On a problem of stochastic reach-avoid set characterization

Peyman Mohajerin Esfahani, Debasish Chatterjee, and John Lygeros

Abstract— We develop a novel framework for formulating a class of stochastic reachability problems with state constraints as a stochastic optimal control problem. Previous approaches to solving these problems are either confined to the deterministic setting or address almost-sure stochastic notions. In contrast, we propose a new methodology to tackle probabilistic specifications that are less specific than almost sure requirements. To this end, we first establish a connection between two stochastic reach-avoid problems and a class of stochastic optimal control problems for diffusions with discontinuous payoff functions. We then derive a weak version of dynamic programming principle (DPP) for the value function. Moreover, based on our DPP, we give an alternate characterization of the value function as the solution to a partial differential equation in the sense of discontinuous viscosity solutions. Finally we validate the performance of the proposed framework on the stochastic Zermelo navigation problem.

I. INTRODUCTION

Reachability is a fundamental concept in the study of dynamical systems, and in view of applications of this concept ranging from engineering, manufacturing, biology, and economics, to name but a few, has been studied extensively in the control theory literature. One particular problem that has turned out to be of fundamental importance in engineering is the so-called "reach-avoid" problem. In the deterministic setting this problem consists of determining the set of initial conditions for which one can find at least one control strategy to steer the system to a target set while avoiding certain obstacles. The set representing the solution to this problem is known as capture basin [1]. This problem finds applications in, air traffic management [26], security of power networks [18]. A direct approach to compute the capture basin is formulated in the language of viability theory in [11], [13]. Related problems involving pursuit-evasion games are solved in, e.g., [2], [22] employing tools from non-smooth analysis, for which computational tools are provided by [13].

An alternative and indirect approach to reachability involves using level set methods defined by value functions that characterize appropriate optimal control problems. Employing dynamic programming techniques for reachability and viability problems in the absence of state-constraints, these value functions can in turn be characterized by solutions to the standard Hamilton-Jacobi-Bellman (HJB) equations

corresponding to these optimal control problems [25], [28]. Numerical algorithms based on level set methods were developed by [31], [32] and have been coded in efficient computational tools by [28], [29]. Extending the scope of this technique, the authors of [20], [7], [27] treat the case of time-independent state constraints and characterize the capture basin by means of a control problem whose value function is continuous.

In the stochastic setting, different probabilistic analogs of reachability problems have been studied extensively: Almost-sure stochastic viability and controlled invariance are treated in [1], [4], [6]; see also the references therein. Methods involving stochastic contingent sets [3], [4], viscosity solutions of second-order partial differential equations [10], [5], [6], and derivatives of the distance function [15] were developed in this context. In [16] the authors developed an equivalence for the invariance problem between a stochastic differential equation and a certain deterministic control system Following the same problem, the authors of [33] studied the differential properties of the reachable set based on the geometrical partial differential equation which is the analogue of the HJB equation for this problem.

Although almost sure versions of reachability specifications are interesting in their own right, they may be too strict a concept in some applications. For example, in the safety assessment context, a common specification involves bounding the probability that undesirable events take place. Motivated by this, in this article we develop a new framework for solving the following stochastic reach-avoid problem: *Given $p \in [0,1]$, a horizon $T > 0$, and two disjoint sets $A, B \subset \mathbb{R}^n$, construct, if possible, a policy such that the controlled processes reaches A prior to entering B within the interval $[0,T]$.* Observe that this is a significantly different problem compared to its almost-sure counterpart referred to above. It is of course immediate that the solution to the above problem is trivial if the initial state is either in B (in which case it is almost surely impossible) or in A (in which case there is nothing to do). However, for generic initial conditions in $\mathbb{R}^n \setminus (A \cup B)$, due to the inherent probabilistic nature of the dynamics, the problem of selecting a policy and determining the probability with which the controlled process reaches the set A prior to hitting B is non-trivial.

