

THE DUBINS TRAVELING SALESPERSON PROBLEM WITH STOCHASTIC DYNAMICS

Ross P. Anderson

Department of Applied Math and Statistics
University of California
Santa Cruz, California 95064
Email: anderson@soe.ucsc.edu

Dejan Milutinović*

Department of Applied Math and Statistics
University of California
Santa Cruz, California 95064
Email: dejan@soe.ucsc.edu

ABSTRACT

Motivated by applications in which a nonholonomic robotic vehicle should sequentially hit a series of waypoints in the presence of stochastic drift, we formulate a new version of the Dubins vehicle traveling salesperson problem. In our approach, we first compute the minimum expected time feedback control to hit one waypoint based on the Hamilton-Jacobi-Bellman equation. Next, minimum expected times associated with the control are used to construct a traveling salesperson problem based on a waypoint hitting angle discretization. We provide numerical results illustrating our solution and analyze how the stochastic drift affects the solution.

NOMENCLATURE

v	Speed of Dubins vehicle
Δx_i	x component of distance from Dubins vehicle to waypoint i
Δy_i	y component of distance from Dubins vehicle to waypoint i
θ	Heading angle of Dubins vehicle
u	Feedback-controlled turning rate
N	Number of waypoints
K	Number of possible final heading angles
θ_k^f	Discrete final heading angle at which the Dubins vehicle should hit the waypoint, $k = 1, \dots, K$
$T(\cdot)$	Minimum expected time to hit the waypoint
$\mu(n)$	Index in $1, \dots, N$ of the n th waypoint to be visited
$c(\cdot)$	Edge length between two nodes

INTRODUCTION

The problem of navigating a nonholonomic robotic vehicle or small aerial vehicle through a series of waypoints in minimum time has recently received much attention. Applications include target surveillance by mobile sensors and small aerial vehicles [1, 2], vehicle routing [3], and environmental sampling [4], for example.

The kinematics of the robotic vehicles is often approximated by that of a Dubins vehicle [5], which is constrained to move at fixed speed in the direction of its heading angle along paths of bounded curvature. The problem of finding the shortest path for the Dubins vehicle through a series of points has been extensively analyzed [6, 7] and has become known as the Dubins Traveling Salesman Problem (Dubins TSP [or DTSP]). Variations on this problem have considered worst-case bounds [8] for the DTSP and include extensions for multiple vehicles [9] and for other dynamical systems [10, 11].

Recent interest has paid particular attention to the stochastic case, in which the waypoint locations arise from a spatial or spatio-temporal stochastic process, and a DTSP is formulated minimize the expected time to hit or service all waypoints [6, 7, 12, 13]. However, in each of these stochastic DTSPs, the stochasticity arises in the locations of the stationary waypoints, while the motion of the vehicle is deterministic. The scenario in which the relative *motion* between the Dubins vehicle and the waypoints is stochastic has alluded attention in literature thus far. Therefore, in this paper, we provide a method for a Dubins vehicle under the influence of stochasticity to plan and execute a closed path (tour) through a series of waypoints before returning to the initial location in minimum (expected) time. In reality, it may be difficult for the Dubins vehicle to follow a prescribed sequence of points due to stochastic drifts. Because the

*Address correspondence to this author.

stochastic drift can influence the tour taken by the Dubins vehicle in practice, our offline approach based on the Hamilton-Jacobi-Bellman equation takes into account the possibility for the vehicle to be “blown” off course by the stochastic effects when approaching a waypoint, and we further include the possibility for a waypoint to be hit at a heading angle that was not intended. Previous work for a Dubins vehicle to hit a single waypoint in minimum expected time in the presence of stochastic drifts can be found in [14], and see [4] for the case of an *a priori* fixed tour.

This paper is organized as follows. In the following section, we formulate the problem and describe how it may be decoupled in terms of a stochastic optimal control problem and a subsequent traveling salesperson optimization problem. Next, we describe how these two stages are linked through an appropriate selection of the state space variables. We then present in more detail the stochastic optimal control problem and the TSP problem. We provide simulations based on our proposed method, and finally, we conclude with a discussion our findings and provide directions for future research.

