On the Curvature-Constrained Traveling Salesman Problem

Jerome Le Ny, *Member, IEEE,* Eric Feron, *Member, IEEE,*
and Emilio Frazzoli, *Member, IEEE*

Abstract

We study the traveling salesman problem for a Dubins vehicle. We prove that this problem is NP-hard, and provide lower bounds on the approximation ratio achievable by some recently proposed heuristics. We also describe new algorithms for this problem based on heading discretization, and evaluate their performance numerically.

I. Introduction

In an instance of the traveling salesman problem (TSP) we are given the distances between any pair of n points. The problem is to find the shortest closed path (tour) visiting every point exactly once. We also call this problem the tour-TSP to distinguish it from the path-TSP, where the requirement that the vehicle must start and end at the same point is removed. This famously intractable problem is often encountered in robotics, and has traditionally been solved in two steps within the common layered controller architectures for mobile robots. At the higher decision-making level, the dynamics of the robot are usually not taken into account and the mission planner might typically chose to solve the TSP for the Euclidean metric (ETSP), i.e., using the Euclidean distances between waypoints. This determines the order in which the waypoints

This work was supported by Air Force - DARPA - MURI award 009628-001-03-132 and Navy ONR award N00014-03-1-0171. Preliminary versions of this work appeared in [1], [2].

J. Le Ny is with the school of Electrical Engineering, University of Pennsylvania, PA 19104, USA jeromel@seas.upenn.edu. E. Frazzoli is with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139, USA frazzoli@mit.edu. E. Feron is with the School of Aerospace Engineering, Georgia Institute ot Technology, Atlanta, GA 30332, USA eric.feron@aerospace.gatech.edu.

should be visited by the robot. The path planner then designs a trajectory joining each pair of successive waypoints and respecting the dynamics of the robot, by solving a succession of shortest path problems. One can thus directly exploit the existing results on the ETSP (or more general combinatorial TSPs) and on the shortest path problem with differential constraints, which are among the most studied problems in optimization and robotics. See e.g. [3]–[8] for a small sample of the available literature.

Even if each problem is solved optimally however, the ad-hoc separation into two successive steps can be inefficient, since the sequence of points chosen by the TSP algorithm is often hard to follow for the physical system. In order to improve the performance of unmanned aerial systems in particular, researchers are now working on integrating the mission planning and path planning stages [9]. In this note we consider the traveling salesman problem for the Dubins vehicle (DTSP), in a *planar environment without obstables*. The Dubins vehicle [8], [20] can only move forward in the plane, at constant speed, and has a limited turning radius. This is a good kinematic model for fixed wing aircraft, which can produce relevant trajectories for the path following controllers of these vehicles. At the same time, we can quickly compute the length of the shortest path between any two configurations of the Dubins vehicle, a critical building block to design efficient algorithms with good performance for the DTSP.

A stochastic version of the DTSP for which the points are distributed randomly and uniformly in the plane was considered in [10]–[13]. Here however, we focus on algorithms and worst-case bounds for the more standard problem where no probability distribution is given for the input. In that case, most of the recently proposed algorithms seem to build on a preliminary solution obtained for the ETSP [14]–[17]. We refer the reader to our more detailed survey of the existing algorithms in [2].

Contributions of this work. In this note, we first prove that the DTSP is NP-hard, thus justifying the work on heuristics and algorithms that approximate the optimal solution. On the negative side, we give some lower bounds on the approximation ratio achievable by recently proposed heuristics. Following a tour based on the ETSP ordering or the ordering of Tang and Özgüner [9] cannot achieve an approximation ratio better than $\Omega(n)$[1]. The same is true for the nearest

[1]We say $f(n) = O(g(n))$ if there exists $c > 0$ such that $f(n) \leq cg(n)$ for all n, and $f(n) = \Omega(g(n))$ if there exists $c > 0$ such that $f(n) \geq cg(n)$ for all n. See section II for the definition of the approximation ratio of an approximation algorithm.

neighbor heuristic, in contrast to the ETSP where it achieves a $O(\log n)$ approximation [18]. Then we propose an algorithm based on heading discretization, which emerges naturally from the previous work on the curvature-constrained shortest path [5]. Its theoretical performance is of the same type as for these shortest-path algorithms and does not improve on the previously mentioned heuristics. However numerical simulations show a significant performance improvement in randomly generated instances over other heuristics when the inter waypoint distances are smaller than the turning radius of the vehicle.

The rest of this paper is organized as follows. We recall some facts about the Dubins paths in Section II and reduce the DTSP to a finite dimensional optimization problem. This section also contains the lower bound on approximation ratios for various heuristics. In Section III we show that the DTSP is NP-hard. Section IV describes our algorithms based on heading discretization. They return in time $O(n^3)$ a tour within $O\left(\min\left(\left(1+\frac{\rho}{\epsilon}\right)\log n, \left(1+\frac{\rho}{\epsilon}\right)^2\right)\right)$ of the optimum, where ρ is the minimum turning radius of the vehicle and ϵ is the minimum Euclidean distance between any two waypoints. Note that throughout the paper, we fix ρ but ϵ is allowed to depend on the problem instance. In particular if the waypoints are sampled in a compact environment we have necessarily $\epsilon = O(1/\sqrt{n})$. Finally, Section V discusses some numerical simulations.