In this article we establish a connection between the stochastic reach-avoid problems and a stochastic optimal control problem with discontinuous payoff functions. In a fashion similar to [9], we propose a weak version of the dynamic programming principle (DPP) which avoids the technical difficulties related to the measurability of value functions. To this end, the proposed approach imposes fairly

The authors are with the Automatic Control Laboratory, ETH Zürich, Physikstrasse 3, 8092 Zürich, Switzerland. Emails: mohajerin,chatterjee,lygeros@control.ee.ethz.ch

mild conditions on the system dynamics and the sets, namely, a non-degeneracy of the diffusion term and an interior cone condition, respectively. In the sequel, we will derive the dynamic programming equation in the sense of viscosity solutions based on our weak DPP. As a by-product of our work, we shall also address the related issue determining whether there exists a policy such that with probability at least p the controlled processes resides in A at time T while avoiding B on the interval $[0, T]$.

In §II we introduce the two stochastic reach-avoid problems dealt with in this article. §III provides a connection between the stochastic reach-avoid problems and a class of stochastic optimal control problems. In §IV we first establish a DPP corresponding to the value function, and characterize it as the (discontinuous) viscosity solution of a partial differential equation. §V presents a connection between §III and §IV and solves the stochastic reach-avoid problem in a "ϵ-conservative" sense. One may observe that this ϵ-precision can be made arbitrarily small. To illustrate the performance of the our techniques, the theoretical results developed in preceding sections are applied to solve the stochastic Zermelo navigation problem in §VI. We conclude with some remarks and directions for future work in §VII.

II. Problem Statement

Consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ whose filtration $\mathbb{F} = (\mathcal{F}_s)_{s \geq 0}$ is generated by a n-dimensional Brownian motion $(W_s)_{s \geq 0}$ adapted to $\mathbb{F}$. Let the natural filtration of the Brownian motion $(W_s)_{s \geq 0}$ be enlarged by its right-continuous completion; — the usual conditions of completeness and right continuity, and $(W_s)_{s \geq 0}$ is a Brownian motion with respect to $\mathbb{F}$ [23, p. 48]. Let $\mathbb{U} \subset \mathbb{R}^m$ be a control set, and let $\mathcal{U}$ denote the set of $\mathbb{F}$-progressively measurable maps into $\mathbb{U}$.[1] The basic object of our study concerns the $\mathbb{R}^n$-valued stochastic differential equation (SDE)

$$\mathrm{d}X_s = f(X_s, u_s)\, \mathrm{d}s + \sigma(X_s, u_s)\, \mathrm{d}W_s, \quad s \geq 0, \quad (1)$$

where $X_0 = x$ given, $f : \mathbb{R}^n \times \mathbb{U} \longrightarrow \mathbb{R}^n$ and $\sigma : \mathbb{R}^n \times \mathbb{U} \longrightarrow \mathbb{R}^{n \times n}$ are measurable maps, $(W_s)_{s \geq 0}$ is the above standard n-dimensional Brownian motion, and $\boldsymbol{u} := (u_s)_{s \geq 0} \in \mathcal{U}$.

Assumption 2.1: We stipulate that

a. $\mathbb{U} \subset \mathbb{R}^m$ is compact;
b. f is continuous and Lipschitz in its first argument uniformly with respect to the second;
c. σ is continuous and Lipschitz in its first argument uniformly with respect to the second.

It is known [8] that under Assumption 2.1 there exists a unique strong solution to the SDE (1). By definition of the filtration $\mathbb{F}$, we see that the control functions $\boldsymbol{u} \in \mathcal{U}$ satisfy the *non-anticipativity condition* [8]—to wit, the increment $W_t - W_s$ is independent of the past history $\{W_y, u_y \mid y \leq s\}$ of the Brownian motion and the control for every $s \in [0, t[$

[1]Recall [17, Chapter IV] that a $\mathbb{U}$-valued process $(y_t)_{t \geq 0}$ is $\mathbb{F}$-progressively measurable if for each $T > 0$ the function $\Omega \times [0, T] \ni (\omega, t) \mapsto y(\omega, t) \in \mathbb{U}$ is measurable, where $\Omega \times [0, T]$ is equipped with $\mathcal{F} \otimes \mathfrak{B}([0, T])$, $\mathbb{U}$ is equipped with $\mathfrak{B}(\mathbb{U})$, and $\mathfrak{B}(S)$ denotes the Borel σ-algebra on a topological space S.

