounded as

$$L_1 \leq N_{\text{strips}} L_{\text{strip}} + (N_{\text{strips}} - 1) L_{\text{u-turn}} + L_{\text{closure}}. \quad (17)$$

It is easy to verify that $N_{\text{strips}} = \lceil \frac{H}{w} \rceil \leq \frac{H}{w} + 1$ and $L_{\text{strip}} \leq W + 2\rho$, where the 2ρ term arises from the fact that the vehicles must start on a strip outside \mathcal{Q} in order to visit a target that might appear on the boundary of \mathcal{Q} in front of it. Bounds on the length of the optimal point-to-point Dubins path from¹⁴ 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 $\kappa \approx 8/3$. Substituting these bounds into Eq. (17) and taking into account Eq. (16),

$$\mathbb{E}[d_1] \leq \frac{W H + 10.38 \rho H}{2 m w} + \frac{W + H + 6.19 \rho}{m}. \quad (18)$$

To calculate $\mathbb{E}[d_2]$, define δ as the smallest distance from the target to any point on the loitering path (see Fig. 5). 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 $w/2$,

$$\mathbb{E}[d_2] = \frac{4\rho}{w} \int_0^{w/2} \sin^{-1} \left(\sqrt{\frac{s}{2\rho}} \right) ds = \frac{4\rho^2}{w} \left[\sqrt{x - x^2} + (-1 + 2x) \sin^{-1}(\sqrt{x}) \right], \quad (19)$$

where $x = w/(4\rho)$. It can be shown that, for all $x \in [0, 1/2]$,

$$\sqrt{1-x} \leq 1 - \frac{x}{2} \quad \text{and} \quad \sin^{-1}(\sqrt{x}) \geq \sqrt{x} + \frac{\sqrt{x^3}}{6}.$$

Substituting these bounds into Eq. (19) and through further algebraic simplification, for $w \leq 2\rho$,

$$\mathbb{E}[d_2] \leq \frac{3}{4}\sqrt{\rho w}. \quad (20)$$

As noted earlier, for $\lambda/m \rightarrow 0^+$, after an initial transient, all targets are generated with the vehicles on the closed loitering path, and an empty target queue. The system time satisfies

$$\bar{T}_{\text{SL}} \leq \frac{\mathbb{E}[d_1] + \mathbb{E}[d_2]}{v}. \quad (21)$$

Therefore, from Equations (21), (18) and (20), for $w \leq 2\rho$, in the limit as $\lambda/m \rightarrow 0^+$,

$$\bar{T}_{\text{SL}} \leq \frac{WH + 10.38\rho H}{2mvw} + \frac{W + H + 6.19\rho}{mv} + \frac{3}{4v}\sqrt{\rho w}. \quad (22)$$

In order to get the least upper bound, the right hand side of Eq. (22) is minimized with respect to w subject to the constraint that $w \leq 2\rho$. The value of w at which \bar{T}_{SL} attains its minimum value is given by Eq. (13) and the minimum value is given by Eq. (14). For a given $\lambda \in \mathbb{R}_+$, as $m \rightarrow +\infty$, $\lambda/m \rightarrow 0^+$. The result in Eq. (15) then follows by taking the limit as $m \rightarrow +\infty$ in Eq. (14). ■

Remark 18 *Theorem 11 and Theorem 17 imply that, in the light load case and for a fixed \mathcal{Q} and ρ , \bar{T}_{opt} belongs to $\Theta(1/(v\sqrt[3]{m}))$.*

Remark 19 *In the light load case, when the turning radius ρ is large compared to the dimensions of the region \mathcal{Q} , Eq. (14) implies that \bar{T}_{opt} belongs to $O(1/(mv))$.*

Simulation Results for the SL policy

Simulations of the SL policy were run for various values of m . For each condition, the system time shown in Fig. 6 is the average of the waiting times of one-hundred thousand target points uniformly generated in the environment. These numerical results 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 $-1/3$.

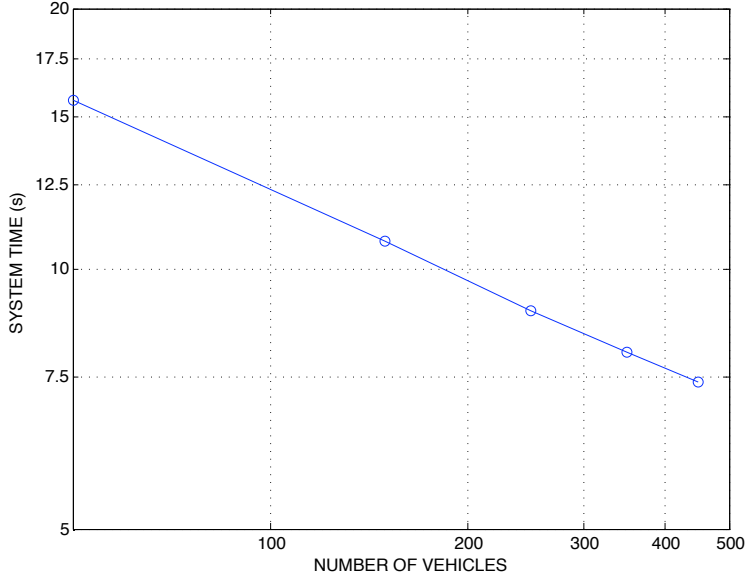


Figure 6. Simulations of the SL policy were run with a square environment of area $A = 9 \times 10^6 \text{ m}^2$, vehicle speed $v = 50 \text{ ms}^{-1}$, minimum-turning radius $\rho = 600 \text{ m}$, and m varying from 50 to 450. The nonholonomic vehicle densities for the scenarios of all data points shown lie within the range $2 \leq d_\rho \leq 18$.

IV. Heavy Load

In this section, performance bounds are studied and algorithms are presented for the case in which the target generation rate is high.

IV.A. Lower Bound

Leveraging the results on the small-time reachable set derived in previous sections, a lower bound on the system time for any policy in the heavy load case can be stated. This lower bound explicitly takes the nonholonomic constraint into account.

