

On Traveling Salesperson Problems for Dubins' vehicle: stochastic and dynamic environments

Ketan Savla

Francesco Bullo

Emilio Frazzoli

Abstract—In this paper we propose some novel planning and routing strategies for Dubins' vehicle, i.e., for a nonholonomic vehicle moving along paths with bounded curvature, without reversing direction. First, we study a stochastic version of the Traveling Salesperson Problem (TSP): given n targets randomly sampled from a uniform distribution in a rectangle, what is the shortest Dubins' tour through the targets and what is its length? We show that the expected length of such a tour is $\Omega(n^{2/3})$ and we propose a novel algorithm that generates a tour of length $O(n^{2/3} \log(n)^{1/3})$ with high probability. Second, we study a dynamic version of the TSP (known as "Dynamic Traveling Repairperson Problem" in the Operations Research literature): given a stochastic process that generates targets, is there a policy that allows a Dubins vehicle to stabilize the system, in the sense that the number of unvisited targets does not diverge over time? If such policies exist, what is the minimum expected waiting period between the time a target is generated and the time it is visited? We propose a novel receding-horizon algorithm whose performance is *almost* within a constant factor from the optimum.

I. INTRODUCTION

The Traveling Salesperson Problem (TSP) with its variations is one of the most widely known combinatorial optimization problems. While extensively studied in the literature, these problems continue to attract great interest from a wide range of fields, including Operations Research, Mathematics and Computer Science. The Euclidean TSP (ETSP) [1], [2] is formulated as follows: given a finite point set P in \mathbb{R}^2 , find the minimum-length tour of P . It is quite natural to formulate this problem in context of Dubins' vehicle, i.e., a nonholonomic vehicle that is constrained to move along paths of bounded curvature, without reversing direction.

The focus of this paper is the analysis of the TSP for Dubins' vehicle; we shall refer to it as DTSP. Exact algorithms, heuristics as well as polynomial-time constant factor approximation algorithms are available for the Euclidean TSP, see [3], [4], [5]. It is known that non-metric versions of the TSP are, in general, not approximable in polynomial time [6]. Furthermore, unlike most other variations of the TSP, it is believed that the DTSP cannot be formulated as a problem on a finite-dimensional graph, thus preventing the use of well-established tools in combinatorial optimization. On the other hand, it is reasonable to believe that exploiting the geometric structure of Dubins' paths one can gain insight into

the nature of the solution, and possibly provide polynomial-time approximation algorithms.

A fairly complete picture is available for the minimum-time point-to-point path planning problem for Dubins' vehicle, see [7] and [8]. However, the DTSP seems not to have been studied that extensively. In [9], we provided some results for the worst case tours of DTSP. A lower bound on the expected cost of a stochastic DTSP visiting randomly generated points was provided in [10]. Here, we shall specifically concentrate on the case when the target points in the environment are generated stochastically according to a uniform probability distribution function. We shall refer to such a problem as stochastic DTSP.

The motivation to study the DTSP arises in robotics and uninhabited aerial vehicles (UAVs) applications, e.g., see [11], [12], [13], [14]. In particular, we envision applying our algorithm to the setting of an UAV monitoring a collection of spatially distributed points of interest. Additionally, from a purely scientific viewpoint, it appears to be of general interest to bring together the work on Dubins' vehicle and that on TSP. UAV applications also motivate us to study the Dynamic Traveling Repairperson Problem (DTRP), in which the aerial vehicle is required to visit a dynamically changing set of targets. This problem was introduced by Bertsimas and van Ryzin in [15] and then decentralized policies achieving the same performances were proposed in [11]. However, as with the TSP, the study of DTRP in context of Dubins' vehicle has eluded attention from the research community.