(In other words, $\boldsymbol{u}$ does not anticipate the future increment of W). We let $(X_s^{t,x;\boldsymbol{u}})_{s \geq t}$ denote the unique strong solution of (1) starting from time t at the state x under the control policy $\boldsymbol{u}$. We denote by $\mathcal{T}$ the collection of all $\mathbb{F}$-stopping times. For $\tau_1, \tau_2 \in \mathcal{T}$ with $\tau_1 \leq \tau_2$ $\mathbb{P}$-a.s. the subset $\mathcal{T}_{[\tau_1, \tau_2]}$ is the collection of all $\mathbb{F}$-stopping times τ that $\tau_1 \leq \tau \leq \tau_2$ $\mathbb{P}$-a.s. and are conditionally independent of $\mathcal{F}_{\tau_1}$ given $(\tau_1, X_{\tau_1}^{t,x;\boldsymbol{u}})$. We also denote by $\mathcal{U}_\tau$ the collection of all processes $\boldsymbol{u} \in \mathbb{U}$ which are conditionally independent of $\mathcal{F}_\tau$ given $(\tau, X_\tau^{t,x;\boldsymbol{u}})$. Measurability on $\mathbb{R}^n$ will always refer to Borel-measurability. In the sequel the complement of a set $S \subset \mathbb{R}^n$ is denoted by S^c.

Definition 2.2 (First entry time): Given a control $\boldsymbol{u}$, the process $(X_s^{t,x;\boldsymbol{u}})_{s \geq t}$, and a measurable set $A \subset \mathbb{R}^n$, we introduce[2] the first entry time to A:

$$\tau_A(t, x) = \inf\{s \geq t \mid X_s^{t,x;\boldsymbol{u}} \in A\}. \quad (2)$$

In view of [19, Theorem 1.6, Chapter 2], $\tau_A(t, x)$ is an $\mathbb{F}$-stopping time.

Fig. 1. The trajectory $X^{(1)}$ hits A prior to B within time $[0, T]$, while $X^{(2)}$ and $X^{(3)}$ do not; all three start from initial state x_0.

Given an initial condition (t, x), we define the set $\mathrm{RA}(t, p; A, B)$ (resp. $\widetilde{\mathrm{RA}}(t, x; A, B)$) as the set of all initial conditions such that there exists an admissible control strategy $\boldsymbol{u} \in \mathcal{U}$ such that with probability more than p the state trajectory $X_s^{t,x;\boldsymbol{u}}$ hits the set A before set B within the time (resp. at the time) horizon T.

Definition 2.3 (Reach-Avoid within the interval $[0, T]$):

$$\mathrm{RA}(t, p; A, B) := \Big\{ x \in \mathbb{R}^n \ \Big| \ \exists \boldsymbol{u} \in \mathcal{U} \ : \ p <$$
$$\mathbb{P}_{t,x}^{\boldsymbol{u}}\Big(\exists s \in [t, T], \ X_s^{t,x;\boldsymbol{u}} \in A \text{ and } \forall r \in [t, s] \ X_r^{t,x;\boldsymbol{u}} \notin B \Big) \Big\}.$$

Definition 2.4 (Reach-Avoid at the terminal time T):

$$\widetilde{\mathrm{RA}}(t, p; A, B) := \Big\{ x \in \mathbb{R}^n \ \Big| \ \exists \boldsymbol{u} \in \mathcal{U} \ :$$
$$p < \mathbb{P}_{t,x}^{\boldsymbol{u}}\Big(X_T^{t,x;\boldsymbol{u}} \in A \text{ and } \forall r \in [t, T] \ X_r^{t,x;\boldsymbol{u}} \notin B \Big) \Big\}.$$

We have suppressed the initial condition in the above probabilities, and will continue doing so in the sequel. A pictorial representation of our problems is in Figure 1.

