

Traveling Salesperson Problems for the Dubins vehicle

Ketan Savla Emilio Frazzoli Francesco Bullo

Abstract—In this paper we study minimum-time motion planning and routing problems for the Dubins vehicle, i.e., a nonholonomic vehicle that is constrained to move along planar paths of bounded curvature, without reversing direction. Motivated by autonomous aerial vehicle applications, we consider the Traveling Salesperson Problem for the Dubins vehicle (DTSP): given n points on a plane, what is the shortest Dubins tour through these points and what is its length? First, we show that the worst-case length of such a tour grows linearly with n and we propose a novel algorithm with performance within a constant factor of the optimum for the worst-case point sets. In doing this, we also obtain an upper bound on the optimal length in the classical point-to-point problem. Second, we study a stochastic version of the DTSP where the n targets are randomly and independently sampled from a uniform distribution. We show that the expected length of such a tour is of order at least $n^{2/3}$ and we propose a novel algorithm yielding a solution with length of order $n^{2/3}$ with probability one. Third and finally, we study a dynamic version of the DTSP: given a stochastic process that generates target points, is there a policy that guarantees that the number of unvisited points does not diverge over time? If such stable policies exist, what is the minimum expected time that a newly generated target waits before being visited by the vehicle? We propose a novel stabilizing algorithms such that the expected wait time is provably within a constant factor from the optimum.

I. INTRODUCTION

In this paper we study a novel class of optimal motion planning problems for a nonholonomic vehicle required to visit collections of points in the plane, where the vehicle is said to *visit* a region in the plane if the vehicle goes to that region and passes through it. This class of problem has two main ingredients. First, the robot model is the so-called Dubins vehicle, namely, a nonholonomic vehicle that is constrained to move along paths of bounded curvature without reversing direction. Second, the objective is to find the shortest path for such vehicle through a given set of target points. Except for the nonholonomic constraint, this task is akin to the classic Traveling Salesperson Problem (TSP) and in particular to the Euclidean TSP (ETSP), in which the shortest path between any two target locations is a straight segment. In summary, the focus of this paper is the analysis and the algorithmic

design of the TSP for the Dubins vehicle; we shall refer to this problem as to the Dubins TSP (DTSP).

A practical motivation to study the DTSP arises naturally in robotics and Uninhabited Aerial Vehicles (UAVs) applications. We envision applying DTSP algorithms to the setting of a UAV monitoring a collection of spatially distributed points of interest. In one scenario, the location of the points of interest might be known and static. Additionally, UAV applications motivate the study of the Dynamic Traveling Repairperson Problem (DTRP), in which the UAV is required to visit a dynamically changing set of targets. Such problems are examples of distributed task allocation problems and are currently generating much interest; e.g., [4] discusses complexity issues related to UAVs assignments problems, [5] considers Dubins vehicles keeping under surveillance multiple mobile targets, [6] considers missions with dynamic threats; other relevant works include [7], [8], [9], [10].

The literature on the Dubins vehicle is very rich and includes contributions from researchers in multiple disciplines. The minimum-time point-to-point path planning problem with bounded curvature was originally introduced by Markov [11] and a first solution was given by Dubins [12]. Modern treatments on point-to-point planning exploit the Pontryagin Minimum Principle [13], carefully account for symmetries in the problem [14], and consider environments with obstacles [15]. The Dubins vehicle is commonly accepted as a reasonably accurate kinematic model for aircraft motion planning problems, e.g., see [16], and its study is included in recent texts [17], [18].

The TSP and its variations continue to attract great interest from a wide range of fields, including operations research, mathematics and computer science. Tight bounds on the asymptotic dependence of the ETSP on the number of targets are given in the early work [19] and in the survey [20]. Exact algorithms, heuristics as well as polynomial-time constant factor approximation algorithms are available for the Euclidean TSP, see [21], [22], [23]. A variation of the TSP with potential robotic applications is the angular-metric problem studied in [24]. The DTRP (without nonholonomic constraints) was introduced in [25]. However, as with the TSP, the study of the DTRP in context of the Dubins vehicle has eluded attention from the research community. Finally, it is worth remarking that, unlike other variations of the TSP, there are no known reductions of the Dubins TSP to a problem on a finite-dimensional graph, thus preventing the use of well-established tools in combinatorial optimization.

The main contributions of this paper are threefold. First, we propose an algorithm for the DTSP through a point set P , called the ALTERNATING ALGORITHM. This algorithm is

To appear in the IEEE Transactions on Automatic Control. This version: January 13, 2008. Preliminary versions of this work appeared as conference papers in [1], [2], [3].

Ketan Savla is with the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology. ksavla@mit.edu. Emilio Frazzoli is with the Aeronautics and Astronautics Department and with the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology. frazzoli@mit.edu. Francesco Bullo is with the Mechanical Engineering Department and with the Center for Control, Dynamical Systems and Computation at the University of California at Santa Barbara. bullo@engineering.ucsb.edu

based on the solution to the ETSP over P and on an alternating heuristic to assign target orientations at each target point. This algorithm performs within a constant factor of the optimal for a certain class of point sets. As an intermediate step in the analysis of the algorithm, we provide an upper bound on the point-to-point minimum length of Dubins optimal paths. Second, we propose an algorithm for the stochastic DTSP, called the RECURSIVE BEAD-TILING ALGORITHM. This algorithm is based on a geometric tiling of the plane, tuned to the Dubins vehicle dynamics, and on a strategy for the vehicle to visit targets from each tile. The RECURSIVE BEAD-TILING ALGORITHM is the first algorithm providing a provable constant-factor approximation to the DTSP optimal solution with probability one. Third, we propose an algorithm for the DTRP in the heavy load case, called the BEAD-TILING ALGORITHM, based on a fixed-resolution version of the RECURSIVE BEAD-TILING ALGORITHM. We show that the performance guarantees for the stochastic DTSP translate into stability guarantees for the performance of the DTRP for the Dubins vehicle in heavy load case. Specifically, we show that the performance of BEAD-TILING ALGORITHM is within a constant factor from the theoretical optimum.

To clarify the contributions of this paper, it is worthwhile to compare our results with the ones existing in the literature. While the problem of flying an aircraft through way-points is a very standard problem in aeronautics (e.g., see [26], [5], [27]), the formal study of the Dubins TSP (algorithmic and performance bounds) was introduced in our early work [1], where a constant-factor approximation algorithm for worst-case point sets for the DTSP was proposed. Subsequently, similar versions of this problem were also considered in [28] and [7]. A simplified version of the problem for a different but closely related kind of vehicle, the Reeds-Shepp vehicle, was considered in [29]. In [2], we introduced the stochastic DTSP and gave the first algorithm yielding, with high probability, a solution with a cost upper bounded by a strictly sub-linear function of the number n of target points. Specifically, it was shown that the lower bound on the stochastic DTSP was of order $n^{2/3}$ and that our algorithm performed asymptotically within a $(\log n)^{1/3}$ factor to this lower bound with high probability. This result was improved in [30] with an algorithm for the stochastic DTSP that asymptotically performs within any $\epsilon(n)$ factor of the optimal with high probability, where $\epsilon(n) \rightarrow +\infty$ as $n \rightarrow +\infty$. In [3] we designed the first algorithm that asymptotically achieves a constant factor approximation to the stochastic DTSP with high probability. Based on our earlier works, this paper presents a comprehensive treatment of the worst-case DTSP, the DTSP in the stochastic setting, and the DTRP.

The paper is organized as follows. Section II contains the problem formulation. The worst-case DTSP is treated in Section III. Section IV contains the treatment of the DTSP for the stochastic setting. Section V contains the treatment of the dynamic version of the DTSP, namely the DTRP. We conclude with a few remarks about future work in Section VI. Some proofs are presented in the appendix.

II. PROBLEM SETUP: FROM THE EUCLIDEAN TO THE DUBINS TRAVELING SALESPERSON PROBLEM

In this section we set up the main problem of the paper and review some basic required notation. A *Dubins vehicle* is a planar vehicle that is constrained to move along paths of bounded curvature, without reversing direction and with bounded speed. Accordingly, we define a *feasible curve for the 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 turning radius. We also let $\text{Length}(\gamma) = \int_0^T \|\gamma'(t)\| dt$ be the length of a differentiable curve $\gamma : [0, T] \rightarrow \mathbb{R}^2$. We represent the vehicle configuration by the triplet $(x, y, \psi) \in \text{SE}(2)$, where (x, y) are the Cartesian coordinates of the vehicle and ψ is its heading.

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 point sets $P \subset \mathcal{Q}$ 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 , with minimum turning radius ρ .

The key objective of this paper is the design of an algorithm that provides a provably good approximation to the optimal solution of the Dubins TSP. To establish what a “good approximation” might be, let us recall what is known about the ETSP. First, given a compact set \mathcal{Q} , there exists [20] a finite constant $\alpha(\mathcal{Q})$ such that, for all $P \in \mathcal{P}_n$,

$$\text{ETSP}(P) \leq \alpha(\mathcal{Q})\sqrt{n}. \quad (1)$$

This upper bound is constructive in the sense that there exist algorithms that generate closed paths through the points P with length of order \sqrt{n} [20]. In the stochastic case, where the n points in P are independently chosen from a distribution φ with compact support $\mathcal{Q} \subset \mathbb{R}^2$, the following deterministic limit holds [19]:

$$\lim_{n \rightarrow +\infty} \frac{\text{ETSP}(P)}{\sqrt{n}} = \beta \int_{\mathcal{Q}} \sqrt{\varphi(q)} dq, \quad \text{with probability 1,}$$

where φ is a probability density function corresponding to the absolutely continuous part of φ , and β is a constant, which has been evaluated as $\beta = 0.712 \pm 0.0001$, e.g., see [31]. The fact that the dependence of the ETSP is sub-linear in n is very important in the study of the DTRP, i.e., the problem in which the point set P is not given a priori, but is generated over time by a point process; see Section V.

Motivated by the results available in the Euclidean case, this paper shows that the DTSP grows with n for the worst-case point sets and with $n^{2/3}$ in the stochastic case (as both lower and upper bounds) with probability one. Additionally, this paper proposes novel algorithms for the DTSP in the worst-case and stochastic settings, whose performances are within a constant factor of the optimal solution in the asymptotic limit as $n \rightarrow +\infty$. Finally, this paper uses these results to introduce the first stabilizing policy for the Dubins DTRP.

We conclude this section with some notation that is the standard concise way to state asymptotic properties. For $f, g :$

$\mathbb{N} \rightarrow \mathbb{R}$, we say that $f \in O(g)$ (respectively, $f \in \Omega(g)$) if there exist $N_0 \in \mathbb{N}$ and $k \in \mathbb{R}_+$ such that $|f(N)| \leq k|g(N)|$ for all $N \geq N_0$ (respectively, $|f(N)| \geq k|g(N)|$ for all $N \geq N_0$). If $f \in O(g)$ and $f \in \Omega(g)$, then we use the notation $f \in \Theta(g)$. Finally, we say that $f \in o(g)$ as $N \rightarrow +\infty$ if $\lim_{N \rightarrow +\infty} f(N)/g(N) = 0$ or, for functions $f, g : \mathbb{R} \rightarrow \mathbb{R}$, we say that $f \in o(g)$ as $x \rightarrow 0$ if $\lim_{x \rightarrow 0} f(x)/g(x) = 0$.