PROBLEM FORMULATION

Here we formulate the problem of navigating a Dubins vehicle in the presence of stochastic drift to visit a series of waypoints and then return to its initial position (e.g., a runway or dock) in minimum expected time.

The Dubins vehicle has position $(x(t), y(t))$ and moves in the direction of its heading angle $\theta(t)$ at fixed speed v relative to the stochastic drift so that its kinematics is described by the equations

$$dx(t) = v \cos(\theta) dt + \sigma dw_x \quad (1)$$

$$dy(t) = v \sin(\theta) dt + \sigma dw_y \quad (2)$$

$$d\theta(t) = u dt, \quad |u| \leq 1 \quad (3)$$

where σ is a known noise intensity of the stochastic drift, and where dw_x and dw_y are mutually independent increments of a standard Wiener process. This type of stochasticity could arise in the model of an uncertain drift field in which the Cartesian components of the drift are assumed to be independent, with zero mean, and with known variance, for example. The feedback-controlled turning rate is u , which is bounded as $u \in [-1, 1]$.

The positions of N waypoints are fixed at (x_i, y_i) , $i = 1, \dots, N$. The evolution equations for the Cartesian components of the distance to a waypoint i , $\Delta x_i(t) = x_i - x(t)$ and $\Delta y_i(t) = y_i - y(t)$, are

$$d\Delta x_i(t) = -v \cos(\theta) dt + \sigma dw_x \quad (4)$$

$$d\Delta y_i(t) = -v \sin(\theta) dt + \sigma dw_y \quad (5)$$

$$d\theta(t) = u dt, \quad |u| \leq 1, \quad (6)$$

where we have chosen to use the symmetry of the Wiener process to revert the sign of the diffusion factor from (1)-(2).




Figure 1. DIAGRAM OF DUBINS VEHICLE AT HEADING ANGLE θ HITTING A SEQUENCE OF TARGETS $\mu(1), \mu(2), \dots$ AT WHICH POINT IT SHOULD RETURN TO ITS INITIAL CONDITION

The goal is for the Dubins vehicle to visit each of the N waypoints exactly once before returning to its initial state $(x(0), y(0), \theta(0))$ in minimum expected time. In addition to the control variable u , we must compute the order in which the waypoints are visited. To this end, we enumerate the waypoints and define a sequence $\mu = (\mu(1), \dots, \mu(N))$, where $\mu(1) = i$ denotes that the first waypoint index is i , $\mu(2)$ is the second waypoint index, and so on. Since the vehicle must return to its original location at the end of its tour, the sequence of waypoints to be visited is $(\mu(1), \dots, \mu(N), N+1)$, where $N+1$ indicates an additional waypoint located at the state $(x(0), y(0), \theta(0))$. To minimize the total expected time to traverse these waypoints under a control u , we can consider the cost function

$$T_{\text{tot}} = \min_{\substack{|u| \leq 1 \\ \mu(1), \dots, \mu(N)}} \mathbb{E} \left[\left(\sum_{n=0}^{N-1} \int_{\tau_{\mu(n)}}^{\tau_{\mu(n+1)}} 1 ds \right) + \int_{\tau_{\mu(N)}}^{\tau_{\mu(N+1)}} 1 ds \right] \quad (7)$$

where $\tau_{\mu(0)} = 0$ and

$$\tau_{\mu(k)} = \min \left\{ t > \tau_{\mu(k-1)} : \sqrt{(\Delta x_{\mu(k)}(t))^2 + (\Delta y_{\mu(k)}(t))^2} \leq r_0 \right\} \quad (8)$$

is the first time that the Dubins vehicles enters into a disc of radius r_0 about the waypoint $\mu(k)$, and where $\tau_{\mu(N+1)}$ is the time that the Dubins vehicle returns to its starting state. Since (7) requires minimization over both the control variable u and the possible permutations to the order assignment variables $(\mu(1), \dots, \mu(N))$, this is a nonstandard optimal control problem.

