

The Markov-Dubins Problem in the Presence of a Stochastic Drift Field

Ross P. Anderson

Efstathios Bakolas

Dejan Milutinović

Panagiotis Tsiotras

Abstract—We consider the problem of navigating a small Dubins-type aerial or marine vehicle to a prescribed destination set in minimum expected time and in the presence of a stochastic drift field induced by local winds or currents. First, we present a deterministic control law that is independent of the local winds/currents and their statistics. Next, by employing numerical techniques from stochastic optimal control, we compute an optimal feedback control strategy that incorporates the stochastic variation in the wind when driving the Dubins vehicle to its destination set in minimum expected time. Our analyses and simulations offer a side-by-side comparison of the optimal deterministic and stochastic optimal feedback control laws for this problem, and they illustrate that the deterministic control can, in many cases, capture the salient features of structure of the stochastic optimal feedback control.

I. INTRODUCTION

We consider the problem of guiding a small aerial or marine vehicle with turning rate constraints to a prescribed terminal position in the presence of a stochastic drift field, which is induced by local winds or currents, in minimum expected time. In particular, it is assumed that the motion of the vehicle may be adequately approximated by a Dubins-type kinematic model [1–3], that is, a unicycle that travels only forward with constant speed and with a prescribed upper bound to the rate of change of the direction of its forward velocity vector. In the absence of the drift field, the vehicle traverses paths of minimal length of bounded curvature, known in the literature as Dubins paths or optimal paths of the Markov-Dubins (MD) problem [1, 4]. This kinematic model is henceforth referred to as the Dubins Vehicle (DV).

Problems characterizing the minimum-time paths of the DV in the absence of drift or in the presence of a deterministic drift field have received considerable attention in the literature. In particular, the characterization of the minimum-time synthesis, that is, a mapping that returns the minimum-time control input given the state vector of the DV, has appeared in [5–7]. The problem of characterizing minimum-time paths of the DV in the presence of a constant drift field was first posed by McGee and Hedrick in [8]. Numerical schemes for the computation of the Dubins-like

paths proposed in [8] have been presented in [9, 10], and the solution of the optimal synthesis for this problem is presented in [10, 11]. A numerical algorithm that computes the minimum-time paths of the DV in the presence of a deterministic, time-varying, yet spatially invariant, drift field appears in [12].

The aforementioned methods address variations and extensions of the MD problem within a completely deterministic optimal control framework. Some recent attempts to address the MD problem within a stochastic control framework can be found in [13, 14]. In particular, Refs. [13, 14] deal with the problem of a DV tracking a target with unpredictable future trajectory using numerical techniques from stochastic optimal control of continuous-time processes [15]. In this work, we develop an optimal *feedback* control that minimizes the expected time required to navigate the DV to its prescribed destination set in the presence of a stochastic drift field.

Our analysis and numerical simulations demonstrate that in many cases, stochastic control laws outperform their deterministic counterparts, which are “blind” to the local winds and their statistics. However, these deterministic control laws can successfully capture the salient features of the structure of the stochastic feedback control, and the similarity between the two control laws suggests that the deterministic optimal control (which has an analytic form) may suffice as a substitute over the stochastic control (which requires the solution of a partial differential equation) in light winds.

This paper is organized as follows. Section II formulates the optimal control problem. Section III presents a deterministic feedback law that provides us with useful insights regarding the optimal stochastic control and a basis with which to compare control laws. Section IV presents the stochastic optimal control and the method used to compute it. Simulation results for both the deterministic and the stochastic controllers are presented in Section V. Finally, Section VI concludes the paper with a summary of remarks.

II. PROBLEM FORMULATION

Here we formulate the problem of controlling the turning rate of a fixed-speed Dubins vehicle (DV) in order to reach a stationary target in the presence of a wind field. The target is fixed at the origin, while the Cartesian components of DV position are $x(t)$ and $y(t)$ (see Fig. 1).

