

The Curvature-Constrained Traveling Salesman Problem For High Point Densities

Jerome Le Ny, Emilio Frazzoli and Eric Feron

Abstract—We consider algorithms for the curvature-constrained traveling salesman problem, when the nonholonomic constraint is described by Dubins’ model. We indicate a proof of the NP-hardness of this problem. In the case of low point densities, i.e., when the Euclidean distances between the points are larger than the turning radius of the vehicle, various heuristics based on the Euclidean Traveling salesman problem are expected to perform well. In this paper we do not put a constraint on the minimum Euclidean distance. We show that any algorithm that computes a tour for the Dubins’ vehicle following an ordering of points optimal for the Euclidean TSP cannot have an approximation ratio better than $\Omega(n)$, where n is the number of points. We then propose an algorithm that is not based on the Euclidean solution and seems to behave differently. For this algorithm, we obtain an approximation guarantee of $O\left(\min\left\{\left(1+\frac{\rho}{\varepsilon}\right)\log n, \left(1+\frac{\rho}{\varepsilon}\right)^2\right\}\right)$, where ρ is the minimum turning radius, and ε is the minimum Euclidean distance between any two points.

I. INTRODUCTION

In order to improve the performance of mobile robotic networks, in particular unmanned aerial systems, researchers are working on integrating the high level problem of mission planning and the low level problem of path planning [1]. Indeed, the current practice of hierarchical control, where the kinematic and dynamic constraints of the vehicles are not taken into account at the mission planning level, can result in poor overall performance, since the sequence of scheduled activities can be very hard to execute on the physical system.

In this paper we consider the traveling salesman problem (TSP): for a given set of points in the plane, design a tour of shortest length going through all the points exactly once. If the length is measured as the sum of the Euclidean distances between consecutive points in the tour, we obtain the Euclidean TSP (ETSP). Here however, we will require the path generated to be feasible for a vehicle with certain dynamic constraints, in particular for a model of unmanned aerial vehicle (UAV). It is not currently known how to solve efficiently this problem if a detailed physical description of the vehicle is used. In theory, this problem could be formulated as an optimal control problem and solved numerically, for example using a mixed-integer linear program [2]. However this procedure is already complex if the sequence of waypoints is specified, and so directly adding a hard

combinatorial problem to it is practical for only a very small number of points. Hence in the past few years, there has been some interest in algorithms for simple kinematic models, in particular the Dubins’ model. The Dubins’ vehicle can only move forward in the plane, at constant speed, and has a limited turning radius. The advantage of this model is that we can generate paths that are almost feasible for fixed wing aircrafts, and therefore provide an order in which to visit the points that is relevant for the low-level controller of the vehicle. At the same time, for this model we have for example an analytical characterization of the shortest path between any two configurations, giving hope that efficient algorithms with good performance can be designed to solve the TSP.

In this paper, we continue the study of the TSP for the Dubins’ vehicle (DTSP). Section II provides a short survey of some recent work on this problem. In section III, we comment on the heuristics based on following a tour optimal for the ETSP and show a lower bound on the best approximation ratio that these algorithms can achieve. In Section IV we show that DTSP is NP-hard, since we are not aware of any previous proof of this result published in the literature. Section V describes an algorithm that is not based on a solution for the ETSP and seems to have a different behavior for high point densities. An approximation ratio of $O\left(\min\left\{\left(1+\frac{\rho}{\varepsilon}\right)\log n, \left(1+\frac{\rho}{\varepsilon}\right)^2\right\}\right)^1$ is shown for this algorithm, where ρ is the minimum turning radius of the vehicle and ε is the minimum Euclidean distance between any two waypoints.