Theorem 20 *For any convex, compact $Q \subset \mathbb{R}^2$ of area A and $\rho, v \in \mathbb{R}^+$, the system time for the Dubins DTRP satisfies*

$$\liminf_{\frac{\lambda}{m} \rightarrow +\infty} \bar{T}_{\text{opt}} \frac{m^3}{\lambda^2} \geq \frac{81 \rho A}{64 v^3}.$$

Proof: Under the actions of any stabilizing policy, the number of outstanding targets approaches a finite steady-state value, \bar{n} , related to the system time by Little's Formula (1). Consider the m -vehicle system in arbitrary configuration. In the following, a bound is computed on the expected time required for a given vehicle i to visit the closest target, in the Dubins distance sense. Call this target the nearest-neighbor, and call such a time t_{nn} . If the vehicle is given a time t to move, the probability that $t_{\text{nn}} > t$ is no less than the probability

that there are no targets reachable within time t . In other words,

$$\Pr(t_{\text{nn}} > t) \geq 1 - \bar{n} \frac{\text{Area}(R_t(g_i))}{A} = 1 - \bar{n} \frac{(vt)^3}{3\rho A}.$$

Defining $c = (\bar{n}v^3)/(3\rho A)$,

$$\begin{aligned} \mathbb{E}[t_{\text{nn}}] &= \int_0^\infty \Pr(t_{\text{nn}} > \xi) d\xi \\ &\geq \int_0^\infty \max\left(0, 1 - \bar{n} \frac{(v\xi)^3}{3\rho A}\right) d\xi \\ &= \int_0^{c^{-1/3}} (1 - c\xi^3) d\xi \\ &= \frac{3}{4v} \left(\frac{3\rho A}{\bar{n}}\right)^{1/3}. \end{aligned}$$

If each vehicle visits its nearest-neighbor target, then in expectation, m targets are visited within time t_{nn} . Assuming, towards a lower bound, that targets are continually serviced at a rate of m/t_{nn} , resulting in a steady-state number of outstanding targets \bar{n}_{nn} ,

$$\frac{m}{t_{\text{nn}}} = \frac{4vm}{3} \left(\frac{\bar{n}_{\text{nn}}}{3\rho A}\right)^{1/3} = \lambda.$$

Substituting Little's formula, i.e., $\bar{n}_{\text{nn}} = \lambda \bar{T}_{\text{nn}}$, and rearranging, the claim follows from $\bar{T}_{\text{opt}} \geq \bar{T}_{\text{nn}}$. ■

Note that the system time depends quadratically on the parameter λ , whereas in the Euclidean case it depends on it linearly. As a consequence, bounded-curvature constraints make the system time more sensitive to increases in the target-generation rate. Perhaps the most striking consequence of the above result is that the lower bound suggests that the system time decreases with the cube of the number of vehicles, whereas in the Euclidean case it decreases with the square of the number of vehicles.

IV.B. The Bead Tiling (BT) Policy

In Savla *et al.*,¹⁴ the single-vehicle Dubins DTRP heavy-load case is analyzed. A policy called the Bead Tiling Algorithm (BT) Policy is shown to have a performance within a constant factor of the optimal for rectangular environments. In the following, the BT policy is adapted for general convex environments, and a method by which the algorithm can be used by a multiple-vehicle system is presented. For the sake of completeness, a modified form of the BT policy from¹⁴ is stated for general convex \mathcal{Q} , and the corresponding upper bound on the system time for the Dubins DTRP is obtained from its analysis.

The BT policy makes use of the shape detailed in Fig. 7, referred to as a “bead”.¹⁴ The algorithm works by tiling the plane with identical beads. In choosing the orientation of the bead rows, we use the same minimum-height bounding box of \mathcal{Q} with width W and height H , as described in the SL policy. The Dubins vehicle visits all beads intersecting \mathcal{Q} in a row-by-row fashion in top-to-bottom sequence, servicing at least one target in every non-empty bead. This process is repeated indefinitely. The size of the beads (and hence the total number intersecting \mathcal{Q}) is scaled with the target generation rate by choosing $\ell = \min\{C_{\text{BTA}}v/\lambda, 4\rho\}$, where

$$C_{\text{BTA}} = \frac{7 - \sqrt{17}}{4} \left(1 + \frac{7\pi\rho H}{3A}\right)^{-1}$$

is a design parameter chosen to maximize the performance of the algorithm. Let \bar{T}_{BT} be the




Figure 7. The bead.

system time for the single-vehicle BT policy. The following theorem is a generalization of the corresponding result from Savla *et al.*¹⁴ for any convex \mathcal{Q} .

Theorem 21 *For any convex, compact $\mathcal{Q} \subset \mathbb{R}^2$ of area A and $\lambda, \rho, v \in \mathbb{R}^+$, the single-vehicle Bead Tiling policy is a stabilizing policy. Moreover, if \mathcal{Q} has a bounding rectangle of width W and height H , the system time for the single-vehicle Bead Tiling policy satisfies*

$$\lim_{\lambda \rightarrow +\infty} \frac{\bar{T}_{\text{BT}}}{\lambda^2} \leq 71 \frac{\rho A}{v^3} \left(1 + \frac{7\pi\rho H}{3A}\right)^3.$$

The BT policy is extended to the m -vehicle Bead Tiling (mBT) policy in the following way. Divide the environment into regions of dominance with lines parallel to the bead rows. Let the area and height of the i -th vehicle’s region be denoted with A_i and H_i . Place the

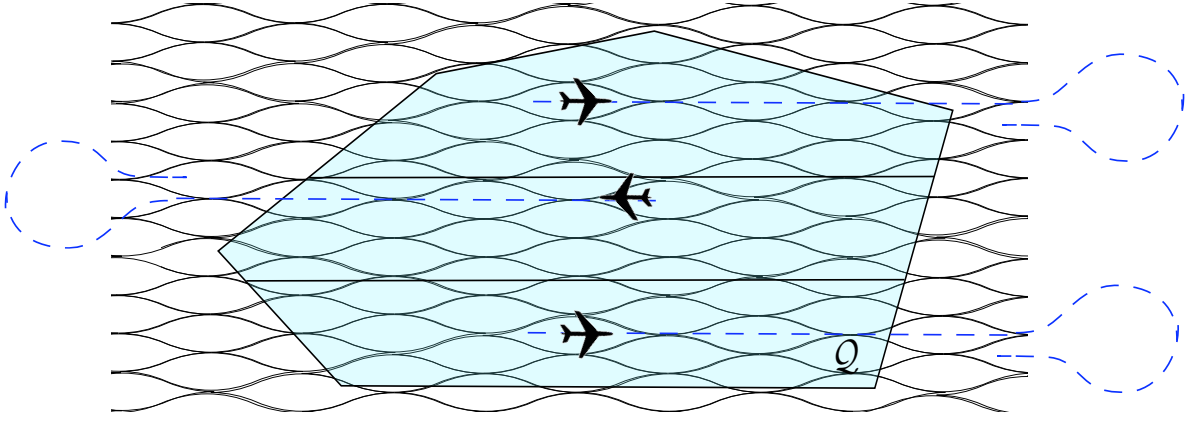


Figure 8. An adaptation of the BT policy for the multiple-vehicle case.

subregion dividers in such a way that

$$A_i + \frac{7}{3}\pi\rho H_i = \frac{1}{m} \left(A + \frac{7}{3}\pi\rho H \right)$$

for all i in I_m . Allocate one subregion to every vehicle and let each vehicle execute the BT policy in its own region. Let \bar{T}_{mBT} be the system time for the mBT policy.