II. HEURISTICS PERFORMANCE

A Dubins vehicle in the plane has its configuration described by its position and heading $(x, y, \theta) \in \mathbb{R}^2 \times S^1$. Its equations of motion are

$$\dot{x} = v_0 \cos(\theta), \quad \dot{y} = v_0 \sin(\theta), \quad \dot{\theta} = \frac{v_0}{\rho} u, \quad \text{with } u \in [-1, 1],$$

where ρ is the minimum turning radius of the vehicle, and u is the available control. Without loss of generality, we assume that the speed v_0 of the vehicle is normalized to 1. The DTSP asks, for a given set of points in the plane, to find the shortest tour through these points that is feasible for a Dubins vehicle. Since we show below that this problem is NP-hard, we focus on the design of approximation algorithms. An α-approximation algorithm (with *approximation ratio* $\alpha \geq 1$) for a minimization problem is an algorithm that produces *on any instance of the problem* with optimum OPT, a feasible solution whose value Z is within a factor α of the optimum, i.e., such that $OPT \leq Z \leq \alpha\, OPT$. In general, α is allowed to depend on the input parameters of the problem, such as the number of points in the DTSP. This definition of approximation ratio, used throughout the paper, corresponds to the worst-case performance of the algorithm [19].

Dubins [20] characterized curvature-constrained shortest paths between an initial and a final configuration. Let P be a feasible path. We call a nonempty subpath of P a C-segment or an S-segment if it is a circular arc of radius ρ or a straight line segment, respectively.

Theorem 1 ([20]): An optimal path between any two configurations in an environment without obstacles is of type CCC or CSC, or a subpath of a path of either of these two types.

We refer to these minimal-length paths as Dubins paths. When a subpath is a C-segment, it can be a left or a right hand turn: denote these two types of C-segments by L and R respectively. The C-segments starting from a configuration $\chi = (x, y, \theta)$ are circular arcs on one of the two circles of radius ρ tangent to the direction θ at (x, y) denoted $\mathcal{C}_L^\chi$ and $\mathcal{C}_R^\chi$ respectively, boundaries of the closed disks denoted $\mathcal{D}_L^\chi$ and $\mathcal{D}_R^\chi$. By Theorem 1, the minimum length path between an initial and a final configuration can be found among the six paths $\{LSL, RSR, RSL, LSR, RLR, LRL\}$. Each of these paths can be explicitly computed and therefore finding the optimum path and length between any two configurations can be done in constant time [20]. Solving the DTSP reduces then to choosing a permutation of the points specifying in which order to visit them, as well as choosing a heading for the vehicle at each of these points. If the Euclidean distances between the waypoints to visit are large with respect to ρ, the DTSP and ETSP behave similarly (see e.g. [11]). Accordingly, researchers have tried in previous work to apply to the DTSP the waypoint ordering optimal for the ETSP [14]–[16], and have concentrated on the choice of headings. Theorem 2 provides a limit on the performance one can achieve using such a technique, which becomes particularly significant when the points are densely distributed with respect to ρ.

Before stating the theorem, we describe two other simple heuristics for the DTSP. The *nearest neighbor heuristic* produces a complete solution for the DTSP, including a waypoint ordering and a heading at each point. We start with an arbitrary point, and choose its heading arbitrarily, fixing an initial configuration. Then at each step, we find a point which is not yet on the path but closest to the last added configuration according to the Dubins metric. This is possible since we also have a complete characterization of the Dubins distance and path between an initial configuration and a final point with free heading [21]. We add this closest point to the path with the associated optimal arrival heading. When all nodes have been added to the path, we add a Dubins path connecting the last obtained configuration and the initial configuration. It is known that the nearest neighbor heuristic achieves a $O(\log n)$ approximation ratio for the ETSP, which is a particular case of the symmetric TSP [18]. The second heuristic, due to Tang and

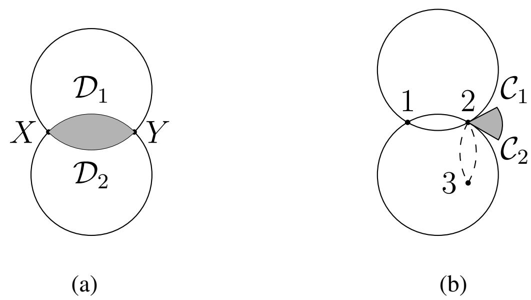

Fig. 1. A path between X and Y contained in the shaded region $\mathcal{D}_1 \cap \mathcal{D}_2$ is called a direct path. On Fig. 1(b), the vehicle cannot pass through points $1, 2$ and 3 using only direct paths. For direct paths $1 \to 2$, we show the range of possible final headings at 2, delimited by the tangent directions to $\mathcal{C}_1$ and $\mathcal{C}_2$. We also delimit the region of direct paths $2 \to 3$.