III. Connection to Stochastic Optimal Control Problem

In this section we establish a connection between the stochastic reach-avoid problems and a stochastic optimal control problem. Consider the value functions $V, \widetilde{V} : [0, T] \times \mathbb{R}^n \to \mathbb{R}$ defined as follows:

$$V(t, x) := \sup_{\boldsymbol{u} \in \mathcal{U}} \mathbb{E}\big[\mathbb{1}_A(X_\tau^{t,x;\boldsymbol{u}})\big], \quad \tau := \tau_{A \cup B} \wedge T, \quad (3\mathrm{a})$$

$$\widetilde{V}(t, x) := \sup_{\boldsymbol{u} \in \mathcal{U}} \mathbb{E}\big[\mathbb{1}_A(X_{\widetilde{\tau}}^{t,x;\boldsymbol{u}})\big], \quad \widetilde{\tau} := \tau_B \wedge T. \quad (3\mathrm{b})$$

where $\wedge$ denotes the minimum operator. Here $\tau_{A \cup B}$ is the hitting time introduced in Definition 2.2, and depends on

[2]By convention, $\inf \emptyset = \infty$.

the initial condition (t, x). Also note that for a measurable function $f : \mathbb{R}^n \to \mathbb{R}$ hereinafter $\mathbb{E}\big[f\big(X_\tau^{t,x;\boldsymbol{u}}\big)\big]$ stands for conditional expectation with initial condition (t, x) given and under the control $\boldsymbol{u}$. For notational simplicity, we drop the initial condition in this section.

In our subsequent work, measurability of the functions V and $\widetilde{V}$ turn out to be irrelevant; see Remark 4.6 for details.

The first result of this section, Proposition 3.2, asserts that $\mathbb{E}\big[\mathbb{1}_A(X_\tau^{t,x;\boldsymbol{u}})\big] = \mathbb{P}_{t,x}^{\boldsymbol{u}}\big(\tau_A < \tau_B, \ \tau_A \leq T\big)$. Since τ_A and τ_B are $\mathbb{F}$-stopping times, it then indicates mapping $(t, x) \mapsto \mathbb{E}\big[\mathbb{1}_A(X_\tau^{t,x;\boldsymbol{u}})\big]$ is well-defined. Furthermore, the similar results can be deduced for the other reachability problem introduced in Definition 2.4 and the value function (3b).

Assumption 3.1: The sets A and B are disjoint and closed.

Proposition 3.2: Consider the system (1), and let $A, B \subset \mathbb{R}^n$ be given. Under Assumptions 2.1 and 3.1 we have

$$\mathrm{RA}(t, p; A, B) = \{x \in \mathbb{R}^n \mid V(t, x) > p\},$$

where the set RA is the set defined in Definition 2.3 and V is the value function defined in (3a).

Proof: In view of Assumption 3.1, the definition of reach-avoid set in 2.3, and [30, Remark 2.3] we can express the set $\mathrm{RA}(t, p; A, B)$ as

$$\mathrm{RA}(t, p; A, B) = \Big\{x \in \mathbb{R}^n \mid \exists \boldsymbol{u} \in \mathcal{U} \ : \\ \mathbb{P}_{t,x}^{\boldsymbol{u}}\big(\tau_A < \tau_B \ \text{ and } \ \tau_A \leq T\big) > p\Big\}. \quad (4)$$

Also, by Assumption 3.1, and the definition of the stopping time τ in (3a), for a given $\boldsymbol{u} \in \mathcal{U}$ we have

$$X_\tau^{t,x;\boldsymbol{u}} \in A \Rightarrow \tau_A \leq \tau \text{ and } \tau \neq \tau_B \Rightarrow T \geq \tau = \tau_A < \tau_B,$$

which means the sample path $X_\cdot^{t,x;\boldsymbol{u}}$ hits the set A before B at the time $\tau \leq T$; let us recall that by Definition 2.2 $\tau_{A \cup B} = \tau_A \wedge \tau_B$. Moreover,

$$X_\tau^{t,x;\boldsymbol{u}} \notin A \Rightarrow \tau \neq \tau_A \Rightarrow \tau = (\tau_B \wedge T) < \tau_A,$$

and this means that the sample path does not succeed in reaching A while avoiding set B within time T. Therefore, the event $\{\tau_A < \tau_B \text{ and } \tau_A \leq T\}$ is equivalent to $\{X_\tau^{t,x;\boldsymbol{u}} \in A\}$, and

$$\mathbb{P}_{t,x}^{\boldsymbol{u}}\big(\tau_A < \tau_B \ \text{ and } \ \tau_A \leq T\big) = \mathbb{E}\big[\mathbb{1}_A(X_\tau^{t,x;\boldsymbol{u}})\big].$$