The contributions of this paper are threefold. First, we propose an algorithm for the stochastic DTSP through a pointset P , called the BEAD-TILING ALGORITHM, based on a smart tiling of the plane, and a strategy for the Dubins' vehicle to service targets from each tile. Second, we obtain an upper bound on the stochastic performance of the proposed algorithm and thus also establish a similar bound on the stochastic DTSP. The upper bound on the performance of BEAD-TILING ALGORITHM belongs to $O(n^{2/3} \log(n)^{1/3})$ whereas we know the lower bound on the achievable performance belongs to $\Omega(n^{2/3})$. Third, we propose an algorithm for DTRP in the heavy load case, called the RECEDING HORIZON BEAD-TILING ALGORITHM, based on a receding horizon version of the BEAD-TILING ALGORITHM. We show that the performance guarantees for the stochastic DTSP translate into stability guarantees for the average performance of the DTRP problem for Dubins' vehicle in heavy load case. Specifically, we show that the performance of RECEDING HORIZON BEAD-TILING ALGORITHM is *almost* within a constant factor of the optimal policy. We contend

Ketan Savla and Francesco Bullo are with the Center for Control, Dynamical Systems and Computation, University of California at Santa Barbara, ketansavla@umail.ucsb.edu, bullo@engineering.ucsb.edu

Emilio Frazzoli is with the Mechanical and Aerospace Engineering Department, University of California at Los Angeles, frazzoli@ucla.edu

that the successful application to the DTRP does indeed demonstrate the significance of the DTSP problem from a control viewpoint.

The paper is organized as follows. In the remainder of the Introduction we establish some basis useful notation. In Section II we review our results on the worst-case Dubins' TSP. In Section III we present the main results of this paper: (i) a novel DTSP algorithm based on a periodic tiling, and (ii) an upper bound on its performance in the stochastic setting. Numerical results are also included. In Section IV we consider the DTRP for Dubins' vehicle and we propose a receding horizon control policy for the heavy load case. Concluding remarks are presented in Section V.

Notation

Here we collect some concepts that will be required in the later sections. A *Dubins' vehicle* is a planar vehicle that is constrained to move along paths of bounded curvature, without reversing direction and maintaining a constant speed. Accordingly, we define a *feasible curve for Dubins' vehicle* or a *Dubins' path*, as a curve $\gamma : [0, T] \rightarrow \mathbb{R}^2$ that is twice differentiable almost everywhere, and such that the magnitude of its curvature is bounded above by $1/\rho$, where $\rho > 0$ is the minimum turn radius. We represent the vehicle *configuration* by the triplet $(x, y, \psi) \in SE(2)$, where (x, y) are the Cartesian coordinates of the vehicle, and ψ is its heading, i.e., $\psi = \text{atan2}(y, x)$ (where atan2 is the four-quadrant version of the arc tangent function).

Let $P = \{p_1, \dots, p_n\}$ be a set of n points in a compact region $\mathcal{Q} \subset \mathbb{R}^2$ and \mathcal{P}_n be the collection of all pointsets $P \subset \mathcal{Q}$ with cardinality n . Let $\text{ETSP}(P)$ denote the cost of the Euclidean TSP over P , i.e., the length of the shortest closed path through all points in P . Correspondingly, let $\text{DTSP}_\rho(P)$ denote the cost of the Dubins' TSP over P , i.e., the length of the shortest closed Dubins' path through all points in P . In what follows, $\rho \in \mathbb{R}_+$ is take constant, and we study the dependence of $\text{DTSP}_\rho : \mathcal{P}_n \rightarrow \mathbb{R}_+$ on n .

For $f, g : \mathbb{N} \rightarrow \mathbb{R}$, we say that $f \in O(g)$ (respectively, $f \in \Omega(g)$) if there exist $N_0 \in \mathbb{N}$ and $k \in \mathbb{R}_+$ such that $|f(N)| \leq k|g(N)|$ for all $N \geq N_0$ (respectively, $|f(N)| \geq k|g(N)|$ for all $N \geq N_0$). If $f \in O(g)$ and $f \in \Omega(g)$, then we use the notation $f \in \Theta(g)$.

II. THE WORST-CASE DTSP

In this section, we review some of our results from [9] where we proposed a simple algorithm, the ALTERNATING ALGORITHM, that gives a sub-optimal tour for the traveling salesperson problem for Dubins' vehicle. We also established a measure of its performance in the worst-case, and of the worst-case cost of the DTSP.

1) *Description of the Algorithm:* The ALTERNATING ALGORITHM works on the following principle: since the optimal path between two configurations of a Dubins' vehicle has been completely characterized in [7], a solution for the Dubins' TSP consists of (i) determining the order in which the Dubins' vehicle visits the given set of points, and (ii) assigning headings for the Dubins' vehicle at the points.