Özgüner [9], only produces a waypoint ordering (in [9], the authors then produce locally optimal headings for this choice of waypoint ordering using a gradient descent algorithm). To construct this ordering, we find the geometric center G of the waypoints, and calculate the orientation angle of each waypoint with respect to G. We then sort the points by increasing values of their orientation to determine the traverse order.

Theorem 2: Any algorithm for the DTSP which always follows the ordering of points that is optimal for the ETSP has an approximation ratio α_n which is $\Omega(n)$, where n is the number of points. If we impose a lower bound ϵ sufficiently small on the minimum Euclidean distance between any two waypoints, then there exist constants C, C', such that the approximation ratio is not better than $\frac{C}{\epsilon + (C'/n)}$. These statements are also true for the nearest neighbor heuristic. Finally, any algorithm following the Tang and Özgüner ordering [9] also has an approximation ratio lower bounded by $\Omega(n)$.

Proof: There are exactly two circles $\mathcal{C}_1, \mathcal{C}_2$ of radius ρ passing through two points X and Y in the plane with Euclidean distance $\|X - Y\| < 2\rho$, see Fig. 1(a). These circles define the boundaries of two closed discs $\mathcal{D}_1$ and $\mathcal{D}_2$. Following [22], we call a path from X to Y a *direct path* if it is contained in $\mathcal{D}_1 \cap \mathcal{D}_2$, and a *detour path* otherwise. It is shown in [22] that a curvature constrained path from X to Y is of length strictly smaller than $\pi\rho$ if and only if it is a direct path. Now consider the configuration of points shown in fig. 2(a), with $\epsilon < \tilde{\epsilon} < \sqrt{2}\rho$. Let n be the number of points, and suppose $n = 4m + 1$, m an integer. For clarity we focus on the path-TSP problem (the extension to the tour-TSP case is easy, by adding a similar path in the reverse direction). The optimal Euclidean path-TSP is shown on Fig. 1(a) as well. Suppose now that a Dubins vehicle tries to follow the points in this order. Then for each sequence of 5 consecutive points the vehicle will have to execute at least two detour paths. For example, if the vehicle follows a direct path between points 1 and 2, it follows from a simple geometric argument that point 3 is in the disc $\mathcal{D}_2$ (see Fig. 1(b)). Moreover it is shown in [22] that the set of possible headings at 2 are directed outside of $\mathcal{D}_2$ as shown on Fig. 1(b), whereas the direct paths from 2 to 3 are contained in $\mathcal{D}_2$. Hence the path from 2 to 3 must be a detour path, i.e., of length greater than $\pi\rho$. The same argument for the points $2, 3, 4$ shows that a direct path between 2 and 3 must be followed by a detour path between 3 and 4. Hence the length of a curvature constrained path through points 1 to 5 in this order is lower bounded by $2\epsilon + 2\pi\rho$. The length of the Dubins path following the ETSP ordering will then be greater than $(n-1)(\epsilon+\pi\rho)/2$.

On the other hand, a Dubins vehicle can simply go through all the points on the top line, execute a U-turn of length at most $6\pi\rho$ (considering CCC paths [23, p.28]), and then go through the points on the lower line, providing an upper bound of $(n-1)\tilde{\epsilon}/2+6\pi\rho$ for the optimal solution. Since this is valid for all $\tilde{\epsilon}$ with $\epsilon < \tilde{\epsilon} < \sqrt{2}\rho$, we deduce that the worst case approximation ratio α_n of the algorithm is at least:

$$\alpha_n \geq \frac{(n-1)(\epsilon+\pi\rho)}{(n-1)\epsilon + 12\pi\rho}. \quad (1)$$

In particular if we choose $\epsilon \leq 1/(n-1)$ for problem instances with n points, we get $\alpha_n = \Omega(n)$.

For the nearest-neighbor heuristic, we use the configuration shown on Fig. 2(b), which includes the first configuration (waypoint and heading) χ chosen by the algorithm. If ϵ is small enough, all the subsequent points are in the interior of the discs $\mathcal{D}_L^\chi$ or $\mathcal{D}_R^\chi$, hence at Dubins distance greater than $\pi\rho$ of the configuration χ (see [21]). The nearest neighbor heuristic chooses one of these points and as $\epsilon \to 0$, it reaches this chosen point by a Dubins path that tends to a circle of radius ρ and with final heading which tends to the initial heading (in the limit $\epsilon \to 0$, a non trivial Dubins path from a point to itself is just a circle of radius ρ passing through the point). Overall it produces a path of length at least $n\pi\rho$ (in fact, this length tends to $n2\pi\rho$ as $\epsilon \to 0$), whereas there is clearly a (Dubins) straight path of length $C_1 + C_2 n\epsilon$ through these points, for some constants C_1 and C_2. Hence again $\alpha_n = \Omega(n)$.

Finally, for the Tang and Özgüner ordering, we close the path of Fig. 2(a) into a cycle to obtain the configuration of Fig. 2(c). The waypoints are on two concentric circles of radius R and $R-\epsilon$, with $R-\epsilon > \rho$. The proof then follows the same steps as for the ETSP ordering.