The DV moves in the direction of its heading angle θ at

R. Anderson is a Graduate Student of Applied Mathematics and Statistics, University of California, Santa Cruz, CA 95064, USA, anderson@soe.ucsc.edu

E. Bakolas is an Assistant Professor in the Department of Aerospace Engineering, University of Texas at Austin, Austin, TX 78712, USA, bakolas@austin.utexas.edu

D. Milutinović is an Assistant Professor in the Department of Applied Mathematics and Statistics, University of California, Santa Cruz, CA 95064, USA, dejan@soe.ucsc.edu (Corresponding author)

P. Tsiotras is a Professor at the School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0150, USA, tsiotras@gatech.edu




Fig. 1. Diagram of a DV in position $[x(t), y(t)]^T$ and moving at heading angle θ in order to converge on a target in minimum time in the presence of wind. The target set \mathcal{T} is shown as a circle of radius δ about the target.

fixed speed v and obeys the equations:

$$dx(t) = v \cos(\theta) dt + dw_x(t, x, y) \quad (1)$$

$$dy(t) = v \sin(\theta) dt + dw_y(t, x, y) \quad (2)$$

$$d\theta(t) = \frac{u}{\rho_{\min}} dt, \quad |u| \leq 1, \quad (3)$$

where $\rho_{\min} > 0$ is the minimum turning radius constraint (in the absence of wind) and u is the control variable, $u \in [-1, 1]$. The motion of the DV is affected by the spatially and/or temporally varying wind field $w(t, x, y) = [w_x(t, x, y), w_y(t, x, y)]^T$, whose increments have been incorporated into the model (1)-(2). In this problem formulation, the model for the wind is unknown. Therefore, we assume that it is described by a stochastic process, and we formulate a stochastic control problem for reaching a target set \mathcal{T} , which is a ball of radius δ around the target. If we introduce the DV-target distance as $r(t) = \sqrt{(x(t))^2 + (y(t))^2}$ then the target set is

$$\mathcal{T} = \{r : r \leq \delta\}, \quad \delta > 0. \quad (4)$$

We avoid complex terminal constraints at the target so that we may examine the stochastic and deterministic control laws side-by-side in a transparent manner.

In order to minimize the time required to reach the target set, we define a cost function

$$J(\mathbf{x}) = \min_{|u| \leq 1} E \left\{ g(\mathbf{x}(T)) + \int_0^T dt \right\}, \quad (5)$$

and upon reaching the target set \mathcal{T} at time T , all motion ceases. In other words, (5) defines the expected value of the minimum time to reach the target set, and we wish to find the turning rate control that solves (5). In (5), $g(\mathbf{x})$ is a terminal cost, which implies that $J(\mathbf{x}) = g(\mathbf{x})$ for any $\mathbf{x} \in \mathcal{T}$ and any u [15]. This could be used to penalize or reward the DV state when hitting the target, but we choose $g(\mathbf{x}) = 0$ in the interest of transparency.

We assume that the wind is a continuous-time stochastic process with respect to the DV position, i.e. $w_x(t, x, y) = w_x(t)$ and $w_y(t, x, y) = w_y(t)$. In other words, there is no explicit relation between a realization of the wind and the DV position, although implicitly this relation may exist. Moreover, in order to focus on the effect of a stochastic wind, rather than perturbations to a known deterministic drift, the Cartesian components of the wind are assumed to evolve

independently. Drawing from the field of estimation, the simplest model to describe an unknown 2D signal suggests that the wind should be modeled as Brownian motion [16]. This choice of modeling has the advantage that an optimal feedback control can be made independent of the exact form of the underlying wind (which may not be known), and instead, it is based on the statistics of the wind. The components $w_x(t)$ and $w_y(t)$ of the wind field in (1)-(3) take the form

$$dw_x(t) = \sigma_W dW_x, \quad dw_y(t) = \sigma_W dW_y, \quad (6)$$

where dW_x and dW_y are mutually independent increments of a unit intensity Wiener process, and where the level of noise intensity σ_W quantifies the uncertainty in the evolution of the wind. In practice, the value of σ_W could be determined from the root mean square of measured wind gusts. Note that although a known deterministic drift could be added to this model, by assuming that $E\{w_x(t)\} = E\{w_y(t)\} = 0$, we omit this possibility for brevity and to keep our results transparent.