This, in view of (4) and arbitrariness of control strategy $\boldsymbol{u} \in \mathcal{U}$ leads to the assertion. ∎

We state the following proposition concerning assertions identical to those of Proposition 3.2 for the reach-avoid problem of Definition 2.4. The proof follows effectively the same approach as that of Proposition 3.2.

Proposition 3.3: Consider the system (1), and let $A, B \subset R^n$ be given. If the set B is closed, then under Assumption 2.1 we have $\widetilde{\mathrm{RA}}(t, p; A, B) = \{x \in \mathbb{R}^n \mid \widetilde{V}(t, x) > p\}$, where the set $\widetilde{\mathrm{RA}}$ is the set defined in Definition 2.4.

IV. ALTERNATIVE CHARACTERIZATION OF REACH-AVOID PROBLEM

The stochastic control problems introduced in (3) are well-known as an exit-time problem. In this section we present an alternative characterization of this class of problems as the (discontinuous) viscosity solution of a partial differential equation. To this end, we generalize the value functions to

$$V(t, x) := \sup_{\boldsymbol{u} \in \mathcal{U}_t} \mathbb{E}\big[\ell\big(X_{\bar{\tau}(t,x)}^{t,x;\boldsymbol{u}}\big)\big], \quad \bar{\tau}(t, x) := \tau_O \wedge T, \quad (5)$$

with the function $\ell : \mathbb{R}^n \to \mathbb{R}$ being bounded measurable, and O a measurable set. Note that τ_O is the stopping time defined in Definition 2.2 that in case of value function (3a) can be considered as $O = A \cup B$. Note once again that measurability of the function V is irrelevant to our work; see Remark 4.6 for details.

Hereafter we shall restrict our control processes to $\mathcal{U}_t$, the set $\mathcal{U}_t$ denotes the collection of all processes $\boldsymbol{u} \in \mathcal{U}$ which are conditionally independent of $\mathcal{F}_t$ given (t, x). In view of independence of the increments of Brownian motion, the restriction of control processes to $\mathcal{U}_t$ is not restrictive, and one can show that the value function in (5) remains the same if $\mathcal{U}_t$ is replaced by $\mathcal{U}$; see, for instance, [24, Theorem 3.1.7, p. 132] and [9, Remark 5.2].

Our objective is to characterize the value function (5) as a (discontinuous) viscosity solution of a suitable Hamilton-Jacobi-Bellman equation. We introduce the set $\mathbb{S} := [0, T] \times \mathbb{R}^n$ and define the lower and upper semicontinuous envelopes of function $V : \mathbb{S} \to \mathbb{R}$:

$$V_*(t, x) := \liminf_{(t', x') \to (t, x)} V(t', x')$$
$$V^*(t, x) := \limsup_{(t', x') \to (t, x)} V(t', x').$$

A. Assumptions and Preliminaries

Assumption 4.1: In addition to Assumption 2.1, we stipulate the following:

a. *(Non-degeneracy)* The controlled processes are uniformly non-degenerate, i.e., there exists $\delta > 0$ such that for all $x \in \mathbb{R}^n$ and $u \in \mathbb{U}$, $\sigma\sigma^\top \geq \delta I$ where $\sigma = \sigma(x, u)$ is the diffusion term in SDE (1).

b. *(Interior Cone Condition)* There exist positive constants h, r an $\mathbb{R}^n$-value bounded map $\eta : \overline{O} \to \mathbb{R}^n$ satisfying

$$\mathrm{B}_{rt}\big(x + \eta(x)t\big) \subset O \qquad \text{for all } x \in \overline{O} \text{ and } t \in (0, h]$$

where $\mathrm{B}_r(x)$ denotes an open ball centered at x and radius r and $\overline{O}$ stands for the closure of the set O.

c. *(Lower Semicontinuity)* The function ℓ in (5) is lower semicontinuous.