For the proposed ordering, each direct path between two points must be followed by a detour path, whereas a Dubins vehicle can turn around each circle once and execute a single maneuver after covering the first circle to reach the second circle. ■

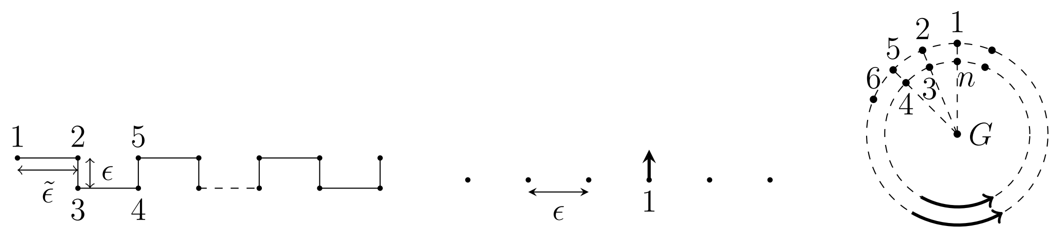

(a) ETSP ordering. (b) Nearest-neighbor heuristic. (c) Tang and Özgüner ordering.

Fig. 2. Waypoint configurations used as counter-examples to previously proposed ordering methods.

Remark 3: Using the notation of the theorem, the "Alternating Algorithm" proposed in [14] and detailed in section V achieves an approximation ratio of $1 + 3.48\frac{\rho}{\epsilon}$ [11] and is based on following the optimal ETSP ordering. Letting $n \to \infty$ in (1), this is not too far from our lower bound $1 + \pi\frac{\rho}{\epsilon}$ on the approximation ratio achievable by this ordering, which is therefore not too conservative.

III. Complexity of the DTSP

It is usually accepted that the DTSP is NP-hard and the goal of this section is to prove this claim rigorously. Note that adding the curvature constraint to the Euclidean TSP could well make the problem easier. [2] Hence, the statement does not follow trivially from the NP-hardness of the ETSP [25], [26]. In the proof of theorem 4, we consider, without loss of generality, the decision version of the problem, which we also call DTSP. That is, given a set of points in the plane and a number $L > 0$, DTSP asks if there exists a tour for the Dubins vehicle visiting all these points exactly once, of length at most L.

Theorem 4: Tour-DTSP and path-DTSP are NP-hard.

Proof: This is a corollary of Papadimitriou's proof of the NP-hardness of ETSP, to which we refer [25]. First recall the Exact Cover Problem: given a family F of subsets of the finite

[2] as in the bitonic TSP [24, p. 364] for example.

set U, is there a subfamily F' of F, consisting of disjoint sets, such that F' covers U? This problem is known to be NP-complete [27]. Papadimitriou described a polynomial-time reduction of Exact Cover to ETSP. That is, given an instance of the Exact Cover problem, we can construct an instance of the Euclidean Traveling Salesman Problem and a number L such that the Exact Cover problem has a solution if and only if the ETSP has an optimal tour of length less than or equal to L. The important fact to observe however, is that if Exact Cover does not have a solution, Papadimitriou's construction gives an instance of the ETSP which has an optimal tour of length $\geq (L+\delta)$, for some $\delta > 0$, and not just $> L$. More precisely, letting $a = 20$ exactly as in his proof, we can take $0 < \delta < \sqrt{a^2+1} - a$.

Now from [14], there is a constant C such that for any instance $\mathcal{P}$ of ETSP with n points and length $ETSP(\mathcal{P})$, the optimal DTSP tour for this instance has length less than or equal to $ETSP(\mathcal{P}) + Cn$. Then if we have n points in the instance of the ETSP constructed as in Papadimitriou's proof, we simply rescale all the distances by a factor $2Cn/\delta$. If Exact Cover has a solution, the ETSP instance has an optimal tour of length no more than $2CnL/\delta$ and so the curvature constrained tour has a length of no more than $2CnL/\delta + Cn$. If Exact Cover does not have a solution, the ETSP instance has an optimal tour of length at least $2CnL/\delta + 2Cn$, and the curvature constrained tour as well. So Papadimitriou's construction, rescaled by $2Cn/\delta$ and using $2CnL/\delta + Cn$ instead of L, where n is the number of points used in the construction, provides a reduction from Exact Cover to DTSP. ■

IV. An Algorithm Based on Heading Discretization

We now propose an algorithm for the DTSP which is inspired from a procedure used for the curvature-constrained shortest path problem, see [5]. It chooses a priori a finite set of possible headings at each point. Suppose, for simplicity, that we choose K headings for each point. We construct a graph with n clusters corresponding to the n waypoints, and each cluster containing K nodes corresponding to the choice of headings. Then, we compute the Dubins distances between configurations corresponding to pairs of nodes in distinct clusters. Finally, we would like to compute a tour through the n clusters which contains exactly one point in each cluster. This problem is called the generalized asymmetric traveling salesman problem, and can be reduced to a standard asymmetric traveling salesman problem (ATSP) over nK nodes [28]. This ATSP can in turn be solved directly using available software such as Helsgaun's implementation of the Lin-Kernighan heuristic [29], or using the $\log n$ approximation algorithm of Frieze et al. [30].

