

On the Dubins Traveling Salesperson Problems: novel approximation algorithms

Ketan Savla¹, Emilio Frazzoli², Francesco Bullo¹

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? Specifically, we study a stochastic version of the DTSP where the n targets are randomly 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 high probability. Additionally, we study a dynamic version of the DTSP: given a stochastic process that generates target points, is there a policy which 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 algorithm 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. 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). Specifically, we study a stochastic version of the DTSP, which we refer to as the stochastic DTSP, where the n targets are randomly sampled from a uniform distribution.

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 interests 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., [1] discusses complexity issues related to UAVs assignments problems, [2] considers Dubins vehicles surveilling multiple mobile targets, [3] considers missions with dynamic threats, other relevant works include [4], [5], [6], [7].

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 [8] and a first solution was given by Dubins [9]. Modern treatments on point-to-point planning exploit the Pontryagin Minimum Principle [10], carefully account for symmetries in the problem [11], and consider environments with obstacles [12]. The Dubins vehicle is commonly accepted as a reasonably accurate kinematic model for aircraft motion planning problems, e.g., see [13], and its study is included in recent texts [14], [15].

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 [16] and in the survey [17]. Exact algorithms, heuristics as well as polynomial-time constant factor approximation algorithms are available for the Euclidean TSP, see [18], [19], [20]. A variation of the TSP with potential robotic applications is the angular-metric problem studied in [21]. The DTRP (without nonholonomic constraints) was introduced in [22]. 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, the Dubins TSP cannot be formulated as a problem on a finite-dimensional graph, thus preventing the use of well-established tools in combinatorial optimization.

The contributions of this paper are threefold. First, we propose an algorithm for the stochastic DTSP through a point set P , called the RECURSIVE BEAD-TILING ALGORITHM, based on a geometric tiling of the plane, tuned to the Dubins vehicle dynamics, and a strategy for the vehicle to service targets from each tile. Second, we obtain an upper bound on the stochastic performance of the proposed algorithm and thus also establish a similar bound on the stochastic DTSP. The upper bound on the performance of the RECURSIVE BEAD-TILING ALGORITHM belongs to $O(n^{2/3})$ with high probability, and it is known that the lower

¹Ketan Savla and Francesco Bullo are with the Center for Control, Dynamical Systems and Computation, University of California at Santa Barbara, ketansavla@umail.ucsb.edu, bullo@engineering.ucsb.edu

²Emilio Frazzoli is with the Mechanical and Aerospace Engineering Department at the University of California at Los Angeles. frazzoli@ucla.edu

bound on the achievable performance belongs to $\Omega(n^{2/3})$. The algorithm we introduce in this paper is the first known algorithm providing a provable constant-factor approximation to the DTSP optimal solution. 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 average 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 literature. While the problem of flying an aircraft through waypoints is a very standard problem in aeronautics, the formal study of the Dubins TSP (algorithmic and performance bounds) was introduced in our early work [23], where a constant-factor approximation algorithm for the worst-case setting of the DTSP was proposed. Subsequently, similar versions of this problem were also considered in [24] and [4]. A simplified version of the problem for a different but closely related kind of vehicle, the Reeds-Shepp vehicle, was considered in [25]. In [26], we introduced the stochastic DTSP and gave the first algorithm yielding, with high probability, a solution with a cost upper bounded by a strictly sublinear 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 [27] 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 this paper we propose the first algorithm that asymptotically achieves a constant factor approximation to the stochastic DTSP with high probability.

Notation

In this section we setup 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 maintaining a constant 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.

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 . For the stochastic DTSP, p_1, \dots, p_n will be assumed to be randomly and independently sampled from a uniform distribution over \mathcal{Q} .

The key objective of this paper is the design of an algorithm that provides a provably good approximation to the optimal solution of the stochastic Dubins TSP. To establish what a “good approximation” might be, let us recall what is known about the ETSP in the stochastic setting. First, given a compact set \mathcal{Q} and a point-set P whose n points are independently chosen from a distribution φ with compact support $\mathcal{Q} \subset \mathbb{R}^2$, the following deterministic limit holds [16]:

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

where $\bar{\varphi}$ 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 [28]. The fact that the dependence of the ETSP is sublinear in n is very important in the study of the DTRP, i.e., the problem in which new locations are continuously added to the set of outstanding points P ; see Section III.