III. THE DTSP IN THE WORST CASE

Here we study the DTSP for worst-case point sets P , i.e., when the points P are chosen in an adversarial fashion.

A. Lower bound

We first give a lower bound on $\text{DTSP}_\rho(P)$ for the worst-case point sets. Given any point set $P \in \mathcal{P}_n$ with $n \geq 2$ and $\rho > 0$, it is immediate to see that $\text{DTSP}_\rho(P) \geq \text{ETSP}(P)$. This bound is improved in the following theorem, whose proof is reported in the appendix.

Theorem 3.1: (Worst-case lower bound on the TSP for the Dubins vehicle) Given $\rho > 0$, there exists a point set $P \in \mathcal{P}_n$, $n \geq 2$, such that

$$\text{DTSP}_\rho(P) \geq \text{ETSP}(P) + 2 \left\lfloor \frac{n}{2} \right\rfloor \pi \rho.$$

Remark 3.2: Theorem 3.1 implies that, for $P \in \mathcal{P}_n$ and in the worst case, $\text{DTSP}_\rho(P) \in \Omega(n)$. \square

B. The Alternating Algorithm

Here we propose a novel algorithm, the ALTERNATING ALGORITHM, that approximates the solution of the DTSP. The underlying principle of the algorithm is the following observation: since the optimal Dubins path between two configurations has been characterized in [12], a solution for the DTSP 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. The algorithm builds on the knowledge of the optimal solution of the ETSP for the same point set, and provides a sub-optimal DTSP tour.

The ALTERNATING ALGORITHM works as follows. 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 (except the n^{th}) edges, and replacing all even-numbered edges with minimum-length Dubins paths preserving the point ordering. In other words, the algorithm consists of the following steps:

- (i) set $(a_1, \dots, a_n) :=$ optimal ETSP ordering of P
- (ii) set $\psi_1 :=$ orientation of segment from a_1 to a_2
- (iii) for $i \in \{2, \dots, n-1\}$, do
 - if i is even, then set $\psi_i := \psi_{i-1}$, else set $\psi_i :=$ orientation of segment from a_i to a_{i+1}
- (iv) if n is even, then set $\psi_n := \psi_{n-1}$, else set $\psi_n :=$ orientation of segment from a_n to a_1
- (v) return the sequence of configurations $\{(a_i, \psi_i)\}_{i \in \{1, \dots, n\}}$.

We illustrate the output of the ALTERNATING ALGORITHM in Figure 1.




Fig. 1. An application of the ALTERNATING ALGORITHM. Left figure: a graph representing the solution of ETSP over a given P . Right figure: a graph representing the solution given by the ALTERNATING ALGORITHM on P where the alternate segments of ETSP are retained.

C. Analysis of the algorithm

In this section we analyze the performance of the ALTERNATING ALGORITHM to obtain an upper bound on $\text{DTSP}_\rho(P)$ and then show that the algorithm performs within a constant factor of the optimal for the worst-case point sets. To obtain an upper bound on the length of the path traversed by the Dubins vehicle while executing the ALTERNATING ALGORITHM, we first obtain an upper bound on the optimal point-to-point problem for the Dubins vehicle.

Problem 3.3: Given an initial configuration $(x_{\text{initial}}, y_{\text{initial}}, \psi_{\text{initial}})$ and a final configuration $(x_{\text{final}}, y_{\text{final}}, \psi_{\text{final}})$, find an upper bound on the length of the shortest Dubins path going from initial to final configuration. \square

To tackle this problem, we introduce some preliminary definitions. Without loss of generality, we assume $(x_{\text{initial}}, y_{\text{initial}}, \psi_{\text{initial}}) = (0, 0, 0)$. Let $\mathcal{C}_\rho : \text{SE}(2) \rightarrow \mathbb{R}_+$ associate to a configuration (x, y, ψ) the length of the shortest Dubins path from $(0, 0, 0)$ to (x, y, ψ) . Define $F_0 :]0, \pi[\times]0, \pi[\rightarrow]0, \pi[$, $F_1 :]0, \pi[\rightarrow \mathbb{R}$ and $F_2 :]0, \pi[\rightarrow \mathbb{R}$ by

$$\begin{aligned} F_0(\psi, \theta) &= 2 \tan^{-1} \left(\frac{\sin(\psi/2) - 2 \sin(\psi/2 - \theta)}{\cos(\psi/2) + 2 \cos(\psi/2 - \theta)} \right), \\ F_1(\psi) &= \psi + \sin \left(\frac{F_0(\psi, \psi/2 - \alpha(\psi))}{2} \right) \\ &\quad + 4 \cos^{-1} \left(\frac{\sin((\psi - F_0(\psi, \psi/2 - \alpha(\psi)))/2)}{2} \right), \\ F_2(\psi) &= 2\pi - \psi + 4 \cos^{-1} \left(\frac{\sin(\psi/2)}{2} \right), \end{aligned} \quad (2)$$

$$F_2(\psi) = 2\pi - \psi + 4 \cos^{-1} \left(\frac{\sin(\psi/2)}{2} \right), \quad (3)$$

where $\alpha(\psi) = \pi/2 - \cos^{-1}(\frac{\sin(\psi/2)}{2})$. The proof of the following result is postponed to the appendix.

Theorem 3.4 (Upper bound on optimal point-to-point length): For $\psi \in [0, 2\pi[$, $(x, y) \in \mathbb{R}^2$, and $\rho > 0$,

$$\mathcal{C}_\rho(x, y, \psi) \leq \sqrt{x^2 + y^2} + \kappa \pi \rho,$$

where $\kappa \in [2.657, 2.658]$ is defined by $\kappa = \frac{1}{\pi} \max\{F_2(\pi), \sup_{\psi \in]0, \pi[} \min\{F_1(\psi), F_2(\psi)\}\}$.

It is a conjecture that $\kappa = 7/3$; we provide some numerical evidence in Appendix D. Next, we let $L_{\text{AA}, \rho}(P)$ denote the length of Dubins path as given by the ALTERNATING

ALGORITHM for a point set P . The following lemma establishes bounds on the performance of the ALTERNATING ALGORITHM.

Lemma 3.5: (Upper bound on the performance of the ALTERNATING ALGORITHM) For any $P \in \mathcal{P}_n$ with $n \geq 2$ and $\rho > 0$,

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

Additionally, if there exists $\eta > 0$ such that $\min_{i,j \in \{1, \dots, n\}, i \neq j} \|p_i - p_j\| \geq \eta \rho$, then for $n \geq 3$

$$L_{AA,\rho}(P) \leq \left(1 + \frac{5\kappa\pi}{6\eta}\right) \text{ETSP}(P).$$

Proof: The first statement follows from Theorem 3.4. The second statement follows from the first by noting that $\min_{i,j \in \{1, \dots, n\}, i \neq j} \|p_i - p_j\| \geq \eta \rho$ implies that $\text{ETSP}(P) \geq n\eta\rho$. ■

Remark 3.6: (i) The first statement of Lemma 3.5 implies that for any point set $P \in \mathcal{P}_n$ with $n \geq 2$ and $\rho > 0$, $\text{ETSP}(P) \leq \text{DTSP}_\rho(P) \leq \text{ETSP}(P) + \kappa \left\lceil \frac{n}{2} \right\rceil \pi \rho$. An important consequence of this result is the following fact: given a point set, for small enough ρ , the order of points in the optimal path for the Euclidean TSP is the same as in the optimal path for the Dubins TSP.

(ii) Theorem 3.1 and Lemma 3.5 imply that there exists a point set $P \in \mathcal{P}_n$ such that $\text{ETSP}(P) + 2 \left\lceil \frac{n}{2} \right\rceil \pi \rho \leq \text{DTSP}_\rho(P) \leq \text{ETSP}(P) + \kappa \left\lceil \frac{n}{2} \right\rceil \pi \rho$, that is, in the worst-case, $\text{DTSP}_\rho(P)$ belongs to $\Theta(n)$.

(iii) The second statement of Lemma 3.5 implies that if the minimal inter-target distance is lower bounded, then $\text{DTSP}_\rho(P)$ is within a constant factor of $\text{ETSP}(P)$. In that case the ALTERNATING ALGORITHM provides $1 + \frac{5\kappa\pi}{6\eta}$ factor approximation to the optimal. Moreover, this constant factor is comparable to the one derived in [7] under the stricter assumption $\eta = 2$. Also, once the conjecture that $\kappa = 7/3$, as stated in Appendix D, is proven, it will decrease the constant factor of $1 + \frac{5\kappa\pi}{6\eta}$ for the ALTERNATING ALGORITHM. However, note that in a bounded environment the minimal inter-target distance must vanish as $n \rightarrow +\infty$ and, therefore, the ALTERNATING ALGORITHM is a constant factor approximation algorithm only for finite point sets with lower bounded inter-target distance. □

Having established bounds on the performance of the ALTERNATING ALGORITHM, we now show that it performs within a constant factor of the optimal for the worst-case point sets.

Theorem 3.7: (Performance of the ALTERNATING ALGORITHM for the worst-case point sets) For $n \geq 2$, $P \in \mathcal{P}_n$ and $\rho > 0$,

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

Furthermore,

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

Proof: The first statement follows from the simple fact that $L_{AA,\rho}(P) \geq \text{DTSP}_\rho(P)$, and from the results in Lemma 3.5 and Theorem 3.1. To prove the second statement, we take the limit as $n \rightarrow +\infty$ in the first statement and we use the bound in equation (1). ■

Remark 3.8: For $P \in \mathcal{P}_n$, Lemma 3.5 implies that $L_{AA,\rho}(P)$ belongs to $O(n)$ and Theorem 3.7 implies that in the worst case, $\text{DTSP}_\rho(P)$ belongs to $\Theta(n)$ and that the ALTERNATING ALGORITHM performs within $\frac{\kappa}{2}$ factor of the optimal for the worst-case point sets. The computational complexity of the ALTERNATING ALGORITHM is of order n . □

IV. STOCHASTIC DTSP

The discussion in the previous section showed that the ALTERNATING ALGORITHM performs well when the points to be visited by the tour are chosen in an adversarial manner. However, this algorithm is not a constant-factor approximation algorithm in the general case. Moreover, this algorithm might not perform very well when dealing with a random distribution of the target points. In particular, we will show that when n points are chosen randomly and independently, 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 RECURSIVE BEAD-TILING ALGORITHM, to visit these points and establish bounds on its performance.

We make the following assumptions: \mathcal{Q} is a rectangle of width W and height H with $W \geq H$; different choices for the shape of \mathcal{Q} affect our conclusions only by a constant. The two axes of the reference frame are parallel to the sides of \mathcal{Q} . In what follows, $P = \{p_1, \dots, p_n\}$ is a random variable, indicating a set of n points randomly and independently generated according to a uniform distribution in \mathcal{Q} .