Let $A = (a_1, \dots, a_n)$ be an ordered set of points that is a permutation of P . Let $\Psi = \{\psi_1, \dots, \psi_n\}$ be a set of heading of the Dubins' vehicle at the n points a_1, \dots, a_n . Therefore the configuration of Dubins' vehicle at a_i is (x_i, y_i, ψ_i) where (x_i, y_i) are the coordinates of a_i , for $i = 1, \dots, n$.

Here is an informal description of ALTERNATING ALGORITHM over P . Compute an optimal ETSP tour of P and label the edges on the tour in order with consecutive integers. A DTSP tour can be constructed by retaining all odd-numbered edges (except the n th one), and replacing all even-numbered edges with minimum-length Dubins' paths preserving the point ordering. The algorithm is formally stated in Table I.

TABLE I
THE ALTERNATING ALGORITHM

Name:	ALTERNATING ALGORITHM
Goal:	To determine an ordering A and a set of headings Ψ for the DTSP through P
Requires:	An algorithm ETSP-ALGO to compute the optimal ETSP ordering of a pointset
<pre> 1: set $A := \text{ETSP-ALGO}(P)$ 2: set $\psi_1 :=$ orientation of segment from a_1 to a_2 3: for $i = 2$ to $n - 1$ do 4: if i is even then 5: set $\psi_i := \psi_{i-1}$ 6: else 7: set $\psi_i :=$ orientation of segment from a_i to a_{i+1} 8: end if 9: end for 10: if n is even then 11: set $\psi_n := \psi_{n-1}$ 12: else 13: set $\psi_n :=$ orientation of segment from a_n to a_1 14: end if </pre>	

2) *Performance of the algorithm:* We now state two results, proved in [9], that characterize the worst-case performance of the ALTERNATING ALGORITHM. Let $L_{AA, \rho}(P)$ be the length of the closed path over P as given by the ALTERNATING ALGORITHM.

Theorem 2.1: (Worst-case performance of the ALTERNATING ALGORITHM) For $n \geq 2$, $\rho > 0$, and $P \in \mathcal{P}_n$,

$$\text{DTSP}_\rho(P) \leq L_{AA, \rho}(P) \leq \text{ETSP}(P) + \kappa\pi\rho \left\lceil \frac{n}{2} \right\rceil,$$

where $\kappa \approx 2.6575$.

From the clear bound $\text{ETSP}(P) \leq \text{DTSP}_\rho(P)$, it follows that the ALTERNATING ALGORITHM provides an $O(n)$ approximation to the DTSP in the general case. Furthermore, the ALTERNATING ALGORITHM provides a constant-factor approximation to large worst-case DTSPs:

Theorem 2.2: For $n \geq 2$ and $\rho > 0$,

$$\begin{aligned} \sup_{P \in \mathcal{P}_n} \text{DTSP}_\rho(P) &\leq \sup_{P \in \mathcal{P}_n} L_{\text{AA},\rho}(P) \\ &\leq \frac{\text{ETSP}(P) + \kappa \lfloor n/2 \rfloor \pi \rho}{\text{ETSP}(P) + 2 \lfloor n/2 \rfloor \pi \rho} \sup_{P \in \mathcal{P}_n} \text{DTSP}_\rho(P). \end{aligned}$$

Furthermore, as $n \rightarrow +\infty$,

$$\sup_{P \in \mathcal{P}_n} \text{DTSP}_\rho(P) \leq \sup_{P \in \mathcal{P}_n} L_{\text{AA},\rho}(P) \leq \frac{\kappa}{2} \sup_{P \in \mathcal{P}_n} \text{DTSP}_\rho(P).$$

III. THE STOCHASTIC DTSP

The discussion in the previous section showed that a simple algorithm, the ALTERNATING ALGORITHM, performs well when the points to be visited by the tour are chosen in an adversarial manner. However, it is reasonable to argue that this algorithm might not perform very well when dealing with a random distribution of the target points. In particular, one can expect that when n points are chosen randomly, the cost of the DTSP increases sub-linearly with n , i.e., that the average length of the path between two points decreases as n increases. In this section, we consider the scenario when n target points are stochastically generated in \mathcal{Q} according to a uniform probability distribution function. We present a novel algorithm, the BEAD-TILING ALGORITHM, to service these points and then establish bounds on its performance.