Theorem 22 *For any convex, compact $\mathcal{Q} \subset \mathbb{R}^2$ of area A and $\lambda, \rho, v \in \mathbb{R}^+$ and $m \in \mathbb{N}$, the mBT policy is a stabilizing policy. Moreover, if \mathcal{Q} has a bounding rectangle of width W and height H , the system time for the mBT policy satisfies*

$$\lim_{\frac{\lambda}{m} \rightarrow +\infty} \bar{T}_{\text{mBT}} \frac{m^3}{\lambda^2} \leq 71 \frac{\rho A}{v^3} \left(1 + \frac{7\pi\rho H}{3A} \right)^3.$$

Proof: The fact that the mBT policy is stabilizing follows from the stabilizing property of the single-vehicle BT policy. Denote the target-generation rate and system time for targets generated in the i -th vehicle's region as λ_i and \bar{T}_i , respectively. Also, let A_i be the area of the i -th vehicle's region, in which case, $\lambda_i = \lambda A_i/A$. Given the definition of the mBT policy, $\lambda/m \rightarrow +\infty$ implies that $\lambda_i \rightarrow +\infty$ for all $i \in I_m$. By Theorem 21, for all $i \in I_m$,

$$\bar{T}_i \leq \frac{71\rho A_i}{v^3} \left(1 + \frac{7\pi\rho H_i}{3A_i} \right)^3 \lambda_i^2, \quad \text{as } \lambda/m \rightarrow +\infty.$$

Substituting $\lambda_i = \lambda A_i/A$, for all $i \in I_m$,

$$\bar{T}_i \leq \frac{71\rho}{A^2 v^3} \left(A_i + \frac{7\pi\rho H_i}{3} \right)^3 \lambda^2, \quad \text{as } \lambda/m \rightarrow +\infty. \quad (23)$$

However, A_i and H_i are selected in the mBT policy so that $A_i + \frac{7\pi\rho H_i}{3} = \frac{1}{m} \left(A + \frac{7\pi\rho H}{3} \right)$.

Substituting this into Eq. (23), for all $i \in I_m$,

$$\bar{T}_i \leq \frac{71\rho A}{m^3 v^3} \left(1 + \frac{7\pi\rho H}{3A}\right)^3 \lambda^2, \quad \text{as } \lambda/m \rightarrow +\infty.$$

Note that \bar{T}_i is independent of i , i.e., the system time in each of the vehicle's regions is the same. The theorem then follows trivially. ■

Remark 23 *Theorem 20 and Theorem 22 imply that the system time for the Dubins DTRP, in the heavy load, is of the order $\lambda^2/(mv)^3$, and that the mBT policy performs within a constant factor of the optimal policy.*

Remark 24 *There exists no stabilizing policy for the Dubins DTRP when the targets are generated in an adversarial worst-case fashion with $\lambda/m \geq v/(\pi\rho)$. This fact is a consequence of the linear lower bound on the length of the Dubins Traveling Salesperson Tour for worst-case point sets as derived in Savla et al.¹⁴*

Simulation results for the m-vehicle Bead Tiling Policy

Simulations of the mBT policy were run for various values of λ and m . For each combination of conditions, an experiment was run for $5 \cdot 10^5$ simulation seconds. The number of outstanding targets was plotted with respect to time to ensure that the data collected represents the performance of the policy once the system has reached steady state. The system times shown in Fig. 9 are the averages of the waiting times of all targets generated and serviced during this period. Moreover, the results are verified for consistency by substituting the known target-generation rate and the measured steady-state number of outstanding targets into Eq. (1) to obtain the system time via an alternative method. Performance results from simulations of the mBT policy, summarized in Fig. 9, validate theoretical predictions that the system time grows quadratically with the target generation rate, and decreases cubically with the number of vehicles. For example, for $m = 1$ and λ increasing from $1/12 \text{ s}^{-1}$ to $1/4 \text{ s}^{-1}$ (Fig. 9, left), note a slope of 1.95 on the log-log plot, implying a power law close to 2. For $\lambda = 1/4 \text{ s}^{-1}$ and m increasing from 1 to 4 (Fig. 9, right), note a slope of -2.91 on the log-log plot, implying a power law close to -3 .

V. Phase Transition

In Section III, two policies were proposed: the Median Circling policy and the Strip Loitering policy. These policies exhibit different modes of behavior. It was shown that the territorial MC policy is optimal as $d_\rho \rightarrow 0^+$ and the gregarious SL policy is constant-factor

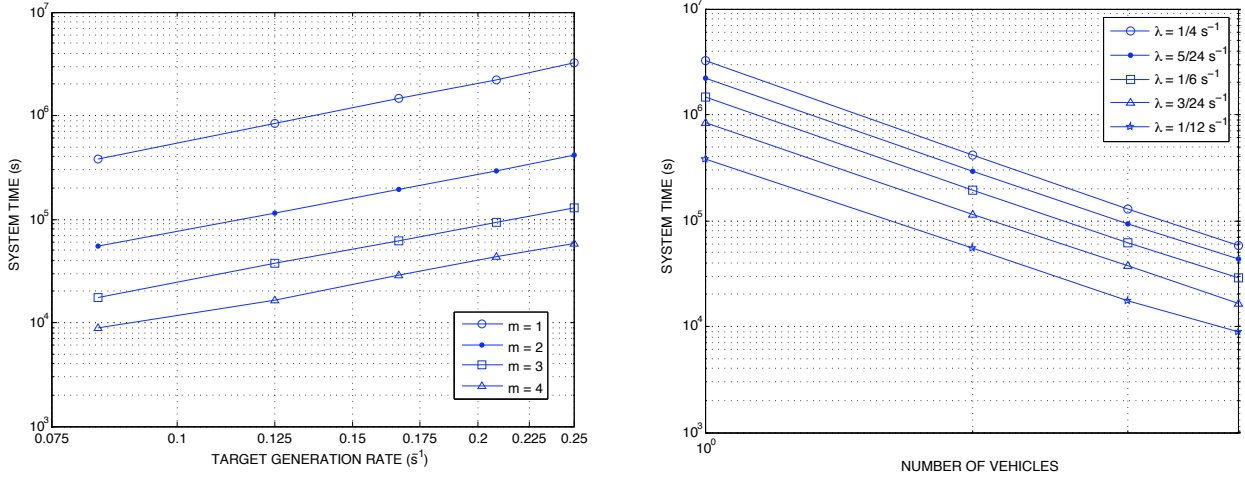


Figure 9. Simulations of the mBT policy were run with a square environment of area $A = 3.6 \times 10^7 \text{ m}^2$, vehicle speed $v = 50 \text{ ms}^{-1}$, minimum-turning radius $\rho = 600 \text{ m}$, number of vehicles m varying from 1 to 4, and target-generation rate λ varying from $1/12 \text{ s}^{-1}$ to $1/4 \text{ s}^{-1}$ in increments of $1/24 \text{ s}^{-1}$. Plots of the system time versus target-generation rate (left), and system time versus the number of vehicles (right) are given.

optimal as $m \rightarrow +\infty$, for constant ρ and Q . This suggests the existence of a phase transition in the optimal policy for the light load scenario as one increases the number of vehicles for a fixed ρ and Q .