II. SURVEY OF RECENT RESULTS

Let us state precisely the Dubins’ traveling salesman problem (DTSP). The configuration of a Dubins’ vehicle in the plane is given by its position and heading $(x, y, \theta) \in \mathbb{R}^2 \times (-\pi, \pi]$. Its equations of motion are

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

Without loss of generality, we will assume that the speed of the vehicle v_0 is normalized to 1. ρ is the minimum turning radius of the vehicle. u is the available control. The DTSP asks, for a given set of points P in the plane, to find the shortest tour through these points that is feasible for a

This work was supported by Air Force - DARPA - MURI award 009628-001-03-132 and Navy ONR award N00014-03-1-0171.

J. Le Ny and E. Frazzoli are with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139-4307, USA jleny@mit.edu, frazzoli@mit.edu.

E. Feron is with the School of Aerospace Engineering, Georgia Tech, Atlanta, GA 30332, USA eric.feron@aerospace.gatech.edu

¹We 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 .

Dubins' vehicle. Given that we show below that this problem is NP-hard, we will focus on the design of approximation algorithms. An α -approximation algorithm ($\alpha \geq 1$) for a minimization problem with optimum OPT is an algorithm that produces *in polynomial time* a feasible solution whose value Z is within a factor α of the optimum, i.e., such that

$$OPT \leq Z \leq \alpha OPT.$$

Dubins [3] 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. We paraphrase the following result from Dubins:

Theorem 1 ([3]): An optimal path between any two configurations is of type CCC or CSC, or a subpath of a path of either of these two types. Moreover, to be optimal, a CCC path must have its middle arc of length greater than $\pi\rho$.

In the following, we will 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. Then we see from Theorem 1 that to find the minimum length path between an initial and a final configuration, it is enough to find the minimum length path among six paths, namely among $\{LSL, RSR, RSL, LSR, RLR, LRL\}$. Each of these paths can be explicitly computed (see for instance [4]) and therefore finding the optimum path and length between any two configurations can be done in constant time.

A. Stochastic Case

Perhaps surprisingly, the stochastic version of the DTSP is currently the best understood case. Let us assume that the points to visit are distributed *uniformly* in a square. Then it has been shown that the expected length of the Dubins' tour scales as $n^{2/3}$, and there is an algorithm that is essentially optimal [5], [6], [7]. The algorithm is reminiscent of Karp's partitioning algorithm [8], but the tiling of the initial square is done using a shape that is more adapted to the dynamics of the Dubins' vehicle. From this body of work, it is clear that the DTSP behaves differently from the ETSP. An essential difference that prevents using directly the classical techniques developed for Euclidean optimization problems [9] is that the DTSP is not scale invariant: if the Euclidean distances between the points are scaled down, the optimal tour length for the Dubins' vehicle is not scaled by the same factor. Note also that it is known that the expected length of the shortest Euclidean tour scales as $n^{1/2}$ [10].

B. Algorithms based on a solution for the Euclidean TSP

If no input distribution is given, fewer results are available. Except for the algorithm in [11], which we will review in the last section, most of the existing algorithms seem to build on a preliminary solution obtained for the Euclidean TSP (ETSP) [12], [13], [14]. When the minimum Euclidean distance between any two points is large compared to the

turning radius of the vehicle, this is a natural idea: as the minimum turning radius tends to zero, the DTSP and ETSP are the same. In [13], the authors impose that no two points in an instance of the problem can be at a distance less than twice the minimum turning radius. Under this hypothesis, they provide a 4.64-approximation algorithm. Their algorithm first solves the underlying ETSP using Christofides' $\frac{3}{2}$ -approximation algorithm [15], and so by using instead Arora's $1 + \varepsilon$ approximation scheme [16] for the ETSP, we get a guaranteed approximation ratio of about 3.1.

Once the order of the points is obtained from the ETSP solution, a feasible path must be obtained. From Dubins' result, the problem is then reduced to fixing the headings at each point. The first algorithm proposed for the DTSP, called the "Alternating Algorithm" [12] follows the edges of the optimal Euclidean tour. All odd-numbered edges are retained (i.e. the subpath is a straight line) as well as the corresponding headings, and the even-numbered edges are replaced with Dubins' paths. [13] instead uses the fact that the optimal path between an initial configuration and a final point with fixed position but free heading is also precisely characterized, of the form CS. Fixing the first heading, one can thus compute successively the optimal path to the next point in the ETSP solution. This fixes the next heading and the algorithm finishes when it is back to the initial point.