Motivated by the Euclidean case, this paper shows that the DTSP grows with $n^{2/3}$ in the stochastic case (as both lower and upper bounds). Additionally, this paper proposes novel algorithms for the DTSP in the stochastic setting, 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 in the 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 $\lim_{x \rightarrow 0} f(x)/g(x) = 0$.

II. STOCHASTIC DTSP

In [23], a simple heuristics, the ALTERNATING ALGORITHM for the Dubins TSP for a given point set was proposed. The length of tour generated by this algorithm was also characterized and it was shown that it belongs to $\Omega(\sqrt{n})$ and $O(n)$. It was also shown that this simple policy 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 this section, we consider the scenario when n target points are stochastically generated in \mathcal{Q} according to a uniform distribution. A novel algorithm was proposed in [26] to service these points in such a way that its tour length grew sub-linearly with the number of points asymptotically with high probability, where an event is said to occur with high probability if the probability of its occurrence approaches 1 as $n \rightarrow +\infty$. Here, we present a novel version of this strategy in the form of the RECURSIVE BEAD-TILING ALGORITHM and characterize 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} . The points $P = (p_1, \dots, p_n)$ are randomly generated according to a uniform distribution in \mathcal{Q} .

A. A lower bound

We begin with a result from [29], that provides a lower bound on the expected length of the stochastic DTSP.

Theorem 2.1: (Lower bound on stochastic DTSP) For all $\rho > 0$, the expected cost of the DTSP for a set P of n uniformly-randomly-generated points in a rectangle of width W and height H satisfies

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

Remark 2.2: Theorem 2.1 implies that $\mathbb{E}[\text{DTSP}_\rho(P)]$ belongs to $\Omega(n^{2/3})$. \square

B. The basic geometric construction

Here we define a useful geometric object and study its properties. Consider two points $p_- = (-\ell, 0)$ and $p_+ = (\ell, 0)$ on the plane, with $\ell \leq \rho$, and construct the region $\mathcal{B}_\rho(\ell)$ as detailed in Figure 1. We refer to such regions as a *bead* of




Fig. 1. 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.

length ℓ . The region $\mathcal{B}_\rho(\ell)$ enjoys the following asymptotic properties as $(\ell/\rho) \rightarrow 0^+$:

(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 2 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 uniformly randomly generated 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 $\nu = n \cdot \mu(\ell(n))$ is a constant, then the limit for large n of the binomial distribution is [30] the Poisson distribution of mean ν , that is,

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

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 $\ell \leq 4\rho$.) Furthermore, the tiling is chosen in such a way that it is aligned with the sides of \mathcal{Q} , see Figure 2.

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 services at least one target in it.

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

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

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 2, 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 services 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 [23].

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. Due to lack of space, we refer the reader to [31] for the missing proofs in this section. 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 [32].

Theorem 2.3 (Targets remaining after recursive phases):

Let $P \in \mathcal{P}_n$ be uniformly randomly 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 $n \rightarrow +\infty$.

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 $h(t_1) < h(t_2)$. 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,

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

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:

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

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

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

In other words, $Y_t = 1$ 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 depends on the targets generated before the t^{th} target, the probability that it takes the value 1 is bounded by

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

regardless of the actual values of $b(1), \dots, b(t-1)$. It is known [32] that if the random variables Y_t satisfy such a condition, the sum $\sum_t Y_t$ is stochastically dominated by a binomially distributed random variable, namely,

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

In particular,

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

where the last inequality follows from Chernoff’s Bound [30]. 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 (1) can be rewritten as

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

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 $v_{i+1}(n) \leq \sum_{t=1}^n Y_t$: the right hand side represents the number of targets that will be visited after the i^{th} phase,




Fig. 2. Sketch of “meta-beads” at successive phases in the recursive bead tiling algorithm. From left to right: phase 1, phase 2 and phase 3.

whereas the left hand side counts the number of beads containing such targets. We have, for all $i \leq i^*(n)$:

$$\Pr[v_{i+1} > \beta_{i+1}n | \mathcal{E}_i] \cdot \Pr[\mathcal{E}_i] \leq \Pr\left[\sum_{t=1}^n Y_t > \beta_{i+1}n\right] \leq \frac{1}{n^2},$$

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

$$\Pr[-\mathcal{E}_{i+1}] \leq \frac{1}{n^2} + \Pr[-\mathcal{E}_i] \leq \frac{i}{n^2}.$$

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,

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

Dealing with conditioning as before, we obtain

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

In other words, the number of targets that are left unvisited after the $(i^*)^{\text{th}}$ phase is bounded by a logarithmic function of n with high probability. ■

In summary, Theorem 2.3 says that after a sufficiently large number of phases, almost all targets will be visited, with high probability. The second key point is to recognize 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 2.4 (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

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

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 2.5 (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

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

Lemma 2.6 (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 BEAD-TILING ALGORITHM can be chosen such that $L_{2j} \leq 2L_{2j+1}$, for all $j \in \mathbb{N}$.