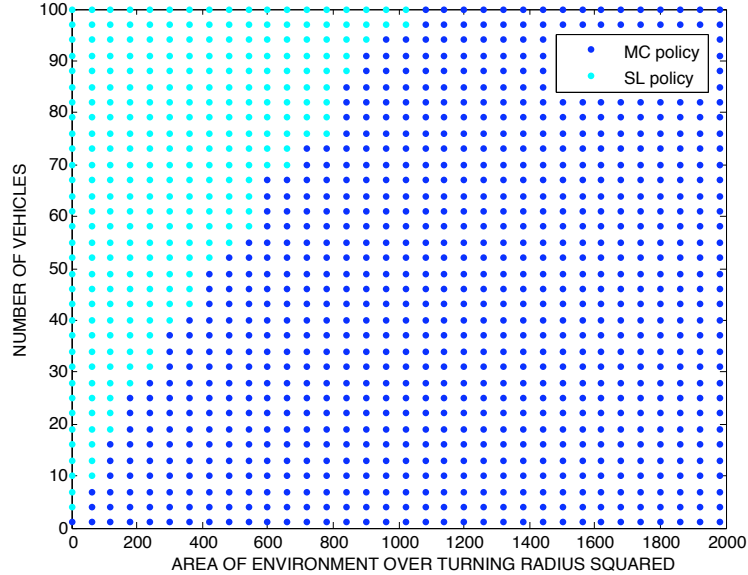


Figure 10. Simulation results aimed at demarcating the phase transition for the case when the environment is a square of area $A = 4 \times 10^8 \text{ m}^2$. The color of the data point signifies which policy had a lower system time for the corresponding values of the parameters.

Simulations of both the MC and SL policy were run with a square environment of area $A = 4 \times 10^8 \text{ m}^2$ for varying values of m and ρ . For each combination of conditions and for each policy, the system time was taken as the average of the waiting times of one-hundred thousand target points uniformly generated in the environment. For the MC policy, $c_2 = 2.91$ is chosen since numerical results show that it gives better performance in terms of c_1 , as

reported in Remark 2. The number of vehicles m is varied between 1 and 100 and the turning radius ρ is varied such that the ratio A/ρ^2 takes values between 1 and about 2000. The results are summarized in Fig. 10 showing which policy performs better for a given value of m and ρ . The color of the data point signifies which policy had a lower system time for the given conditions. The plot illustrates that the transition curve between the two policies could be captured by a straight line of the form $m = kA/\rho^2$ for some constant $k > 0$, thereby highlighting the significance of the nonholonomic density, $d_\rho = \rho^2 m/A$ in capturing the transition. Moreover, the slope of the approximating line gives an estimate of this critical value of the nonholonomic density to be $d_\rho^{\text{sq-sim}} \approx 0.0925$.




Figure 11. The MC (left) and SL (right) policies with an infinite number of vehicles on an unbounded domain.

It is desirable to study the fundamental factors driving the transition, ignoring its dependence on the shape of \mathcal{Q} . Towards this end, envision an infinite number of vehicles operating on the unbounded plane, where the system is characterized by the vehicle density. Denote the inverse of the density with a , the area per vehicle. Illustrations of the MC and SL policies operating on an unbounded domain are shown in Fig. 11. In this case, the configuration p^* yielding the minimum of the function $\mathcal{H}(p)$ is that in which the Voronoi partition induced by p^* is a network of regular hexagons,⁴⁹ each of area a . Each generator p_i is the median of its own Voronoi region, and it is known that if \mathcal{Q} is a regular hexagon of area a then $\mathcal{H}_1^*(\mathcal{Q}) \approx 0.377\sqrt{a}$. The system time of the policy satisfies $\bar{T}_{\text{MC}} \leq 0.377\sqrt{a} + 3.76\rho$. This expression is obtained by using the value of c_1 from Remark 2 in Theorem 13. 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 $\bar{T}_{\text{SL}} \leq L/2 + 3\sqrt{\rho w}/4$. The area per vehicle is related to the design parameters by $Lw = a$. Substituting gives $\bar{T}_{\text{SL}} \leq a/2w + 3\sqrt{\rho w}/4$. Minimizing with respect to w , with the constraint that $w \leq 2\rho$,

gives $w = \min\{(4a/3\sqrt{\rho})^{2/3}, 2\rho\}$, yielding

$$\bar{T}_{\text{SL}} \leq \begin{cases} 1.238(\rho a)^{1/3} & \text{for } d_\rho \geq 0.471, \\ a/4\rho + 1.06\rho & \text{otherwise,} \end{cases} \quad (24)$$

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 $d_\rho^{\text{unbd-theory}} = 1/x^2 \approx 0.0587$. This is a lower critical density than implied by the simulation results on a square environment, 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.

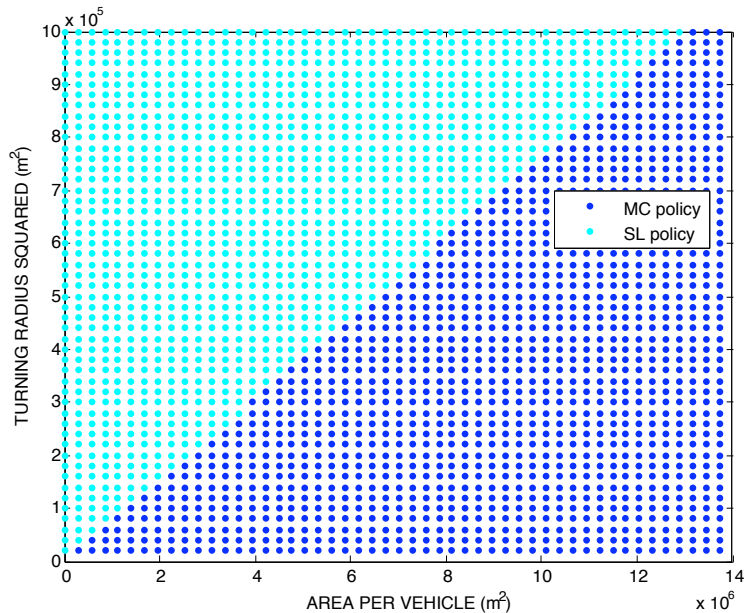


Figure 12. Simulation results aimed at demarcating the phase transition for the case when the environment is an unbounded domain. The color of the data point signifies which policy had a lower system time for the corresponding values of the parameters.

To further verify the above result, extensive simulations of the two policies operating in an unbounded domain were run for various values of a and ρ^2 , and Fig. 12 shows which policy has the lower system time for the given set of conditions. For the MC policy, this reduces to simulating a single vehicle operating on a regular hexagon. In the simulations of the MC policy, trials were run varying the loitering radius (by varying c_2) for a given set of a and ρ^2 , and the MC policy's system time was taken as the lowest value obtained. Avoiding excessive detail, it should be noted that the optimal c_2 is also a function of d_ρ and that for all conditions under which the MC policy is dominant, the optimal c_2 is one.