We found this method to perform very well in practice. The type of approximation result one can expect with such an a priori discretization of the headings is the same as in the shortest path case, that is, the path obtained is close to the optimum if the optimum is "robust". This means that the discretization must be sufficiently fine so that the algorithm can find a Dubins path close to the optimum which is of the same type, as defined in section II (see [7]). A limitation of the algorithm is that we must solve a TSP over nK points instead of n.

A. *Performance Bound for* $K = 1$.

Suppose that we choose $K = 1$ in the previous paragraph. Our algorithm can then be described as follows:

1) Fix the headings at all points, say to 0, or by choosing them randomly uniformly in $[-\pi, \pi)$, independently for each point.
2) Compute the $n(n-1)$ Dubins distances between all pairs of points.
3) Construct a complete graph with one node for each point and edge weights given by the Dubins distances.
4) We obtain a *directed* graph where the edges satisfy the triangle inequality. Compute an exact or approximate solution for the asymmetric TSP on this graph.

Next we derive an upper bound on the approximation ratio provided by this algorithm. Let us first introduce some results and notation that will be used in this derivation. We denote the Euclidean distance between two locations $X = (x, y)$ and $X' = (x', y')$ by $E(X, X')$. The Dubins distance between two configurations $\chi = (X, \theta)$ and $\chi' = (X', \theta')$ is denoted $D(\chi, \chi')$ (note that $D(\chi, \chi') \neq D(\chi', \chi)$). The first theorem is taken from [11, theorem 3.4].

Theorem 5: There exist a constant $\kappa \in [2.657, 2.658]$ such that for any two configurations $\chi = (X, \theta)$ and $\chi' = (X', \theta')$ we have

$$D(\chi, \chi') \leq E(X, X') + \kappa\pi\rho.$$

From the trivial bound $D(\chi, \chi') \geq E(X, X')$, we obtain the following corollary of theorem 5.

Corollary 6: Consider two choices of headings $\theta, \hat{\theta}$ at point X and $\theta', \hat{\theta}'$ at point X', with corresponding configurations $\chi, \hat{\chi}, \chi', \hat{\chi}'$. Then we have

$$D(\hat{\chi}, \hat{\chi}') \leq \left(1 + \frac{\kappa\pi\rho}{E(X, X')}\right) D(\chi, \chi').$$

Now denote by $\{\hat{\theta}_i\}_{i=1}^n$ the headings fixed in the first step of the algorithm, and by $\{\hat{\chi}_i = (X_i, \hat{\theta}_i)\}_{i=1}^n$ the corresponding configurations at the waypoint $\{X_i = (x_i, y_i)\}_{i=1}^n$. Let ϵ be the minimum Euclidean distance between any two waypoints

$$\epsilon = \min_{i \neq j} E(X_i, X_j).$$

As in the previous corollary, since $D(\hat{\chi}_j, \hat{\chi}_i) \geq E(X_j, X_i) = E(X_i, X_j)$, we have

$$\max_{i \neq j} \frac{D(\hat{\chi}_i, \hat{\chi}_j)}{D(\hat{\chi}_j, \hat{\chi}_i)} \leq 1 + \frac{\kappa\pi\rho}{\epsilon}. \tag{2}$$

With this bound on the arc distances, we can use a modified version of Christofides algorithm, also due to Frieze et al. [30], to obtain a $\frac{3}{2}\left(1 + \frac{\kappa\pi\rho}{\epsilon}\right)$ approximation for the ATSP in step 4. The time complexity of the first three steps of our algorithm is $O(n^2)$. To solve the ATSP, we can run the two algorithms of Frieze et al. [30] and choose the tour with minimum length, thus obtaining an approximation ratio of $\min\left(\log n, \frac{3}{2}\left(1 + \frac{\kappa\pi\rho}{\epsilon}\right)\right)$. This step solving the ATSP runs in time $O(n^3)$, so overall the running time of our algorithm is $O(n^3)$. The following theorem then describes the approximation ratio of our algorithm.

Theorem 7: Given a set of n points in the plane, the algorithm described above with $K = 1$ returns a Dubins traveling salesman tour with length within a factor

$$\min\left(\left(1 + \frac{\kappa\pi\rho}{\epsilon}\right)\log n, \frac{3}{2}\left(1 + \frac{\kappa\pi\rho}{\epsilon}\right)^2\right)$$

of the length of the optimum tour. The running time of this algorithm is $O(n^3)$.