The same idea is used in [14], where the authors also extend the method to computing the optimal subpaths between three consecutive points of the Euclidean solution. Obviously this cannot do worse than with two points. They also improve empirically on that technique by using a receding horizon version of the algorithm, using only the first Dubins' path in the three-point solution and recomputing at the middle waypoint. More generally, one could try to compute subpaths for a larger number of consecutive points, using if possible an analytical solution similar to the two point case, or numerically, using for example a mixed-integer formulation.

III. PERFORMANCE LIMIT FOR ETSP-BASED ALGORITHMS

There is however a limit on the performance one can achieve using the technique described in the previous paragraph, computing first an ordering based on the optimal solution for the Euclidean TSP. This limit is significant in particular in the case where the points are densely distributed in the plane. Then algorithms based on the Euclidean metric are not necessarily a good choice any more, returning sequences that will require too many maneuvers from the vehicle. The metric used to compute the distances is clearly an important element, and for example the angular-metric TSP [17], which also aims at modeling a kinematic constraint, has little to do with the ETSP. It is also interesting to note that the solution obtained for the stochastic DTSP did not involve the ETSP solution.

Theorem 2: Any algorithm for the DTSP following the ordering of points that is optimal for the ETSP has an approximation ratio R_n which is $\Omega(n)$. If we impose a lower

bound ε sufficiently small on the minimum Euclidean distance between any two waypoints, then there exist constants C, C' , independent of n , such that the approximation ratio is not better than $\frac{C}{\varepsilon + (C'/n)}$.

Proof: Let us call the configuration of points shown on fig. 1 a chain. Let n be the number of points, and suppose $n = 4m$. For clarity we focus on the path-TSP problem but extension to the tour-TSP case is easy, by adding a similar chain in the reverse direction. The optimal Euclidean path-TSP is shown on the figure as well. Suppose now that a Dubins' vehicle tries to follow this order of points, and suppose ε is sufficiently small. Then it should be clear that for each sequence of 4 consecutive points, two on the upper line and two on the lower line, the Dubins' vehicle will have to execute a maneuver of length at least C , where C is a constant of order $2\pi\rho$. For instance, if the vehicle tries to go through the two top points without a large maneuver, it will exit the second point with a heading almost horizontal and will have to make a large turn to catch the third point on the lower line. Hence the length of the Dubins' path will be greater than mC .

On the other hand, a Dubins vehicle can simply go through all the points on the top line, execute a U-turn of length C' of order $2\pi\rho$, and then go through the points on the lower line, providing an upper bound of $2n\varepsilon$ for the optimal solution. So we deduce that the worst case approximation ratio of the algorithm is at least:

$$R_n \geq \frac{\frac{n}{4}C}{2n\varepsilon + C'}.$$

But we can choose ε as small as we want and thus $R_n = \Omega(n)$. ■

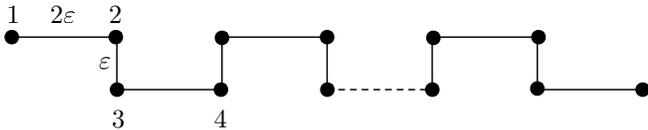


Fig. 1. A chain and the associated Euclidean optimal path.

Remark 3: In [1], the authors determine the order of the points without using the solution of the Euclidean TSP. Instead, they construct the geometric center of the waypoints, calculate the orientations of the waypoints with respect to that center, and traverse the points in order of increasing orientations. It is easy to adapt Theorem 2 to this case, using a “circular chain” for example.