We thus propose to solve the problem in two stages. The first stage is comprised of an optimal control problem for a Dubins vehicle to hit a single waypoint in minimum expected time in the presence of the stochastic drift. This allows us to simultaneously compute the optimal control and the expected time to a waypoint using that optimal control. Once these times have been computed, we construct a TSP to choose the optimal sequence of waypoints. In this second stage, the cost to go from one waypoint to the next (i.e., the edge time length) is based on the expected time under the stochastic optimal control. However, as we will see in the following section, the two stages must remain linked




Figure 2. PATH THROUGH FIVE WAYPOINTS USING THE OPTIMAL MINIMUM-TIME CONTROL TO HIT EACH WITHOUT REGARD TO FINAL HEADING ANGLE (RETURN PATH AND STOCHASTIC EFFECTS NOT SHOWN).

in terms of the heading angles at which the Dubins vehicle hits the waypoints.

STATE SPACE SELECTION

In Appendix A, we briefly describe the stochastic optimal control problem for a Dubins vehicle to hit a single waypoint i in minimum time. This results in a feedback control law $u(\Delta x_i, \Delta y_i, \theta)$ for the turning rate of the vehicle. When this control is applied sequentially to drive the Dubins vehicle through multiple waypoints, the resulting trajectory is not necessarily the minimum-time trajectory for the sequence of waypoints. For example, in Figure 2, it is seen that the heading angle at which the vehicle hits its first waypoint from the south determines the behavior of the vehicle through the next four. If the first waypoint were approached from the west direction, the trajectory would be much shorter, because the Dubins vehicle could continue straightforward to the second waypoint and so on.

Since the heading angle of the Dubins vehicle as it approaches one waypoint determines its state when it leaves that waypoint, we consider separate optimal control problems for different final heading angles, and these heading angles are incorporated into the TSP.

In this work, the optimal sequence of waypoints is computed offline and is not recomputed if the Dubins vehicle misses the waypoint due to stochastic effects, i.e., if it is “blown” off course. Consequently, by incorporating the stochasticity into the edge lengths, the time for traversing all waypoints $(\mu(1), \dots, \mu(N), N+1)$ includes the chance that the Dubins vehicle will miss a waypoint and have to circle back around. If, on the other hand, the Dubins vehicle hits the waypoint but at a heading angle that was not computed by our optimization, it can continue on to the next waypoint, although this may change the time required for the entire tour. To take this into account, we also consider an upper bound to the minimum expected total time based on the scenario in which the Dubins vehicle hits each of its waypoints at a costly heading angle.




Figure 3. (A) EXPECTED TIME FOR A DUBINS VEHICLE TO HIT A WAYPOINT WITH FREE FINAL HEADING ANGLE WHEN CURRENT $\theta = -144^\circ$. (B) EXPECTED TIME FOR A DUBINS VEHICLE TO HIT A WAYPOINT AT -180° WHEN CURRENT $\theta = -144^\circ$.

STOCHASTIC OPTIMAL CONTROL PROBLEM

Given the kinematics (4)-(6), we wish to minimize the expected time to drive the Dubins vehicle from any initial state to a waypoint i with a fixed final heading angle θ^f . Then the state vector $\mathbf{x}(t) = [\Delta x_i, \Delta y_i, \theta]^\top$ is independent of any other waypoints, and the minimum expected time to hit that waypoint is

$$T(\Delta x_i, \Delta y_i, \theta | \theta^f) = \min_{|u| \leq 1} \mathbb{E} \left(\int_0^{\tau_i} 1 ds \right) \quad (9)$$

where the stopping time to the waypoint is (cf. (8))

$$\tau_i = \min \left\{ t > 0 : \sqrt{(\Delta x_i(t))^2 + (\Delta y_i(t))^2} \leq r_0, |\theta - \theta^f| \leq \varepsilon_\theta \right\}. \quad (10)$$