Let us define the DV-target distance $r(t) = \sqrt{(x(t))^2 + (y(t))^2}$, and let φ be the angle between the vehicle velocity vector and the line-of-sight to the target, given by $\varphi = \tan^{-1}(y/x) - \theta + \pi$ and mapped to lie in $\varphi \in (-\pi, \pi]$ (see Fig. 1). Based on (1)-(3) and (6) it can be shown using Itô's lemma that the relative DV-target system coordinates obey

$$dr(t) = \left(-v \cos(\varphi) + \frac{\sigma_W^2}{2r} \right) dt + \sigma_W dW_0 \quad (7)$$

$$d\varphi(t) = \left(\frac{v}{r} \sin(\varphi) - \frac{u}{\rho_{\min}} \right) dt + \frac{\sigma_W}{r} dW_{\perp}, \quad (8)$$

where $|u| \leq 1$, and where dW_0 and dW_{\perp} are mutually independent increments of Wiener processes aligned with the direction of DV motion θ . Note the appearance of a positive bias $\sigma_W^2/2r$ in the relation for $r(t)$, which is a consequence of the random process included in our analysis.

In the limiting case where $\sigma_W \rightarrow 0$, the problem is reduced to one that ignores the presence of stochastic winds. Section III develops an optimal feedback control that drives the DV to the target in minimum time when the stochastic wind vanishes ($\sigma_W \rightarrow 0$). In this deterministic case, the cost function is the same as (5), but without the expectation operator. We shall see later on that this deterministic optimal control, when applied to the DV in the presence of stochastic winds, will capture the salient features of the stochastic optimal feedback control.

III. DETERMINISTIC CASE

Before addressing the problem of characterizing the optimal stochastic feedback control laws that drive the DV to its target in the presence of a stochastic wind field, we shall briefly discuss a method for designing deterministic feedback controllers for the same problem. The proposed control scheme, which is based on analytic arguments, will give us significant insights for the subsequent analysis and will illustrate some interesting patterns of the synthesis of the

stochastic optimal control problem. In particular, we propose a deterministic control law that is completely independent of any information about the distribution of the winds. In other words, we design a feedback control law under the assumption that the local winds are modeled by (7)-(8) with $\sigma_W \rightarrow 0$. Therefore, our deterministic control law is “blind” to the presence and the statistics of the actual local winds. This approach will give us two navigation laws that are similar to the pure pursuit strategy from missile guidance [17], which is, in turn, a control strategy that forces the velocity vector of the controlled object (the DV in our case) to point towards its destination at every instant of time.

Note that when applying a feedback law that imitates the pure pursuit strategy, the DV will not be able to instantaneously change its motion in order to point its velocity vector toward the target, in the presence of winds. This happens for two reasons. The first reason is because the rate at which the DV can rotate its velocity vector is bounded by the turning rate constraint (3). This is true even in the absence of winds. The second reason has to do with the fact that, by hypothesis, the pure pursuit law does not account for the local winds, and, consequently, even if the DV were able to rotate its forward velocity vector $[v \cos \theta, v \sin \theta]^T$ arbitrarily, it would be this forward velocity vector that points toward the target rather than the inertial velocity $[\dot{x}, \dot{y}]^T$.

The proposed pure pursuit-like navigation law takes the following state-feedback form

$$u(\varphi) = \begin{cases} +1 & \text{if } \varphi \in (0, \pi], \\ 0 & \text{if } \varphi = 0, \\ -1 & \text{if } \varphi \in (-\pi, 0). \end{cases} \quad (9)$$

One important observation is that the control law (9) does not essentially depend on the distance $r(t)$ of the DV from the target but only on the angle φ . We shall refer to the state feedback control law given in (9) as the *geometric pure pursuit* (GPP for short) law. Note that the GPP law drives the DV to the line $\mathcal{S}_0 := \{(r, \varphi) : \varphi = 0\}$, which behaves as a “switching surface.” In the absence of wind, once the DV reaches \mathcal{S}_0 , it would travel along \mathcal{S}_0 until it reaches the target (such that $r = 0$ at the final time T) with the application of the control input $u = 0$. Therefore, the GPP law is a bang-off control law with one switching at most, that is, a control law which is necessarily a control sequence $\{\pm 1, 0\}$.