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 2.7 (Path length for the algorithm): Let $P \in \mathcal{P}_n$ be uniformly randomly generated in the rectangle of width W and height H . For any $\rho > 0$, with high probability

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

Proof: For simplicity we let $L_{\text{RBTA},\rho}(P) = L_{\text{RBTA}} \cdot$ 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 2.6, 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 2.5, yielding

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

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

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

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 high probability. From Theorem 2.3, we know that with high probability there will be at most $24 \log_2 n$ unvisited targets after the $\lceil \log_2 n \rceil$ recursive phases. From [23] we know that, with high probability, the length of a ALTERNATING ALGORITHM tour through these points satisfies

$$L''_{\text{RBTA}} \leq \kappa \lceil 12 \log_2 n \rceil \pi \rho + o(\log_2 n).$$

Remark 2.8: Theorems 2.1 and 2.7 imply that, with high probability, the RECURSIVE BEAD-TILING ALGORITHM is $\frac{32}{\sqrt[3]{3}} \left(1 + \frac{7}{3} \pi \frac{\rho}{W} \right)$ -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 . \square

III. THE DTRP FOR DUBINS VEHICLE

We now turn our attention to the Dynamic Traveling Reairperson Problem (DTRP) that was introduced by Bertsimas and van Ryzin in [22]. 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 , is 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 will 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

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

that is, if the Dubins vehicle is able to service targets at a rate that is, on average, at least as fast as the rate at which new targets are generated. Let T_j be the time that the j^{th} target spends within the set D , i.e., the time elapsed from the time the j^{th} target is generated to the time it is serviced. If the system is stable, then we can write the balance equation (known as Little's formula [33]):

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

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.,

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

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 [29] a lower bound for the system time, and then present a novel approximation algorithm providing an upper bound on the performance that holds with high probability.

Theorem 3.1: (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

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

Remark 3.2: Theorem 3.1 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 [22]. \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

$$C_{\text{BTA}} = \frac{7 - \sqrt{17}}{4} \left(1 + \frac{7}{3} \pi \frac{\rho}{W}\right)^{-1}. \quad (2)$$

- (ii) Traverse all non-empty beads once, visiting one target per non-empty bead.
 (iii) 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 2.4.

Theorem 3.3: (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:

$$\lim_{\lambda \rightarrow +\infty} \frac{T_{\text{DTRP}}}{\lambda^2} \leq \lim_{\lambda \rightarrow +\infty} \frac{T_{\text{BTA}}}{\lambda^2} \leq 70.5 \rho W H \left(1 + \frac{7}{3} \pi \frac{\rho}{W}\right)^3.$$

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

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

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 2.4. As a consequence, targets in B will be visited at a rate no smaller than

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

In summary, the expected time T_B between the appearance of a target in B and its servicing 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 [33], 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

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

Since equation (3) 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 (3) is a constant that depends on problem parameters ρ , W , and H , and on the design parameter C_{BTA} , as defined in equation (2). Stability of the queue is established by noting that $C_{\text{BTA}} < \left(1 + \frac{7}{3} \pi \frac{\rho}{W}\right)^{-1}$. Additionally, the choice of C_{BTA} in equation (2) minimizes the right hand side of (3) yielding the numerical bound in the statement. \blacksquare

Remark 3.4: The achievable performance of the BEAD-TILING ALGORITHM provides a $55.7 \left(1 + \frac{7}{3} \pi \frac{\rho}{W}\right)^3$ -factor approximation to the lower bound established in Theorem 3.1. 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 worst-case DTSP derived in [23]. \square

IV. CONCLUSIONS

In this paper, we have studied the TSP problem for vehicles that follow paths of bounded curvature in the plane. For the stochastic setting, we have obtained upper bounds that are within a constant factor of the lower bound established in literature [29]; the upper bounds are constructive in the sense that they are achieved by novel algorithm. Similar analysis has been done for a vehicle modeled as a double integrator in [34]. The same paper extends the results to the three dimensional case too. It is interesting to compare our results with the Euclidean setting (i.e., the setting in which curves do not have curvature constraints). The results are summarized in the following table, where $d \in \mathbb{N}$ is the dimension of the space.