Proof: Call OPT the optimal value of the DTSP, σ^* the corresponding optimal permutation specifying the order of the waypoints, and $\{\chi_i^*\}_{i=1}^n$ the optimal configurations. We have

$$OPT = \sum_{i=1}^{n-1} D(\chi^*_{\sigma^*(i)}, \chi^*_{\sigma^*(i+1)}) + D(\chi^*_{\sigma^*(n)}, \chi^*_{\sigma^*(1)}) =: L(\{\chi_i^*\}_{i=1}^n, \sigma^*).$$

Considering the permutation σ^* for the graph problem (where the edge weights are the distances $\{D(\hat{\chi}_i, \hat{\chi}_j)\}$) and $\hat{\sigma}^*$ the optimal permutation for the graph problem, we have

$$L(\{\hat{\chi}_i\}, \hat{\sigma}^*) \leq L(\{\hat{\chi}_i\}, \sigma^*) \leq \left(1 + \frac{\kappa\pi\rho}{\epsilon}\right) L(\{\chi_i^*\}, \sigma^*),$$

where the last inequality follows from corollary 6. We do not obtain the optimal permutation for the ATSP on the graph in general, instead we use the approximation algorithm mentioned above. Calling $\hat{\sigma}$ the permutation obtained, we have:

$$\begin{aligned} L(\{\hat{\chi}_i\}, \hat{\sigma}) &\leq \min\left(\log n, \frac{3}{2}\left(1 + \frac{\kappa\pi\rho}{\epsilon}\right)\right) L(\{\hat{\chi}_i\}, \hat{\sigma}^*) \\ &\leq \left[\min\left(\left(1 + \frac{\kappa\pi\rho}{\epsilon}\right)\log n, \frac{3}{2}\left(1 + \frac{\kappa\pi\rho}{\epsilon}\right)^2\right)\right] OPT. \end{aligned}$$

■

We note that the bound provided is in fact worse than the bound available for the "Alternating Algorithm" (AA), see remark 3. It is the same if we assume that the ATSP in step (4) is solved to optimality. However, experiments suggest that it performs in fact significantly better in general than the AA when the waypoint distances are smaller than the turning radius (see section V).

B. On the ρ/ϵ Term in Approximation Ratios

Note that every approximation ratio mentioned so far contains a ρ/ϵ term, where ϵ is the minimum Euclidean distance between any two waypoints. This term is particularly problematic for densely distributed waypoints, although it seems in general too pessimistic for points in general positions. Now the naive algorithm which considers all possible waypoint permutations and runs for each permutation the algorithm of Lee et al. to determine the headings [22] provides a tour within 5.03 of the optimum and runs in time $\Omega(n \cdot n!)$. We have also obtained an algorithm that runs faster than the naive algorithm, in time $O(n^2 \cdot 2^n)$, and achieves an approximation ratio of $O(\log n)$ [31]. It would be interesting to understand if one can design a polynomial-time algorithm that provides an approximation ratio free from the ρ/ϵ factor.

V. Numerical Simulations

Fig. 3 presents simulation results comparing the practical performance of different algorithms proposed for the DTSP. Points are generated randomly and uniformly in a 10×10 square. All the TSP tours (symmetric and asymmetric) are computed using the software LKH [29]. We compare the performance of the Alternating Algorithm (AA) [14], the nearest neighbor heuristic (NN) of section II, and the algorithm described in section IV. AA is based on the optimal Euclidean ordering. It computes the optimal ETSP, and keeps every other edge in the loop. These edges are straight lines followed by the vehicle, which must then connect the end of a straight path with the start of the next one by a Dubins path. Compared to our algorithm with the headings chosen randomly in step 1, we see that the performance of AA is similar for low point densities. In fact, AA clearly outperforms the randomized heading version of the algorithm if the waypoints are very sparsely distributed and the optimal tour tends to become the same as the optimal Euclidean tour. However, in scenarios with waypoints densely distributed with respect to the vehicle turning radius, which arise for example for UAVs in urban environments, or loitering weapons flying at high speed [17], large performance gains can be obtained by our algorithms over the AA.

Moreover for the choice of heading, even with just one discretization level, we can use the same heading as AA, and compute the solution of the ATSP using this heading (see step 4 in paragraph IV-A). This clearly always performs at least as well as AA, and the figure shows the significant performance improvement due to only changing the ordering, as the number of points increases. Also shown on the figure are the performance curves for an increasing number of discretization levels. With 5 discretization levels, a tour through 100 points can be computed on a standard laptop in about one minute, requiring the solution of an ATSP with 500 points, well below the limits of state-of-the-art TSP software. The asymptotic lower bound shown is taken from [12]. In particular, it is known that for the specific case of waypoints uniformly distributed in a rectangle, as in our experiments, the average length of the DTSP scales as $n^{2/3}$ [11], [12]. In fact in this case Savla et al. [11] have proposed an algorithm that is not based on the ETSP ordering and returns a tour whose length is within a constant factor of the optimum (in table I, the theoretical upper bound for this algorithm with our problem parameters is $193\, n^{2/3}$ but numerical experiments seem to suggest a smaller constant [11]).

Empirically, we observed the growth rates shown in table I for the different algorithms. With 10 discretization levels, the rate is close to optimal on such random inputs. Note also the good asymptotic performance of the nearest neighbor heuristic, which moreover is much easier to compute than the other heuristics. Although not used for the comparisons on Fig. 3, it is worth in practice including the headings of this nearest neighbor heuristic as part of the set of headings used in the discretization of section IV. Finally, fig. 4 shows an example of tours through 50 points computed using the Alternating Algorithm and our algorithm using 10 discretization levels.