It is important to highlight that the GPP law turns out to be the time-optimal control law of the MD problem with free terminal heading for the majority (but not all) of boundary conditions (see Fig. 2), when there are no winds [18, 19]. However, there are still initial configurations for which the navigation law (9) does not give us a satisfactory answer to the steering problem, especially when the DV is close to the target with a relatively large $|\varphi|$. In particular, it can be shown [18, 19] that if the DV starts at time $t = 0$ in any point that belongs to one of the two regions \mathcal{C}_+ and \mathcal{C}_- , defined by (see Fig. 2)

$$\mathcal{C}_+ = \{(r, \varphi) : r \leq 2\rho_{\min} \sin(-\varphi), \varphi < 0\} \quad (10)$$

$$\mathcal{C}_- = \{(r, \varphi) : r \leq 2\rho_{\min} \sin(\varphi), \varphi > 0\}, \quad (11)$$




Fig. 2. Time-optimal partition of the control input space and state feedback control law of the MD problem with free terminal heading in the absence of winds. One can use this control strategy as a feedback law for the case of a stochastic drift field. Control sequences for an initial state in each time-optimal partition are indicated in red background.




Fig. 3. Level sets of the minimum time-to-go function of the MD problem with free terminal heading.

then the target cannot be reached by means of the GPP law without the presence of a stochastic drift. Therefore, in order to complete the design of a feedback control law for any possible state of the DV, we need to consider the optimal synthesis of the MD problem [18, 19]. It turns out that the boundaries of \mathcal{C}_+ and \mathcal{C}_- , denoted, respectively, by \mathcal{S}_- and \mathcal{S}_+ (the choice of the subscript notation will become apparent shortly later), correspond to two new “switching surfaces” along which the DV travels all the way to the target. In particular, when the DV starts in the interior of \mathcal{C}_+ (respectively, \mathcal{C}_-), then the minimum-time control action is $u = 1$ (respectively, $u = -1$), which may appear to be counterintuitive, since its effect is to increase $|\varphi|$ rather to decrease it. The control input remains constant until the DV reaches the “switching surface” \mathcal{S}_- (respectively, \mathcal{S}_+), where the control switches to $u = -1$ (respectively, $u = +1$), and subsequently, the DV travels along \mathcal{S}_- (respectively, \mathcal{S}_+) all the way to the target driven by $u = -1$ (respectively, $u = +1$). The net effect is that when the DV starts in regions \mathcal{C}_\pm , the DV must first distance itself from the target so that its minimum turning radius ρ_{\min} is sufficient to drive it to the target. Note that in this case the control law is bang-bang with one switching at most, that is, a control sequence $\{\pm 1, \mp 1\}$. The situation is illustrated in Fig. 2 for $\rho_{\min} = 1$.

The GPP law given in (9), therefore, needs to be updated appropriately to account for the previous remarks. In partic-

ular, the new feedback control law is given by

$$u(r, \phi) = \begin{cases} +1 & \text{if } (r, \phi) \in \Sigma_+, \\ 0 & \text{if } (r, \phi) \in \Sigma_0, \\ -1 & \text{if } (r, \phi) \in \Sigma_-, \end{cases} \quad (12)$$

where,

$$\begin{aligned} \Sigma_+ &:= \{(r, \phi) : \phi \in (0, \pi)\} \cap (\text{int}\mathcal{C}_-)^c \cup \text{int}\mathcal{C}_+ \\ \Sigma_- &:= \{(r, \phi) : \phi \in (-\pi, 0)\} \cap (\text{int}\mathcal{C}_+)^c \cup \text{int}\mathcal{C}_- \\ \Sigma_0 &:= \{(r, \phi) : \phi = 0\}. \end{aligned}$$

We henceforth refer to the state feedback law (12) as the *optimal pure pursuit* (OPP) law. Note that in the absence of winds, the OPP law is the optimal control law of the MD problem with free final heading. Figure 3 illustrates the level sets of the minimum time-to-go function, which can be computed analytically by using standard optimal control techniques (Maximum Principle) and geometric tools, as shown in [20].