We assume that the environment \mathcal{Q} is a rectangle of width W and height H ; different choices for the shape of \mathcal{Q} affect our conclusions only by a constant. In what follows we select a reference frame whose two axes are parallel to the sides of \mathcal{Q} . Let n target points be generated stochastically according to uniform distribution in the region \mathcal{Q} . Let $\mathcal{P} = (p_1, \dots, p_n)$ be the locations of these target points.

A. A lower bound

First, we summarize a result from [10], that provides a lower bound on the expected length of the stochastic DTSP.

Theorem 3.1: (Lower bound on stochastic DTSP) For all $\rho > 0$, the expected cost of a stochastic DTSP visiting a set P of n uniformly-randomly-generated points in a rectangle of width W and height H satisfies the following inequality:

$$\lim_{n \rightarrow +\infty} \frac{\mathbb{E}[\text{DTSP}_\rho(P)]}{n^{2/3}} \geq \frac{3}{4} (3\rho W H)^{1/3}.$$

B. A constructive upper bound

In this section, we design a novel algorithm that computes a Dubins' path through a pointset in the square \mathcal{Q} . We will show that the proposed algorithm provides a $O(\log(n)^{1/3})$ approximation to the optimal DTSP with high probability. We start by describing some useful geometric objects.

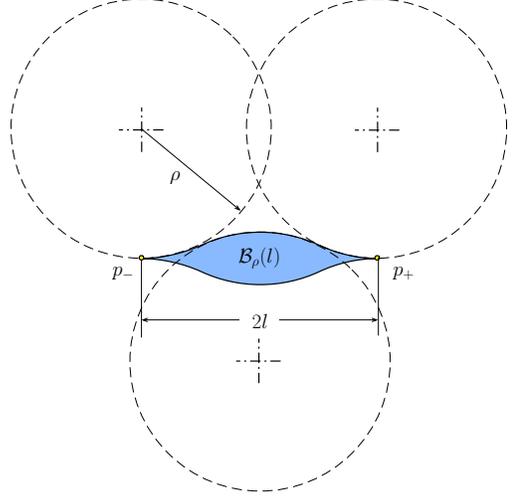


Fig. 1. Construction of the “bead” $\mathcal{B}_\rho(l)$. The figure shows how the upper half of the boundary is constructed, the bottom half is symmetric.

1) *The basic geometric construction:* Consider two points $p_- = (-l, 0)$ and $p_+ = (l, 0)$ on the plane, with $l \leq \rho$, and construct the region $\mathcal{B}_\rho(l)$ as detailed in Figure 1. In the following, we will refer to such regions as *beads*. The region $\mathcal{B}_\rho(l)$ enjoys the following asymptotic properties as the $(l/\rho) \rightarrow 0^+$:

(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

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

(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}_ρ . 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 facts are 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 $l \in (0, \rho]$. (Recall that a tiling of the plane is a collection of set whose intersection has measure zero and whose union covers the plane.) This tiling has the following critical property, adapted from [16].

Proposition 3.2: Given the number n of uniformly-randomly-generated points 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), let

$$l_n = \sqrt[3]{\frac{6\rho W H \log n}{n}}. \quad (1)$$

Then, the maximum number of targets in any single bead $\mathcal{B}_\rho(l_n)$ is $3e \log n$ with high probability.

3) *The BEAD-TILING ALGORITHM* : We here design an algorithm, that we will call the BEAD-TILING ALGORITHM, that calculates a Dubins' path through a pointset in the rectangle \mathcal{Q} . The basic idea is to exploit an appropriate beads-based tiling and the properties of the beads. In what follows we shall tacitly assume that n is sufficiently large so that $l_n \in (0, \rho]$.

BEAD-TILING ALGORITHM: Given n targets, compute 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 2 (the cusps of the beads are aligned with the longer side). Next, compute the Dubins' tour with the following properties:

- 1) it visits all non-empty beads once,
- 2) it visits all rows¹ 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.