	Simple vehicle	Dubins vehicle
Length of TSP tour (worst case)	$\Theta(n^{1-\frac{1}{d}})$ [17]	$\Theta(n)$ ($d = 2, 3$) [23]
Exp. Length of TSP tour (stochastic)	$\Theta(n^{1-\frac{1}{d}})$ [17]	$\Theta(n^{1-\frac{1}{2d-1}})$ w.h.p. ($d = 2, 3$)
System time for DTRP	$\Theta(\lambda^{d-1})$ [22] ($d = 1$)	$\Theta(\lambda^{2(d-1)})$ ($d = 2, 3$)

Remarkably, the differences between the various TSP bounds play a crucial role when studying the DTRP problem;

e.g., stable policies exist only when the TSP cost grows strictly sub-linearly with n . For the DTRP problem we have proposed the novel BEAD-TILING ALGORITHM and shown its stability for a uniform target-generation process with intensity λ . It is known that the system time for the DTRP problem for Dubins vehicle belongs to $\Omega(\lambda^2)$ and based on the new policy, we have shown that the system time belongs to $O(\lambda^2)$. Thus the system time of the DTRP problem for Dubins vehicle belongs to $\Theta(\lambda^2)$. This result differs from the result in the Euclidean case, where it is known that the system time belongs to $\Theta(\lambda)$. As a consequence, bounded-curvature constraints make the system much more sensitive to increases in the target generation rate.

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 serviced 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 and AFOSR MURI Award F49620-02-1-0325. The authors would like to thank John J. Enright and Prof. João Hespanha for helpful discussions.

REFERENCES

- [1] 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.
- [2] 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.
- [3] 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.
- [4] S. Rathinam, R. Sengupta, and S. Darbha, "A resource allocation algorithm for multi-vehicle systems with non holonomic constraints," Tech. Rep. UCB-ITS-RR-2005-2, Institute of Transportation Studies, University of California at Berkeley, May 2005. To appear in *IEEE Transactions on Automation Sciences and Engineering*.
- [5] 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.
- [6] A. E. Gil, 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.
- [7] 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.
- [8] A. A. Markov, "Some examples of the solution of a special kind problem on greatest and least quantities," *Soobshch. Kharkovsk. Mat. Obshch.*, vol. 1, pp. 250–276, 1887. (in Russian).
- [9] 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.
- [10] 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.
- [11] A. M. Shkel and V. J. Lumelsky, "Classification of the Dubins set," *Robotics and Autonomous Systems*, vol. 34, pp. 179–202, 2001.
- [12] 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.
- [13] 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.
- [14] S. M. LaValle, *Planning Algorithms*. Cambridge, UK: Cambridge University Press, 2006.
- [15] U. Boscaiu and B. Piccoli, *Optimal Syntheses for Control Systems on 2-D Manifolds*. Mathématiques et Applications, New York: Springer Verlag, 2004.
- [16] 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.
- [17] 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.
- [18] 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.
- [19] S. Arora, "Nearly linear time approximation scheme for Euclidean TSP and other geometric problems," in *Proc. 38th IEEE Annual Symposium on Foundations of Computer Science*, (Miami Beach, FL), pp. 554–563, Oct. 1997.
- [20] S. Lin and B. W. Kernighan, "An effective heuristic algorithm for the traveling-salesman problem," *Operations Research*, vol. 21, pp. 498–516, 1973.
- [21] 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.
- [22] 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.
- [23] 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.
- [24] 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.
- [25] S. Darbha, "Combinatorial motion planning for a collection of Reeds-Shepp vehicles," tech. rep., ASEE/AFOSR SFFP, AFRL, Eglin, Aug. 2005.
- [26] 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.
- [27] 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.
- [28] 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.
- [29] 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.
- [30] R. Motwani and P. Raghavan, *Randomized Algorithms*. Cambridge, UK: Cambridge University Press, 1995.
- [31] K. Savla, E. Frazzoli, and F. Bullo, "Traveling Salesperson Problems for the Dubins vehicle," *IEEE Transactions on Automatic Control*, June 2006. Submitted.
- [32] 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.
- [33] L. Kleinrock, *Queueing Systems. Volume I: Theory*. New York: John Wiley, 1975.
- [34] K. Savla, F. Bullo, and E. Frazzoli, "On Traveling Salesperson Problems for a double integrator," in *IEEE Conf. on Decision and Control*, (San Diego, CA), Dec. 2006. To appear.