On Multiple UAV Routing with Stochastic Targets: Performance Bounds and Algorithms

J. J. Enright* and E. Frazzoli†

University of California – Los Angeles, Los Angeles, CA 90095

K. Savla‡ and F. Bullo§

University of California – Santa Barbara, Santa Barbara, CA 93106

In this paper we consider the following problem. A number of Uninhabited Aerial Vehicles (UAVs), modeled as vehicles moving at constant speed along paths of bounded curvature, must visit stochastically-generated targets in a convex, compact region of the plane. Targets are generated according to a spatio-temporal Poisson process, uniformly in the region. It is desired to minimize the expected waiting time between the appearance of a target, and the time it is visited. We present partially centralized algorithms for UAV routing, assigning regions of responsibility to each vehicle, and compare their performance with respect to asymptotic performance bounds, in the light and heavy load limits. Simulation results are presented and discussed.

I. Introduction

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

Variations of problems falling in this class have been studied in a number of papers in the recent past, e.g., [1–4]. In these papers, the problem is set up in such a way that the location of targets is known a priori; a strategy is computed that attempts to optimize the cost of servicing the known targets. In [5] a stable receding-horizon strategy is proposed, but its performance in not characterized. In [6], we addressed the case in which new targets are generated continuously by a stochastic process: we provided algorithms for minimizing the expected waiting time between the appearance of a target and the time it is serviced by one of the vehicles. A limitation of the results presented in [6] is the fact that omni-directional vehicles were considered in the problem formulation: as such, the results are not applicable to many vehicles of interest, including aircraft and car-like robots.

In [7] we presented the initial results of our work aimed at extending the results available in the literature to address non-holonomic vehicle dynamics. In particular, we considered paths with bounded curvature, which provide a good approximation of feasible trajectories for aircraft, and provided performance bounds and routing algorithms for the single-vehicle case. In this paper, we improve our previous results and extend them to the multiple-vehicle case. The main contributions of this paper are: (i) The establishment of a new

*Graduate Research Assistant, Mechanical and Aerospace Engineering Department, 420 Westwood Plaza, Los Angeles, CA 90095. AIAA Student Member.

†Assistant Professor, Mechanical and Aerospace Engineering Department, 420 Westwood Plaza, Los Angeles, CA 90095. AIAA Member.

‡Graduate Research Assistant, Electrical Engineering Department, 2338 Engineering Bldg II, UCSB, Santa Barbara, CA 93106-5070, USA

§Associate Professor, Mechanical and Environmental Engineering Department, 2338 Engineering Bldg II, UCSB, Santa Barbara, CA 93106-5070, USA.

lower bound for the achievable performance in light load conditions (i.e., when targets are generated very rarely), for the single- and multiple-vehicle cases; (ii) The extension to the multiple-vehicle case of lower bounds and upper bounds in heavy load (i.e., when targets are generated very often) first derived in [7, 8]; (iii) The design of partially-centralized task allocation and routing strategies achieving performance levels that are provably within a constant factor from the optimum, both in the light load and heavy load cases.

The paper is structured as follows. In Section II we introduce some notation and formulate the problem we wish to address. In Section III we address the case in which new targets appear very rarely, and provide lower bounds and constructive upper bounds on the achievable performance. In Section IV we investigate the same objectives in the case in which the target generation rate is very high. In Section V, we draw some conclusions and discuss some directions for future work.

II. Problem Formulation

The basic version of the problem we wish to study in this paper is known as the Dynamic Traveling Repairperson Problem (DTRP), and was introduced by Bertsimas and van Ryzin in [9]. The m-vehicle version of the problem, m-DTRP, was first studied by the same authors in [10]. Our problem is different from the original single- or multiple-vehicle DTRP since we consider a vehicle that is constrained to move at unit speed along paths of bounded curvature, i.e., we impose a non-holonomic constraint on the vehicle's dynamics. In the remainder of the section, we define the details of the problem and its components.

Let the environment $\mathcal{Q} \subset \mathbb{R}^2$ be a convex, compact set with unit area, and let $\|\cdot\|$ denote the Euclidean norm in $\mathbb{R}^2$. Consider $m \geq 1$ UAVs, modeled as nonholonomic vehicles constrained to move at constant speed along a path with bounded curvature, and let $1/\rho$ be the maximum curvature. Without loss of generality, we will assume that the vehicles move at unit speed. In other words, let the configuration $g_i \in SE(2)$ of the i-th vehicle ($1 \leq i \leq m$) be given in coordinates by $g_i = (x_i, y_i, \theta_i)$, where x_i, y_i are respectively the projections of the vehicle's position along inertially fixed orthogonal axes, and θ_i is the orientation of the vehicle's longitudinal axis with respect to the $y = 0$ axis; then the dynamics of the vehicle are described by the differential equations

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

The vehicles have unlimited range and target-serving capacity. In the following, we will indicate by $p_i = (x_i, y_i)$ the position of the i-th vehicle. Moreover, we will indicate by $g = \{g_1, g_2, \ldots, g_m\} \in SE(2)^m$ the configuration of the m-vehicle system. We note that the above kinematic model of an airplane's dynamics is very common in the literature on UAV motion planning; the model is very similar to the one studied in [11], with the difference that the vehicle we consider is constrained to move at constant speed. Results in terms of minimum-length paths for Dubins' vehicle hold for our model, where they assume an additional connotation of being minimum-time paths as well. In the following, we will often use the expressions "Dubins vehicle" and "Dubins paths" to indicate a vehicle modeled by (1) and paths that are feasible with respect to the same model.

Information on outstanding targets—the demand—at time t is summarized as a finite set of target positions $D(t) \subset \mathcal{Q}$, with $n(t) := \mathrm{card}(D(t))$. Targets are generated, and inserted into D, according to a homogeneous (i.e., time-invariant) spatio-temporal Poisson process, with time intensity $\lambda > 0$, and uniform spatial density. In other words, given a set $\mathcal{S} \subseteq \mathcal{Q}$, the expected number of targets generated in $\mathcal{S}$ within the time interval $[t, t']$ is

$$\begin{aligned} &\mathrm{E}\left[\mathrm{card}(D(t') \cap \mathcal{S}) - \mathrm{card}(D(t) \cap \mathcal{S})\right] = \\ &\quad = \lambda (t' - t) \mathrm{Area}(\mathcal{S}). \end{aligned}$$

(Strictly speaking, the above equation holds in the case in which targets are not being removed from the queue D.) Servicing of a target $e_j \in D$, and its removal from the set D, is achieved when the UAV moves to the target's position.