IV. STOCHASTIC CASE

This section describes the value iteration computation for the optimal feedback control corresponding to the kinematic model (7)-(8) and cost functional (5) and presents the resulting control policies. When discretizing a state space for value iteration in stochastic optimal control problems, it is important that the chosen spatial and temporal step sizes accurately scale in the same way as the stochastic process. To take this into account, we employ the *Markov chain approximation method* [15], which constructs a discrete-time and discrete-state approximation to the cost function in the form of a controlled Markov chain that is “locally-consistent” with the process under control. Once the control is computed, it is valid for any initial state, including locations near the target. We first review the method as tailored for this problem before presenting the computed optimal control.

A. Markov Chain Approximation Method and Value Iteration

Denote by \mathcal{L}^u the differential operator associated with the controlled stochastic process (7)-(8), which, for the sake of brevity, we write in terms of the mean drift $\mathbf{b}(\mathbf{x}) \in \mathbb{R}^2$, the diffusion $\mathbf{a}(\mathbf{x}) \in \mathbb{R}^{2 \times 2}$, and the state vector $\mathbf{x} = [r, \phi]^T$, as $\mathcal{L}^u = \sum_{i=1}^2 b_i(\mathbf{x}) \frac{\partial}{\partial x_i} + \frac{1}{2} \sum_{i,j=1}^2 a_{ij}(\mathbf{x}) \frac{\partial^2}{\partial x_i \partial x_j}$. The state \mathbf{x} is in a domain $\mathbb{X} = \{\mathbf{x} \mid \delta \leq r < r_{\max}, -\pi \leq \phi \leq \pi\}$, which is semi-periodic because $[r, \pi]^T = [r, -\pi]^T$. It follows that the domain boundary is composed of two disjoint segments, i.e., $\partial\mathbb{X} = \{\mathbf{x} : r = \delta \text{ or } r = r_{\max}\}$.

It can be shown [15] that a sufficiently smooth $J(\mathbf{x})$ given by (5) satisfies

$$\mathcal{L}^u J(\mathbf{x}) + 1 = 0, \quad (13)$$

so that the stochastic Hamilton-Jacobi-Bellman equation for the minimum cost $V(\mathbf{x})$ over all control sequences is $\inf_u [\mathcal{L}^u V(\mathbf{x}) + 1] = 0$. This PDE has mixed boundary conditions on $\partial\mathbb{X}$. At $r = r_{\max}$, we can use reflecting boundary conditions $(\nabla V(\mathbf{x}))^T \hat{\mathbf{n}} = 0$ with the boundary normals $\hat{\mathbf{n}}$. For the part of boundary $r = \delta$ that belongs

to the target set \mathcal{T} , we have to use an absorbing boundary condition with $V(\mathbf{x}) = g(\mathbf{x}) = 0$.

The state transition probabilities $p(\mathbf{y} \mid \mathbf{x}, u)$ from the state \mathbf{x} to the state $\mathbf{y} \in \mathbb{X}$ under the control u appear as coefficients in the finite-difference approximations of the operator \mathcal{L}^u in (13). Using the so-called up-wind approximations for derivatives, the finite-difference discretization for $J(\cdot)$ with step sizes h_r and h_ϕ is

$$\begin{aligned} J^h(r, \phi) &= \Delta t_u + p(r + h_r, \phi \mid r, \phi, u) J^h(r + h_r, \phi) \\ &\quad + p(r - h_r, \phi \mid r, \phi, u) J^h(r - h_r, \phi, \gamma) \\ &\quad + p(r, \phi + h_\phi \mid r, \phi, u) J^h(r, \phi + h_\phi) \\ &\quad + p(r, \phi - h_\phi \mid r, \phi, u) J^h(r, \phi - h_\phi), \end{aligned}$$