IV. COMPLEXITY OF THE DTSP

It is usually accepted that the DTSP is NP-hard, but, to the authors' knowledge, no proof of this result has been published so far. Note that adding the curvature constraint to the Euclidean TSP could well make the problem easier, as in the bitonic TSP [18, p. 364]², and so the statement does not follow trivially from the NP-hardness of the Euclidean

TSP, which was shown by Papadimitriou [19] and by Garey, Graham and Johnson [20].

Deducing the result directly from the result for the ETSP or the asymmetric TSP does not seem easy. Instead, we can go back to Papadimitriou's proof and exploit his construction. A sketch of the proof is provided below. As a side remark, it is not clear that DTSP is in NP. In fact, there is a difficulty even in the case of ETSP with distances that are not rounded, because evaluating the length of a tour involves computing many square roots. In the case of DTSP, given a permutation of the points and a set of headings, deciding whether the tour has length less than a bound L might require computing trigonometric functions and square roots accurately, even if the headings are restricted to be rational multiples of π . Hence in the following we concentrate on the NP-hardness result.

Let us first define an instance of the decision version of the Dubins TSP, which with a slight abuse of notation we also call DTSP.

Dubins traveling salesman problem (DTSP): Given a set of points in the plane with rational coordinates and a rational number $L > 0$, does there exist a tour for the Dubins' vehicle visiting all these points exactly once, of length at most L ?

We also consider the two cases tour-DTSP and path-DTSP, depending on the presence of the requirement that the vehicle must start and end at the same point or not. Euclidean path-TSP and tour-TSP are both NP-hard.

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

Proof: [sketch] This can be seen as a corollary of Papadimitriou's proof of the NP-hardness of Euclidean TSP (ETSP), to which we refer [19]. First recall the Exact Cover Problem: given a family F of subsets of the finite 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 [21]. Papadimitriou gives a polynomial-time reduction of Exact Cover to Euclidean TSP. 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 problem has an optimal tour of length less than or equal to L . In fact, it can be observed that in the case where Exact Cover does not have a solution, Papadimitriou's construction, possibly rescaled, gives an instance of the ETSP which has an optimal tour of length $\geq (L + 1)$, not just $> L$.

Now it is intuitively clear, and more precisely proved in [12], that 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 $ETSP(\mathcal{P}) + Cn$. This bound is obtained using the permutation of points corresponding to the optimal Euclidean tour, and we can take $C = \frac{3\pi}{2}\rho$. Of course the length of the curvature constrained tour is always greater than or equal to the length of the Euclidean tour.

Using this, if we have n points in the instance of the ETSP constructed as in Papadimitriou's proof, we can construct

²The authors would like to acknowledge Vincent Blondel for this example

a new ETSP instance by simply scaling all the distances by a factor $2Cn$. Then if Exact Cover has a solution, the ETSP instance has an optimal tour of no more than $2CnL$ and so the curvature constrained tour has a length of no more than $2CnL + Cn$. If Exact Cover does not have a solution, the ETSP instance has an optimal tour of at least $2CnL + 2Cn$, and the curvature constrained tour as well. So Papadimitriou's construction, rescaled by $2Cn$ and using $2CnL + Cn$ instead of L , where n is the number of points used in the construction, provides a reduction from Exact Cover to DTSP. ■

The reduction works basically because, as the Euclidean distances between the points increase, the additional length Cn of the curvature constrained tour becomes negligible compared to the length of the Euclidean tour. Theorem 4 thus shows that DTSP is at least as hard as ETSP. However it does not capture the potential additional difficulty due to the problem of optimizing the headings at each point. This later problem is essentially a continuous optimization problem however, and it is not clear that the standard complexity theory for discrete problems is the tool we should use to capture this aspect. Note that there is an important literature on complexity theory for certain continuous problems, see [22].