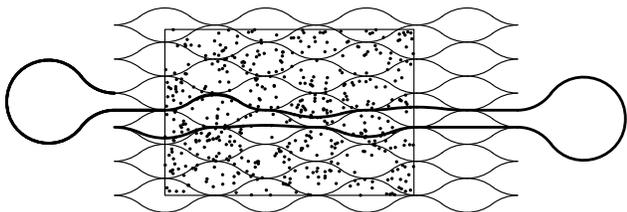


Fig. 2. Sketch of the aligned periodic tiling and of the BEAD-TILING ALGORITHM

Next, we let $L_{\text{BTA},\rho}(P)$ denote the length of the tour designed by the BEAD-TILING ALGORITHM through P with a minimum turn radius ρ . To characterize this length, we start by studying the path length needed to visit all non-empty beads once.

Lemma 3.3: Consider a pointset $P \in \mathcal{P}_n$ and a periodic tiling of the plane into beads equal to $\mathcal{B}_\rho(l_n)$. Take a pointset $\tilde{P} \subset P$ such that each bead, with a nonempty intersection with \mathcal{Q} , contains at most one point. Then, as $n \rightarrow +\infty$ and as $\rho \rightarrow +\infty$,

$$\text{DTSP}_\rho(\tilde{P}) = O\left(\rho^{4/3} \left(\frac{n}{\log n}\right)^{2/3}\right).$$

Proof: Let us first compute the length of a pass, in either direction. The number of beads traversed will be no more than

$$\left\lceil \frac{\max\{W, H\}}{2l_n} \right\rceil = \left\lceil c_1 \left(\frac{n}{\rho \log n}\right)^{1/3} \right\rceil,$$

¹Here, by row we mean a maximal string of beads with non-empty intersection with \mathcal{Q} .

where $c_1 = \frac{\max\{W, H\}}{2\sqrt[3]{6WH}}$ is a constant. Hence, the total path length per pass will be bounded by:

$$L_{\text{pass}} \leq \max\{W, H\} + 2l_n + o\left(\frac{l_n^2}{\rho^2}\right),$$

as $(l_n/\rho) \rightarrow 0^+$. Applying a result from [9], the cost of a u-turn, i.e., the length of the path needed to reverse direction and move to the next row of beads, is bounded by

$$L_{\text{u-turn}} \leq \frac{7}{3}\pi\rho + \frac{w(l_n)}{2} = \frac{7}{3}\pi\rho + \frac{l_n^2}{4\rho} + o\left(\frac{l_n^3}{\rho^3}\right).$$

The total number of passes will be at most

$$N_{\text{pass}} = \left\lceil \frac{2 \min\{W, H\}}{w(l_n)} \right\rceil \leq \frac{2 \min\{W, H\}}{l_n^2/(2\rho) + o(l_n^3/\rho^3)} + 1.$$

The cost of closing the tour is bounded by a constant, say

$$L_{\text{closure}} \leq \min\{W, H\} + \frac{w(l_n)}{2} + \frac{7}{3}\pi\rho.$$

In summary, the total path length will be bounded by

$$\text{DTSP}_\rho(\tilde{P}) = N_{\text{pass}}(L_{\text{pass}} + L_{\text{u-turn}}) + L_{\text{closure}}.$$

Neglecting higher-order terms, this can be simplified to

$$\begin{aligned} \text{DTSP}_\rho(\tilde{P}) \approx & \frac{4\rho WH}{l_n^2} + W + H + \frac{14\pi\rho}{3} \\ & + \min\{W, H\} \left(1 + \frac{8\rho}{l_n} + \frac{28\pi\rho^2}{3l_n^2}\right). \end{aligned}$$

Recalling our selection of $l_n = \sqrt[3]{6\rho WH \log n/n}$ from (1), we obtain the desired result. ■

Based on the results obtained so far, we are now ready to state an upper bound on the length of the path traveled by Dubins' vehicle to service all the targets while executing the BEAD-TILING ALGORITHM.