where the transition probabilities multiplying $J^h(\cdot)$ are

$$\begin{aligned} p(r \pm h_r, \phi \mid r, \phi, u) &= \Delta t_u \left(\frac{\max[0, (\mp v \cos(\phi)) \pm \sigma_W^2/2r]}{h_r} + \sigma_W^2/2h_r^2 \right) \\ p(r, \phi \pm h_\phi \mid r, \phi, u) &= \Delta t_u \left(\frac{\max[0, (\pm v \sin(\phi)/r \mp u/\rho_{\min})]}{h_\phi} + \sigma_W^2/2r^2 h_\phi^2 \right), \end{aligned}$$

and where “max” is a result of the up-wind approximation, and Δt_u is a state- and control-dependent interpolation interval of the piecewise constant chain [15], given by

$$\begin{aligned} \Delta t_u(\mathbf{x}, u) &= (|-v \cos(\phi) + \sigma_W^2/2r|/h_r \\ &\quad + |v \sin(\phi)/r - u/\rho_{\min}|/h_\phi + \frac{\sigma_W^2}{h_r^2} + \frac{\sigma_W^2}{r^2 h_\phi^2})^{-1}. \end{aligned}$$

The dynamic programming equation for the Markov chain used for value iteration is as follows [15]:

$$V^h(\mathbf{x}) = \min_{|u| \leq 1} \left\{ \Delta t_u(\mathbf{x}, u) + \sum_{\mathbf{y}} p(\mathbf{y} \mid \mathbf{x}, u) V^h(\mathbf{y}) \right\} \quad (14)$$

for all $\mathbf{x} \in \mathbb{X} \setminus \partial\mathbb{X}$. For the reflective part of the boundary, $r = r_{\max}$, we use, instead of (14), $V^h(\mathbf{x}) = \sum_{\mathbf{y}} p(\mathbf{y} \mid \mathbf{x}) V^h(\mathbf{y})$ [15, pp. 143], where $p(\mathbf{y} \mid \mathbf{x}) = 1$ for $\mathbf{y} = [r_{\max} - h_r, \phi]^T$ and $\mathbf{x} = [r_{\max}, \phi]^T$; otherwise, $p(\mathbf{y} \mid \mathbf{x}) = 0$. Finally, for those states $\mathbf{x} \in \mathcal{T}$ in the target set (4), we impose the terminal condition $V^h(\mathbf{x}) = g(\mathbf{x}) = 0$.

Equation (14) along with the reflective transition equation and terminal boundary condition equation are used in the method of value iteration until the cost converges (the interested reader may find a proof of convergence in [15]).

B. Optimal Control in a Stochastically-varying Wind

Here we describe the stationary optimal control computed for the stochastic wind. We chose parameters as $\rho_{\min} = 1$, $r_{\max} = 10$, $v = 1$, $h_r = 0.02$, and $h_\phi = 0.025$, with all units in meters and seconds.

The structure of the optimal control law for the discrete-time Markov chain that approximates the continuous-time control problem is seen in Fig. 4(a) for $\sigma_W = 0.1$ and Fig. 4(b) for $\sigma_W = 0.5$. As in the deterministic case (Fig. 2), the control is composed of bang-bang regions instructing


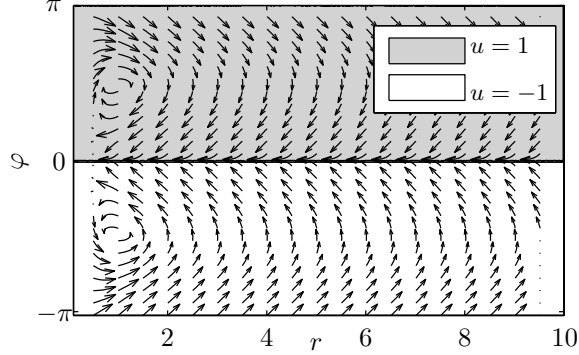

 (a) Control optimal for approximating Markov chain for $\sigma_W = 0.1$

 (b) With more variation in the wind ($\sigma_W = 0.5$), the optimal control yields the GPP model (9).