V. A RANDOMIZED ALGORITHM

In this section, we describe an algorithm which does not use an ordering of the waypoints obtained from a solution for the ETSP and we provide a simple analysis of its approximation ratio. We do not assume any bound on the minimum Euclidean distance between any two points of an instance. This algorithm was described in [11], although we refine somewhat the analysis here.

A first version of the algorithm can be described as follows:

- 1) Fix the headings at all points to be 0.
- 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 approximate solution for this asymmetric TSP.

In step 4 of the algorithm, we have obtained a metric asymmetric TSP (ATSP) because once the headings at two points A and B are fixed, the length of the Dubins' path from A to B is in general different from the length of the Dubins' path from B to A . Unfortunately there are less available results for the asymmetric TSP than for symmetric case. The smallest approximation ratio currently known seems to be $0.842 \log n$ [23], improving upon the $\log n$ of Frieze et al. [24] and the $0.999 \log n$ in [25]. For the path version of the problem, there is only a $O(\log n)$ approximation [26].

There is however some additional information that we can exploit once the headings have been fixed, which differentiate the problem from a general ATSP. We have the following:

Lemma 5: Let \hat{d}_{ij} denote the Dubins' distance from point $X_i = (x_i, y_i)$ to point $X_j = (x_j, y_j)$ in the plane after the headings $\{\theta_k\}_{k=1}^n$ have been fixed. Then:

$$\max_{i,j} \frac{\hat{d}_{ij}}{\hat{d}_{ji}} \leq 1 + \frac{4\pi\rho}{\varepsilon},$$

where ε is, for the rest of this section, the minimum Euclidean distance between any two waypoints.

Proof: This result is based on the following stronger fact:

$$\hat{d}_{ij} \leq \hat{d}_{ji} + 4\pi\rho.$$

To show this, consider the Dubins' path from the configuration (X_i, θ_i) to (X_j, θ_j) . We construct a feasible path (not necessarily optimal) for the Dubins' vehicle, to return from (X_j, θ_j) to (X_i, θ_i) . The Dubins' vehicle moves along an initial circle (of radius ρ), a straight line or a middle circle, and a final circle. In the computation of the forward Dubins' path, there are two possible initial circles, and two possible final circles to consider. In every possible case, we can maneuver the Dubins' vehicle on these circles, to return from configuration (X_j, θ_j) to (X_i, θ_i) . The length of the intermediate part is never increased. It can be verified in the different cases that the additional length of a return path constructed this way is never more than $4\pi\rho$ plus the length of the forward path. An example for a RSL path is provided on Fig. 2.

Now the bound of the lemma is obtained since $\hat{d}_{ij} \geq \varepsilon$. ■

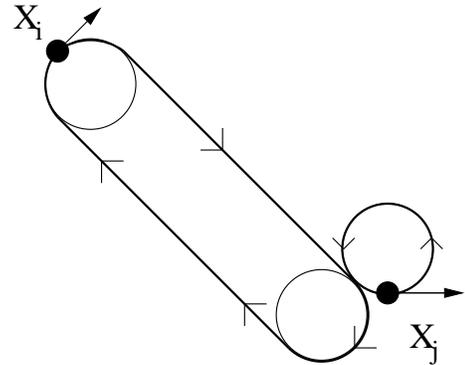


Fig. 2. A feasible return path for a RSL Dubins' path.

With this bound on the arc distances, we can use a modified version of Christofides' algorithm to obtain a $\frac{3}{2} \left(1 + \frac{4\pi\rho}{\varepsilon}\right)$ approximation for the ATSP in step 4 [24].

The complexity of the three first steps is $O(n^2)$. For simplicity, we will consider only the algorithms of Frieze et al. to solve the ATSP. That is, we run both algorithms and choose the tour with minimum length, thus obtaining an approximation ratio of $\min\left(\log n, \frac{3}{2} \left(1 + \frac{4\pi\rho}{\varepsilon}\right)\right)$. The algorithm for solving the ATSP runs in $O(n^3)$, so overall the running time of our algorithm is $O(n^3)$.