Algorithm	Dubins tour average length
Alternating Algorithm	$5.3 \times n^{0.94}$
Nearest Neighbor Heuristic	$9.9 \times n^{0.69}$
Randomized Headings + ATSP	$9.3 \times n^{0.75}$
AA + ATSP with 1 discretization level	$8.2 \times n^{0.71}$
5 discretization levels	$6.7 \times n^{0.7}$
10 discretization levels	$6.6 \times n^{0.68}$

TABLE I

EXPERIMENTAL GROWTH OF THE AVERAGE DUBINS TOUR LENGTHS FOR THE DIFFERENT ALGORITHMS, WHEN THE WAYPOINTS ARE DISTRIBUTED RANDOMLY AND UNIFORMLY IN A 10-BY-10 SQUARE.

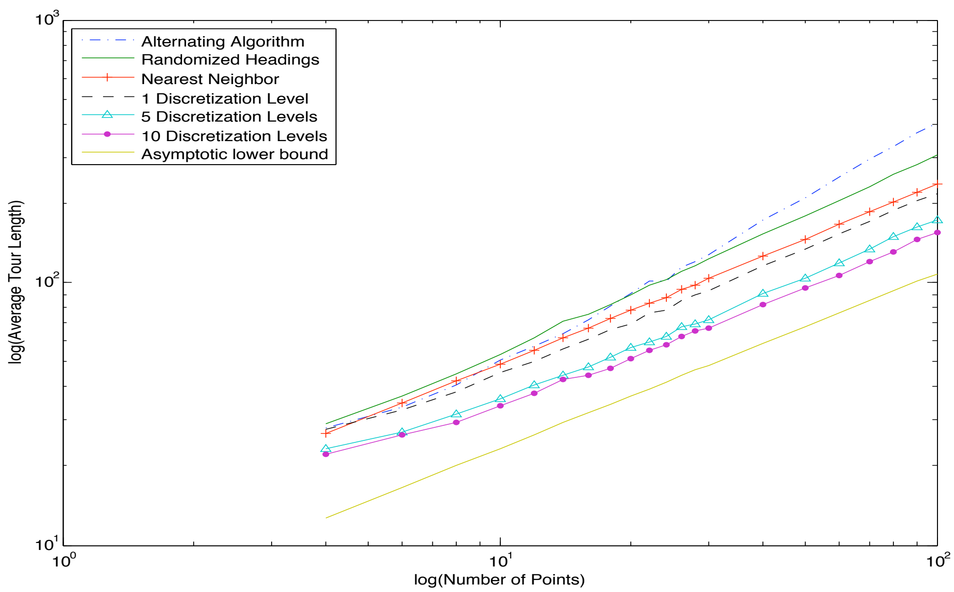

Fig. 3. Average tour length vs. number of points in a 10×10 square, on a log-log scale. The average is taken over 30 experiments for each given number of points. Except for the randomized heading choice, the heading of the AA is always included in the set of discrete headings in our algorithm (paragraph IV-A). Hence the difference between the AA curve and the 1 discretization level curve shows the performance improvement obtained by only changing the waypoint ordering.

REFERENCES

[1] J. Le Ny and E. Feron, "An approximation algorithm for the curvature-constrained traveling salesman problem," in *Proceedings of the 43rd Annual Allerton Conference on Communications, Control and Computing*, September 2005.

[2] J. Le Ny, E. Feron, and E. Frazzoli, "The curvature-constrained traveling salesman problem for high point densities," in

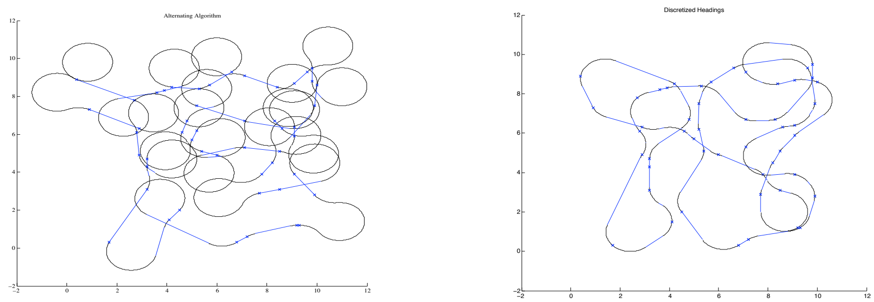

(a) Alternating Algorithm (b) 10 discretization levels.

Fig. 4. Dubins tours through 50 points randomly distributed in a 10×10 square. The turning radius of the vehicle is 1.

Proceedings of the 46th IEEE Conference on Decision and Control, 2007, pp. 5985 – 5990.

[3] E. Lawler, J. Lenstra, A. R. Kan, and D. Shmoys, *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization*. Wiley, 1985.

[4] S. Arora, "Polynomial time approximation schemes for Euclidian traveling salesman and other geometric problems," *Journal of the ACM*, vol. 45, no. 5, pp. 753–782, 1998.

[5] P. Jacobs and J. Canny, "Planning smooth paths for mobile robots," in *Nonholonomic Motion Planning*, Z. Li and J. Canny, Eds. Kluwer Academic, 1992, pp. 271–342.