A. Lower bound

We begin with a result from [32] that provides a lower bound on the expected length of the stochastic DTSP.

Theorem 4.1 (Lower bound on stochastic DTSP): Let $P \in \mathcal{P}_n$ be uniformly, randomly and independently generated in the rectangle of width W and height H . For any $\rho > 0$,

$$\liminf_{n \rightarrow +\infty} \frac{\text{E}[\text{DTSP}_\rho(P)]}{n^{2/3}} \geq \frac{3}{4} \sqrt[3]{3\rho WH}.$$

Remark 4.2: Theorem 4.1 implies that $\text{E}[\text{DTSP}_\rho(P)]$ belongs to $\Omega(n^{2/3})$. □

B. The basic geometric construction

Here we define a useful geometric object and study its properties. Consider two points p_- and p_+ on the plane, with $\ell = \|p_+ - p_-\|_2 \leq 4\rho$, and construct the region $\mathcal{B}_\rho(\ell)$ as detailed in Figure 2. We refer to such a region as a *bead* of length ℓ . The region $\mathcal{B}_\rho(\ell)$ enjoys the following asymptotic properties as $(\ell/\rho) \rightarrow 0^+$:




Fig. 2. Construction of the “bead” $\mathcal{B}_\rho(\ell)$. The figure shows how the upper half of the boundary is constructed, the bottom half is symmetric.

(P1) Its maximum “thickness” is

$$w(\ell) = 4\rho \left(1 - \sqrt{1 - \frac{\ell^2}{16\rho^2}} \right) = \frac{\ell^2}{8\rho} + \rho \cdot o\left(\frac{\ell^3}{\rho^3}\right).$$

(P2) Its area is

$$\text{Area}(\mathcal{B}_\rho(\ell)) = \frac{\ell w(\ell)}{2} = \frac{\ell^3}{16\rho} + \rho^2 \cdot o\left(\frac{\ell^4}{\rho^4}\right).$$

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

$$\text{Length}(\gamma_p) \leq 4\rho \arcsin\left(\frac{\ell}{4\rho}\right) = \ell + \rho \cdot o\left(\frac{\ell^3}{\rho^3}\right).$$

These facts are verified using elementary planar geometry. Finally, the bead has the property that the plane can be periodically tiled¹ by identical copies of $\mathcal{B}_\rho(\ell)$, for any $\ell \in]0, 4\rho]$. This fact is illustrated in Figure 3 below.

Next, we study the probability of targets belonging to a given bead. Consider a bead B entirely contained in \mathcal{Q} and assume n points are randomly and independently generated according to a uniform distribution in \mathcal{Q} . The probability that the i^{th} point is sampled in B is

$$\mu(\ell) = \frac{\text{Area}(\mathcal{B}_\rho(\ell))}{\text{Area}(\mathcal{Q})}.$$

Furthermore, the probability that exactly k out of the n points are sampled in B has a binomial distribution, i.e., indicating with n_B the total number of points sampled in B ,

$$\Pr[n_B = k | n \text{ samples}] = \binom{n}{k} \mu^k (1 - \mu)^{n-k}.$$

If the bead length ℓ is chosen as a function of n in such a way that $\nu = n \cdot \mu(\ell)$ is a constant, then the limit for large n of the binomial distribution is [33] the Poisson distribution of mean ν , that is,

$$\lim_{n \rightarrow +\infty} \Pr[n_B = k | n \text{ samples}] = \frac{\nu^k}{k!} e^{-\nu}.$$

¹A tiling of the plane is a collection of sets whose intersection has measure zero and whose union covers the plane.

C. The Recursive Bead-Tiling Algorithm

In this section, we design a novel algorithm that computes a Dubins path through a point set in \mathcal{Q} . The proposed algorithm consists of a sequence of phases; during each phase, a Dubins tour (i.e., a closed path with bounded curvature) is constructed that “sweeps” the set \mathcal{Q} . We begin by considering a tiling of the plane such that $\text{Area}(\mathcal{B}_\rho(\ell)) = WH/(2n)$; in such a case, $\mu(\ell(n)) = 1/(2n)$, $\nu = 1/2$, and

$$\ell(n) = 2\left(\frac{\rho WH}{n}\right)^{\frac{1}{3}} + o(n^{-\frac{1}{3}}), \quad (n \rightarrow +\infty).$$

(Note that this implies that n must be large enough in order that $\ell(n) \leq 4\rho$.) Furthermore, the tiling is chosen in such a way that it is aligned with the sides of \mathcal{Q} , see Figure 3. In the first phase of the algorithm, a Dubins tour is constructed with the following properties:

- (i) it visits all non-empty beads once,
- (ii) it visits all rows² in sequence top-to-down, alternating between left-to-right and right-to-left passes, and visiting all non-empty beads in a row,
- (iii) when visiting a non-empty bead, it visits at least one target in it.

In order to visit the targets outstanding after the first phase, a second phase is initiated. Instead of considering single beads, we now consider “meta-beads” composed of two beads each, as shown in Figure 3, and proceed in a way similar to the first phase, i.e., a Dubins tour is constructed with the following properties:

- (i) the tour visits all non-empty meta-beads once,
- (ii) it visits all (meta-bead) rows in sequence top-to-down, alternating between left-to-right and right-to-left passes, and visiting all non-empty meta-beads in a row,
- (iii) when visiting a non-empty meta-bead, it visits at least one target in it.

This process is iterated $\lceil \log_2 n \rceil$ times, and at each phase, meta-beads composed of two neighboring meta-beads from the previous phase are considered; in other words, the meta-beads at the i^{th} phase are composed of 2^{i-1} neighboring beads. After the last recursive phase, the leftover targets are visited using the ALTERNATING ALGORITHM.

D. Analysis of the algorithm

In this section, we calculate an upper bound on the length of Dubins path as given by the RECURSIVE BEAD-TILING ALGORITHM. By comparing this upper bound with the lower bound established earlier, we will conclude that the algorithm provides a constant factor approximation to the optimal stochastic DTSP with high probability. We begin with a key result about the number of outstanding targets after the execution of the $\lceil \log_2 n \rceil$ recursive phases; the proof of this result is based upon techniques similar to those developed in [34].

Theorem 4.3 (Targets remaining after recursive phases): Let $P \in \mathcal{P}_n$ be uniformly, randomly and independently

²A row is a maximal sequence of horizontally-aligned beads with non-empty intersection with \mathcal{Q} .




Fig. 3. Sketch of “meta-beads” at successive phases in the recursive bead tiling algorithm. From left to right: phase 1, phase 2 and phase 3. Note that for phase 2 (and for all subsequent even-numbered phases), the vehicle will have to visit every row of meta-beads twice, once to visit targets in the meta-beads with the darker shade and once to visit targets in the meta-beads with the lighter shade.

generated in \mathcal{Q} . The number of unvisited targets after the last recursive phase of the RECURSIVE BEAD-TILING ALGORITHM over P is less than $24 \log_2 n$ with high probability, i.e., with probability approaching one as $1 - \frac{\log_2 2n}{n^2}$.

Proof: Associate a unique identifier to each bead, let $b(t)$ be the identifier of the bead in which the t^{th} target is sampled, and let $h(t) \in \mathbb{N}$ be the phase at which the t^{th} target is visited. Without loss of generality, assume that targets within a single bead are visited in the same order in which they are generated, i.e., if $b(t_1) = b(t_2)$ and $t_1 < t_2$, then $h(t_1) < h(t_2)$. Note that we assume here that only one target per bead is visited at each phase. The resultant analysis will give an upper bound on the path length for the RECURSIVE BEAD-TILING ALGORITHM. Let $v_i(t)$ be the number of beads that contain unvisited targets at the inception of the i^{th} phase, computed after the insertion of the t^{th} target. Furthermore, let m_i be the number of i^{th} phase meta-beads (i.e., meta-beads containing 2^{i-1} neighboring beads) with a non-empty intersection with \mathcal{Q} . Clearly, $v_i(t) \leq v_i(n)$, $m_i \leq 2m_{i+1}$, and $v_1(n) \leq n \leq m_1/2$ with certainty. The t^{th} target will not be visited during the first phase if it is sampled in a bead that already contains other targets. In other words,

$$\Pr [h(t) \geq 2 | v_1(t-1)] = \frac{v_1(t-1)}{m_1} \leq \frac{v_1(n)}{2n} \leq \frac{1}{2}.$$

Similarly, the t^{th} target will not be visited during the i^{th} phase if (i) it has not been visited before the i^{th} pass, and (ii) it belongs to a meta-bead that already contains other targets not visited before the i^{th} phase:

$$\begin{aligned} & \Pr [h(t) \geq i+1 | (v_i(t-1), v_{i-1}(t-1), \dots, v_1(t-1))] \\ &= \Pr [h(t) \geq i+1 | h(t) \geq i, v_i(t-1)] \\ & \quad \cdot \Pr [h(t) \geq i | (v_{i-1}(t-1), \dots, v_1(t-1))] \\ &\leq \frac{v_i(t-1)}{m_i} \Pr [h(t) \geq i | (v_{i-1}(t-1), \dots, v_1(t-1))] \\ &= \prod_{j=1}^i \frac{v_j(t-1)}{m_j} \leq \prod_{j=1}^i \frac{2^{j-1} v_j(n)}{2n} = \left(\frac{2^{\frac{i-3}{2}}}{n} \right)^i \prod_{j=1}^i v_j(n). \end{aligned}$$

Given a sequence $\{\beta_i\}_{i \in \mathbb{N}} \subset \mathbb{R}_+$ and given a fixed $i \geq 1$, define a sequence of binary random variables

$$Y_t(i) = \begin{cases} 1, & \text{if } h(t) \geq i+1 \text{ and } v_i(t-1) \leq \beta_i n, \\ 0, & \text{otherwise.} \end{cases}$$

In other words, $Y_t(i) = 1$ if the t^{th} target is not visited during the first i phases even though the number of beads still containing unvisited targets at the inception of the i^{th} phase is less than $\beta_i n$. Even though the random variable $Y_t(i)$ depends on the targets generated before the t^{th} target, the probability that it takes the value 1 is bounded by

$$\Pr [Y_t(i) = 1 | b(1), b(2), \dots, b(t-1)] \leq 2^{\frac{i(i-3)}{2}} \prod_{j=1}^i \beta_j =: q_i,$$

regardless of the actual values of $b(1), \dots, b(t-1)$. It is known [34] that if the random variables $Y_t(i)$ satisfy such a condition, the sum $\sum_t Y_t(i)$ is stochastically dominated by a binomially distributed random variable, namely,

$$\Pr \left[\sum_{t=1}^n Y_t(i) > k \right] \leq \Pr [B(n, q_i) > k],$$

where $B(n, q_i)$ denotes a binomially distributed random variable with parameters n and q_i . In particular,

$$\Pr \left[\sum_{t=1}^n Y_t(i) > 2nq_i \right] \leq \Pr [B(n, q_i) > 2nq_i] < 2^{-nq_i/3}, \quad (4)$$

where the last inequality follows from Chernoff's Bound [33]. Now, it is convenient to define $\{\beta_i\}_{i \in \mathbb{N}}$ by

$$\beta_1 = 1, \quad \beta_{i+1} = 2q_i = 2^{\frac{i(i-3)}{2}+1} \prod_{j=1}^i \beta_j = 2^{i-2} \beta_i^2,$$

which leads to $\beta_i = 2^{1-i}$. In turn, this implies that equation (4) can be rewritten as

$$\Pr \left[\sum_{t=1}^n Y_t(i) > \beta_{i+1} n \right] < 2^{-\beta_{i+1} n/6} = 2^{-\frac{n}{3 \cdot 2^i}},$$

which is less than $1/n^2$ for $i \leq i^*(n) := \lfloor \log_2 n - \log_2 \log_2 n - \log_2 6 \rfloor \leq \log_2 n$. Note that $\beta_i \leq 12 \frac{\log_2 n}{n}$, for all $i > i^*(n)$.