A static feedback control policy for the system is a map $\pi : SE(2)^m \times 2^{\mathcal{Q}} \to [-1/\rho, 1/\rho]^m$, assigning a control input to each vehicle, as a function of the current state of the system, i.e., $\omega_i(t) = \pi_i(g(t), D(t))$. (The subscript i denotes the component of the output vector, and of the control policy, that are relevant to the i-th vehicle.) In this paper we will also consider policies that compute a control input for the vehicles based on a snapshot of the target configuration at a certain time in the past, at which certain computations are made. Let $\mathcal{T}_\pi = \{t_1, t_2, \ldots, t_i, \ldots\}$ be a strictly increasing sequence of times at which such computations are started: with some abuse of terminology, we will say that π is a *receding horizon strategy* if it is based on the most recent target data available, i.e., $\omega_i(t) = \pi_i(g(t), D_{\mathrm{rh}}(t))$, with

$$D_{\mathrm{rh}}(t) = D(\max\{t_{\mathrm{rh}} \in \mathcal{T}_\pi : t_{\mathrm{rh}} < t\}).$$

Note that the computation time schedule $\mathcal{T}_\pi$ is an additional output of the policy π, but we do not formalize this in order not to make the notation exceedingly cumbersome.

The (receding horizon) policy π is stable if, under its action,

$$n_\pi := \lim_{t \to +\infty} \mathrm{E}\left[n(t) | \dot{p} = \pi(p, D_{\mathrm{rh}})\right] < +\infty,$$

that is, if the UAV is able to service targets at a rate that is—on average—at least as fast as the rate at which new targets are generated.

Let T_j be the time that the j-th target spends within the set D, i.e., the time elapsed from the time e_j is generated to the time it is serviced. If the system is stable, then we can write the balance equation (known as Little's formula [12])

$$n_\pi = \lambda T_\pi,$$

where $T_\pi := \lim_{j \to +\infty} \mathrm{E}\left[T_j\right]$ is the steady-state system time under the policy π. Our objective is to minimize the steady-state system time, over all possible static feedback control policies, i.e.,

$$T^* = \inf_\pi T_\pi.$$

In the following, we are interested in designing control policies that provide constant-factor approximations of the optimal achievable performance. In particular, as done in [9], we will analyze the asymptotic cases of light load, as $\lambda \to 0$, and heavy load, as $\lambda \to \infty$; in the light load case, we will derive lower and upper bounds on the achievable system time T^*. The upper bound is given in terms of system times achieved by an explicit, readily implementable algorithm, and is therefore constructive. In the heavy load case, we derive a lower bound for the system time, and present an approximation algorithm providing an upper bound on the performance that holds with high probability.

III. The light load case

In this section we will derive bounds on the achievable system time T^* in light-load conditions, i.e., when $\lambda \to 0^+$, and design routing policies whose performance is provably within a constant factor from the lower bound, and hence from the optimum. Before we address the core of the problem, we present some results on a problem that will prove to be crucial to the derivation of our bounds, that is, the analysis of the length of Dubins paths (for a single vehicle) from an initial configuration to a given point in the plane.

A. Some results on minimal-length Dubins paths

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 [11]; a further classification is given in [13]. Our purpose in this section is to study the length of optimal paths given different boundary conditions. Namely, we will consider the problem of steering a Dubins vehicle from a given initial configuration to a point in the plane; the difference with the original problem posed by Dubins is that the final heading at the target point is not constrained a priori.

In other words, let $L_\rho(q) : \mathbb{R}^2 \to \mathbb{R}$ be the minimum length of a path satisfying (1), steering a Dubins vehicle from the identity in $SE(2)$ (i.e., $g(0) = (0, 0, 0)$) to a point q in the plane, without any constraints on the final heading, i.e., to any configuration in the set $\{q\} \times S^1 \subset SE(2)$. The characteristics of paths of minimal length with such boundary conditions were studied in [14], 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 upper bounded by $2\pi\rho$, 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 $q = 0$.)

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. In other words, having defined $\mathcal{D}_\rho^+ = \{q \in \mathbb{R}^2 : \|q - (0, \rho)\| < \rho\}$, $\mathcal{D}_\rho^- = \{q \in \mathbb{R}^2 : \|q - (0, -\rho)\| < \rho\}$, if $q \in \mathcal{D}_\rho^+ \cup \mathcal{D}_\rho^-$ the optimal path is of type CC, otherwise it is of type CL.

Let us state a few direct consequences of the above results, which we believe are useful in gaining an insight on the nature of the problem we are discussing.

Proposition 1 *The minimum length $L_\rho(q)$ of a Dubins path steering a vehicle from $g_0 = (0, 0, 0) \in SE(2)$ to a point $q = (x, y) \in \mathbb{R}^2$ is given by:*

$$L_\rho(q) = \begin{cases} \sqrt{d_{\mathrm{c}}^2(q) - 2\rho d_{\mathrm{c}}(q)} + \rho\left(\theta_{\mathrm{c}}(q) - \arccos\dfrac{\rho}{d_{\mathrm{c}} - \rho}\right) & \text{for } q \notin \mathcal{D}_\rho^+ \cup \mathcal{D}_\rho^-; \\ \rho\left(2\pi - \alpha(q) + \arcsin\dfrac{d_{\mathrm{c}}(q)\sin\theta_{\mathrm{c}}(q)}{d_{\mathrm{f}}(q)} + \arcsin\dfrac{\rho\sin\alpha(q)}{d_{\mathrm{f}}(q)}\right) & \text{otherwise.} \end{cases} \tag{2}$$

where $d_{\mathrm{c}}(q) = \sqrt{x^2 + (|y| - \rho)^2}$ and $\theta_{\mathrm{c}}(q) = \operatorname{atan2}(x, (\rho - |y|))$ are polar coordinates of the point q with respect to the center of either $\mathcal{D}_\rho^+$ or $\mathcal{D}_R^-$, whichever is the closest, $d_{\mathrm{f}}(q) = \sqrt{x^2 + (|y| + \rho)^2}$ is the distance of q from the other center, and

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

Proof: The proof is based on the results in [14] mentioned in the text, and elementary planar geometry. ■

The facts stated in the following remarks are consequences of Proposition 1.

Remark 2 *The level sets of the function L_ρ are segments of well-studied curves. More precisely, level sets of L_ρ are:*

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

We 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., [15]). A depiction of such level sets is provided in Figure 1.