The (stochastic) Hamilton-Jacobi-Bellman (HJB) equation for the minimum expected time to hit the waypoint is

$$0 = \min_{|u(\mathbf{x})| \leq 1} \left\{ \partial_{\mathbf{x}} T^\top \begin{bmatrix} -v \cos \theta \\ -v \sin \theta \\ u \end{bmatrix} + \frac{\sigma^2}{2} (\partial_{\Delta x_i}^2 T + \partial_{\Delta y_i}^2 T) + 1 \right\} \quad (11)$$

with boundary condition

$$T(\Delta x_i, \Delta y_i, \theta | \theta^f) = 0 \quad \text{if } \sqrt{(\Delta x_i)^2 + (\Delta y_i)^2} \leq r_0 \quad \text{and } |\theta - \theta^f| \leq \varepsilon_\theta \quad (12)$$

To solve this partial differential equation, we employ the Markov Chain Approximation method [15], which ensures that the behavior of the spatially-discretized time to the waypoint function $T(\Delta x_i, \Delta y_i, \theta | \theta^f)$ is locally consistent the stochastic kinematics (4)-(6). Transition probabilities for the controlled Markov chain we use to compute the control using the method are in




Figure 4. ILLUSTRATION OF $N = 2$ TARGETS (CLUSTERS), EACH WITH $K = 6$ NODES, I.E., POSSIBLE TERMINAL HEADING ANGLES (PIE WEDGES). THE TSP CONSTRUCTS A TOUR STARTING WITH THE MARKED INITIAL CONDITION (I.E., FROM ONLY ONE INITIAL HEADING), VISITING EACH CLUSTER ONCE, AND RETURNS TO THE INITIAL CONDITION. THIS RESULTS IN A TSP WITH 13 NODES.

Appendix B. The minimum expected time to hit the waypoint i at heading angle θ^f is computed offline for K discrete values of uniformly-spaced final headings, $\theta_1^f, \dots, \theta_K^f \in [0, 2\pi]$. In Figure 3, we show an example of the minimum expected time for the Dubins vehicle to hit one waypoint when the incident heading angle is either free or fixed.

TRAVELING SALESPERSON PROBLEM

Once the times (9) have been computed for K final heading angles, we construct a graph of $NK + 1$ nodes, K for each waypoint, and one extra node for the final position (i.e., the initial state). The K nodes for each waypoint are organized into clusters (see Fig. 4). The goal for this section is to find a tour that begins with the initial node, visits each of the N clusters exactly once, i.e., hits only one heading angle for each waypoint, and then returns to the initial node in minimum expected time.

The directed edge length c between two nodes i and j in different clusters, with initial and final heading angles k_1 and k_2 , respectively, are based on the expected minimum time to travel that route (9) and are chosen as

$$c((i, k_1) \rightarrow (j, k_2)) = T(x_j - x_i, y_j - y_i, \theta_{k_1}^f | \theta_{k_2}^f),$$

$$i, j = 1, \dots, N, i \neq j, k_1, k_2 = 1, \dots, K, \quad (13)$$

while the edge weights leaving and returning to the final node ($N + 1$) are

$$c((N + 1) \rightarrow (j, k_2)) = T(x_j - x_0, y_j - y_0, \theta_0 | \theta_{k_2}^f)$$

$$j = 1, \dots, N, k_2 = 1, \dots, K. \quad (14)$$

$$c((i, k_1) \rightarrow (N + 1)) = T(x_0 - x_i, y_0 - y_i, \theta_{k_1}^f | \theta_0)$$

$$i = 1, \dots, N, k_1 = 1, \dots, K. \quad (15)$$

The edge weights between any two nodes that are part of the same cluster are set to be infinite. Our algorithm is similar to the heading angle discretization algorithm proposed in [8],




Figure 5. (A) MINIMUM-TIME PATH ASSUMING FIXED FINAL HEADING ANGLE (RETURN PATH NOT SHOWN). STOCHASTICITY IS NOT SHOWN FOR ILLUSTRATIVE PURPOSES. (B) MINIMUM-TIME PATH UNDER WORST-CASE FINAL HEADING ANGLE

but, here, the edge lengths incorporate the stochasticity of the times required to travel the route. This problem to find the minimum expected time route through nodes with edge lengths (13)-(15) is the generalized asymmetric traveling salesperson problem (GATSP) [16]. This NP-hard problem is generally difficult to solve, but may be cast as a standard TSP with $NK + 1$ nodes through a clever transformation [17] that does not change the optimal tour. This allows us to efficiently solve the TSP with the so-called Helsgaun implementation of the Lin-Kernighan heuristic [18] with a computational complexity of approximately $O((NK)^{2.2})$.