To analyze the performance of this algorithm, we will use the following bound.

Lemma 6: Let d_{ij} be the length of the Dubins' path between two configurations (X_i, θ_i) and (X_j, θ_j) and let ε_{ij} be the Euclidean distance between X_i and X_j . If we perturb the initial and final headings by angles less than $\delta \in (-\pi, \pi]$, and call \hat{d}_{ij} the new Dubins' length, we have

$$\hat{d}_{ij} \leq \left(1 + 2\rho \max \left\{ \frac{3|\delta| + \pi |\sin \frac{\delta}{2}|}{\varepsilon_{ij}}, \right. \right. \quad (1)$$

$$\left. \frac{|\delta| + 4 \arccos \left(1 - |\sin \frac{\delta}{2}|/2 \right)}{\pi\rho} \right\} \right) d_{ij}. \quad (2)$$

Proof: The bound follows from the discussion in section 4.8 of [27]³. Let us mention that the two terms in the max on the right-hand side correspond to the cases where the initial Dubins' path is a CSC path with opposite initial and final turning directions and a CCC path respectively. Perturbations of a CSC path with identical initial and final turning directions are dominated by the first term. Moreover, for the second term, the fact that a CCC optimal path must have minimum length $\pi\rho$ was used. ■

By fixing the headings a priori, we can make an error up to $|\delta| = \pi$ with respect to the optimal heading at each point. Hence if d_{ij} are the Dubins' distances between the points if the optimal headings were selected, and \hat{d}_{ij} are the distances after we fix the headings in step 1, we obtain from the lemma:

$$\hat{d}_{ij} \leq \left(1 + \max \left\{ \frac{8\pi\rho}{\varepsilon}, \frac{14}{3} \right\} \right) d_{ij} := C_1 d_{ij}. \quad (3)$$

Call OPT the optimal value of the DTSP and σ^* the corresponding optimal permutation specifying the order of the waypoints. We have $OPT = \sum_{i=1}^{n-1} d_{\sigma^*(i)\sigma^*(i+1)} + d_{\sigma^*(n)\sigma^*(1)} := L(\{d_{ij}\}, \sigma^*)$. Considering the permutation σ^* for the graph problem (where the edge weights are the distances $\{d_{ij}\}$) and $\hat{\sigma}^*$ the optimal permutation for the graph problem, we have

$$L(\{\hat{d}_{ij}\}, \hat{\sigma}^*) \leq L(\{d_{ij}\}, \sigma^*) \leq C_1 L(\{d_{ij}\}, \sigma^*).$$

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:

$$\begin{aligned} L(\{\hat{d}_{ij}\}, \hat{\sigma}) &\leq \min \left(\log n, \frac{3}{2} \left(1 + \frac{4\pi\rho}{\varepsilon} \right) \right) L(\{d_{ij}\}, \hat{\sigma}^*) \\ &\leq \left(C_1 \min \left(\log n, \frac{3}{2} \left(1 + \frac{4\pi\rho}{\varepsilon} \right) \right) \right) OPT. \end{aligned}$$

Therefore, we obtain with the specific assignment of headings (or any assignment in fact) an approximation guaranteed to be within a factor

$$\left(1 + \max \left\{ \frac{8\pi\rho}{\varepsilon}, \frac{14}{3} \right\} \right) \min \left(\log n, \frac{3}{2} \left(1 + \frac{4\pi\rho}{\varepsilon} \right) \right)$$

³It appears that eq. (4.38) and (4.48) in [27] contain a typographical error, namely a ρ is missing.

of the optimum. More succinctly, for the case of high point densities which is of most interest for us, this bound is of the order $O\left(\min\left(\frac{\rho}{\varepsilon} \log n, \left(\frac{\rho}{\varepsilon}\right)^2\right)\right)$.

A. Randomized Version

Instead of fixing all the headings to 0 in step 1 of the algorithm, it is perhaps more natural to choose them randomly and independently in $(-\pi, \pi]$ for each point.