Remark 3 *The function L_ρ is:*

- *Not continuous on the "top half" of the boundary of $\mathcal{D}_\rho^+$ and $\mathcal{D}_\rho^-$, 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 x axis;*
- *Continuous and differentiable everywhere else.*

Remark 4 *The length of the optimal path to reach a point inside one of the circles $\mathcal{D}_\rho$ is at least $\pi\rho$; in particular, one can verify that*

$$\sup_{q \in \mathcal{D}_\rho^+} L_\rho(q) = \rho\left(2\pi + 2\arctan\frac{\sqrt{15}}{9} - \arccos\frac{7}{8}\right) \approx 6.5906\rho, \tag{3}$$

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

$$q_{\text{sup}} = \left(\frac{\rho\sqrt{15}}{8}, \frac{\rho}{8}\right) \in \partial\mathcal{D}_\rho^+. \tag{4}$$

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

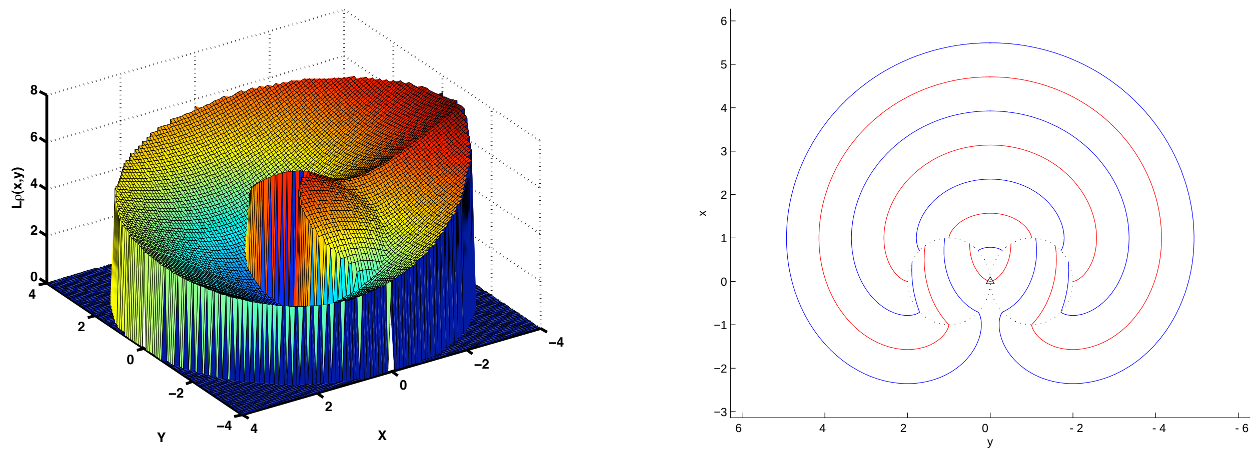

Figure 1. A 3D plot of the function L_ρ (left) and its level sets L_ρ, for $\rho = 1$, shown at increments of $\pi/4$ (right). The initial vehicle configuration is represented by the black triangle, and is at coordinates $(0, 0)$ heading "up". Note that the positive direction of the x axis is "up" and of the y axis is "left."

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

Proposition 5 *For any $q = (x, y) \in \mathbb{R}^2$, the following holds:*

$$\|q\| + \rho(\theta(q) - \sin\theta(q)) \leq L_\rho(q) \leq \|q\| + 2\pi\rho, \tag{5}$$

where $\theta(q) = \arctan |y|/x$.

Proof: Let us prove the lower bound 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 $Ł_\rho(q) \geq \pi\rho$ within $\mathcal{D}_\rho^+ \cup \mathcal{D}_\rho^-$, which proves that the bound holds in this case.

If $q \notin \mathcal{D}_\rho^+ \cup \mathcal{D}_\rho^-$, we proceed in the following way. Define the point q' as depicted in Figure 2, i.e, by translating the segment from the origin to q normally to itself, until it is tangent to the minimum-radius circle $\partial\mathcal{D}_\rho$. Because of the triangle inequality $L_\rho(q') \leq L_\rho(q)$. Since $L_\rho(q') = \rho\theta(q) + \|q\| - \rho \sin\theta(q)$, the bound holds in this case as well.

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

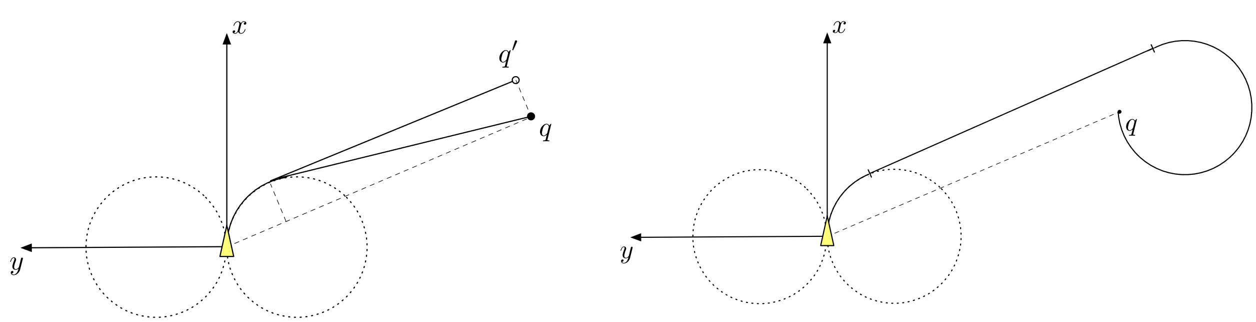

Figure 2. Constructions for the lower bound (left) and upper bound (right) bound on L_ρ.

We now present the last result of this section, which was established through numerical optimization, based on the functions presented in Proposition 1.

Proposition 6 *The function $L_\rho(q)$ satisfies the following inequality, for all $q \in \mathbb{R}^2$:*

$$L_\rho(q) \leq c_1\rho + \|q - (0, c_2\rho)\|, \tag{6}$$

with $c_1 \approx 3.75595$, and $c_2 \approx 2.91801$.

B. A lower bound on the system time

In order to study optimal policies in the light load case, we need to introduce a problem from geometric optimization. Given a set $\mathcal{Q} \subset \mathbb{R}^d$ and a set of points $p = \{p_1, p_2, \ldots, p_m\} \in \mathcal{Q}^m$, 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

$$H_m(p, \mathcal{Q}) := \mathrm{E}\left[\min_{i\in\{1,\ldots,m\}} \|p_i - q\|\right] = \sum_{i=1}^{m} \int_{\mathcal{V}_i(p)} \|p_i - q\| dq, \tag{7}$$

where $\mathcal{V}(p) = \{\mathcal{V}_1(p), \mathcal{V}_2(p), \ldots, \mathcal{V}_m(p)\}$ is the Voronoi partition of the set $\mathcal{Q}$ generated by the points p. In other words, $q \in \mathcal{V}_i(p)$ if $\|q - p_i\| \leq \|q - p_k\|$, for all $k \in \{1, \ldots, m\}$. The set $\mathcal{V}_i$ is referred to as the Voronoi cell of the generator p_i. The function H_m is known in the locational optimization literature as the *continuous Weber function* or the *continuous multi-median function*; see [16, 17] and references therein.

The m-median of the set $\mathcal{Q}$ is the global minimizer

$$p_m^*(\mathcal{Q}) = \operatorname*{argmin}_{p\in\mathcal{Q}^m} H_m(p, \mathcal{Q}).$$

We let $H_m^*(\mathcal{Q}) = H_m(p_m^*(\mathcal{Q}), \mathcal{Q})$ be the global minimum of H_m. It is straightforward to show that the map $p \mapsto H_1(p, \mathcal{Q})$ is differentiable and strictly convex on $\mathcal{Q}$. Therefore, it is a simple computational task to compute $p_1^*(\mathcal{Q})$. It is convenient to refer to $p_1^*(\mathcal{Q})$ as the median of $\mathcal{Q}$. On the other hand, the map $P \mapsto H_m(P, \mathcal{Q})$ with $m > 1$ is differentiable (whenever $(p_1, \ldots, p_m)$ are distinct) but not convex, thus making the solution of the continuous m-median problem hard in the general case. In fact, it is known [16, 18] that the discrete version of the m-median problem is NP-hard for $d \geq 2$. Gradient algorithms for the continuous m-median problems can be designed [19] by means of the equality

$$\frac{\partial H_m(p, \mathcal{Q})}{\partial p_i} = \int_{\mathcal{V}_i(p)} \frac{p_i - q}{\|p_i - q\|}\, dq. \tag{8}$$

We will not pursue further the issue of computation of the m-median and of the corresponding $H_m^*(\mathcal{Q})$, but will assume that these values are available.