The above formulation will find a total expected minimum tour length based on the scenario where the Dubins vehicle hits its waypoint at the intended heading angle. However, since the waypoint may not be hit at its prescribed heading angle due to stochastic dynamics, we also wish to find the minimum expected time tour in the scenario where the actual hitting angle is adversarial, i.e., the most costly. To this end, we construct a TSP with $N + 1$ nodes, where the edge weights are

$$c(i \rightarrow j) = \max_{\theta_{k_1}^f, \theta_{k_2}^f} T(x_j - x_i, y_j - y_i, \theta_{k_1}^f | \theta_{k_2}^f),$$

$$i, j = 1, \dots, N, i \neq k. \quad (16)$$

This formulation results in a standard TSP that we solve using [18].

RESULTS

In this section, we provide examples of the solution using our approach and analyze the effect of stochastic dynamics on the heading angle. In all cases, $\sigma = 0.2$, $\nu = 1$, $r_0 = 0.1$, and $K = 36$ so that each θ^f corresponds to 10° intervals in $[0, 2\pi]$. The formulated TSP problems were solved using [18] and required approximately 7.5 [s] for $NK + 1 = 253$ nodes, for example.

In Figure 5, we show an example path generated by our proposed method assuming fixed heading angle and under a worst-




Figure 6. MINIMUM-TIME PATH FOR THE DUBINS VEHICLE THROUGH FIVE TARGETS ASSUMING A FIXED FINAL HEADING (IN THE DIRECTION OF THE ARROW HEADS). THE PATH WITH STOCHASTIC EFFECTS IS OVERLAID ON THE TRAJECTORY WITHOUT STOCHASTICITY.

case heading angle. In the former case, the optimal path is a straight line through the seven waypoints, while the worst-case scenario results in circuitous paths through the waypoints similarly to that in Figure 2. Figure 6 shows a more complex path through five waypoints including the effects of the stochasticity. It is seen that in the stochastic case, the Dubins vehicle missed waypoint 5 and had to circle around. However, the chosen sequence of waypoints accounts for this possibility since the edge weights in the TSP were based on the expected minimum time.

We find that the heading angle of the Dubins vehicle when hitting the waypoint is only minimally different from the prescribed hit heading angle θ^f , despite the relatively large influence of the stochastic motion ($\sigma = 0.2$) compared to the speed of the vehicle ($v = 1$). Figure 7(a) shows the distribution in the difference in the intended and actual heading angles when hitting a series of randomly-placed waypoints. This implies that the stochastic optimal control performs well in directing the Dubins vehicle heading angle to point toward the waypoint as it approaches, and, consequently, the TSP assuming that the intended heading angle is the final heading angle should generally be adequate. However, based on the multiple peaks seen in Figure 7(b), the Dubins vehicle may have to make several attempts (i.e., circles around) to hit the target. Although this possibility is included in our expected minimum times, it suggests that a control strategy that is able to re-plan a sequence of waypoints upon missing one may be useful in a stochastic scenario.

Next, to compare the fixed final heading case to the worst-case heading scenario, we examine the expected times and optimal waypoint sequence based on the complexity of the trajectories in terms of the number of targets and the distance between targets. Figure 8(a) shows the difference between the minimum expected tour times for the fixed final heading case and the worst-case heading angle based on the distance between waypoints and the number of waypoints. To compare the waypoint order and




Figure 7. (A) DISTRIBUTION OF THE DIFFERENCE IN INTENDED HEADING ANGLES AND ACTUAL HEADING ANGLES UPON HITTING RANDOMLY-POSITIONED WAYPOINTS. (mean=-0.08, var=0.04.) (B) HISTOGRAM OF TIME TO HIT A FIXED WAYPOINT, WHERE MULTIPLE PEAKS INDICATE THE OCCURRENCE OF MISSING THE WAYPOINT.