The following is an intuitive explanation for this fact: the MC policy is dominant when the Euclidean cost is dominant and a choice of one for c_2 is the closest the Dubins vehicle can come to replicating the optimal policy of the Euclidean vehicle (waiting at the median). For the SL policy, the unbounded domain scenario reduces to simulating a single vehicle responsible for targets appearing in the region of the strip in front of it, as dictated by the chosen w and L discussed above. Trials were run varying w (and hence $L = a/w$) as a fraction of 2ρ . Experiments suggest that the optimal w is 2ρ for all conditions near the phase transition. This agrees with the theoretical analysis above regarding the unbounded domain case. Also, once a target appears, the vehicle is allowed to take a minimum-time path rather than the service method given in the formal description of the SL policy. This practice lends itself to a fairer comparison between the two policies. For each set of conditions, for each policy, and for each trial varying design parameters (c_2 for the MC policy and w, L for the SL policy), the system time was taken as the average of the waiting times of one-hundred thousand target points uniformly generated in the vehicle’s region of responsibility. Figure 12 provides a stronger justification for approximating the phase transition with a line, and hence the nonholonomic density. The slope of the transition line in Fig. 12 implies a critical nonholonomic density of $d_\rho^{\text{unbd-sim}} \approx 0.0759$. This is within a factor of 1.3 from $d_\rho^{\text{unbd-theory}}$, further favoring the MC policy. This discrepancy is ascribed to the conservative use of a strict upper bound on $L_\rho(g, q)$ in Theorem 13 rather than taking its expected value over a given domain, as is effectively done through simulation. To gain further intuition on the nature of this critical condition, consider the area per vehicle $a^{\text{crit}} = \rho^2/d_\rho^{\text{crit}}$, which yields $a^{\text{unbd-theory}} \approx 5.42\pi\rho^2$ and $a^{\text{unbd-sim}} \approx 4.19\pi\rho^2$. In other words, the transition occurs when each vehicle is responsible for a region of area 5.42 or 4.19 times that of a minimum turning-radius disk. Note that the value of $d_\rho^{\text{unbd-sim}}$ is within a factor of 0.821 from $d_\rho^{\text{sq-sim}}$. The relative closeness between the values of d_ρ obtained via different methods gives a compelling guideline in the form of d_ρ and its critical value that can be used by a system architect to decide upon the optimal strategy in the light load scenario for given problem parameters.

VI. Conclusions

In this paper, several routing problems in stochastic time-varying environments were introduced, for vehicles that move with constant forward speed along paths of bounded curvature in the plane. For each of them we derived fundamental limits on the achievable performance in terms of system time, as a function of problem parameters such as the size of the fleet, the minimum turning-radius, the target-generation rate, the vehicle speed and the size of the environment. In addition, for each problem, we presented routing algorithms that provably approximate the optimal performance.

Specifically, it has been shown that, in light load, the optimal system time is proportional to the inverse cube root of the number of vehicles, contrasting with the holonomic case where it is proportional to the inverse square root. This implies that the nonholonomic motion constraints inhibit the added performance gained per vehicle. 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. In heavy load, it was shown that the system time is of order proportional to the square of the target appearance rate, and inversely proportional to the cube of the number of vehicles, whereas it is known to be proportional to the target appearance rate and inversely proportional to the square of the number of vehicles in the holonomic case. This implies that the system time in the nonholonomic case is more sensitive to the target generation rate as well as the size of the vehicle fleet and vehicle speed than the holonomic case. A phase transition was recognized in the optimal algorithm and numerical experiments were performed towards demarcating the parameter values at which this phase transitions occur. This led to a set of guidelines available to system architects and control algorithm designers aiming to maximize the performance of their autonomous-vehicle systems. The routing algorithm can be chosen according to the size of the vehicle team. On other hand, depending on the conditions of the operating environment, a designer might use the analysis provided to decide if the additional cost of another vehicle is worth the added value in performance it will bring.

In the future, various extensions of the problem considered in this paper will be studied. For example, nonholonomic dynamics will be combined with constraints on the information available to the vehicles by limiting communication and sensing capabilities. Another important extension is the inclusion of obstacles or no-fly zones restricting the motion of the UAVs. The focus of this paper was on the analysis of fundamental performance bounds, and on the design of algorithms achieving provable approximations of the optimal performance. Avenues for future work include further work on efficient implementations of the proposed algorithms, including for example the design of decentralized strategies for the individual vehicles which give rise to a group behavior that is near optimal for the system as a whole. A game-theoretic approach might lend itself naturally to such a completely distributed control paradigm. The understanding of phase transitions gained from this work can aid in studying similar phenomena in biological systems. Finally, experimental validation of the proposed algorithms on fixed-wing UAV platform is currently being pursued using commercial off-the-shelf autopilots.

A. The Dubins distance function

In this section, a brief exposition about the Dubins distance function $L_\rho(g, q)$ is given. Since the Dubins distance is invariant under transformation of coordinates from the global frame to the Dubins frame, without loss of generality, attention can be focused on $L_\rho(g_0, q)$, where $g_0 = (0, 0, 0)$ and q is the coordinate of the destination point in the Dubins frame. Define $\mathcal{D}_\rho^+ := \{q \in \mathbb{R}^2 : \|q - (0, \rho)\| < \rho\}$, $\mathcal{D}_\rho^- := \{q \in \mathbb{R}^2 : \|q - (0, -\rho)\| < \rho\}$. If $q \in \mathcal{D}_\rho^+ \cup \mathcal{D}_\rho^-$ the optimal path is of type CC , otherwise it is of type CL . The following proposition gives a closed form expression for $L_\rho(g_0, q)$.

Proposition 25 *The minimum length $L_\rho(g_0, q)$ of a Dubins path steering a vehicle from $g_0 = (0, 0, 0) \in \text{SE}(2)$ to a point $q = (x, y) \in \mathbb{R}^2$ is given by:*

$$L_\rho(g_0, q) = \begin{cases} \sqrt{d_c^2(q) - \rho^2} + \rho \left(\theta_c(q) - \arccos \frac{\rho}{d_c} \right), & \text{for } q \notin \mathcal{D}_\rho^+ \cup \mathcal{D}_\rho^-; \\ \rho \left(2\pi - \alpha(q) + \arcsin \frac{d_c(q) \sin \theta_c(q)}{d_f(q)} + \arcsin \frac{\rho \sin \alpha(q)}{d_f(q)} \right), & \text{otherwise,} \end{cases}$$

where $d_c(q) = \sqrt{x^2 + (|y| - \rho)^2}$ and $\theta_c(q) = \text{atan2}(x, \rho - |y|)$ are polar coordinates of the point q with respect to the center of either \mathcal{D}_ρ^+ or \mathcal{D}_ρ^- , whichever is the closest, $d_f(q) = \sqrt{x^2 + (|y| + \rho)^2}$ is the distance of q from the other center, and

$$\alpha(q) = \arccos \frac{5\rho^2 - d_f(q)^2}{4\rho^2}.$$

Proof: The proof is based on the results in Thomaschewski⁴⁶ mentioned in the paper, and elementary planar geometry. ■

The facts stated in the following remarks are consequences of Proposition 25.

Remark 26 *The level sets of the function $L_\rho(g_0, q)$ are segments of well-studied curves. More precisely, level sets of $L_\rho(g_0, q)$ are:*

- Arcs of circle involutes, for $q \notin \mathcal{D}_\rho^+ \cup \mathcal{D}_\rho^-$;
- Arcs of epicycloids, for $q \in \mathcal{D}_\rho^+ \cup \mathcal{D}_\rho^-$.