It is interesting to investigate how the optimal value of the multi-median function depends on the number of generators, i.e., how $H_m^*(\mathcal{Q})$ depends on the integer m, for a fixed environment $\mathcal{Q}$. If the environment $\mathcal{Q}$ is a square of side L, it is known that $H_1^*(\mathcal{Q}) = c_{\text{square}} L$, with $c_{\text{square}} \approx 0.52$. Simple scaling arguments let one conclude that in the case of a square environment,

$$\frac{c_{\text{square}} L}{\lceil \sqrt{m} \rceil} \leq H_m^*(\mathcal{Q}) \leq \frac{c_{\text{square}} L}{\lfloor \sqrt{m} \rfloor}.$$

Such arguments can be generalized to general shapes for the environment, and show that $H_m^*(\mathcal{Q})$ is $\Theta(1/\sqrt{m})$ for a given $\mathcal{Q}$.

We start the analysis of the light case stating the following bound. The result was first proven by [9] for holonomic vehicles, and as such is trivially valid for non-holonomic vehicles. A short proof is reported for completeness.

Theorem 7 *The system time T^* for the problem stated in Section II satisfies:*

$$T^* \geq H_m^*(\mathcal{Q}) \tag{9}$$

Proof: Let us consider the i-th target, generated at time t_i, at a random position $e_i \in \mathcal{Q}$. The time necessary to service the i-th target is lower bounded by minimum time taken by the vehicle to move from its configuration at time t_i to the target's position e_i. The configuration of the vehicle at time t_i is in general unknown, since it depends on the chosen control policy, and on the history of generated targets. If we assumed that the vehicle is always in such a location that it minimizes the *a priori* expected Euclidean distance to a randomly-generated target, we would get a lower bound on the expected service time: such a point is the median of the set $\mathcal{Q}$. In other words,

$$T^* \geq \min_{g \in SE(2)^m} \mathrm{E}\left[\min_{i \in \{1,\ldots,m\}} \mathrm{DubinsDistance}(g_i, q)\right] \geq \min_{p \in \mathcal{Q}^m} \mathrm{E}\left[\min_{i \in \{1,\ldots,m\}} \|p_i - q\|\right] = H_m^*(\mathcal{Q}).$$

■

Theorem 7 holds for any policy, and any value of λ. However, it is most useful in the light-load case and as such it is reported in this section. However, the theorem provides little insight into the specifics of the routing problem in light load for Dubins vehicles. In particular, Dubins vehicles *cannot* wait for the generation of targets at a single location. As a consequence, we need to characterize the configuration of the agents at the appearance of new targets in terms of Dubins paths, that we will call *loitering patterns*. In general, optimal loitering patterns will have to be computed based on the shape of the region assigned to a certain agent. However, we will concentrate on *circular loitering patterns*; the rationale for doing so is that this allows us to provide algorithms and bounds that are independent of the particular shape of the environment. Furthermore, it seems unlikely that UAVs in the field will be able to compute optimal loitering patterns as their assigned regions change in real time; on the other hand, determining the location of the center, and the radius of a circular loitering patterns are much easier tasks.

We will also consider the case in which the agents group into $l \leq m$ teams, that share Voronoi regions. In other words, to each team we associate a single generator p_i, $i \in \{1, \ldots, m\}$. Each team will be composed of m/l agents; for simplicity, we assume that such a number is an integer. Non-integer values can be understood in a time-averaged sense, as agents join different teams over time.

We have the following:

Theorem 8 *Under the assumption that the m agents are grouped into l teams of m/l agents each, and all agents execute circular loitering patterns while waiting for targets, the system time T^* for the problem stated in Section II satisfies:*

$$T^* \geq H_l^*(\mathcal{Q}) + \rho\left(\frac{\pi l}{2m} + \frac{m}{\pi l}\left(\cos\frac{\pi l}{m} - 1\right)\right). \tag{10}$$

In particular, if we impose the constraint that $m = l$ (no teaming up allowed), the bound takes the form:

$$T^* \geq H_m^*(\mathcal{Q}) + \rho\left(\frac{\pi}{2} - \frac{2}{\pi}\right) \approx H_m^*(\mathcal{Q}) + 0.9342\rho. \tag{11}$$

Proof: The proof of this theorem is similar to the proof of Theorem 7, with the difference that now we use a non-trivial bound for the Dubins distance. More specifically, we use the lower bound in Proposition 5. In other words, while we assume that the *position* of each vehicle is such that it minimizes the expected Euclidean distance to a random target, we model the fact that the heading must change over time in order to maintain the agents close to the desired point. The assumption of circular loitering patterns lets us consider a uniform distribution of the agents' heading over time. We get:

$$\begin{aligned} T^* &\geq \min_{g \in SE(2)^m} \mathrm{E}\left[\min_{i \in \{1,\ldots,m\}} \mathrm{DubinsDistance}(g_i, q)\right] \\ &\geq \min_{p \in \mathcal{Q}^l} \mathrm{E}\left[\min_{i \in \{1,\ldots,l\}} \|p_i - q\|\right] + \mathrm{E}\left[\min_{j \in \{1,\ldots,m/l\}} (|\theta_j(q)| - \sin|\theta_j(q)|)\right] \\ &\geq H_m^*(\mathcal{Q}) + \frac{m}{\pi l}\int_0^{\pi l/m} \theta - \sin\theta \, d\theta, \end{aligned}$$

which yields the stated result upon integration. ■

Some comments are in order at this point. The lower bound shows that when ρ is very small, compared to $H_l^*(\mathcal{Q})$, the result for the holonomic case (Theorem 7) is recovered, according to intuition. However, should ρ be large compared to $H_l^*(\mathcal{Q})$, which can happen when many agents are available, and acting in separate teams, it might be advantageous to group agents to reduce the cost penalties induced by non-holonomic constraints.

C. A constructive upper bound

Recall Proposition 6. Its significance is the following: if a vehicle moves along a circular path $\gamma : \mathbb{R} \to SE(2)$ of radius $c_2\rho$ around a point p, then for all $t \in \mathbb{R}$, $\mathrm{DubinsDistance}(\gamma(t), q) \leq c_1\rho + \|q - p\|$.