Theorem 3.4: (*Upper bound on the length of the total path*) Let $P \in \mathcal{P}_n$ be uniformly randomly generated in a rectangle. For all $\rho > 0$, there exists $\delta > 0$ such that the following inequality holds with high probability:

$$\lim_{n \rightarrow +\infty} \frac{\mathbb{E}[L_{\text{BTA},\rho}(P)]}{n^{2/3} \log(n)^{1/3}} < \delta, \quad \text{w.h.p.}$$

Proof: By Proposition 3.2 we know that each bead contains at most order $\log(n)$ targets. Hence, at most order $\log(n)$ tours through each bead are necessary. The proof follows from the upper bound in Lemma 3.3. ■

C. Simulations

In this section we present the results of the BEAD-TILING ALGORITHM and the ALTERNATING ALGORITHM. We summarize the result in Figure 3. The points are stochastically generated according to a uniform distribution in a square with $A = 25$. The minimum turning radius for the Dubins' vehicle, i.e., $\rho = 1$. Each data point in the upper sequence of points in the logarithmic plot in Figure 3 represents the mean of lengths of Dubins' path as given by the BEAD-TILING ALGORITHM, taken over 10 instances of the experiment for the corresponding value of n on a logarithmic scale,

whereas each data point in the lower sequence of points represents the corresponding quantity for the ALTERNATING ALGORITHM. The solid curve in the plot represents the function $\log(\beta_1 n^{2/3} \log(n)^{1/3})$, for $\beta_1 \approx 115$. The dashed curve in the plot represents the function $\log(\beta_2 n)$, for $\beta_2 \approx 3.5$. The fact that all the dots for BEAD-TILING ALGORITHM lie below the solid line is consistent with our results for the BEAD-TILING ALGORITHM. The nature of these two curves indicates that for high values of n , the BEAD-TILING ALGORITHM will outperform the ALTERNATING ALGORITHM. This is consistent with our asymptotic characterizations of the two algorithms.

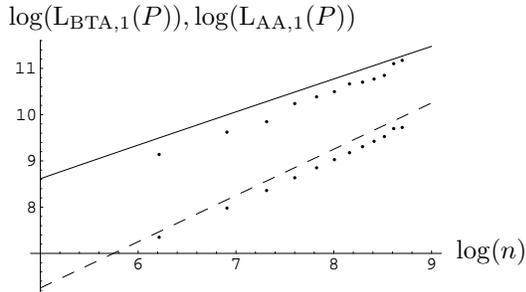


Fig. 3. Numerical experimental results of the BEAD-TILING ALGORITHM and the ALTERNATING ALGORITHM. The solid and dashed curves are the functions $\log(\beta_1 n^{2/3} \log(n)^{1/3})$, for $\beta_1 \approx 115$, and $\log(\beta_2 n)$, for $\beta_2 \approx 3.5$, respectively. The upper and lower sequence of points are the average $L_{BTA,1}(P)$ and the average $L_{AA,1}(P)$ over 10 random instances of $P \in \mathcal{P}_n$, respectively.

IV. THE DTRP FOR DUBINS' VEHICLE

We now turn our attention to a related problem which is known as the Dynamic Traveling Repairperson Problem (DTRP), and was introduced by Bertsimas and van Ryzin in [15]. Our problem is different from the single-vehicle-DTRP in [15] since we consider here a *Dubins'* vehicle for targets servicing task, i.e., we impose the same non-holonomic constraint on the vehicle dynamics that we have been considering so far in this paper.

A. Model and problem statement

In this subsection we describe in some detail the vehicle and sensing model and the DTRP definition. The key aspect of the DTRP is that the aerial vehicle is required to visit a dynamically growing set of targets, generated by some stochastic process. We assume that the Dubins' vehicle has unlimited range and target-servicing capacity. To simplify notations, we also assume that the Dubins' vehicle moves constantly at a unit speed.

Information on the outstanding targets – the demand – at time t is summarized as a finite set of target positions $D(t) \subset \mathcal{Q}$, with $n(t) := \text{card}(D(t))$. 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

$$E[\text{card}(D(t') \cap \mathcal{S}) - \text{card}(D(t) \cap \mathcal{S})] = \lambda(t' - t) \text{Area}(\mathcal{S}).$$

(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 Dubins' vehicle is a map $\Phi : \text{SE}(2) \times 2^{\mathcal{Q}} \rightarrow [-1/\rho, 1/\rho]$, assigning a control input to each vehicle, as a function of the current state of the system. 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 $\mathcal{T}_\Phi = \{t_1, t_2, \dots, t_i, \dots\}$ be a strictly increasing sequence of times at which such computations are started: with some abuse of terminology, we will say that Φ is a *receding horizon strategy* if it is based on the most recent target data available - $D_{\text{th}}(t)$, with

$$D_{\text{th}}(t) = D(\max\{t_{\text{th}} \in \mathcal{T}_\Phi : t_{\text{th}} < t\}).$$