Let \mathcal{E}_i be the event that $v_i(n) \leq \beta_i n$. Note that if \mathcal{E}_i is true, then $v_{i+1}(n) \leq \sum_{t=1}^n Y_t(i)$: the right hand side represents the number of targets that will be visited after the i^{th} phase, whereas the left hand side counts the number of

beads containing such targets. We have, for all $i \leq i^*(n)$:

$$\begin{aligned} & \Pr \left[v_{i+1}(n) > \beta_{i+1}n \mid \mathcal{E}_i \right] \cdot \Pr[\mathcal{E}_i] \\ & \leq \Pr \left[\sum_{t=1}^n Y_t(i) > \beta_{i+1}n \right] \leq \frac{1}{n^2}, \end{aligned}$$

that is, $\Pr[\neg\mathcal{E}_{i+1} \mid \mathcal{E}_i] \cdot \Pr[\mathcal{E}_i] \leq \frac{1}{n^2}$, and thus (recall that \mathcal{E}_i is true with certainty):

$$\begin{aligned} & \Pr[\neg\mathcal{E}_{i+1}] \\ & = \Pr[\neg\mathcal{E}_{i+1} \mid \mathcal{E}_i] \cdot \Pr[\mathcal{E}_i] + \Pr[\neg\mathcal{E}_{i+1} \mid \neg\mathcal{E}_i] \cdot \Pr[\neg\mathcal{E}_i] \\ & \leq \frac{1}{n^2} + \Pr[\neg\mathcal{E}_i] \leq \frac{i}{n^2}. \end{aligned}$$

In other words, for all $i \leq i^*(n)$, $v_i(n) \leq \beta_i n$ with high probability.

Let us now turn our attention to the phases such that $i > i^*(n)$. The total number of targets visited after the $(i^*)^{\text{th}}$ phase is dominated by a binomial variable $B(n, 12 \log_2 n/n)$; in particular,

$$\begin{aligned} & \Pr \left[v_{i^*+1}(n) > 24 \log_2 n \mid \mathcal{E}_{i^*} \right] \cdot \Pr[\mathcal{E}_{i^*}] \\ & \leq \Pr \left[\sum_{t=1}^n Y_t(i) > 24 \log_2 n \right] \\ & \leq \Pr \left[B(n, 12 \log_2 n/n) > 24 \log_2 n \right] \leq 2^{-12 \log_2 n}. \end{aligned}$$

Dealing with conditioning as before, we obtain

$$\Pr[v_{i^*+1}(n) > 24 \log_2 n] \leq \frac{1}{n^{12}} + \Pr[\neg\mathcal{E}_{i^*}] \leq \frac{1}{n^{12}} + \frac{\log_2 n}{n^2}. \quad (5)$$

In other words, the number of unvisited targets after the $(i^*)^{\text{th}}$ phase is bounded by a logarithmic function of n with high probability. Equation (5) also shows that this probability approaches one as $1 - \frac{\log_2 2n}{n^2}$. ■

In summary, Theorem 4.3 says that after a sufficiently large number of phases, almost all targets will be visited, with high probability. A simple application of the Borel-Cantelli Lemma [35] to the upper bound in equation (5) gives the following corollary.

Corollary 4.4: With probability one, the number of unvisited targets after the last recursive phase of the RECURSIVE BEAD-TILING ALGORITHM over P is less than $24 \log_2 n$ asymptotically.

We also observe that (i) the length of the first phase is of order $n^{2/3}$ and (ii) the length of each phase is decreasing at such a rate that the sum of the lengths of the $\lceil \log_2 n \rceil$ recursive phases remains bounded and proportional to the length of the first phase. (Since we are considering the asymptotic case in which the number of targets is very large, the length of the beads will be very small; in the remainder of this section we will tacitly consider the asymptotic behavior as $\ell/\rho \rightarrow 0^+$.)

Lemma 4.5 (Path length for the first phase): Consider a tiling of the plane with beads of length ℓ . For any $\rho > 0$ and for any set of target points, the length L_1 of a path visiting once and only once each bead with a non-empty intersection with a rectangle \mathcal{Q} of width W and length H satisfies

$$L_1 \leq \frac{16\rho WH}{\ell^2} \left(1 + \frac{7}{3}\pi \frac{\rho}{W} \right) + \rho \cdot o\left(\frac{\rho}{\ell}\right).$$

Proof: A path visiting each bead once can be constructed by a sequence of passes, during which all beads in a row are visited in a left-to-right or right-to-left order. In each row, there are at most $\lceil W/\ell \rceil + 1$ beads with a non-empty intersection with \mathcal{Q} . Hence, the cost of each pass is at most:

$$L_1^{\text{pass}} \leq W + 2\ell + \rho \cdot o\left(\frac{\ell^2}{\rho^2}\right).$$

Two passes are connected by a U-turn maneuver, in which the direction of travel is reversed, and the path moves to the next row, at distance equal to one half the width of a bead. Since the length of the shortest path to reverse the heading of a Dubins vehicle with co-located initial and final points is $(7/3)\pi\rho$, the length of the U-turn satisfies

$$L_1^{\text{U-turn}} \leq \frac{7}{3}\pi\rho + \frac{1}{2}w(\ell) \leq \frac{7}{3}\pi\rho + \frac{\ell^2}{16\rho} + \rho \cdot o\left(\frac{\ell^3}{\rho^3}\right).$$

The total number of passes, i.e., the total number of rows of beads with non-empty intersection with \mathcal{Q} , satisfies

$$N_1^{\text{pass}} \leq \left\lceil \frac{2H}{w(\ell)} \right\rceil + 1 \leq \frac{16\rho H}{\ell^2} + 2 + o\left(\frac{\rho}{\ell}\right).$$

A simple upper bound on the cost of closing the tour is given by

$$\begin{aligned} L_1^{\text{close}} & \leq (W + 2\ell) + (H + 2w(\ell)) + 2\pi\rho \\ & = W + H + 2\pi\rho + 2\ell + \rho \cdot o(\ell/\rho). \end{aligned}$$

In summary, the total length of the path followed during the first phase is

$$\begin{aligned} L_1 & \leq N_1^{\text{pass}} (L_1^{\text{pass}} + L_1^{\text{U-turn}}) + L_1^{\text{close}} \\ & \leq \left(\frac{16\rho H}{\ell^2} + 2 + o\left(\frac{\rho}{\ell}\right) \right) \\ & \quad \cdot \left(W + 2\ell + \frac{7}{3}\pi\rho + \frac{\ell^2}{16\rho} + \rho \cdot o\left(\frac{\ell^2}{\rho^2}\right) \right) \\ & \quad + W + H + 2\pi\rho + 2\ell + \rho \cdot o(\ell/\rho) \\ & \leq \frac{16\rho WH}{\ell^2} \left(1 + \frac{7}{3}\pi \frac{\rho}{W} \right) + \rho \cdot o\left(\frac{\rho}{\ell}\right). \end{aligned}$$

Based on this calculation, we can estimate the length of the paths in generic phases of the algorithm. Since the total number of phases in the algorithm depends on the number of targets n , as does the length of the beads ℓ , we will retain explicitly the dependency on the phase number. ■

Lemma 4.6 (Path length at odd-numbered phases): Consider a tiling of the plane with beads of length ℓ . For any $\rho > 0$ and for any set of target points, the length L_{2j-1} of a path visiting once and only once each meta-bead with a non-empty intersection with a rectangle \mathcal{Q} of width W and length H at phase number $(2j-1)$, $j \in \mathbb{N}$, satisfies

$$\begin{aligned} L_{2j-1} & \leq 2^{5-j} \left[\frac{\rho WH}{\ell^2} \left(1 + \frac{7}{3}\pi \frac{\rho}{W} \right) + \rho \cdot o\left(\frac{\rho}{\ell}\right) \right] \\ & \quad + 32 \frac{\rho H}{\ell} + \rho \cdot o\left(\frac{\rho}{\ell}\right) + 2^j \left[3\ell + \rho \cdot o\left(\frac{\ell}{\rho}\right) \right]. \end{aligned}$$

Proof: During odd-numbered phases, the number of beads in a meta-bead is a perfect square and the considerations made

in the proof of Lemma 4.5 can be readily adapted. The length of each pass satisfies

$$L_{2j-1}^{\text{pass}} \leq (W + 2^j \ell) \left[1 + o\left(\frac{\ell}{\rho}\right) \right].$$

The length of each U-turn maneuver is bounded as

$$\begin{aligned} L_{2j-1}^{\text{U-turn}} &\leq \frac{7}{3} \pi \rho + 2^{j-2} w(\ell) \\ &\leq \frac{7}{3} \pi \rho + 2^{j-2} \left[\frac{\ell^2}{8\rho} + \rho \cdot o\left(\frac{\ell^3}{\rho^3}\right) \right], \end{aligned}$$

from which

$$L_{2j-1}^{\text{pass}} + L_{2j-1}^{\text{U-turn}} = W + \frac{7}{3} \pi \rho + o\left(\frac{\ell}{\rho}\right) + 2^j \left[\ell + \rho \cdot o\left(\frac{\ell}{\rho}\right) \right].$$

The number of passes satisfies:

$$N_{2j-1}^{\text{pass}} \leq 2^{5-j} \left[\frac{\rho H}{\ell^2} + o\left(\frac{\rho}{\ell}\right) \right] + 2.$$

Finally, the cost of closing the tour is bounded by

$$L_{2j-1}^{\text{close}} \leq W + H + 2\pi\rho + 2^j [\ell + \rho \cdot o(\ell/\rho)].$$

Therefore, a bound on the total length of the path is

$$\begin{aligned} L_{2j-1} &= N_{2j-1}^{\text{pass}} (L_{2j-1}^{\text{pass}} + L_{2j-1}^{\text{U-turn}}) + L_{2j-1}^{\text{close}} \\ &\leq 2^{5-j} \left[\frac{\rho W H}{\ell^2} \left(1 + \frac{7\pi\rho}{3W} \right) + \rho \cdot o\left(\frac{\rho}{\ell}\right) \right] \\ &\quad + 32 \frac{\rho H}{\ell} + \rho \cdot o\left(\frac{\rho}{\ell}\right) + 2^j \left[3\ell + \rho \cdot o\left(\frac{\ell}{\rho}\right) \right]. \end{aligned}$$

Lemma 4.7 (Path length at even-numbered phases):

Consider a tiling of the plane with beads of length ℓ . For any $\rho > 0$, a rectangle \mathcal{Q} of width W and length H and any set of target points, paths in each phase of the RECURSIVE BEAD-TILING ALGORITHM can be chosen such that $L_{2j} \leq 2L_{2j+1}$, for all $j \in \mathbb{N}$.