Based on this observation, consider the following control policy, which we call the offset median (OM) policy. Let assume that all agents act independently, without teaming up: in this case, $l = m$. Let p^* be the m-median of $\mathcal{Q}$, and define the *loitering station* for the i-th agent as a circular trajectory of radius $c_2\rho$ centered at p_i^*. In the offset median policy, each agent visits all targets in its own Voronoi region $\mathcal{V}_i(p^*)$ in a greedy fashion: in other words, it always pursues the closest target in a Dubins' distance sense. When no targets are available, it returns to its loitering station; the direction in which the orbit is followed is inconsequential, and can be chosen in such a way that the station is reached in minimum time. We refer to such an algorithm as "partially centralized," meaning that the location of loitering stations (and hence the Voronoi partition) is decided in a centralized fashion, but each agent in its own area acts independently from the others.

Theorem 9 *An upper bound on the system time of the Offset Median policy in light load is*

$$T_{\mathrm{OM}} \leq T^* + c_1\rho \ \text{as} \ \lambda \to 0 \tag{12}$$

Proof: Consider a generic initial condition for all the agents' configurations in $\mathcal{Q}$ and for the outstanding target positions $D(0)$, with $n_0 = \mathrm{card}(D(0))$. An upper bound to the time needed to service all of the initial targets is $n_0(\mathrm{diam}(\mathcal{Q}) + 2\pi\rho)$. When there are no targets outstanding in the target set D, all vehicles move at unit speed toward the assigned loitering station, which is reached in at most $\mathrm{diam}(\mathcal{Q}) + 2\pi\rho$ units of time.

The time needed to service the initial targets and go to the median is hence bounded by $t_{\mathrm{ini}} \leq (n_0 + 1)\,[\mathrm{diam}(\mathcal{Q}) + 2\pi\rho]$. The probability that at the end of this initial phase the number of targets is reduced to zero is

$$\begin{aligned} P\left[n(t_{\mathrm{ini}}) = 0\right] &= \exp(-\lambda t_{\mathrm{ini}}) \\ &\geq \exp\left(-\lambda(n_0+1)(\mathrm{diam}(\mathcal{Q}) + 2\pi\rho)\right), \end{aligned}$$

that is, $P\left[n(t_{\mathrm{ini}}) = 0\right] \to 1^-$ as $\lambda \to 0$. As a consequence, after an initial transient, all targets will be generated with the vehicles in the assigned loitering station, and an empty demand queue.

After the initial transient, when the next target arises, say the j-th target at location e_j, the expected time for servicing will be $T_j \leq \min_{i\in\{1,\ldots,m\}} \|e_j - p_i^*\| + c_1\rho$. The system time can hence be computed as

$$T_{\mathrm{OM}} = \lim_{j\to+\infty} \mathrm{E}\left[T_j\right] = H_m^*(\mathcal{Q}) + c_1\rho.$$

■

In other words, we have shown that the system time achieved by the OM policy is within a constant additive factor $(c_1 - \pi/2 + 2/\pi)\rho \approx 2.8218\rho$ from the optimal—if teaming is not allowed. The additive factor, which can be considered as a penalty due to the non-holonomic constraints imposed on the vehicle's dynamics, depends linearly on the minimum turn radius ρ. At this time, neither the lower bound nor the upper bound are known to be tight. As indicated by the lower bound, teaming strategies are likely able to further reduce the upper bound; such strategies are currently under investigation.

In Figure 3 we show simulation results that confirm our theoretical predictions. When the minimum turn radius is very small, the performance of the OM policy approximates the lower bound valid for a vehicle without kinematics constraints, i.e., as $\rho \to 0$, $T_{\mathrm{OM}} \to H^*(\mathcal{Q})$. As ρ increases, the penalty associated to the bounded curvature constraints dominates, and the system time increases linearly with ρ. For a fixed minimum turning radius, numerical experiments show that the lower bound (11) is approached asymptotically as the dimensions of the environment increase. Furthermore, the upper bound captures well the penalty associated to turning in small environment. Note that, while the effect of the non-holonomic constraints decreases as the dimensions of the environment increase, it never vanishes, and remains noticeable in the form of the offset $\rho(\pi/2 - 2/\pi)$.

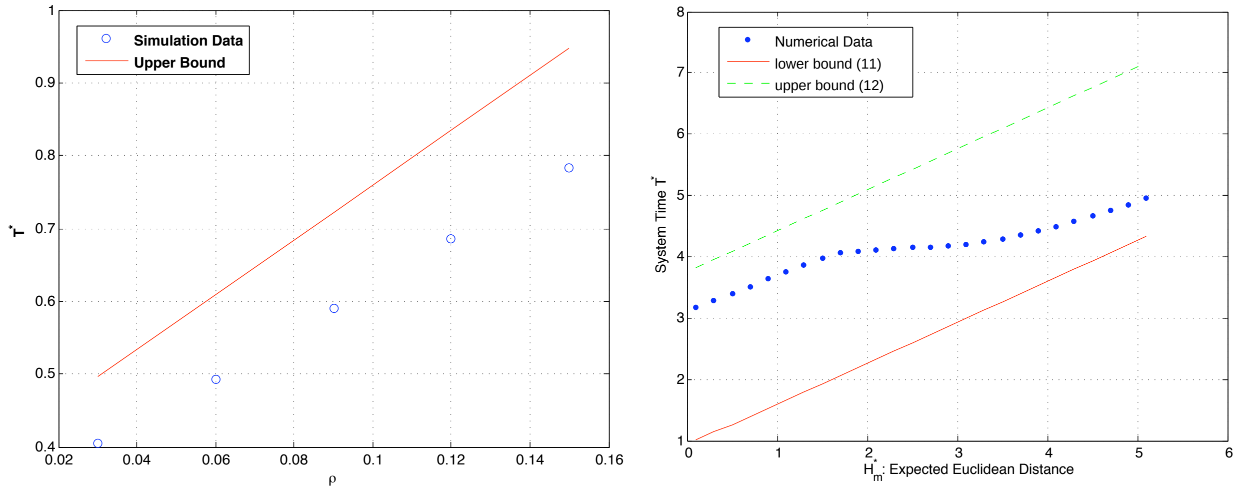

Figure 3. (Performance of the OM policy in light load, as a function of the minimum turning radius (left) and environment size (right), compared to the bounds derived in the text.

IV. The heavy load limit

In this section, we turn our attention to the heavy-load limit, in which $\lambda \to \infty$. In the heavy load case, the nature of optimal control policies is related to the well-known Traveling Salesperson Problem (TSP). We will first discuss some well-known results for the Euclidean version of the TSP, then derive a lower bound on the asymptotic cost of TSP problems for bounded-curvature vehicles. Based on this result, we will provide a lower bound on the system time in the heavy load limit.