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

$$n_\Phi := \lim_{t \rightarrow +\infty} E[n(t) | \dot{p} = \Phi(p, D_{\text{th}})] < +\infty,$$

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 [17])

$$n_\Phi = \lambda T_\Phi,$$

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

$$T^* = \inf_{\Phi} T_\Phi.$$

B. Lower and constructive upper bounds

In what follows, we are interested in designing a control policy that provide a constant-factor approximation of the optimal achievable performance. Consistent with the theme of the paper, we shall consider the case of *heavy load*, i.e., the problem as $\lambda \rightarrow +\infty$. We shall review a known lower bound for the system time, and present a novel approximation algorithm providing an upper bound on the performance that holds with high probability.

We start by summarizing a result from [10], that provides a lower bound on the system time for any policy in the heavy load case.

Theorem 4.1: The system time T^* for the DTRP problem, satisfies the following lower bound for the heavy load case:

$$\lim_{\lambda \rightarrow +\infty} \frac{T^*}{\lambda^2} \geq \frac{81}{64} \rho W H.$$

Note that the system time depends quadratically on the parameter λ , whereas in the Euclidean case it depends only linearly on it.

The bound derived in Theorem 3.4 can be directly used to derive a constructive upper bound on the system time. We propose a simple strategy, that we call the RECEDING HORIZON BEAD-TILING ALGORITHM (RH-BTA), based on an iterative invocation the BEAD-TILING ALGORITHM. The strategy consists of the following two steps:

- 1) at time t_0 , execute the BEAD-TILING ALGORITHM for all the outstanding targets, and
- 2) update the target list and iterate.

Theorem 4.2: The RECEDING HORIZON BEAD-TILING ALGORITHM is a stable policy for the stochastic DTRP problem in heavy load. The performance of the RECEDING HORIZON BEAD-TILING ALGORITHM provides the following upper bound on the system time: for any $\epsilon > 0$,

$$\lim_{\lambda \rightarrow +\infty} \frac{T^*}{\lambda^{2+\epsilon}} \leq 9.88^3 \rho W H \left(1 + \frac{7}{3} \frac{\pi \rho}{\max\{W, H\}} \right)^3.$$

Note that the achievable performance of the RECEDING HORIZON BEAD-TILING ALGORITHM provides an *almost* constant-factor approximation to the lower bound established in Theorem 4.1 in the sense that the exponent of λ in the last equation can be selected arbitrarily close to 2. The ratio between the constants of the upper bound and lower bound is still significant. We believe that the lower bound is exceedingly optimistic: the large value of the approximation factor may be due to 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.

Finally, note that there exists no stable policy for the DTRP when the targets are generated in an adversarial worst-case fashion with high intensity. This fact is a consequence of the linear lower bound on the worst-case DTSP in Theorem 2.2.

V. CONCLUSIONS

Here and in the companion paper [9], we have studied the TSP problem for vehicles that follow paths of bounded curvature in the plane. We have obtained lower and upper bounds in the worst-case and stochastic settings; the upper bounds are constructive in the sense that they are achieved by two novel algorithms. It is interesting to compare our results with the Euclidean setting (i.e., the setting in which curves do not have curvature constraints). For a given compact set and a pointset P of n points, it is known [1], [2] that the ETSP(P) belongs to $\Theta(\sqrt{n})$. This is true for both stochastic and worst-case settings. In this paper, we showed that, given a fixed $\rho > 0$, the *stochastic* DTSP $_{\rho}(P)$ belongs to $\Omega(n^{2/3})$ and to $O(n^{2/3} \log(n)^{1/3})$. In the companion paper [9], we have showed that the *worst-case* DTSP $_{\rho}(P)$ belongs to $\Theta(n)$.

Remarkably, the differences between these various bounds play a crucial role when studying the DTRP problem; e.g., stable policies exist only when the TSP cost grows strictly sub-linearly with n . For the DTRP problem we have proposed the novel receding-horizon policy RH-BTA and shown its stability for a uniform target-generation process with intensity λ . Based on this policy, we have shown that the system time for the DTRP problem for Dubins' vehicle belongs to $\Omega(\lambda^2)$ and $O(\lambda^{2+\epsilon})$ for any $\epsilon > 0$. This

result differs from the result in the Euclidean case, where it is known that the system time belongs to $\Theta(\lambda)$. As a consequence, bounded-curvature constraints make the system much more sensitive to increases in the target generation rate.