Proof: Consider a generic meta-bead B_{2j+1} traversed in the $(2j+1)^{\text{th}}$ phase, and let l_3 be the length of the path segment within B_{2j+1} . The same meta-bead is traversed at most twice during the $(2j)^{\text{th}}$ phase; let l_1, l_2 be the lengths of the two path segments of the $(2j)^{\text{th}}$ phase within B_{2j+1} . By convention, for $i \in \{1, 2, 3\}$, we let $l_i = 0$ if the i^{th} path does not intersect B_{2j+1} . Without loss of generality, the order of target points can be chosen in such a way that $l_1 \leq l_2 \leq l_3$, and hence $l_1 + l_2 \leq 2l_3$. Repeating the same argument for all non-empty meta-beads, we prove the claim. ■

Finally, we can summarize these intermediate bounds into the main result of this section. We let $L_{\text{RBTA},\rho}(P)$ denote the length of the Dubins path computed by the RECURSIVE BEAD-TILING ALGORITHM for a point set P .

Theorem 4.8: (Path length for the RECURSIVE BEAD-TILING ALGORITHM) Let $P \in \mathcal{P}_n$ be uniformly, randomly and independently generated in the rectangle of width W and height H . For any $\rho > 0$, with probability one,

$$\begin{aligned} \limsup_{n \rightarrow +\infty} \frac{\text{DTSP}_{\rho}(P)}{n^{2/3}} &\leq \limsup_{n \rightarrow +\infty} \frac{L_{\text{RBTA},\rho}(P)}{n^{2/3}} \\ &\leq 24 \sqrt[3]{\rho W H} \left(1 + \frac{7}{3} \pi \frac{\rho}{W} \right). \end{aligned}$$

Proof: For simplicity we let $L_{\text{RBTA},\rho}(P) = L_{\text{RBTA}}$. Clearly, $L_{\text{RBTA}} = L'_{\text{RBTA}} + L''_{\text{RBTA}}$, where L'_{RBTA} is the path length of the first $\lceil \log_2 n \rceil$ phases of the RECURSIVE BEAD-TILING ALGORITHM and L''_{RBTA} is the length of the path required to visit all remaining targets. An immediate consequence of Lemma 4.7 is that

$$L'_{\text{RBTA}} = \sum_{i=1}^{\lceil \log_2(n) \rceil} L_i \leq 3 \sum_{j=1}^{\lceil \log_2(n)/2 \rceil} L_{2j-1}.$$

The summation on the right hand side of this equation can be expanded using Lemma 4.6 yielding

$$\begin{aligned} L'_{\text{RBTA}} &\leq 3 \left\{ \left[\frac{\rho W H}{\ell^2} \left(1 + \frac{7\pi\rho}{3W} \right) + \rho \cdot o\left(\frac{\rho^2}{\ell^2}\right) \right] \sum_{j=1}^{\lceil \log_2(n)/2 \rceil} 2^{5-j} \right. \\ &\quad \left. + \left(32 \frac{\rho H}{\ell} + \rho \cdot o\left(\frac{\rho}{\ell}\right) \right) \left\lceil \frac{\log_2 n}{2} \right\rceil \right. \\ &\quad \left. + [3\ell + \rho \cdot o(\ell/\rho)] \sum_{j=1}^{\lceil \log_2(n)/2 \rceil} 2^j \right\}. \end{aligned}$$

Since $\sum_{j=1}^k 2^{-j} \leq \sum_{j=1}^{+\infty} 2^{-j} = 1$, and $\sum_{j=1}^k 2^j = 2^{k+1} - 2 \leq 2^{k+1}$, the previous equation can be simplified to

$$\begin{aligned} L'_{\text{RBTA}} &\leq 3 \left\{ 32 \left[\frac{\rho W H}{\ell^2} \left(1 + \frac{7\pi\rho}{3W} \right) + \rho \cdot o\left(\frac{\rho}{\ell}\right) \right] \right. \\ &\quad \left. + \left(32 \frac{\rho H}{\ell} + \rho \cdot o\left(\frac{\rho}{\ell}\right) \right) \left\lceil \frac{\log_2 n}{2} \right\rceil \right. \\ &\quad \left. + [3\ell + \rho \cdot o(\ell/\rho)] \cdot (4\sqrt{n}) \right\}. \end{aligned}$$

$$\begin{aligned} L'_{\text{RBTA}} &\leq 3 \left\{ 32 \left[\frac{\rho W H}{\ell^2} \left(1 + \frac{7\pi\rho}{3W} \right) + \rho \cdot o\left(\frac{\rho}{\ell}\right) \right] \right. \\ &\quad \left. + \left(32 \frac{\rho H}{\ell} + \rho \cdot o\left(\frac{\rho}{\ell}\right) \right) \left\lceil \frac{\log_2 n}{2} \right\rceil \right. \\ &\quad \left. + [3\ell + \rho \cdot o(\ell/\rho)] \cdot (4\sqrt{n}) \right\}. \end{aligned}$$

Recalling that $\ell = 2(\rho W H/n)^{1/3} + o(n^{-1/3})$ for large n , the above can be rewritten as

$$L'_{\text{RBTA}} \leq 24 \sqrt[3]{\rho W H n^2} \left(1 + \frac{7}{3} \pi \frac{\rho}{W} \right) + o(n^{2/3}).$$

Now it suffices to show that L''_{RBTA} is negligible with respect to L'_{RBTA} for large n with probability one. From Corollary 4.4, we know that asymptotically, with probability one, there will be at most $24 \log_2 n$ unvisited targets after the $\lceil \log_2 n \rceil$ recursive phases. Lemma 3.5 would then imply that, with probability one, the length of a ALTERNATING ALGORITHM tour through these points asymptotically satisfies

$$L''_{\text{RBTA}} \leq \kappa [12 \log_2 n] \pi \rho + o(\log_2 n). \quad \blacksquare$$

Next, we state a result for the concentration of $\text{DTSP}_\rho(P)$ around its mean, which will let us compare the lower bound in Theorem 4.1 with the upper bound in Theorem 4.8.

Lemma 4.9 (Concentration around the mean): Let $P \in \mathcal{P}_n$ be uniformly, randomly and independently generated in the rectangle of width W and height H . For any $\rho > 0$, with probability one,

$$|\text{DTSP}_\rho(P) - \mathbb{E}[\text{DTSP}_\rho(P)]| \in O(\sqrt{n \log n}).$$

Proof: The proof presented here closely follows the one for the Long Common Sub-sequence Problem in Chapter 1 of [36]. We use Doob's method to construct a martingale from the random variable $\text{DTSP}_\rho(P)$. First let $\mathcal{F}_k = \sigma(p_1, \dots, p_k)$, that is, \mathcal{F}_k is the sigma-field generated by the first k elements of $P = \{p_1, \dots, p_n\}$, and then we set

$$d_i = \mathbb{E}[\text{DTSP}_\rho(P) | \mathcal{F}_i] - \mathbb{E}[\text{DTSP}_\rho(P) | \mathcal{F}_{i-1}].$$

The sequence $\{d_i\}$ can be easily checked to be a martingale-difference sequence adapted to the increasing sequence of sigma-fields $\{\mathcal{F}_i\}$. Moreover, d_i 's are related to the original variables via the following relation:

$$\text{DTSP}_\rho(P) - \mathbb{E}[\text{DTSP}_\rho(P)] = \sum_{i=1}^n d_i.$$

Consider a new sequence of independent random variables $\{\hat{p}_i\}$ with the same distribution as the original $\{p_i\}$. Accordingly, define $\hat{P}_i := \{p_1, \dots, p_{i-1}, \hat{p}_i, p_{i+1}, \dots, p_n\}$. Since \mathcal{F}_i has no information about \hat{p}_i , we have

$$\mathbb{E}[\text{DTSP}_\rho(P) | \mathcal{F}_{i-1}] = \mathbb{E}[\text{DTSP}_\rho(\hat{P}_i) | \mathcal{F}_i],$$

and this representation then lets us rewrite the expression for d_i in terms of a single conditional expectation:

$$d_i = \mathbb{E}[\text{DTSP}_\rho(P) - \text{DTSP}_\rho(\hat{P}_i) | \mathcal{F}_i].$$

From Theorem 3.4, one can easily check that

$$|\text{DTSP}_\rho(P) - \text{DTSP}_\rho(\hat{P}_i)| \leq 2 \text{diam}(\mathcal{Q}) + 2\kappa\pi\rho =: c.$$

Since conditional expectations cannot increase the upper bound, we have $|d_i| \leq c$ for all $i \in \{1, \dots, n\}$. Finally, by Azuma's Inequality, we have the useful tail bound:

$$\Pr[|\text{DTSP}_\rho(P) - \mathbb{E}[\text{DTSP}_\rho(P)]| \geq t] \leq 2 \exp(-t^2/(2nc^2)).$$

A straightforward application of the Borel-Cantelli Lemma with $t = \sqrt{2c^2n(\log n)(1+\epsilon)}$, where ϵ is some positive constant, gives us the desired result. ■

Remark 4.10: Lemma 4.9 implies that, with probability one,

$$\lim_{n \rightarrow +\infty} \left(\frac{\text{DTSP}_\rho(P)}{n^{2/3}} - \frac{\mathbb{E}[\text{DTSP}_\rho(P)]}{n^{2/3}} \right) = 0.$$

This statement together with Theorems 4.1 and 4.8 implies that, with probability one, the RECURSIVE BEAD-TILING ALGORITHM is a $(32/\sqrt[3]{3}) \left(1 + \frac{7}{3}\pi \frac{\rho}{W}\right)$ factor approximation (with respect to n) to the optimal DTSP and that $\text{DTSP}_\rho(P)$ belongs to $\Theta(n^{2/3})$. The computational complexity of the RECURSIVE BEAD-TILING ALGORITHM is of order n . □




Fig. 4. Plot of $\log(L_{\text{RBTA}, \rho}(P))$ vs. $\log(n)$

E. Numerical Results

In this section we present numerical results for the RECURSIVE BEAD-TILING ALGORITHM. The results are summarized in the form of a logarithmic plot in Figure 4. The points comprising the set P are randomly and independently generated according to a uniform distribution in a rectangle of width $W = 10$ and height $H = 8$. The minimum turning radius for the Dubins vehicle is $\rho = 1$. Each point represents the mean of Dubins path length as given by the RECURSIVE BEAD-TILING ALGORITHM, taken over 10 instances of the experiment for the corresponding values of n . The lower solid line represents the function $\log(C_l n^{2/3})$ where C_l is the value of the quantity $\frac{3}{4}\sqrt[3]{3\rho WH}$ corresponding to the lower bound in Theorem 4.1. Similarly, the upper solid line represents the function $\log(C_u n^{2/3})$, with C_u being the value of $24\sqrt[3]{\rho WH} \left(1 + \frac{7}{3}\pi \frac{\rho}{W}\right)$ corresponding to the upper bound in Theorem 4.8. From the simulations we gather the following qualitative observations. First, the lower bound to $\text{DTSP}_\rho(P)$ established in Theorem 4.1 is fairly conservative when considered as a lower bound to $L_{\text{RBTA}, \rho}(P)$. Second, the upper bound to $L_{\text{RBTA}, \rho}(P)$ established in Theorem 4.8 becomes less conservative and the data conforms more accurately with the $2/3$ exponent as n grows.