Consider an optimal tour, and two successive points i, j in this tour. The Dubins' path between these two points has length d_{ij} , and the optimal headings are θ_i and θ_j . Following [27] in the derivation of the bound (1), we know that if we make an error of $\delta \in (-\pi, \pi]$ on θ_i , the difference in path length is bounded by:

$$\begin{aligned} \Delta d &\leq \rho \max \left\{ 3|\delta| + \pi \left| \sin \left(\frac{\delta}{2} \right) \right|, \right. \quad (4) \\ &\quad \left. |\delta| + 4 \arccos \left(1 - \left| \sin \frac{\delta}{2} \right|/2 \right) \right\}. \end{aligned}$$

This leads to the inequality (1), with an additional factor 2 when taking into account the error on θ_j as well. Now (4) is derived for a change from θ_i to $\theta_i + \delta$ in the initial heading. Of course, we do not know the optimal θ_i so the natural idea is to choose $(\theta_i + \delta)$ uniformly in $(-\pi, \pi]$, in which case the error δ is distributed uniformly in $(-\pi, \pi]$ as well. This implies that Δd becomes a random variable whose expectation is bounded by:

$$\begin{aligned} E[\Delta p] &\leq \rho \int_{-\pi}^{\pi} \max \left\{ 3|\delta| + \pi \left| \sin \frac{\delta}{2} \right|, \right. \\ &\quad \left. |\delta| + 4 \arccos \left(1 - \left| \sin \frac{\delta}{2} \right|/2 \right) \right\} \frac{d\delta}{2\pi} \\ &\leq 6.79\rho. \end{aligned}$$

Replacing the corresponding expression in (1), we obtain as a final upper bound:

$$E[\hat{d}_{ij}] \leq \left(1 + \frac{13.58\rho}{\varepsilon_{ij}} \right) d_{ij}. \quad (5)$$

It is also possible to refine the bound using the fact that a CCC path has length at least $\pi\rho$ as in lemma 6, but for our purpose ε and ρ will be of the same order and (5) is then enough. It improves by a factor of roughly 2 on the bound obtained earlier when all headings were fixed deterministically, although now the bound is in expectation. If we repeat the previous analysis, we obtain the following result:

Theorem 7: There is a randomized polynomial-time algorithm that, given a set of n points in the plane, returns a Dubins' traveling salesman tour with expected length within a factor

$$\left(1 + \frac{13.58\rho}{\varepsilon} \right) \min \left(\log n, \frac{3}{2} \left(1 + \frac{4\pi\rho}{\varepsilon} \right) \right)$$

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

Of course, in practice, one should try to run the randomized algorithm several times and choose the shortest tour obtained, in order to increase the probability that this tour has a length less than the expected value.

VI. CONCLUSION

We have presented several results on the Dubins' traveling salesman problem, including a proof of its NP-hardness. We argued that, although much of the recent literature studying this problem has focused on a large separation between the input waypoints, the most challenging problem is probably to devise good algorithms for the case where the points are densely packed. We obtained a lower bound on the approximation ratio that any algorithm which tries to follow the order of points optimal for the Euclidean tour can achieve. In general, guided by the existing results for the stochastic case, we expect the Dubins' TSP and Euclidean TSP to have significantly different characteristics. Finally we presented an algorithm for the DTSP that is not based on the ETSP solution. Although the theoretical approximation guarantee that we provide in this paper is not necessarily very strong, our technique has several interesting practical features. First, the algorithm is fast and very simple and experimentally tends to have a good behavior for high point density. We can therefore run it simultaneously with other existing algorithms and take the best solution. Moreover, note that one could obtain the headings using any existing algorithm, such as the alternating algorithm for example, and then recompute the ordering of the points using the ATSP solution. This has the potential of improving the practical performance of any algorithm, but was so far not used apparently, even in experimental studies. Finally, by computing the ordering of the points after the headings, we are not subject to the counter examples of the type presented in section III. We are currently pursuing this idea to potentially remove the ρ/ε term in the approximation ratio.