Recall that a circle involute is the curve traced by a point fixed to a line, as the line rolls without slipping on a circle; an epicycloid is the curve traced by a point fixed to a circle, as the circle rolls without slipping on another circle. For further details on these families of curves see, e.g., Lawrence.⁵⁷ A depiction of such level sets is provided in Fig. 13.

Remark 27 *The function $L_\rho(g_0, q)$ is:*




Figure 13. The level sets of the function $L_\rho(g_0, q)$, for $\rho = 1$, shown at increments of $\pi/4$. The initial vehicle configuration is represented by the black triangle, and is at coordinates $(0, 0)$ heading *up*. Note that the positive direction of the x axis is *up* and of the y axis is *left*.

- *not continuous on the top half of the boundary of \mathcal{D}_ρ^+ and \mathcal{D}_ρ^- , i.e., for $\{(x, y) \in \mathcal{D}_\rho^+ \cup \mathcal{D}_\rho^- : x \geq 0\}$;*
- *continuous, but not differentiable on the bottom half of the boundary of the same set, i.e., for $\{(x, y) \in \mathcal{D}_\rho^+ \cup \mathcal{D}_\rho^- : x < 0\}$, and on the negative part of the x -axis, i.e., for $\{(x, y) : x < 0, y = 0\}$;*
- *continuous and differentiable everywhere else.*

Remark 28 *The length of the optimal path to reach a point inside one of the circles \mathcal{D}_ρ is at least $\pi\rho$; furthermore, one can verify that*

$$\sup_{q \in \mathcal{D}_\rho^+} L_\rho(g_0, q) = \rho \left(2\pi + 2 \arctan \frac{\sqrt{15}}{9} - \arccos \frac{7}{8} \right) \approx 6.5906\rho,$$

and that such a supremum is attained in the limit as q approaches, from within \mathcal{D}_ρ^+ , the point

$$q_{\text{sup}} = \left(\frac{\rho\sqrt{15}}{8}, \frac{\rho}{8} \right) \in \partial\mathcal{D}_\rho^+.$$

Because of symmetry, a similar result holds within \mathcal{D}_ρ^- .

As is apparent from the preceding discussion, the function L_ρ is fairly complicated. The next result establishes lower and upper bounds on $L_\rho(g_0, q)$ based on the Euclidean norm of q .

Proposition 29 *For any $q = (x, y) \in \mathbb{R}^2$, the following holds:*

$$\|q\| + \rho(\theta(q) - \sin \theta(q)) \leq L_\rho(g_0, q) \leq \|q\| + 2\pi\rho, \quad (25)$$

where $\theta(q) = \text{atan2}(|y|, x)$.

Proof: The lower bound is proved first. If $q \in \mathcal{D}_\rho^+ \cup \mathcal{D}_\rho^-$, then $\|q\| \leq 2\rho \sin \theta(q)$, and hence

$$\|q\| + \rho(\theta(q) - \sin \theta(q)) \leq \rho(\theta(q) + \sin \theta(q)) \leq \pi\rho.$$

However $L_\rho(g_0, q) \geq \pi\rho$ within $\mathcal{D}_\rho^+ \cup \mathcal{D}_\rho^-$, which proves that the bound holds in this case.

If $q \notin \mathcal{D}_\rho^+ \cup \mathcal{D}_\rho^-$, proceed in the following way. Define the point q' as depicted in Fig. 14, i.e., by translating the segment from the origin to q normally to itself, until it is tangent to the minimum-radius circle $\partial\mathcal{D}_\rho$. Since $L_\rho(g_0, q') = \rho\theta(q) + \|q\| - \rho \sin \theta(q)$, the bound holds in this case as well.




Figure 14. Constructions for the lower bound (left) and upper bound (right) on L_ρ .

The proof for the upper bound is constructive: consider a path of type *CLC* in which the first subpath has arc length $\rho\theta(q)$, the linear subpath has length $\|q\|$, and the final subpath has arc length $\rho(2\pi - \theta(q))$. Such a path terminates at the desired point q , and has the required length. ■

B. Technical lemmas for Theorem 10

In this section, the technical lemmas used in the proof of Theorem 10 are presented.

Lemma 30 Consider $A_1, A_2, A_3 \in \mathbb{R}_+$ such that $A_3 > A_2 > A_1 \geq 0$, then the average Dubins distance satisfies the following monotonic relationship:

$$\frac{\int_{R_{A_2}(g_i) \setminus R_{A_1}(g_i)} L_\rho(g_i, q) dq}{A_2 - A_1} \leq \frac{\int_{R_{A_3}(g_i) \setminus R_{A_1}(g_i)} L_\rho(g_i, q) dq}{A_3 - A_1}.$$

Proof: Let $\bar{L} := \frac{\int_{R_{A_2}(g_i) \setminus R_{A_1}(g_i)} L_\rho(g_i, q) dq}{A_2 - A_1}$. One can then write

$$\begin{aligned} \frac{\int_{R_{A_3}(g_i) \setminus R_{A_1}(g_i)} L_\rho(g_i, q) dq}{A_3 - A_1} &= \frac{\int_{R_{A_3}(g_i) \setminus R_{A_2}(g_i)} L_\rho(g_i, q) dq + \int_{R_{A_2}(g_i) \setminus R_{A_1}(g_i)} L_\rho(g_i, q) dq}{A_3 - A_1} \\ &\geq \frac{\int_{R_{A_3}(g_i) \setminus R_{A_2}(g_i)} L_\rho(g_i, q) dq + \bar{L}(A_2 - A_1)}{A_3 - A_1}. \end{aligned}$$

For any $q \in R_{A_3}(g_i) \setminus R_{A_2}(g_i)$, $L_\rho(g_i, q) \geq \bar{L}$. Therefore,

$$\frac{\int_{R_{A_3}(g_i) \setminus R_{A_1}(g_i)} L_\rho(g_i, q) dq}{A_3 - A_1} \geq \bar{L}. \quad \blacksquare$$

Lemma 31 *The function defined in Eq. (8) is convex.*

Proof: Consider $A_1, A_3 \in \mathbb{R}^+$ such that $A_3 > A_1$ and $\alpha \in (0, 1)$. Define $A_2 := \alpha A_1 + (1 - \alpha)A_3$. Convexity is shown by proving that for any $\alpha \in (0, 1)$, $f(A_2) \leq \alpha f(A_1) + (1 - \alpha)f(A_3)$ (the proof for the case when $\alpha = 0$ or $\alpha = 1$ is trivial). One can write

$$\begin{aligned} f(A_1) &= \int_{R_{A_1}(g_i)} L_\rho(g_i, q) dq, \\ f(A_2) &= \int_{R_{A_1}(g_i)} L_\rho(g_i, q) dq + \int_{R_{A_2}(g_i) \setminus R_{A_1}(g_i)} L_\rho(g_i, q) dq, \\ f(A_3) &= \int_{R_{A_1}(g_i)} L_\rho(g_i, q) dq + \int_{R_{A_3}(g_i) \setminus R_{A_1}(g_i)} L_\rho(g_i, q) dq. \end{aligned}$$