A. The Euclidean Traveling Salesperson Problem

The Euclidean TSP (ETSP) is formulated as follows: given a set D of n points in $\mathbb{R}^d$, find the minimum-length tour of D. Let ETSP(D) denote the minimum length of a tour through all the points in D; by convention, ETSP$(\emptyset) = 0$. The asymptotic behavior of stochastic ETSP problems for large n exhibits the following interesting property. Assume that the locations of the n targets are independent random variables, uniformly distributed in a compact set $\mathcal{Q}$; in [20] it is shown that there exists a constant $\beta_{\mathrm{TSP},2}$ such that, almost surely,

$$\lim_{n \to +\infty} \frac{\mathrm{TSP}(D)}{\sqrt{n}} = \beta_{\mathrm{TSP},2}. \tag{13}$$

In other words, the optimal cost of stochastic ETSP tours approaches a deterministic limit, and grows as the square root of the number of points in D; the current best estimate of the constant in (13) is $\beta_{\mathrm{TSP},2} = 0.7120 \pm 0.0002$, see [21, 22].

B. The Traveling Salesperson Problem for a Dubins vehicle

While the ETSP has attracted a great deal of interest from the scientific community, its bounded-curvature counterpart (which we will call DTSP) has not been studied extensively. In [23] we did some initial work, mainly in terms of upper bounds for worst-case tours. Here we extend to the multiple-vehicle case a result from [7], which can be seen as a first step in the search of deterministic bounds similar to those available for the ETSP.

Theorem 10 (From [7]) *The expected cost of a stochastic DTSP visiting a set D of n randomly-generated points in $\mathcal{Q}$ satisfies the following inequality:*

$$\lim_{n \to \infty} \frac{\mathrm{DTSP}(D, \rho)}{n^{2/3}} \geq \frac{3}{4}(3\rho)^{1/3} \tag{14}$$

Proof: Choose a random point $p_i \in D$ as the initial position of the vehicle on the tour, and choose the heading randomly. We would like to compute a bound on the expected distance, according to the metric induced by the length of Dubins' paths, to the closest next point in the tour; let us call such distance δ^*.

To this purpose, consider the set R_δ of points that are reachable from a Dubins' vehicle with an arc of length $\delta \leq \rho$. It can be verified that the area of such a set is

$$\text{Area}(R_\delta) = \frac{\delta^3}{3\rho}. \tag{15}$$

In other words, the area of the set R_δ decreases faster than the area of a circle of radius δ as $\delta \to 0$.

Given a distance δ, the probability that $\delta^* > \delta$ is no less than the probability that there is no other target reachable with a path of length at most δ; in other words,

$$\Pr[\delta^* > \delta] \geq 1 - n\frac{\text{Area}(R_\delta)}{\text{Area}(\mathcal{Q})} = 1 - n\frac{\delta^3}{3\rho}.$$

In terms of expectation, defining $c = n/(3\rho)$,

$$\begin{aligned}
\mathrm{E}\left[\delta^*\right] &= \int_0^\infty \Pr(\delta^* > \xi)\, d\xi \\
&\geq \int_0^\infty \max\{0, 1 - n\frac{\xi^3}{3\rho}\}\, d\xi \\
&= \int_0^{c^{-1/3}} (1 - c\xi^3)\, d\xi \\
&= \frac{3}{4}\left(\frac{3\rho}{n}\right)^{1/3}
\end{aligned}$$

The expected total tour length will be no smaller than n times the expected length of the shortest path between two points, i.e.,

$$\mathrm{E}\left[\text{DTSP}(D, \rho)\right] \geq \frac{3}{4}\left(3\rho n^2\right)^{1/3}.$$

Dividing both sides by $n^{2/3}$ and taking the limit as $n \to \infty$, we get the desired result. ■

C. Lower bound on the system time

At this point we can state the desired result in terms of a lower bound on the system time for any policy in the heavy load case, in the multiple-vehicle case.

Theorem 11 *The system time for the problem stated in Section II, satisfies the following inequality, for $\lambda \to \infty$:*

$$T^* \geq \frac{81}{64}\rho\frac{\lambda^2}{m^3} \quad (\lambda \to \infty) \tag{16}$$

Proof: Let us assume that a stabilizing policy is available. In such a case, the number of outstanding targets approaches a finite steady-state value, n^*, related to the system time by Little's formula, i.e., $n^* = \lambda T^*$. In order for the policy to be stabilizing, the time needed, on average, to service m targets must be no greater than the average time interval in which m new targets are generated. Since there are m vehicles, the average time needed for them to service one target each, in parallel, is no greater than the expected minimum distance (in the Dubins' sense) from an arbitrarily placed vehicle to the closest target; in other words, we can write the stability condition $\mathrm{E}\left[\delta^*(n^*)\right] \leq m/\lambda$. A bound on the expected value of δ^* has been computed in the proof of Theorem 10, yielding

$$\frac{3}{4}\left(\frac{3\rho}{n^*}\right)^{1/3} \leq \mathrm{E}\left[\delta^*(n^*)\right] \leq m\lambda.$$

Using Little's formula $n^* = \lambda T^*$, and rearranging, we get the desired result. ■

Note that the system time depends quadratically on the parameter λ, whereas in the Euclidean case it depends only linearly on it. As a consequence, bounded-curvature constraints make the system time much more sensitive to increases in the target generation rate. However, 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 agents.

D. An upper bound on the system time

A tight upper bound on the DTRP for a Dubins' vehicle is not yet available, as the DTSP problem is still largely unexplored. Very recently, a new algorithm for the stochastic DTSP was discovered [8], which provides a tour with cost $O(n^{2/3}(\log n)^{1/3})$. While this cannot provide a tight bound, but rather an approximation factor growing as $(\log n)^{1/3}$, it guarantees a sub-linear increase of the cost with the number of targets and, as shown in the following, ensures stability of the DTRP for a Dubins' vehicle. Before proceeding, let us describe the algorithm introduced in [8], called the *"Bead Tiling Algorithm."*

1. *The basic geometric construction*

Consider two points $p_- = (-l, 0)$ and $p_+ = (l, 0)$ on the plane, with $l \leq 2\rho$, and construct the region $\mathcal{B}_\rho(l)$ as detailed in Figure 4. In the following, we will refer to such regions as *beads*. The region $\mathcal{B}_\rho(l)$ enjoys the

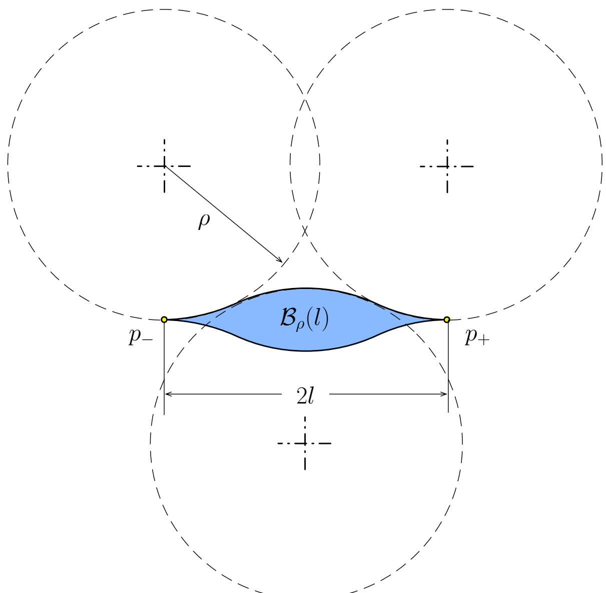

Figure 4. Construction of the "bead" $\mathcal{B}_\rho(l)$. The figure shows how the upper half of the boundary is constructed, the bottom half is symmetric.