V. THE DTRP FOR THE DUBINS VEHICLE

We now turn our attention to the Dynamic Traveling Repairperson Problem (DTRP) that was introduced by Bertsimas and van Ryzin in [25]. When compared with previous work, the novel feature of the following work is the focus on the Dubins vehicle.

A. Model and problem statement

In this subsection we describe the vehicle and sensing model and the DTRP definition. The key aspect of the DTRP is that the Dubins 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 and that it moves at a unit speed with minimum turning radius $\rho > 0$.

Information about the outstanding targets representing the demand at time t is described by a finite set of 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 inside the rectangle \mathcal{Q} of width W and height H . 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

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

(Strictly speaking, the above equation holds when targets are not being removed from the queue D .) Servicing of a target and its removal from the set D are achieved when the Dubins vehicle moves to the target position.

A feedback control policy for the Dubins vehicle is a map Φ assigning a control input to the vehicle as a function of its configuration and of the current outstanding targets. We also consider policies that compute a control input based on a snapshot of the outstanding target configurations at certain time sequences. Let $\mathcal{T}_\Phi = \{t_k\}_{k \in \mathbb{N}}$ be a strictly increasing sequence of times at which such computations are started: with some abuse of terminology, we say that Φ is a *receding horizon strategy* if it is based on the most recent target data $D_{\text{rh}}(t)$, where

$$D_{\text{rh}}(t) = D(\max\{t_{\text{rh}} \in \mathcal{T}_\Phi \mid t_{\text{rh}} \leq t\}).$$

The (receding horizon) policy Φ is a stable policy for the DTRP if, under its action,

$$n_\Phi = \limsup_{t \rightarrow +\infty} \mathbb{E}[n(t) \mid \dot{p} = \Phi(p, D_{\text{rh}})] < +\infty,$$

that is, if the Dubins vehicle is able to visit 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 the j^{th} target is generated to the time it is visited. If the system is stable, then we can write the balance equation (known as Little's Formula [37]):

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

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

$$T_{\text{DTRP}} = \inf\{T_\Phi \mid \Phi \text{ is a stable control policy}\}.$$

B. Lower and constructive upper bounds

In what follows, we design a control policy that provides a constant-factor approximation of the optimal achievable performance. Consistently with the theme of the paper, we consider the case of *heavy load*, i.e., the problem as the time intensity $\lambda \rightarrow +\infty$. We first review from [32] a lower bound for the system time, and then present a novel approximation algorithm providing an upper bound on the performance.

Theorem 5.1: (Lower bound on the system time for the DTRP) For any $\rho > 0$, the system time T_{DTRP} for the DTRP in a rectangle of width W and height H satisfies

$$\liminf_{\lambda \rightarrow +\infty} \frac{T_{\text{DTRP}}}{\lambda^2} \geq \frac{81}{64} \rho W H.$$

Remark 5.2: Theorem 5.1 implies that the system time for the Dubins vehicle depends quadratically on the time intensity

λ , whereas in the Euclidean case it depends only linearly on it, e.g., see [25]. \square

We now propose a simple strategy, the BEAD-TILING ALGORITHM, based on the concepts introduced in the previous section. The strategy consists of the following steps:

- (i) Tile the plane with beads of length $\ell := \min\{C_{\text{BTA}}/\lambda, 4\rho\}$, where

$$C_{\text{BTA}} = \frac{7 - \sqrt{17}}{4} \left(1 + \frac{7}{3} \pi \frac{\rho}{W}\right)^{-1}. \quad (6)$$

- (ii) Update D to contain information of all (and only) the outstanding targets.
- (iii) Visit all non-empty beads once, visiting one target per non-empty bead.
- (iv) Repeat step (ii).

The following result characterizes the system time for the closed loop system induced by this algorithm and is based on the bound derived in Lemma 4.5.

Theorem 5.3: (System time for the BEAD-TILING ALGORITHM) For any $\rho > 0$ and $\lambda > 0$, the BEAD-TILING ALGORITHM is a stable policy for the DTRP and the resulting system time T_{BTA} satisfies:

$$\limsup_{\lambda \rightarrow +\infty} \frac{T_{\text{DTRP}}}{\lambda^2} \leq \limsup_{\lambda \rightarrow +\infty} \frac{T_{\text{BTA}}}{\lambda^2} \leq 71 \rho W H \left(1 + \frac{7}{3} \pi \frac{\rho}{W}\right)^3.$$

Proof: Consider a generic bead B , with non-empty intersection with \mathcal{Q} . Target points within B will be generated according to a Poisson process with rate λ_B satisfying

$$\lambda_B = \lambda \frac{\text{Area}(B \cap \mathcal{Q})}{WH} \leq \lambda \frac{\text{Area}(B)}{WH} = \frac{C_{\text{BTA}}^3}{16\rho W H \lambda^2} + o\left(\frac{1}{\lambda^2}\right).$$

The vehicle will visit B at least once every L_1 time units, where L_1 is the bound on the length of a path through all beads, as computed in Lemma 4.5. As a consequence, targets in B will be visited at a rate no smaller than

$$\mu_B = \frac{C_{\text{BTA}}^2}{16\rho W H \lambda^2} \left(1 + \frac{7}{3} \pi \frac{\rho}{W}\right)^{-1} + o\left(\frac{1}{\lambda^2}\right).$$

In summary, the expected time T_B between the appearance of a target in B and its visit by the vehicle is no more than the system time in a queue with Poisson arrivals at rate λ_B , and deterministic service rate μ_B . Such a queue is called a $M/D/1$ queue in the literature [37], and its system time is known to be

$$T_{M/D/1} = \frac{1}{\mu_B} \left(1 + \frac{1}{2} \frac{\lambda_B}{\mu_B - \lambda_B}\right).$$

Using the computed bounds on λ_B and μ_B , and taking the limit as $\lambda \rightarrow +\infty$, we obtain

$$\begin{aligned} \limsup_{\lambda \rightarrow +\infty} \frac{T_B}{\lambda^2} &\leq \limsup_{\lambda \rightarrow +\infty} \frac{T_{M/D/1}}{\lambda^2} \\ &\leq \frac{16\rho W H}{C_{\text{BTA}}^2} \left(1 + \frac{7}{3} \pi \frac{\rho}{W}\right)^{-1} \left(1 + \frac{1}{2} \frac{C_{\text{BTA}}}{\left(1 + \frac{7}{3} \pi \frac{\rho}{W}\right)^{-1} - C_{\text{BTA}}}\right). \end{aligned} \quad (7)$$

Since equation (7) holds for *any* bead intersecting \mathcal{Q} , the bound derived for T_B holds for all targets and is therefore

a bound on T_{BTA} . The expression on the right hand side of (7) is a constant that depends on problem parameters ρ , W , and H , and on the design parameter C_{BTA} , as defined in equation (6). Stability of the queue is established by noting that $C_{\text{BTA}} < (1 + 7/3 \pi \rho/W)^{-1}$. Additionally, the choice of C_{BTA} in equation (6) minimizes the right hand side of (7) yielding the numerical bound in the statement. ■

Remark 5.4: The achievable performance of the BEAD-TILING ALGORITHM provides a constant-factor approximation to the lower bound established in Theorem 5.1. Also, there exists no stable policy for the DTRP when the targets are generated in an adversarial worst-case fashion with $\lambda \geq (\pi\rho)^{-1}$. This fact is a consequence of the linear lower bound on the DTSP for worst-case point sets derived in Theorem 3.1. □

VI. CONCLUSIONS

In this paper, we have studied the TSP for vehicles that follow paths of bounded curvature in the plane. For the worst-case and the stochastic settings, we have obtained upper bounds that are within a constant factor of the lower bound; the upper bounds are constructive in the sense that they are achieved by novel algorithms. It is interesting to compare our results with the Euclidean setting (i.e., the setting in which vehicle paths do not have curvature constraints). For a given compact set and a point set P of n points, it is known [19], [20] 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 *worst-case* DTSP $_{\rho}(P)$ belongs to $\Theta(n)$ and the *stochastic* DTSP $_{\rho}(P)$ belongs to $\Theta(n^{2/3})$ with high probability.

Remarkably, the differences between these various bounds play a crucial role when studying the DTRP; e.g., stable policies exist only when the TSP cost grows strictly sub-linearly with n . For the DTRP we have proposed the novel policy and shown its stability for a uniform target-generation process with intensity λ . It is known [32] that the system time for the DTRP for the Dubins vehicle belongs to $\Omega(\lambda^2)$; the policy proposed in this paper shows that the system time belongs to $O(\lambda^2)$. Thus, the system time of the DTRP for the Dubins vehicle belongs to $\Theta(\lambda^2)$. This result differs from the result in the Euclidean case, where it is known [25] that the system time belongs to $\Theta(\lambda)$. Therefore, our analysis rigorously establishes the following intuitive fact: bounded-curvature constraints make the system much more sensitive to increases in the target generation rate.

Future directions of research include finding a *single* algorithm which would provide constant factor approximation to the DTSP for the worst case *as well as* the stochastic setting. It is also interesting to consider the *non-uniform* stochastic DTSP when the points to be visited are sampled according to a non-uniform probability distribution. Other avenues of future research are to use the tools developed in this paper to study Traveling Salesperson Problems for other dynamical vehicles, study centralized and decentralized versions of the DTRP and general task assignment and surveillance problems for multi-Dubins (and other dynamical) vehicles.

ACKNOWLEDGMENT

This material is based upon work supported in part by ONR YIP Award N00014-03-1-0512, AFOSR MURI Award F49620-02-1-0325 and NSF Awards 0705451 and 0705453. The authors thank Professor João Hespanha and Jerome Le Ny for interesting conversations. The authors would also like to thank anonymous referees for several useful and insightful comments.