Now use Lemma 30 to derive the relation

$$\int_{R_{A_2}(g_i) \setminus R_{A_1}(g_i)} L_\rho(g_i, q) dq \leq (1 - \alpha) \int_{R_{A_3}(g_i) \setminus R_{A_1}(g_i)} L_\rho(g_i, q) dq,$$

and, in turn,

$$\begin{aligned}
f(A_2) &= \int_{R_{A_1}(g_i)} L_\rho(g_i, q) dq + \int_{R_{A_2}(g_i) \setminus R_{A_1}(g_i)} L_\rho(g_i, q) dq \\
&\leq \int_{R_{A_1}(g_i)} L_\rho(g_i, q) dq + (1 - \alpha) \int_{R_{A_3}(g_i) \setminus R_{A_1}(g_i)} L_\rho(g_i, q) dq \\
&= \alpha \int_{R_{A_1}(g_i)} L_\rho(g_i, q) dq + (1 - \alpha) \left(\int_{R_{A_1}(g_i)} L_\rho(g_i, q) dq + \int_{R_{A_3}(g_i) \setminus R_{A_1}(g_i)} L_\rho(g_i, q) dq \right) \\
&= \alpha f(A_1) + (1 - \alpha) f(A_3).
\end{aligned}$$

Therefore, f is convex. ■

Acknowledgments

The authors thank Marco Pavone for insightful feedback during the development of this work. This material is based upon work supported in part by the Michigan/AFRL Collaborative Center on Control Sciences, AFOSR Award FA 8650-07-2-3744, by NSF Grants no. 0705451 and 0705453, and by ONR Award N00014-07-1-0721. 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

- ¹Psaraftis, H., “Dynamic Vehicle Routing Problems,” *Vehicle Routing: Methods and Studies*, edited by B. Golden and A. Assad, Studies in Management Science and Systems, Elsevier, 1988.
- ²Bertsimas, D. J. and van Ryzin, G. J., “A Stochastic and Dynamic Vehicle Routing Problem in the Euclidean Plane,” *Operations Research*, Vol. 39, 1991, pp. 601–615.
- ³Bertsimas, D. J. and van Ryzin, G. J., “Stochastic and Dynamic Vehicle Routing in the Euclidean Plane with Multiple Capacitated Vehicles,” *Operations Research*, Vol. 41, No. 1, 1993, pp. 60–76.
- ⁴Bertsimas, D. J. and van Ryzin, G. J., “Stochastic and Dynamic Vehicle Routing with General Inter-arrival and Service Time Distributions,” *Advances in Applied Probability*, Vol. 25, 1993, pp. 947–978.
- ⁵Dubins, L., “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, 1957, pp. 497–516.
- ⁶LaValle, S. M., *Planning algorithms*, Cambridge University Press, Cambridge, UK, 2006.
- ⁷Boscain, U. and Piccoli, B., *Optimal Syntheses for Control Systems on 2-D Manifolds*, Mathématiques et Applications, Springer Verlag, New York, 2004.
- ⁸Pallottino, L. and Bicchi, A., “On optimal cooperative conflict resolution for air traffic management systems,” *IEEE Trans. on Intelligent Transportation Systems*, Vol. 1, No. 4, 2000, pp. 221–231.
- ⁹Tomlin, C., Mitchell, I., and Ghosh, R., “Safety verification of conflict resolution manoeuvres,” *IEEE*

Trans. on Intelligent Transportation Systems, Vol. 2, No. 2, June 2001, pp. 110–120.

¹⁰Chandler, P., Rasmussen, S., and Pachter, M., “UAV cooperative path planning,” *AIAA Conf. on Guidance, Navigation, and Control*, 2000.

¹¹Beard, R. W., McLain, T. W., Goodrich, M. A., and Anderson, E. P., “Coordinated Target assignment and intercept for unmanned air vehicles,” *IEEE Trans. on Robotics and Automation*, Vol. 18, No. 6, 2002, pp. 911–922.

¹²Schumacher, C., Chandler, P. R., Rasmussen, S. J., and Walker, D., “Task allocation for wide area search munitions with variable path length,” *Proc. of the American Control Conference*, Denver, CO, 2003, pp. 3472–3477.

¹³Le Ny, J., Feron, E., and Frazzoli, E., “On the Curvature-Constrained Traveling Salesman Problem,” *IEEE Trans. on Automatic Control*, 2008, To appear.

¹⁴Savla, K., Frazzoli, E., and Bullo, F., “Traveling Salesperson Problems for the Dubins vehicle,” *IEEE Trans. on Automatic Control*, Vol. 53, No. 6, 2008, pp. 1378–1391.

¹⁵Drezner, Z., editor, *Facility Location: A Survey of Applications and Methods*, Springer Series in Operations Research, Springer Verlag, New York, 1995.

¹⁶Enright, J., Frazzoli, E., Savla, K., and Bullo, F., “On Multiple UAV Routing with Stochastic Targets: Performance Bounds and Algorithms,” *Proc. of the AIAA Conf. on Guidance, Navigation, and Control*, San Francisco, CA, August 2005.

¹⁷Savla, K., Bullo, F., and Frazzoli, E., “The coverage problem for loitering Dubins vehicles,” *IEEE Conf. on Decision and Control*, New Orleans, LA, 2007.

¹⁸Enright, J., Savla, K., and Frazzoli, E., “Coverage Control for Teams of Nonholonomic Agents,” *IEEE Conf. on Decision and Control*, 2008.

¹⁹Cortés, J., Martínez, S., Karatas, T., and Bullo, F., “Coverage control for mobile sensing networks,” *IEEE Transactions on Robotics and Automation*, Vol. 20, No. 2, 2004, pp. 243–255.

²⁰Choset, H., “Coverage of known spaces: the boustrophedon cellular decomposition,” *Autonomous Robots*, Vol. 9, No. 3, 2000, pp. 247–253.

²¹Batalin, M. and Sukhatme, G. S., “Coverage, exploration and deployment by a mobile robot and communication network,” *Proceedings of the 2nd International Workshop on Information Processing in Sensor Networks*, Palo Alto, CA, 2003, pp. 376–391.

²²Guo, Y. and Balakrishnan, M., “Complete coverage control for nonholonomic mobile robots in dynamic environments,” *Int’l Coinference on Robotics and Automation*, Orlando, FL, 2006, pp. 1704–1709.

²³Liu, Y., Cruz, J., and Sparks, A. G., “Coordinated networked uninhabited aerial vehicles for persistent area denial,” *IEEE Conf. on Decision and Control*, Paradise Island, Bahamas, 2004, pp. 3351–3356.

²⁴Moore, B. and Passino, K., “Distributed Balancing of AAVs for Uniform Surveillance Coverage,” *IEEE Conference on Decision and Control*, 2005, pp. 7060–7065.