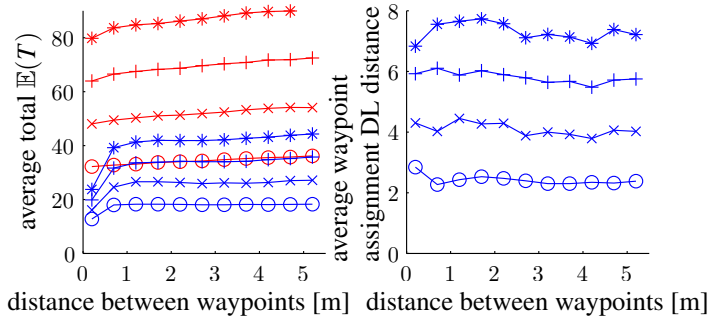


Figure 8. (LEFT) MEAN TIME TO TRAVERSE ALL WAYPOINTS AND RETURN TO THE INITIAL STATE UNDER THE FIXED-HEADING TSP (BLUE) AND THE WORST-CASE HEADING (RED), AVERAGED OVER 100 RANDOM TARGET POSITIONS BASED ON THE DISTANCE BETWEEN TARGETS. THE MARKERS \circ , \times , $+$, AND $*$ ARE FOR TWO THROUGH FIVE TARGETS, RESPECTIVELY. (RIGHT) DAMERAU-LEVENSHTEIN (DL) DISTANCE BETWEEN TARGET ASSIGNMENT SEQUENCES GENERATED BY THE FIXED-HEADING TSP AND WORST-CASE HEADING TSP, AVERAGED OVER 100 RANDOM TARGET POSITIONS BASED ON THE DISTANCE BETWEEN TARGETS.

heading angle (i.e., the nodes selected by the TSP algorithm) generated by the fixed final heading case to the worst-case heading angle TSP, we examine the Damerau-Levenshtein (DL) distance [19] between the two sequences, which corresponds to the minimum number of operations (e.g., substitution, transposition, etc.) to change from one sequence to the other. It is seen in Figure 8(b) that the DL distance is largely independent of the distance between waypoints, but increases with number of waypoints. Consequently, the change in the DTSP solution due to stochastic dynamics appears to be independent of the Cartesian location of the waypoints.

DISCUSSION

This work considers the problem of navigating a Dubins vehicle on a tour through a series of waypoints in the presence of a stochastic drift. First, a stochastic optimal control problem is formulated to determine the minimum expected time required to visit one waypoint. These expected times are used as costs for a secondary traveling salesperson optimization problem in which the order of nodes to be visited is chosen based on a hitting angle discretization. The resulting tours (including waypoint ordering, incident heading angles, and paths) anticipate the effects of the stochasticity on the motion of the vehicle.

Future work will consider ways to compute a control for hitting a waypoint with maximum probability, or with minimal variance in the expected time. Moreover, the case where the waypoints are not stationary will be examined.

ACKNOWLEDGMENT

This work is supported by NSF GRFP No. DGE-0809125. The authors wish to thank anonymous reviewers for their feedback.