[6] J.-P. Laumond, P. Jacobs, M. Taix, and R. Murray, "A motion planner for nonholonomic mobile robots," *IEEE Transactions on Robotics and Automation*, vol. 10, no. 5, pp. 577–593, October 1994.

[7] P. Agarwal and H. Wang, "Approximation algorithms for curvature-contrained shortest paths," *SIAM Journal on Computing*, vol. 30, no. 6, pp. 1739–1772, 2001.

[8] S. LaValle, *Planning Algorithms*. Cambridge University Press, 2006.

[9] Z. Tang and Ü. Özgüner, "Motion planning for multitarget surveillance with mobile sensor agents," *IEEE Transactions on Robotics*, vol. 21, pp. 898–908, 2005.

[10] K. Savla, F. Bullo, and E. Frazzoli, "On traveling salesperson problems for Dubins' vehicle: stochastic and dynamic environments," in *IEEE conference on decision and control*, Seville, Spain, December 2005, pp. 4530 – 4535.

[11] K. Savla, E. Frazzoli, and F. Bullo, "Traveling salesperson problems for the Dubins vehicle," *IEEE Transactions on Automatic Control*, vol. 53, no. 9, July 2008.

[12] J. Enright and E. Frazzoli, "UAV routing in a stochastic, time-varying environment," in *Proceedings of the IFAC World Congress*, 2005.

[13] S. Itani and M. Dahleh, "On the stochastic TSP for the Dubins vehicle," in *Proceedings of the American Control Conference*, 2007, pp. 443–448.

[14] K. Savla, E. Frazzoli, and F. Bullo, "On the point-to-point and traveling salesperson problems for Dubins' vehicle," in *Americal Control Conference*, Portland, OR, June 2005, pp. 786 – 791.

[15] S. Rathinam, R. Sengupta, and S. Darbha, "A resource allocation algorithm for multi-vehicle systems with non-holonomic constraints," *IEEE Transactions on Automation Science and Engineering*, pp. 98 – 104, 2006.

[16] X. Ma and D. Castañón, "Receding horizon planning for Dubins traveling salesman problems," in *IEEE Conference on Decision and Control*, 2006, pp. 5453 – 5458.

[17] R. J. Kenefic, "Finding good Dubins tours for UAVs using particle swarm optimization," *Journal of Aerospace Computing, Information, and Communication*, vol. 5, pp. 47–56, February 2008.

[18] D. Rosenkrantz, R. Stearns, and P. Lewis II, "An analysis of several heuristics for the traveling salesman problem," *SIAM Journal on Computing*, vol. 6, pp. 563–581, 1977.

[19] V. Vazirani, *Approximation Algorithms*. Springer, 2001.

[20] L. Dubins, "On curves of minimal length with a constraint on average curvature and with prescribed initial and terminal positions and tangents," *American Journal of Mathematics*, vol. 79, pp. 497–516, 1957.

[21] J. D. Boissonnat and X. N. Bui, "Accessibility region for a car that only moves forward along optimal paths," INRIA, Sophia-Antipolis, France, Tech. Rep. 2181, 1994.

[22] J.-H. Lee, O. Cheong, W.-C. Kwon, S. Y. Shin, and K.-Y. Chwa, "Approximation of curvature-constrained shortest paths through a sequence of points," in *ESA '00: Proceedings of the 8th Annual European Symposium on Algorithms*. London, UK: Springer-Verlag, 2000, pp. 314–325.

[23] X.-N. Bui, P. Souères, J.-D. Boissonnat, and J.-P. Laumond, "The shortest path synthesis for non-holonomic robots moving forwards," INRIA, Tech. Rep., 1993.

[24] T. Cormen, C. Leiserson, R. Rivest, and C. Stein, *Introduction to algorithms*. MIT Press, 2001.

[25] C. Papadimitriou, "The Euclidean traveling salesman problem is NP-complete," *Theoretical Computer Science*, vol. 4, pp. 237–244, 1977.

[26] M. Garey, R. Graham, and D. Johnson, "Some NP-complete geometric problems," in *Proceedings of the 8th annual ACM symposium on Theory of computing*, 1976, pp. 10 – 22.

[27] R. Karp, "Reducibility among combinatorial problems," *Complexity of Computer Computations*, pp. 85–103, 1972.

[28] A. Behzad and M. Modarres, "New efficient transformation of the generalized traveling salesman problem into traveling salesman problem," in *Proceedings of the 15th International Conference of Systems Engineering*, Las Vegas, 2002.

[29] K. Helsgaun, "An effective implementation of the Lin-Kernighan traveling salesman heuristic," *European Journal of Operational Research*, vol. 126, no. 1, pp. 106–130, 2000.

[30] A. Frieze, G. Galbiati, and F. Maffioli, "On the worst-case performance of some algorithms for the asymmetric traveling salesman problem," *Networks*, vol. 12, pp. 23–39, 1982.

[31] J. Le Ny, "Performance optimization for unmanned vehicle systems," Ph.D. dissertation, Massachusetts Institute of Technology, September 2008.