In the future, we plan to study centralized and decentralized versions of the DTRP and general task assignment and surveillance problems for various non-holonomic vehicles.

ACKNOWLEDGEMENT

This material is based upon work supported in part by ONR YIP Award N00014-03-1-0512 and AFOSR MURI Award F49620-02-1-0325. The authors would like to thank John J. Enright for helpful discussions.

REFERENCES

- [1] J. Beardwood, J. Halton, and J. Hammersly, "The shortest path through many points," in *Proceedings of the Cambridge Philosophy Society*, vol. 55, pp. 299–327, 1959.
- [2] J. M. Steele, "Probabilistic and worst case analyses of classical problems of combinatorial optimization in Euclidean space," *Mathematics of Operations Research*, vol. 15, no. 4, p. 749, 1990.
- [3] D. Applegate, R. Bixby, V. Chvátal, and W. Cook, "On the solution of traveling salesman problems," in *Documenta Mathematica, Journal der Deutschen Mathematiker-Vereinigung*, (Berlin, Germany), pp. 645–656, Aug. 1998. Proceedings of the International Congress of Mathematicians, Extra Volume ICM III.
- [4] S. Arora, "Nearly linear time approximation scheme for Euclidean TSP and other geometric problems," in *Proc. 38th IEEE Annual Symposium on Foundations of Computer Science*, (Miami Beach, FL), pp. 554–563, Oct. 1997.
- [5] S. Lin and B. W. Kernighan, "An effective heuristic algorithm for the traveling-salesman problem," *Operations Research*, vol. 21, pp. 498–516, 1973.
- [6] S. Sahni and T. Gonzalez, "P-complete approximation problems," *Journal of the Association of Computing Machinery*, vol. 23, no. 3, pp. 555–565, 1976.
- [7] L. E. Dubins, "On curves of minimal length with a constraint on average curvature and with prescribed initial and terminal positions and tangents," *American Journal of Mathematics*, vol. 79, pp. 497–516, 1957.
- [8] A. M. Shkel and V. J. Lumelsky, "Classification of the Dubins set," *Robotics and Autonomous Systems*, vol. 34, pp. 179–202, 2001.
- [9] K. Savla, E. Frazzoli, and F. Bullo, "On the point-to-point and traveling salesperson problems for Dubins' vehicle," in *American Control Conference*, (Portland, OR), June 2005. 786-791.
- [10] J. J. Enright and E. Frazzoli, "UAV routing in a stochastic time-varying environment," in *IFAC World Congress*, (Prague, Czech Republic), July 2005. Electronic Proceedings.
- [11] E. Frazzoli and F. Bullo, "Decentralized algorithms for vehicle routing in a stochastic time-varying environment," in *IEEE Conf. on Decision and Control*, (Paradise Island, Bahamas), pp. 3357–3363, Dec. 2004.
- [12] Z. Tang and Ü. Özgüner, "Motion planning for multi-target surveillance with mobile sensor agents," *IEEE Transactions on Robotics*, Jan. 2005. To appear.
- [13] R. W. Beard, T. W. McLain, M. A. Goodrich, and E. P. Anderson, "Coordinated target assignment and intercept for unmanned air vehicles," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 9, pp. 911–922, 2002.
- [14] S. Darbha, "Combinatorial motion planning for a collection of Reeds-Shepp vehicles," tech. rep., ASEE/AFOSR SFFP, AFRL, Eglin, Aug. 2005.
- [15] D. J. Bertsimas and G. J. van Ryzin, "A stochastic and dynamic vehicle routing problem in the Euclidean plane," *Operations Research*, vol. 39, pp. 601–615, 1991.
- [16] F. Xue and P. R. Kumar, "The number of neighbors needed for connectivity of wireless networks," *Wireless Networks*, vol. 10, no. 2, pp. 169–181, 2004.
- [17] R. C. Larson and A. R. Odoni, *Urban Operations Research*. Englewood Cliffs, NJ: Prentice Hall, 1981.