 Fig. 4. Dubins vehicle optimal turning rate control policy $u(r, \varphi)$.

the DV to turn left or right, which aids in our comparison of the stochastic feedback control to the deterministic OPP control. With $\sigma_W = 0.1$, the optimal control is comprised of four regions, two directing the target to turn left, and others instructing a turn to the right. The reader should note the similarity between Fig. 4(a) and the OPP control illustrated in Fig. 2. In particular, the structure of the regions \mathcal{C}_- and \mathcal{C}_+ have changed somewhat as a consequence of the stochastic variation of the wind. In Fig. 4(b), a higher noise intensity of $\sigma_W = 0.5$ causes the control to return to GPP (9), and this control strategy remains optimal for even larger σ_W .

V. PERFORMANCE COMPARISON

This section provides a comparison of performance of the proposed feedback control laws. As an example, Fig. 5 shows two DV's approaching a target in the presence of the stochastic wind. When entering the region near the target where the two control laws differ (cf. Figs. 2 and 4(a)), the DV applying the OPP control law (12) is instructed to loop around twice before reaching the target.

We next examine the mean hitting time, i.e., the average time required for the DV to hit the target set as a function of its initial state. Figure 6 compares the mean hitting time under (7)-(8), using both the OPP control law (12) and the optimal control shown in Fig. 4(a). It is seen that the mean times under OPP are greater in regions near the target, although in a small subset of these regions, the stochastic optimal control law has higher standard deviation.





 Fig. 5. Two DVs, initially located at $(x, y, \theta) = (-5, 0, \pi/6)$ attempt reach the target set \mathcal{T} centered about $(0, 0)$ in minimum time in a stochastic wind ($\sigma_W = 0.1$). The blue DV applies the stochastic optimal control for $\sigma_W = 0.1$, while the red DV applies the optimal pure pursuit (OPP) strategy. The right panel shows a closeup of the target.

 (a) Mean time $E(T)$ to hit the target based on the starting location in state space (r, φ) with OPP control (top) and the stochastic optimal control $u(r, \varphi)$ (bottom).

 (b) Left: difference in mean hitting time $E(T_{\text{OPP}}) - E(T_{\text{stoch}})$. Red implies OPP has higher $E(T)$. Right: difference in standard deviation $\text{std}(T_{\text{OPP}}) - \text{std}(T_{\text{stoch}})$. Orange and red regions correspond to higher OPP standard deviation; dark blue regions indicate higher standard deviation under stochastic control.

Fig. 6. Comparison of distributions of time required to hit the target under both OPP control (12) and the stochastic optimal control.

Since the stochastic control is specific to the intensity σ_W of the wind, we also show how the expected minimum hitting time changes when a control computed for one value of $\sigma_W = \sigma_W^{(\text{control})}$ performs against a wind of intensity $\sigma_W^{(\text{actual})}$ in Fig. 7. Although a larger $\sigma_W^{(\text{actual})}$ leads to longer DV paths, a suitable $\sigma_W^{(\text{control})}$ mitigates this effect.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we have addressed the problem of driving a small vehicle with Dubins-type kinematics to a prescribed destination with free final heading in the presence of a




Fig. 7. Time to reach the target, initially at a distance of 5 [m], in a stochastic wind with simulated intensity $\sigma_W^{(\text{actual})}$ using a control computed for $\sigma_W^{(\text{control})}$. Each difference in noise intensity is an average of the $E(T)$ from all combinations of $\sigma_W^{(\text{control})}$ and $\sigma_W^{(\text{actual})}$, further averaged over 1000 simulations for each pair of noise intensities.

stochastic drift field in minimum expected time. We have proposed two approaches to this problem. The first one, which was based on analytic techniques, was to employ a deterministic feedback control law that is similar to the pure pursuit law from the field of missile guidance. The second approach was to tackle the problem computationally by employing numerical tools from stochastic optimal control theory. Our side-by-side comparison and simulations have revealed that although, in general, stochastic control laws that explicitly account for the drift field outperform the analytic deterministic laws that are suboptimal in stochastic winds, the latter ones in many cases capture the structure of the stochastic feedback control. Future work includes the extension of the techniques presented herein to problems with a more realistic model of the drift field, including spatially-correlated winds.

VII. ACKNOWLEDGMENTS

The authors thank the reviewers for their helpful comments. This work is supported by NSF GRFP 0809125.