REFERENCES

- [1] 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), pp. 786–791, June 2005.
- [2] K. Savla, F. Bullo, and E. Frazzoli, "On traveling salesperson problems for Dubins' vehicle: stochastic and dynamic environments," in *IEEE Conf. on Decision and Control and European Control Conference*, (Seville, Spain), pp. 4530–4535, Dec. 2005.
- [3] K. Savla, E. Frazzoli, and F. Bullo, "On the Dubins Traveling Salesperson Problems: Novel approximation algorithms," in *Robotics: Science and Systems II (Proceedings of the Second RSS Conference, August 2006, Philadelphia PA)* (G. S. Sukhatme, S. Schaal, W. Burgard, and D. Fox, eds.), Cambridge, MA: MIT Press, 2007. Available electronically at <http://arxiv.org/abs/cs.RO/0603010>.
- [4] P. R. Chandler, M. Pachter, D. Swaroop, J. M. Fowler, J. K. Howlett, S. Rasmussen, C. Schumacher, and K. Nygard, "Complexity in UAV cooperative control," in *American Control Conference*, (Anchorage, AK), pp. 1831–1836, 2002.
- [5] Z. Tang and Ü. Özgüner, "Motion planning for multi-target surveillance with mobile sensor agents," *IEEE Transactions on Robotics*, vol. 21, no. 5, pp. 898–908, 2005.
- [6] 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. 6, pp. 911–922, 2002.
- [7] S. Rathinam, R. Sengupta, and S. Darbha, "A resource allocation algorithm for multi-vehicle systems with non holonomic constraints," *IEEE Transactions on Automation Sciences and Engineering*, vol. 4, no. 1, pp. 98–104, 2007.
- [8] 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.
- [9] A. E. Gil, S. Ganapathy, K. M. Passino, and A. Sparks, "Cooperative scheduling of tasks for networked uninhabited autonomous vehicles," in *IEEE Conf. on Decision and Control*, (Maui, Hawaii), pp. 522–527, Dec. 2003.
- [10] A. Richards, J. Bellingham, M. Tillerson, and J. How, "Coordination and control of multiple UAVs," in *AIAA Conf. on Guidance, Navigation and Control*, (Monterey, CA), Aug. 2002.
- [11] A. A. Markov, "Some examples of the solution of a special kind problem in greatest and least quantities," *Soobshch. Kharkovsk. Mat. Obshch.*, vol. 1, pp. 250–276, 1887. (in Russian).
- [12] 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.
- [13] J.-D. Boissonnat, A. Cérézo, and J. Leblond, "Shortest paths of bounded curvature in the plane," *Journal of Intelligent and Robotic Systems*, vol. 11, pp. 5–20, 1994.
- [14] A. M. Shkel and V. J. Lumelsky, "Classification of the Dubins set," *Robotics and Autonomous Systems*, vol. 34, pp. 179–202, 2001.
- [15] J.-D. Boissonnat and S. Lazard, "A polynomial-time algorithm for computing a shortest path of bounded curvature amidst moderate obstacles," *International Journal of Computational Geometry & Applications*, vol. 13, pp. 189–229, 2003.
- [16] C. Tomlin, I. Mitchell, and R. Ghosh, "Safety verification of conflict resolution manoeuvres," *IEEE Transactions on Intelligent Transportation Systems*, vol. 2, no. 2, pp. 110–120, 2001.
- [17] S. M. LaValle, *Planning Algorithms*. Cambridge, UK: Cambridge University Press, 2006.
- [18] U. Boschain and B. Piccoli, *Optimal Syntheses for Control Systems on 2-D Manifolds*. Mathématiques et Applications, New York: Springer Verlag, 2004.

- [19] 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.
- [20] 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.
- [21] 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.
- [22] S. Arora, "Nearly linear time approximation scheme for Euclidean TSP and other geometric problems," in *IEEE Symposium on Foundations of Computer Science*, (Miami Beach, FL), pp. 554–563, Oct. 1997.
- [23] S. Lin and B. W. Kernighan, "An effective heuristic algorithm for the traveling-salesman problem," *Operations Research*, vol. 21, pp. 498–516, 1973.
- [24] A. Aggarwal, D. Coppersmith, S. Khanna, R. Motwani, and B. Schieber, "The angular-metric traveling salesman problem," *SIAM Journal on Computing*, vol. 29, no. 3, pp. 697–711, 1999.
- [25] 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.
- [26] G. Yang and V. Kapila, "Optimal path planning for unmanned air vehicles with kinematic and tactical constraints," in *IEEE Conf. on Decision and Control*, (Las Vegas, NV), pp. 1301–1306, 2002.
- [27] D. H. Walker, T. W. McLain, and J. K. Howlett, "Coordinated UAV target assignment using distributed tour calculations," in *Theory and Algorithms for Cooperative Systems (Conference on Cooperative Control and Optimization, Destin, FL, Nov 2003)* (D. Grundel, R. Murphey, and P. M. Pardalos, eds.), pp. 535–564, Singapore: World Scientific Publishing, 2004.
- [28] J. Le Ny and E. Feron, "An approximation algorithm for the curvature-constrained traveling salesman problem," in *Allerton Conference on Communications, Control and Computing*, (Monticello, IL), Sept. 2005.
- [29] S. Darbha, "Combinatorial motion planning for a collection of Reeds-Shepp vehicles," tech. rep., ASEE/AFOSR SFFP, AFRL, Eglin, Aug. 2005.
- [30] S. Itani and M. A. Dahleh, "On the average length of the stochastic TSP with Dubins' metric," tech. rep., LIDS, Massachusetts Institute of Technology, 2005.
- [31] G. Percus and O. C. Martin, "Finite size and dimensional dependence of the Euclidean traveling salesman problem," *Physical Review Letters*, vol. 76, no. 8, pp. 1188–1191, 1996.
- [32] 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.
- [33] R. Motwani and P. Raghavan, *Randomized Algorithms*. Cambridge, UK: Cambridge University Press, 1995.
- [34] Y. Azar, A. Z. Broder, A. R. Karlin, and E. Upfal, "Balanced allocations," *SIAM Journal on Computing*, vol. 29, no. 1, pp. 180–200, 1999.
- [35] W. Feller, *An Introduction to Probability Theory and Its Applications*, vol. 1. New York: John Wiley, 1968.
- [36] J. M. Steele, *Probability Theory and Combinatorial Optimization*. Philadelphia, PA: SIAM, 1987.
- [37] L. Kleinrock, *Queueing Systems. Volume I: Theory*. New York: John Wiley, 1975.
- [38] K. Savla, *Multi UAV Systems with Motion and Communication Constraints*. PhD thesis, Electrical and Computer Engineering Department, University of California at Santa Barbara, Santa Barbara, Aug. 2007. Available electronically at <http://ccdc.mee.ucsb.edu>.

APPENDIX

We first present a sketch of the proof of Theorem 3.1; a more detailed proof is presented in [38].

A. Proof of Theorem 3.1

We first describe the construction of the set $P \in \mathcal{P}_n$ for which the statement holds true. Let C_r be a circle of radius $r < \rho$ with center at the origin. For $i \in \{1, \dots, n\}$, define the i^{th} point b_i by

$$b_i = (r \cos(2\pi i/n), r \sin(2\pi i/n)).$$

This definition ensures that $b_i \neq b_j$ for $i \neq j$. Let $P(r) = \{b_1, \dots, b_n\}$. Let $F = (f_1, \dots, f_n)$ be the (possibly, suboptimal) order of points which the Dubins vehicle will go through while executing any algorithm (not necessarily the optimal algorithm) over $P(r)$. Let τ denote the closed path followed by the Dubins vehicle. Let D_r be a closed disk of radius r with center at the origin.

$$\text{Length}(\tau) = \text{Length}(\tau \text{ inside } D_r) + \text{Length}(\tau \text{ outside } D_r).$$

Replacing τ by segments to join the n points in $P(r)$ gives the following inequality:

$$\begin{aligned} & \text{Length}(\tau \text{ inside } D_r) \\ & \geq \text{Length}(\tau \text{ inside } D_r \text{ replaced by segments}) \\ & = \sum_{i=1}^{n-1} \|f_i - f_{i+1}\| + \|f_n - f_1\| \\ & \quad - \text{Length}(\tau \text{ outside } D_r \text{ replaced by segments}). \end{aligned}$$

Therefore, the total length of τ can be lower bounded as follows:

$$\begin{aligned} \text{Length}(\tau) & \geq \sum_{i=1}^{n-1} \|f_i - f_{i+1}\| + \|f_n - f_1\| \\ & \quad + \text{Length}(\tau \text{ outside } D_r) \\ & \quad - \text{Length}(\tau \text{ outside } D_r \text{ replaced by segments}) \end{aligned} \quad (8)$$

Let v be the number of point-to-point paths contained in that part of τ which lies outside D_r . Since the length of the longest segment lying entirely in D_r is $2r$, $\text{Length}(\tau \text{ outside } D_r \text{ replaced by segments})$ will then be upper bounded by $2vr$. Also, $\sum_{i=1}^{n-1} \|f_i - f_{i+1}\| + \|f_n - f_1\|$ is lower bounded by the length of the ETSP tour over $P(r)$. This together with equation (8) gives that

$$\text{Length}(\tau) \geq \text{ETSP}(P(r)) + \text{Length}(\tau \text{ outside } D_r) - 2vr. \quad (9)$$

From [12] it follows that under the minimum radius of curvature constraint, for its optimality, τ is composed of line segments and arcs of circle of radius ρ . Let $\{\zeta_1, \dots, \zeta_v\}$ denote the angular displacements of the vehicle as it travels along τ outside D_r along its v point-to-point sections. Then,

$$\text{Length}(\tau \text{ outside } D_r) \geq \sum_{i=1}^v \zeta_i \rho. \quad (10)$$

From equations (9) and (10) it follows that

$$\text{Length}(\tau) \geq \text{ETSP}(P(r)) + \sum_{i=1}^v \zeta_i \rho - 2vr. \quad (11)$$

Now, we use the fact that as $r \rightarrow 0$, $\zeta_i \rightarrow 2\pi$ for all i . By taking the limit in (11) as $r \rightarrow 0^+$, we obtain

$$\text{Length}(\tau) \geq \text{ETSP}(P(r)) + 2\pi v \rho. \quad (12)$$

The inequality (12) holds true for any algorithm over the set P . Therefore, it holds true for the optimal algorithm when v attains its minimum value of $\lfloor n/2 \rfloor$. Substituting this value of v in (12) we obtain the desired lower bound.

B. Dubins classification of optimal curves

Following [12], the minimum length feasible curve for the Dubins vehicle is either (i) an arc of a circle of radius ρ , followed by a line segment, followed by an arc of a circle of radius ρ , or (ii) a sequence of three arcs of circles of radius ρ , or (iii) a sub-path of a path of path type (i) or (ii). To specify the *type* of these minimum length feasible curves for the Dubins path we follow the notations used in [14]. Three elementary motions are considered: turning to the left, turning to the right (both along a circle of radius ρ), and straight line motion S . Three operators are introduced: L_v (for left/counterclockwise turn of length $v > 0$), R_v (for right/clockwise turn of length $v > 0$), S_v (for straight motion of length $v > 0$). The operators L_v , R_v , and S_v , transform an arbitrary configuration $(x, y, \psi) \in \text{SE}(2)$ into its corresponding image point in $\text{SE}(2)$ by

$$\begin{aligned} &(x + \sin(\psi + v) - \sin \psi, y - \cos(\psi + v) + \cos \psi, \psi + v), \\ &(x - \sin(\psi - v) + \sin \psi, y + \cos(\psi - v) - \cos \psi, \psi - v), \\ &(x + v \cos \psi, y + v \sin \psi, \psi), \end{aligned}$$

respectively. Thus, the *Dubins set* \mathcal{D} which is the domain for the type of the minimum length feasible curve for a Dubins vehicle between a given initial and final configuration is given by $\mathcal{D} = \{LSL, RSR, RSL, LSR, RLR, LRL\}$. One may refer to [12] for a detailed discussion on the construction of these path types between a given initial and final configuration. One may note that there are sets of initial and final configurations for which all the path types may not be feasible between those configurations.