REFERENCES

- [1] Tang, Z., and Ozguner, U., 2005. "Motion planning for multitarget surveillance with mobile sensor agents". *IEEE Transactions on Robotics*, **21**(5), Oct., pp. 898–908.
- [2] Beard, R., McLain, T., and Goodrich, M., 2002. "Coordinated target assignment and intercept for unmanned air vehicles". *Proceedings 2002 IEEE International Conference on Robotics and Automation*(May), pp. 2581–2586.
- [3] Bullo, F., Frazzoli, E., Pavone, M., Savla, K., and Smith, S. L., 2011. "Dynamic Vehicle Routing for Robotic Systems". *Proceedings of the IEEE*, **99**(9), pp. 1482–1504.
- [4] Anderson, R.P., Dinolov, G., Milutinovic, D., and Moore, A., 2012. "Maximally-informative ocean modeling system (ROMS) navigation of an AUV in uncertain ocean currents". In *ASME Dynamic Systems and Control Conference*.
- [5] Dubins, L. E., 1957. "On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents". *American Journal of Mathematics*, **79**(3), pp. 497–516.
- [6] Savla, K., Frazzoli, E., and Bullo, F., 2005. "On the point-to-point and traveling salesperson problems for Dubins' vehicle". In *Proceedings of the 2005 American Control Conference*, pp. 786–791.
- [7] Savla, K., Frazzoli, E., and Bullo, F., 2008. "Traveling Salesperson Problems for the Dubins Vehicle". *IEEE Transactions on Automatic Control*, **53**(6), pp. 1378–1391.
- [8] Le Ny, J., Feron, E., and Frazzoli, E., 2012. "On the Dubins Traveling Salesman Problem". *IEEE Transactions on Automatic Control*, **57**(1), Jan., pp. 265–270.
- [9] Frazzoli, E., and Bullo, F., 2004. "Decentralized algorithms for vehicle routing in a stochastic time-varying environment". In *Proceedings of the 43rd IEEE Conference on Decision and Control*, pp. 3357–3364.
- [10] Enright, J. J., and Frazzoli, E., 2006. "The Traveling Salesman Problem for the Reeds-Shepp Car and the Differential Drive Robot". In *Proceedings of the 45th IEEE Conference on Decision and Control*, Ieee, pp. 3058–3064.
- [11] Itani, S., Frazzoli, E., and Dahleh, M. A., 2008. "Travelling Salesperson Problem for dynamic systems". In *Proceedings of the 17th IFAC World Congress*, Vol. 0, pp. 13318–13323.
- [12] Enright, J. J., and Frazzoli, E., 2005. "UAV routing in a stochastic, time-varying environment". In *Proceedings of the 16th IFAC World Congress*.
- [13] Itani, S., and Dahleh, M. a., 2007. "On the Stochastic TSP for the Dubins Vehicle". In *Proceedings of the 2007 American Control Conference*, Ieee, pp. 443–448.
- [14] Anderson, R. P., Bakolas, E., Milutinović, D., and Tsiotras, P. "Optimal feedback guidance of a small aerial vehicle in a stochastic wind". *AIAA Journal of Guidance, Navigation, and Control*.
- [15] Kushner, H. J., and Dupuis, P., 2001. *Numerical Methods for Stochastic Control Problems in Continuous Time*, 2nd ed. Springer.
- [16] Laporte, G., Asef-Vaziri, A., and Sriskandarajah, C., 1996. "Some Applications of the Generalized Travelling Salesman Problem". *The Journal of the Operational Research Society*, **47**(12), pp. 1461–1467.
- [17] Behzad, A., and Modarres, M., 2002. "A new efficient transformation of the generalized traveling salesman problem into traveling salesman problem". In *Proceedings of the 15th International Conference of Systems Engineering*.
- [18] Helsgaun, K., 2000. "An effective implementation of the Lin-Kernighan traveling salesman heuristic". *European Journal of Operational Research*, **126**(1), pp. 106–130.
- [19] Damerau, F. J., 1964. "A technique for computer detection and correction of spelling errors". *Communications of the ACM*, **7**(3), pp. 171–176.