Note that if the set A in §III is open, then $\ell(\cdot) = \mathbb{1}_A(\cdot)$ satisfies Assumption 4.1.c. The interior cone condition in Assumption 4.1.b. concerns shapes of the set O; figure 2 illustrate two typical scenarios:

Let us define the function $J : \mathbb{S} \times \mathbb{U} \to \mathbb{R}$:

$$J\big(t, x; \boldsymbol{u}\big) := \mathbb{E}\Big[\ell\big(X_{\bar{\tau}(t,x)}^{t,x;\boldsymbol{u}}\big)\Big], \quad (6)$$

Fig. 2. Interior cone condition holds at every point of the boundary.

where the stopping time $\bar{\tau}$ is defined in (5) and depends on the initial condition (t, x). In the following proposition, we establish continuity of $\bar{\tau}(t, x)$ and lower semicontinuity of $J(t, x, \boldsymbol{u})$ with respect to (t, x).

Proposition 4.2: Consider the system (1), and suppose that Assumptions 2.1 and 4.1 hold. Then for any strategy $\boldsymbol{u} \in \mathcal{U}$, the function $(t, x) \mapsto \bar{\tau}(t, x)$ is continuous $\mathbb{P}$-a.s. Moreover, the function $(t, x) \mapsto J(t, x, \boldsymbol{u})$ defined in (6) is uniformly bounded and lower semicontinuous:

$$J(t, x; \boldsymbol{u}) \leq \liminf_{(t', x') \to (t, x)} J(t', x'; \boldsymbol{u}).$$

Proof: Here we briefly sketch the proof in words and refer to [30, Proposition 4.3] for more details. Based on the Assumptions 4.1.a. and 4.1.b., it is a well-known property of non-degenerate processes that the set of sample paths that hit the boundary of O and do not enter the set is negligible. Hence, it can be observed that this property in conjunction with almost-sure continuity of $(t, x) \mapsto X_s^{t,x;\boldsymbol{u}}$ uniformly with respect to s leads to almost-sure continuity of the exit

time function $(t, x) \mapsto \bar{\tau}(t, x)$. Furthermore, according to the continuity of the exit time map and the process trajectories, one can infer that the lower semicontinuity of payoff function J in (6) is inherited to the value function V in (3a). Note also that the function J is bounded since ℓ is. ∎

Remark 4.3: As a consequence of Proposition 4.2, given $(t, x, \boldsymbol{u}) \in \mathbb{S} \times \mathcal{U}$ the function

$$\Omega \ni \omega \mapsto J\big(\theta(\omega), X_{\bar{\tau}(\omega)}^{t,x;\boldsymbol{u}}(\omega); \boldsymbol{u}\big) \in \mathbb{R}$$

is $\mathbb{P}$-measurable. ◁

Proposition 4.4 (Strong Markov Property): Consider the system (1) satisfying Assumption 2.1. Then, for a stopping time $\theta \in \mathcal{T}_{[t,T]}$ and an admissible control $\boldsymbol{u} = \mathbb{1}_{[t,\theta]}\boldsymbol{u}_1 + \mathbb{1}_{(\theta,T]}\boldsymbol{u}_2$, where $\boldsymbol{u}_1 \in \mathcal{U}_t$ and $\boldsymbol{u}_2 \in \mathcal{U}_\theta$,

$$\mathbb{E}\Big[\ell\big(X_{\bar{\tau}(t,x)}^{t,x;\boldsymbol{u}}\big) \,\Big|\, \mathcal{F}_\theta\Big] = \mathbb{1}_{\{\bar{\tau}(t,x)<\theta\}}\ell\big(X_{\bar{\tau}(t,x)}^{t,x;\boldsymbol{u}_1}\big)$$
$$+ \mathbb{1}_{\{\bar{\tau}(t,x)\geq\theta\}}J\big(\theta, X_\theta^{t,x;\boldsymbol{u}_1}, \boldsymbol{u}_2\big) \quad \mathbb{P}\text{-a.s.}$$