In the remaining part of the paper we will need to frequently use the curves of type LRL and RLR starting with the initial configuration $(0, 0, 0)$ and the final configuration $(0, 0, \psi)$. We introduce some additional notations to facilitate presentation of the same. We introduce notations for the path type LRL . For $\psi \neq 0$, let $C_{p_1}(\psi)$ be a circle with center $O_{C_{p_1}} := (0, \rho)$ and radius ρ , and let $C_{p_2}(\psi)$ be a circle with center $O_{C_{p_2}} := (-\rho \sin \psi, \rho \cos \psi)$ and radius ρ . Note that $\psi \neq 0$ implies that $C_{p_1}(\psi) \cap C_{p_2}(\psi)$ is either a point or 2 points. Then let $C_{m_1}(\psi)$ and $C_{m_2}(\psi)$ be two circles with radius ρ that are tangent to both $C_{p_1}(\psi)$ and $C_{p_2}(\psi)$, see Figure 5 and Figure 6.




Fig. 5. LRL curves returning to the origin for $\psi \in [0, \pi]$. Fig. 6. LRL curves returning to the origin for $\psi \in]\pi, 2\pi[$.

By construction, $C_{p_1}(\psi)$ intersects $C_{m_1}(\psi)$ and $C_{m_2}(\psi)$ at one point each: let $P_1(\psi)$ be the first of these two points that is reached moving left from the origin O along $C_{p_1}(\psi)$. Without

loss of generality, assume $P_1(\psi) \in C_{m_1}(\psi)$. Let $O_{C_{m_1}}$ be the center of C_{m_1} . Let $P_2(\psi) = C_{m_1}(\psi) \cap C_{p_2}(\psi)$. In order to remove ambiguity, we shall pick that heading of the tangent line to a circle at a given point which is consistent with the orientation of that circle to be *the* orientation of the tangent to that circle at that point. Let the orientation of the Dubins vehicle at P_1 be along the orientation of the tangent to C_{p_1} at P_1 . Similarly, let the orientation of the Dubins vehicle at P_2 be along the orientation of the tangent to C_{p_2} at P_2 . Let the vehicle configuration at P_1 and P_2 be denoted by $J_{p_1}, J_{p_2} \in \text{SE}(2)$, respectively. Let t_1, t_2, t_3 be such that $L_{t_1}(0, 0, 0) = J_{p_1}$, $R_{t_2}(J_{p_1}) = J_{p_2}$ and $L_{t_3}(J_{p_2}) = (0, 0, \psi)$. Let $LRL_O(\psi)$ and $RLR_O(\psi)$ be the minimum length curves of types LRL and RLR respectively from the configuration $(0, 0, 0)$ to the configuration $(0, 0, \psi)$.

For $\psi \neq 0$, we define *forbidden cones* $V_1, V_2 : [0, 2\pi[\rightarrow \mathbb{R}^2$ to be the open, positive cones with symmetry axes $(d, \psi/2)_{d \in \mathbb{R}_+}$ and $(d, \pi + \psi/2)_{d \in \mathbb{R}_+}$, respectively, and half angle for both of them given by $\alpha(\psi) = \pi/2 - \cos^{-1}(\frac{\sin(\psi/2)}{2})$. We also write $V_1^c(\psi) = \mathbb{R}^2 \setminus V_1(\psi)$ and $V_2^c(\psi) = \mathbb{R}^2 \setminus V_2(\psi)$.

C. Proof of Theorem 3.4

We begin with some preliminary results. To keep the presentation simple in this section, we either sketch the proofs or omit them altogether. We refer to [38] for detailed proofs.

We start by providing bounds on the Dubins path length starting and ending at the same points. One can prove that for $d = 0$, the minimal length feasible curve for the Dubins vehicle is of type LRL or RLR .

Lemma A.1 (Optimal path length returning to the origin): Let $d = 0$ and $\theta \in [0, 2\pi[$.

- (i) if $\psi \in]0, \pi]$, then $LRL_O(\psi)$ is the optimal path and $C_\rho(0, \theta, \psi) = \rho\psi + 4\rho \cos^{-1}\left(\frac{\sin(\psi/2)}{2}\right)$,
- (ii) if $\psi \in]\pi, 2\pi[$, then $RLR_O(\psi)$ is the optimal path and $C_\rho(0, \theta, \psi) = \rho(2\pi - \psi) + 4\rho \cos^{-1}\left(\frac{\sin(\psi/2)}{2}\right)$.

Therefore, for all $\psi \in [0, 2\pi[$ and $\rho > 0$,

$$C_\rho(0, 0, \psi) \leq C_\rho(0, 0, \pi) = \frac{7}{3}\pi\rho.$$

Next, we start to analyze the general case where $(x, y) \neq (0, 0)$. In what follows, we let $(d, \theta) = \text{polar}(x, y)$ be the polar coordinates of $(x, y) \neq (0, 0)$ and, with a slight abuse of notation, we let $C_\rho(d, \theta, \psi) = C_\rho(x, y, \psi)$.

Lemma A.2: (Upper bound on the optimal length via LRL_O and RLR_O) For $\psi \in]0, 2\pi[$, and $(d, \theta) = \text{polar}(x, y)$,

- (i) if $(x, y) \in V_1^c(\psi)$, then $C_\rho(d, \theta, \psi) \leq d + \text{Length}(LRL_O(\psi))$,
- (ii) if $(x, y) \in V_2^c(\psi)$, then $C_\rho(d, \theta, \psi) \leq d + \text{Length}(RLR_O(\psi))$.

Therefore, for $(d, \theta) = \text{polar}(x, y)$ and $(x, y) \in \left(\bigcup_{\psi \in]0, \pi]} V_1^c(\psi)\right) \cup \left(\bigcup_{\psi \in]\pi, 2\pi[} V_2^c(\psi)\right)$,

$$C_\rho(d, \theta, \psi) \leq d + C_\rho(0, \theta, \psi) \leq d + \frac{7}{3}\pi\rho.$$

It now remains to obtain a bound on $C_\rho(d, \theta, \psi)$ when $(x, y) \in V_1(\psi)$ or $(x, y) \in V_2(\psi)$ where $(d, \theta) = \text{polar}(x, y)$.

To this effect, let the vehicle start moving at time $t = 0$ at unit speed along C_{p_1} in the counterclockwise direction and keep updating the parameters d, θ, ψ as if the coordinate system was moving along with the vehicle. Consequently $V_1(\psi)$ keeps shrinking and there is a time instant $t = t^*$ when the final configuration is such that $(x, y) \notin V_1(\psi)$. This construction along with Lemma A.2 gives the following result.

Lemma A.3: For $\psi \in]0, \pi[$, $(x, y) \in V_1(\psi)$, $(d, \theta) = \text{polar}(x, y)$ and $\rho > 0$,

$$C_\rho(d, \theta, \psi) \leq d + \rho F_1(\psi),$$

where $F_1(\psi)$ is as defined in equation (2).

From the definition, it follows that for $(x, y) \neq (0, 0)$, $(x, y) \in V_1(\psi)$ implies $(x, y) \in V_2^c(\psi)$. This observation along with Lemma A.1 and Lemma A.2 leads to the next lemma.

Lemma A.4: For $\psi \in]0, \pi[$, $(x, y) \in V_1(\psi)$, $(d, \theta) = \text{polar}(x, y)$ and $\rho > 0$,

$$C_\rho(d, \theta, \psi) \leq d + \rho F_2(\psi),$$

where $F_2(\psi)$ is as defined in equation (3).

Combining Lemma A.3 and Lemma A.4, one gets the following result.

Lemma A.5: For $\psi \in]0, \pi[$, $(x, y) \in V_1(\psi)$, $(d, \theta) = \text{polar}(x, y)$ and $\rho > 0$,

$$\begin{aligned} C_\rho(d, \theta, \psi) &\leq d + \rho \max\{F_2(\pi), \sup_{\psi \in]0, \pi[} \min\{F_1(\psi), F_2(\psi)\}\} \\ &= d + \kappa \pi \rho. \end{aligned}$$

Similarly, one can prove that for $\psi \in]\pi, 2\pi[$, $(x, y) \in V_2(\psi)$, $(d, \theta) = \text{polar}(x, y)$ and $\rho > 0$, we have $C_\rho(d, \theta, \psi) \leq d + \kappa \pi \rho$. Combining this with Lemma A.2 and Lemma A.5, we can state that for $\psi \in]0, 2\pi[$, $(x, y) \in \mathbb{R}^2$, $(d, \theta) = \text{polar}(x, y)$ and $\rho > 0$,

$$C_\rho(d, \theta, \psi) \leq d + \kappa \pi \rho. \quad (13)$$

It now remains to prove a similar bound on $C_\rho(d, \theta, 0)$ for which we state the following lemma.

Lemma A.6: For $(x, y) \in \mathbb{R}^2$, $(d, \theta) = \text{polar}(x, y)$ and $\rho > 0$,

$$C_\rho(d, \theta, 0) \leq d + 2\pi \rho.$$

Lemma A.6 combined with equation (13) completes the proof for Theorem 3.4. It is easy to check that for $\psi \in]0, \pi[$, $F_1(\psi)$ is a monotonically increasing function of ψ and $F_2(\psi)$ is a monotonically decreasing function of ψ . Therefore, there exists a unique ψ^* such that $F_1(\psi^*) = F_2(\psi^*)$. By numerical calculations one can find that $\kappa \simeq 2.6575$.

D. Numerical Results

The length of the optimal Dubins path, $C_\rho(d, \theta, \psi)$, was calculated for numerous sets of final configurations (d, θ, ψ) starting with an initial configuration of $(0, 0, 0)$ and the corresponding values of the quantity $\frac{C_\rho(d, \theta, \psi) - d}{\pi \rho}$ were evaluated for each of the instances. The results suggest that these values are bounded by $\frac{7}{3}$. Moreover, it appears that the value of $\frac{7}{3}$ is achieved only when the Dubins vehicle makes a transition

from a state of the form $(0, 0, 0)$ to a state of the form $(0, 0, \pi)$ according to our setup. Hence, we conjecture that the value of κ , which is the upper bound for the quantity $\frac{C_\rho(d, \theta, \psi) - d}{\pi \rho}$, is $\frac{7}{3}$.