following asymptotic properties as the $(l/\rho) \to 0^+$:

(P1) The maximum "thickness" of the region is equal to

$$w(l) = 4\rho\left(1 - \sqrt{1 - \frac{l^2}{4\rho^2}}\right) = \frac{l^2}{2\rho} + o\left(\frac{l^3}{\rho^3}\right).$$

(P2) The area of $\mathcal{B}_\rho(l)$ is equal to

$$\text{Area}[\mathcal{B}_\rho(l)] = l w(l) = \frac{l^3}{2\rho} + o\left(\frac{l^4}{\rho^4}\right).$$

(P3) For any $p \in \mathcal{B}_\rho$, there is at least one Dubins' path γ_p through the points $\{p_-, p, p_+\}$, entirely contained within $\mathcal{B}_\rho$. The length of any such path is at most

$$\text{Length}(\gamma_p) \leq 4\rho \arcsin\left(\frac{l}{2\rho}\right) = 2l + o\left(\frac{l^2}{\rho^2}\right).$$

These properties can be verified using elementary planar geometry.

2. *Periodic tiling of the plane*

An additional property of the geometric shape introduced above is that the plane can be periodically tiled by identical copies of $\mathcal{B}_\rho(l)$, for any $l \in (0, 2\rho]$. (Recall that a tiling of the plane is a collection of set whose intersection has measure zero and whose union covers the plane.)

3. *The* Bead-Tiling Algorithm

We here design an algorithm, that we will call Bead-Tiling Algorithm, that calculates a Dubins' path through a pointset in a rectangular environment $\mathcal{Q}$ of width W and height H (or equivalently, in a general environment contained in a rectangle with the stated dimensions). The basic idea is to exploit an appropriate beads-based tiling and the properties of the beads. First of all, choose the dimension of the beads as follows:

$$l_n = \sqrt[3]{\frac{6\rho W H \log n}{n}}. \tag{17}$$

In what follows we shall tacitly assume that n is sufficiently large so that $l_n \in (0, 2\rho]$.

> Bead-Tiling Algorithm: Given n targets, compute a a periodic tiling of the plane based on bead $\mathcal{B}_\rho(l_n)$ and aligned with the sides of $\mathcal{Q}$ as shown in Figure 5 (the cusps of the beads are aligned horizontally). Next, compute the Dubins tour with the following properties:
>
> 1. it visits all non-empty beads once,
> 2. it visits all rows[a] in sequence top-to-down, alternating between left-to-right and right-to-left passes, and visiting all non-empty beads in a row,
> 3. when visiting a non-empty bead, it services at least one target in it.
>
> Iterate until all targets are visited.

It is a consequence of bead's property (P3) that there exists a Dubins path visiting at least one target in any non-empty bead.

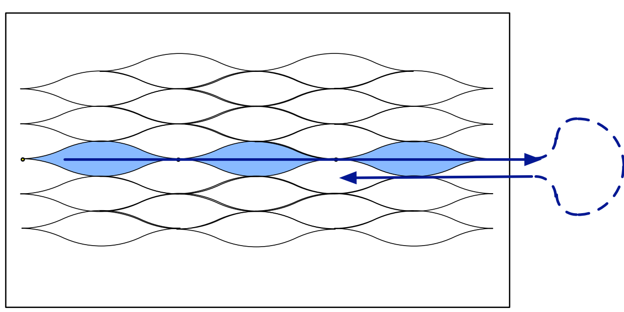

Figure 5. Sketch of the aligned periodic tiling and of the Bead-Tiling Algorithm

The length L_{BTA} of the path generated by the Bead-Tiling Algorithm, provides an upper bound to the length of the optimal Dubins TSP tour, DTSP. The following theorem is proven in [8]:

[a] A row is a maximal string of beads with non-empty intersection with $\mathcal{Q}$.

Theorem 12 (Adapted from [8]) *Let $P \in \mathcal{P}_n$ be uniformly randomly generated in a rectangle of sides W and H. The following inequality holds with high probability:*

$$\lim_{n\to+\infty} \frac{\mathrm{E}\left[\mathrm{DTSP}(P)\right]}{n^{2/3}(\log n)^{1/3}} \le \lim_{n\to+\infty} \frac{\mathrm{E}\left[\mathrm{L}_{\mathrm{BTA},\rho}(P)\right]}{n^{2/3}(\log n)^{1/3}} < 9.88 \sqrt[3]{\rho W H} \left(1 + \frac{7}{3}\pi \frac{\rho}{\max\{W,H\}}\right), \quad w.h.p.$$

Remark 13 *According to the upper bound in Theorem 12, the shape of the environment affects the cost of* DTSP*tours: for equal areas, "thin" environments result in shorter tours. This is in contrast with the corresponding results for the* ETSP*. The intuition is that, in the* Bead-Tiling Algorithm*, a thin environment will require fewer direction reversals, whose cost is directly proportional to ρ.*

4. An upper bound on the system time: the single-vehicle case

The bound derived in Theorem 12 can be directly used to derive a bound on the system time for a single vehicle. A simple strategy, that we call *Receding-Horizon Bead-Tiling Algorithm (RH-BTA)* would entail servicing all the outstading targets at some initial time t_0 using the Bead-Tiling Algorithm, with the following modifications:

1. The size of the beads is chosen in such a way that the number of beads is equal to the number of outstanding targets.
2. Completes a single tour of all the beads.
3. Update the target list and iterate.

Such a strategy lets us state the following upper bound on the system time:

Proposition 14 *The system time for the problem stated in Section II satisfies the following inequality:*

$$\lim_{\lambda\to+\infty} T^* \le 16\rho W H \left(\frac{e}{e-1}\right)\left(1 + \frac{7}{3}\pi\frac{\rho}{\max\{W,H\}}\right)^3 \lambda^2$$

Proof: (Sketch) Let us refer to the times at which the *RH-BTA* is initiated, i.e., the times at which a snapshot of the outstanding targets is taken and used to compute a DTSP tour, as the strictly increasing sequence $\mathcal{T}_{\mathrm{RH-BTA}} = \{t_1, \ldots, t_i, t_{i+1}, \ldots\}$. Note that the time schedule $\mathcal{T}_{\mathrm{RH-BTA}}$ is a sequence of stochastic variables depending on the target generation process. Since each vehicle is guaranteed to be able to service only one target per bead, at the end of a tour, in each region there will be a remainder of "old" unserviced targets—in addition to the targets generated while the vehicle was completing the tour. The number of unserviced old targets is upper bounded by the number of empty beads at the start of the tour: for large $n(t_i)$, the expected number of empty cells is $n(t_i)/e$. In order to show this, let X_i be a stochastic variable taking the value 1 if the i-th cell is empty, and 0 otherwise. If there are n randomly generated targets, and n cells of equal area, the expected value of X_i is

$$\mathrm{E}\left[X_i\right] = \left(1 - \frac{1}{n}\right)^n.$$

Thus, the expected number of empty cells, for large n, satisfies

$$\lim_{n\to\infty} \frac{1}{n}\mathrm{E}\left[\sum_{i=1}^{n} X_i\right] = \lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n}\mathrm{E}\left[X_i\right] = \lim_{n\to\infty}\left(1 - \frac{1}{n}\right)^n = \frac{1}{e}.$$