REFERENCES

- [1] Z. Tang and Ü. Özgüner, "Motion planning for multitarget surveillance with mobile sensor agents," *IEEE Transactions on Robotics*, vol. 21, pp. 898–908, 2005.
- [2] T. Schouwenaars, B. De Moor, E. Feron, and J. How, "Mixed integer programming for multi-vehicle path planning," in *Proceedings of the European Control Conference*, 2001.
- [3] 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.
- [4] A. Shkel and V. Lumelsky, "Classification of the Dubins' set," *Robotics and Autonomous Systems*, vol. 34, pp. 179–202, 2001.
- [5] K. Savla, F. Bullo, and E. Frazzoli, "On the stochastic traveling salesperson problem for Dubins' vehicle," in *IEEE conference on decision and control*, Seville, Spain, December 2005.
- [6] J. Enright and E. Frazzoli, "UAV routing in a stochastic, time-varying environment," in *Proceedings of the IFAC World Congress*, 2005.
- [7] S. Itani and M. Dahleh, "On the average length of the stochastic TSP with Dubins' metric," LIDS, Massachusetts Institute of Technology, Tech. Rep., 2005.
- [8] R. Karp, "Probabilistic analysis of partitioning algorithms for the traveling salesman problem in the plane," *Mathematics of Operations Research*, vol. 2, pp. 209–224, 1977.
- [9] J. Steele, *Probability Theory and Combinatorial Optimization*. SIAM, 1987.
- [10] J. Beardwood, J. Halton, and J. Hammersley, "The shortest path through many points," *Proceedings of the Cambridge Philosophical Society*, vol. 55, pp. 299–327, 1959.
- [11] 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.
- [12] K. Savla, E. Frazzoli, and F. Bullo, "On the point-to-point and traveling salesperson problems for Dubins' vehicle," in *American Control Conference*, Portland, OR, June 2005.
- [13] 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*, 2006.
- [14] X. Ma and D. Castañón, "Receding horizon planning for Dubins traveling salesman problems," in *IEEE Conference on Decision and Control*, 2006.
- [15] N. Christofides, "Worst-case analysis of a new heuristic for the travelling salesman problem," Carnegie-Mellon Univ., Technical report, CSIA, 1976.
- [16] 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.
- [17] 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.
- [18] T. Cormen, C. Leiserson, R. Rivest, and C. Stein, *Introduction to algorithms*. MIT Press, 2001.
- [19] C. Papadimitriou, "The Euclidean traveling salesman problem is NP-complete," *Theoretical Computer Science*, vol. 4, pp. 237–244, 1977.
- [20] M. Garey, R. Graham, and D. Johnson, "Some NP-complete geometric problems," in *Proceedings of the eighth annual ACM symposium on Theory of computing*, 1976.
- [21] R. Karp, "Reducibility among combinatorial problems," *Complexity of Computer Computations*, pp. 85–103, 1972.
- [22] L. Blum, F. Cucker, M. Shub, and S. Smale, *Complexity and real computation*. Springer, 1998.
- [23] H. Kaplan, M. Lewenstein, N. Shafir, and M. Sviridenko, "Approximation algorithms for asymmetric TSP by decomposing directed regular multidigraphs," in *Proceedings of IEEE FOCS*, 2003, pp. 56–67.
- [24] 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.
- [25] M. Bläser, "A new approximation algorithm for the asymmetric TSP with triangle inequality," in *Proceedings of ACM-SIAM SODA*, 2002, pp. 638–645.
- [26] C. Chekuri and M. Pál, "An $O(\log n)$ approximation ratio for the asymmetric traveling salesman path problem," University of Illinois at Urbana-Champaign, Tech. Rep., 2006.
- [27] 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.