²⁵Fekete, S. P., Mitchell, J. S. B., and Weinbrecht, K., “On the continuous Weber and k -median problems,” *Proceedings of the Sixteenth Annual Symposium on Computational Geometry (Hong Kong, 2000)*, ACM, New York, 2000, pp. 70–79.

²⁶Agarwal, P. K. and Sharir, M., “Efficient algorithms for geometric optimization,” *ACM Computing Surveys*, Vol. 30, No. 4, 1998, pp. 412–458.

²⁷Richards, A., Bellingham, J., Tillerson, M., and How, J., “Coordination and Control of Multiple UAVs,” *Proc. of the AIAA Conf. on Guidance, Navigation, and Control*, Monterey, CA, 2002.

- ²⁸Bertuccelli, L. F., Alighanbari, M., and How, J. P., “Robust Planning For Coupled Cooperative UAV Missions,” *IEEE Conf. on Decision and Control*, Atlantis, Paradise Island, Bahamas, December 2004.
- ²⁹Earl, M. and D’Andrea, R., “Iterative MILP methods for vehicle control problems,” *IEEE Trans. on Robotics*, Vol. 21, December 2005, pp. 1158–1167.
- ³⁰Li, W. and Cassandras, C., “A Cooperative Receding Horizon Controller for Multivehicle Uncertain Environments,” *IEEE Trans. on Automatic Control*, Vol. 51, No. 2, 2006, pp. 242–257.
- ³¹Ma, X. and Castanon, D. A., “Receding Horizon Planning for Dubins Traveling Salesman Problems,” *IEEE Conference on Decision and Control*, San Diego, CA, Dec 2006.
- ³²Ousingsawat, J. and Campbell, M. E., “Optimal Cooperative Reconnaissance Using Multiple Vehicles,” *AIAA Journal of Guidance, Control, and Dynamics*, Vol. 30, No. 1, January-February 2007, pp. 122–132.
- ³³Murphey, R., “Target-based weapon target assignment problems,” *Nonlinear Assignment Problems: Algorithms and Applications*, edited by P. Pardalos and L. Pitsoulis, Kluwer Academic Publisher, 1999, pp. 39–53.
- ³⁴Arslan, G., Marden, J., and Shamma, J., “Autonomous Vehicle-Target Assignment: a game theoretic formulation,” *ASME Journal of Dynamic Systems, Measurement and Control*, Vol. 129, No. 5, September 2007, pp. 584–596.
- ³⁵Bethke, B., Valenti, M., and How, J. P., “UAV Task Assignment,” *IEEE Robotics and Automation Magazine*, March 2008.
- ³⁶Alighanbari, M., Bertuccelli, L. F., and How, J. P., “A Robust Approach to the UAV Task Assignment Problem,” *IEEE Conf. on Decision and Control*, San Diego, CA, December 2006.
- ³⁷Smith, S. L. and Bullo, F., “Monotonic Target Assignment for Robotic Networks,” *IEEE Trans. on Automatic Control*, June 2007, Submitted.
- ³⁸Shamma, J., *Cooperative Control of Distributed Multi-Agent Systems*, John Wiley & Sons, 2007.
- ³⁹Kish, B. A., Jacques, D. R., and Pachter, M., “Optimal Control of Sensor Threshold for Autonomous Wide-Area-Search Munitions,” *AIAA Journal of Guidance, Control, and Dynamics*, Vol. 30, No. 5, September-October 2007, pp. 1239–1248.
- ⁴⁰Vicsek, T., Czirók, A., Ben-Jacob, E., and Cohen, I. Shochet, O., “Novel type of phase transition in a system of self-driven particles,” *Physical Review Letters*, Vol. 75, No. 6, 1995, pp. 1226–1229.
- ⁴¹Spears, W., Spears, J., Hamann, J., and Heil, R., “Distributed, physics-based control of swarms of vehicles,” *Autonomous Robots*, Vol. 17, No. 2-3, 2004, pp. 137–162.
- ⁴²Buhl, J., Sumpter, D. J., Couzin, I. D., Hale, J., Despland, E., Miller, E., and Simpson, S. J., “From disorder to order in marching locusts,” *Science*, Vol. 312, 2006, pp. 1402–1406.
- ⁴³Yechout, T. R., Morris, S. L., Bossert, D. E., and Hallgreen, W. F., *Introduction to Aircraft Flight Mechanics: Performance, Static Stability, Dynamic Stability, and Classical Feedback Control*, AIAA Education Series, AIAA, Reston, VA, 2003.
- ⁴⁴Little, J., “A proof of the queueing formula $l = \lambda w$,” *Operations Research*, Vol. 9, 1961, pp. 383–387.
- ⁴⁵Shkel, A. M. and Lumelsky, V. J., “Classification of the Dubins Set,” *Robotics and Autonomous Systems*, Vol. 23, 2001, pp. 179–202.
- ⁴⁶Thomaschewski, B., “Dubins’ problem for free terminal direction,” Preprint M09/2001, Technical University of Ilmenau, Ilmenau, Germany, 2001.
- ⁴⁷Rathinam, S., Sengupta, R., and Darbha, S., “A Resource Allocation Algorithm for Multivehicle Systems With Nonholonomic Constraints,” *tase*, Vol. 4, No. 1, 2007, pp. 98–104.

⁴⁸Okabe, A., Boots, B., Sugihara, K., and Chiu, S., *Spatial tessellations: Concepts and Applications of Voronoi diagrams*, Wiley Series in Probability and Statistics, John Wiley & Sons, Chichester, UK, 2nd ed., 2000.

⁴⁹Zemel, E., “Probabilistic Analysis of Geometric Location Problems,” *Annals of Operations Research*, Vol. 1, No. 3, October 1984.

⁵⁰Papadimitriou, C. H., “Worst-Case and Probabilistic Analysis of a Geometric Location Problem,” *SIAM Journal on Computing*, Vol. 10, No. 3, August 1981.

⁵¹Megiddo, N. and Supowit, K. J., “On the complexity of some common geometric location problems,” *SIAM Journal on Computing*, Vol. 13, No. 1, 1984, pp. 182–196.

⁵²Masuyama, S., Ibaraki, T., and Hasegawa, T., “The computational complexity of the m-center problem on the plane,” *Trans. IECE of Japan*, Vol. 64, No. 2, 1981, pp. 57–64.

⁵³Arsie, A., Savla, K., and Frazzoli, E., “Efficient Routing of Multiple Vehicles with No Communications,” *IEEE Trans. on Automatic Control*, 2009, To appear.

⁵⁴Larson, R. C. and Odoni, A. R., *Urban Operations Research*, Prentice-Hall, Englewood Cliffs, NJ, 1981.

⁵⁵Boyd, S. and Vandenberghe, L., *Convex Optimization*, Cambridge University Press, 2004.

⁵⁶*Unmanned Aircraft Systems Roadmap: 2005–2030*, Office of the Secretary of Defense, 2005.

⁵⁷Lawrence, J., *A Catalog of Special Plane Curves*, Dover, New York, NY, 1972.