Informally, assuming that λ and the number of outstanding targets n are large enough, we can write

$$\mathrm{E}\left[n(t_{i+1})\right] \le \lambda\,\mathrm{E}\left[t_{i+1} - t_i\right] + \frac{n(t_i)}{e} \le \lambda\beta_{\mathrm{Pass}}(\rho, W, H) n(t_i)^{2/3} + \frac{n(t_i)}{e},$$

where

$$\beta_{\mathrm{Pass}} := \sqrt[3]{16\rho W H}\left(1+\frac{7}{3}\pi\frac{\rho}{\max\{W,H\}}\right)$$

has been derived in [8].

The above equation describes a dynamical system, whose equilibrium points satisfy the inequality

$$\left(1-\frac{1}{e}\right)n^*(\lambda) \leq \lambda\beta_{\mathrm{Pass}}(\rho,W,H)n^*(\lambda)^{2/3}.$$

Note that the above bound describes a stable equilibrium point, as for all values of $n(t_i)$ not satisfying the above inequality, the expected value of $n(t_{i+1})$ will be strictly smaller than $n(t_i)$. Rearranging, and using Little's formula, we get the desired result. ■

5. An upper bound on the system time: the multiple-vehicle case

Let us turn now to the multiple-vehicle case. Note that the system time depends on (i) the area of the region assigned to a vehicle, and (ii) on the shape of the region. In particular, the system time is minimized, for a given area, when one of its dimensions is maximized. This suggest the following strategy for partitioning the environment $\mathcal{Q}$, that we call the *Strip-tiling* strategy. Let $\mathcal{Q}'$ be the smallest rectangle that contains $\mathcal{Q}$; and let W', H' be its width and height, with $W' \geq H'$. Partition $\mathcal{Q}$ into m strips of width W' and height H'/m, and assign each strip to a vehicle. Let each vehicle execute the *RH-BTA* strategy inside the assigned region. Then, the following holds:

Proposition 15 *The system time for the problem stated in Section II satisfies the following inequality:*

$$\lim_{\lambda\to+\infty} T^* \leq 16\rho W' H' \left(\frac{e}{e-1}\right)^3 \left(1+\frac{7}{3}\pi\frac{\rho}{W'}\right)^3 \frac{\lambda^2}{m^3}$$

Note that the achievable performance of the *Strip-tiling* algorithm provides a constant-factor approximation to the lower bound established in Theorem 11. The factor is approximately 63.3: the ratio between upper bound and lower bound is consequently still significant. We believe that the lower bound is exceedingly optimistic: the large value of the approximation factor may depend from the lack of a tight lower bound. On the other hand, the *RH-BTA* algorithm is the first polynomial-time algorithm to provide such a guarantee.

V. Conclusions

In this paper, we have considered the problem of steering a number of UAVs, modeled as Dubins vehicles, in order to minimize the expected waiting time between the appearance of randomly-generated targets and the time they are visited by one of the UAVs. We have proposed control policies that achieve a system time that is provably within a constant additive factor from the optimal, in the light load case. In the heavy load case, we have developed a lower bound on the system time, showing that it depends at least quadratically on the target generation rate, and linearly on the minimum turn radius, and decreases with the cube of the number of vehicles. An algorithm was provided, providing a constant factor approximation to the optimal solution. Future research directions will include the search for tighter bounds, and the extension of the present work to the fully decentralized control case.

Acknowledgments

The research leading to this document was partially supported by AFOSR, through grant F49620-02-1-0325. 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 organization.

References

[1]Schumacher, 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.

[2]Beard, 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.

[3]Richards, 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.

[4]Gil, A. E., Passino, K. M., and Sparks, A., "Cooperative Scheduling of Tasks for Networked Uninhabited Autonomous Vehicles," *Proc. IEEE Conf. on Decision and Control*, Maui, Hawaii, 2003, pp. 522–527.

[5]Li, W. and Cassandras, C. G., "Stability Properties of a Cooperative Receding Horizon Controller," *Proc. IEEE Conf. on Decision and Control*, Maui, Hawaii, 2003, pp. 492–497.

[6]Frazzoli, E. and Bullo, F., "Decentralized Algorithms for Vehicle Routing in a Stochastic Time-Varying Environment," *Proc. IEEE Conf. on Decision and Control*, Paradise Island, Bahamas, December 2004.

[7]Enright, J. J. and Frazzoli, E., "UAV Routing in a Stochastic, Time-Varying Environment," *Proc. of the IFAC World Congress*, Prague, Czech Republic, July 2005.

[8]Savla, K., Bullo, F., and Frazzoli, E., "The stochastic travelling salesperson problem for Dubins' vehicle," *Proc. IEEE Conf. on Decision and Control*, 2005, To appear.

[9]Bertsimas, 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.

[10]Bertsimas, 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.

[11]Dubins, 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.

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

[13]Shkel, A. M. and Lumelsky, V. J., "Classification of the Dubins Set," *Robotics and Autonomous Systems*, Vol. 23, 2001, pp. 179–202.

[14]Thomaschewski, B., "Dubins' problem for free terminal direction," Preprint M09/2001, Technical University of Ilmenau, Ilmenau, Germany, 2001.

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

[16]Agarwal, P. K. and Sharir, M., "Efficient algorithms for geometric optimization," *ACM Computing Surveys*, Vol. 30, No. 4, 1998, pp. 412–458.

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

[18]Megiddo, 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.

[19]Cortés, 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.

[20]Beardwood, J., Halton, J., and Hammersley, J., "The Shortest Path Through Many Points," *Proceedings of the Cambridge Philoshopy Society*, Vol. 55, 1959, pp. 299–327.

[21]Percus, G. and Martin, O. C., "Finite size and dimensional dependence of the Euclidean traveling salesman problem," *Physical Review Letters*, Vol. 76, No. 8, 1996, pp. 1188–1191.

[22]Johnson, D. S., McGeoch, L. A., and Rothberg, E. E., "Asymptotic Experimental Analysis for the Held-Karp Traveling Salesman Bound," *Proc. 7th Annual ACM-SIAM Symposium on Discrete Algorithms*, 1996, pp. 341–350.

[23]Savla, K., Frazzoli, E., and Bullo, F., "On the point-to-point and traveling salesperson problems for Dubins' vehicles," *Proc. of the American Control Conference*, 2005.