REFERENCES

- [1] 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, no. 3, pp. 497–516, 1957.
- [2] R. Isaacs, "Games of pursuit," RAND Corporation, Santa Monica, CA, RAND Report P-257, 1951.
- [3] V. S. Patsko and V. L. Turova, "Numerical study of the "homicidal chauffeur" differential game with the reinforced pursuer," *Game Theory and Applications*, vol. 12, pp. 123–152, 2007.
- [4] H. J. Sussmann, "The Markov-Dubins problem with angular acceleration control," in *Proceedings of 36th IEEE Conference on Decision and Control*, San Diego, CA, Dec. 1997, pp. 2639–2643.
- [5] T. Pecsvaradi, "Optimal horizontal guidance law for aircraft in the terminal area," *IEEE Transactions on Automatic Control*, vol. 17, no. 6, pp. 763–772, 1972.
- [6] X. N. Bui and J. D. Boissonnat, "The shortest paths synthesis for nonholonomic robots moving forwards," Institut National de Recherche en Informatique et en Automatique, Sophia-Antipolis, France, Research Note 2153, 1993.
- [7] V. S. Patsko, S. G. Pyatko, and A. A. Fedotov, "Three-dimensional reachability set for a nonlinear control system," *Journal of Computer and Systems Sciences International*, vol. 42, no. 3, pp. 320–328, 2003.
- [8] T. G. McGee and J. K. Hedrick, "Optimal path planning with a kinematic airplane model," *Journal of Guidance, Control, and Dynamics*, vol. 30, no. 2, pp. 629–633, 2007.
- [9] L. Tchy and C. A. Woosley, "Minimum-time path-planning for unmanned aerial vehicles in steady uniform winds," *Journal of Guidance, Control, and Dynamics*, vol. 32, no. 6, pp. 1736–1746, 2009.
- [10] E. Bakolas and P. Tsiotras, "Time-optimal synthesis for the Zermelo-Markov-Dubins problem: the constant wind case," in *Proceedings of American Control Conference*, Baltimore, MD, June 30–July 2, 2010, pp. 6163–6168.
- [11] E. Bakolas, "Optimal steering for kinematic vehicles with applications to spatially distributed agents," Ph.D. dissertation, School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA, 2011.
- [12] R. L. McNeely, R. V. Iyer, and P. R. Chandler, "Tour planning for an unmanned air vehicle under wind conditions," *Journal of Guidance, Control, and Dynamics*, vol. 30, no. 5, pp. 1299–1306, 2007.
- [13] R. P. Anderson and D. Milutinović, "Dubins vehicle tracking of a target with unpredictable trajectory," in *Proc. 4th ASME Dynamic Systems and Control Conference*, Arlington, VA, Oct. 31–Nov. 2, 2011.
- [14] R. Anderson and D. Milutinović, "A stochastic approach to Dubins feedback control for target tracking," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, San Francisco, CA, Sep. 25–30, 2011, pp. 3917–3922.
- [15] H. J. Kushner and P. Dupuis, *Numerical Methods for Stochastic Control Problems in Continuous Time*, 2nd ed. Springer, 2001.
- [16] C. Gardiner, *Stochastic Methods: A Handbook for the Natural and Social Sciences*, 4th ed. Springer, 2009.
- [17] E. Bakolas and P. Tsiotras, "Feedback navigation in an uncertain flow-field and connections with pursuit strategies," *Journal of Guidance, Control, and Dynamics*, 2012, in press.
- [18] J. D. Boissonnat and X. N. Bui, "Accessibility region for a car that only moves forward along optimal paths," Institut National de Recherche en Informatique et en Automatique, Sophia-Antipolis, France, Research Note 2181, 1994.
- [19] B. Thomaschewski, "Dubins' problem for the free terminal direction," *Preprint*, pp. 1–14, 2001.
- [20] E. Bakolas and P. Tsiotras, "Optimal synthesis of the asymmetric sinistral/dextral Markov-Dubins problem," *Journal of Optimization Theory and Applications*, vol. 150, no. 2, pp. 233–250, 2011.