Appendix A: SOCP for Single Waypoint with Free Final Heading

Given the kinematics (4)-(6), we wish to minimize the expected time to drive the Dubins vehicle from any initial state to a waypoint i . The minimum expected time to hit that waypoint is

$$T(\Delta x_i, \Delta y_i, \theta) = \min_{|u| \leq 1} \mathbb{E} \left(\int_0^{\tau_i} 1 ds \right) \quad (17)$$

where the stopping time to the waypoint is (cf. (8))

$$\tau_i = \min \left\{ t > 0 : \sqrt{(\Delta x_i(t))^2 + (\Delta y_i(t))^2} \leq r_0 \right\}. \quad (18)$$

The (stochastic) Hamilton-Jacobi-Bellman (HJB) equation for the minimum expected time to hit the waypoint is (11), but with with boundary condition

$$T(\Delta x_i, \Delta y_i, \theta | \theta^f) = 0 \quad \text{if } \sqrt{(\Delta x_i)^2 + (\Delta y_i)^2} \leq r_0.$$

Appendix B: Markov Chain Approximation Transition Probabilities

The iterative formula for value iteration on the optimal time to a waypoint function is

$$T^h(\mathbf{x} | \theta_k^f) = \min_{|u| \leq 1} \left\{ \sum_{j=1}^3 p(\mathbf{x} \pm e_j h | \mathbf{x}, u) T^h(\mathbf{x} \pm e_j h | \theta_k^f) \Delta t^h(\mathbf{x}, u) \right\},$$

where e_j are the standard basis vectors (for Δx_i , Δy_i , or θ) and h is the step size in the respective direction. Here, the Markov chain transition probabilities $p(\cdot | \cdot)$ are given by

$$\begin{aligned} p(\Delta x_i \pm h, \Delta y_i, \theta | \Delta x_i, \Delta y_i, \theta, u) &= \Delta t^h(\mathbf{x}, u) \left[\frac{\sigma^2}{2h^2} + \frac{(-v \cos \theta)^\pm}{h} \right] \\ p(\Delta x_i, \Delta y_i \pm h, \theta | \Delta x_i, \Delta y_i, \theta, u) &= \Delta t^h(\mathbf{x}, u) \left[\frac{\sigma^2}{2h^2} + \frac{(-v \sin \theta)^\pm}{h} \right] \\ p(\Delta x_i, \Delta y_i, \theta \pm h | \Delta x_i, \Delta y_i, \theta, u) &= \Delta t^h(\mathbf{x}, u) \left[\frac{(u)^\pm}{h} \right] \\ f^\pm &= \max\{0, \pm f\}, \end{aligned} \quad (19)$$

and the interpolation time interval is

$$\Delta t^h(\mathbf{x}, u) = \left[\frac{2\sigma^2}{h^2} + \frac{|-v \cos \theta|}{h} + \frac{|-v \sin \theta|}{h} + \frac{|u|}{h} \right]^{-1}. \quad (20)$$

The computational domain in this problem is semi-periodic; the points on the boundary $\theta = \pi$ and $\theta = -\pi$ are identical, i.e., $(\Delta x_i, \Delta y_i, \pi) = (\Delta x_i, \Delta y_i, -\pi)$ for any Δx_i and Δy_i . For those states \mathbf{x} along the boundary of the computational domain, we apply reflective boundary conditions [15]:

$$\begin{aligned} T^h(\Delta x_i, \Delta y_i, \theta | \theta_k^f) &= T^h(\Delta x_i - h, \Delta y_i, \theta | \theta_k^f) \text{ if } \Delta x_i = \Delta x_i^{(\max)} \\ T^h(\Delta x_i, \Delta y_i, \theta | \theta_k^f) &= T^h(\Delta x_i + h, \Delta y_i, \theta | \theta_k^f) \text{ if } \Delta x_i = \Delta x_i^{(\min)} \\ T^h(\Delta x_i, \Delta y_i, \theta | \theta_k^f) &= T^h(\Delta x_i, \Delta y_i + h, \theta | \theta_k^f) \text{ if } \Delta y_i = \Delta y_i^{(\max)} \\ T^h(\Delta x_i, \Delta y_i, \theta | \theta_k^f) &= T^h(\Delta x_i, \Delta y_i - h, \theta | \theta_k^f) \text{ if } \Delta y_i = \Delta y_i^{(\min)} \end{aligned}$$

For those states corresponding to a successful waypoint hit (12), we apply absorbing boundary conditions [15], which resemble (12), but with $T(\cdot | \theta_k^f)$ replaced by $T^h(\cdot, \theta_k^f)$.