Proof: By Definition 2.2, one has

$$\mathbb{1}_{\{\theta\leq\bar{\tau}(t,x)\}}\bar{\tau}(t,x) = \mathbb{1}_{\{\theta\leq\bar{\tau}(t,x)\}}\big(\bar{\tau}(\theta, X_\theta^{t,x;\boldsymbol{u}})+\theta-t\big) \quad \mathbb{P}\text{-a.s.}$$

Hereafter one can follow effectively the same computations as in the proof of [9, Proposition 5.1] to arrive at the assertion. ∎

B. Dynamic Programming Principle

The following Theorem provides a dynamic programming principle (DPP) for the exit time problem introduced in (5).

Theorem 4.5 (Exit Time Problem DPP): Consider the system (1), and suppose that Assumptions 2.1 and 4.1 hold. Then for every $(t, x) \in \mathbb{S}$ and for all stopping times $\theta \in \mathcal{T}_{[t,T]}$,

$$V(t, x) \leq \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\mathbb{1}_{\{\bar{\tau}(t,x)<\theta\}}\ell\big(X_{\bar{\tau}(t,x)}^{t,x;\boldsymbol{u}}\big)$$
$$+ \mathbb{1}_{\{\bar{\tau}(t,x)\geq\theta\}}V^*\big(\theta, X_\theta^{t,x;\boldsymbol{u}}\big)\Big], \quad (7)$$

and

$$V(t, x) \geq \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\mathbb{1}_{\{\bar{\tau}(t,x)<\theta\}}\ell\big(X_{\bar{\tau}(t,x)}^{t,x;\boldsymbol{u}}\big)$$
$$+ \mathbb{1}_{\{\bar{\tau}(t,x)\geq\theta\}}V_*\big(\theta, X_\theta^{t,x;\boldsymbol{u}}\big)\Big], \quad (8)$$

where V is the value function defined in (5).

Proof: The proof follows techniques developed in [9] in conjunction with the strong Markov property in Proposition 4.4. To wit, one first constructs a control strategy for a small enough open cylinder around an operating point, assembles an appropriate covering for the set $\mathbb{S}$ with such open cylinders, and uses this covering to arrive at a control strategy which satisfies the required conditions within ϵ precision, $\epsilon > 0$ being pre-assigned and arbitrary. For the detailed proof we refer to [30, Proposition 4.7]. ∎

Remark 4.6: The dynamic programming principles in (7) and (8) are introduced in a weaker sense than the standard DPP for stochastic optimal control problems [21]. To wit, note that one does not have to verify the measurability of the value function V defined in (5) to apply our DPP. ◁

C. Dynamic Programming Equation

Our objective in this subsection is to demonstrate how the DPP derived in §IV-B characterizes the value function V as a (discontinuous) viscosity solution to an appropriate HJB equation. For the general theory of viscosity solutions we refer to [14] and [21].

Definition 4.7 (Dynkin Operator): Given $u \in \mathbb{U}$, we denote by $\mathcal{L}^u$ the Dynkin operator/infinitesimal generator associated to the controlled diffusion (1) as

$$\mathcal{L}^u \Phi(t,x) := \partial_t \Phi(t,x) + \langle f(x,u), \partial_x \Phi(t,x) \rangle$$
$$+ \frac{1}{2}\text{Tr}[\sigma\sigma^\top(x,u)\partial_x^2\Phi(t,x)],$$

where Φ is a real-valued function smooth on the interior of $\mathbb{S}$, with $\partial_t\Phi$ and $\partial_x\Phi$ denoting the partial derivatives with respect to t and x respectively, and $\partial_x^2\Phi$ denoting the Hessian matrix with respect to x.

Theorem 4.8 (Exit Time DPE): Consider the system (1), and suppose that Assumptions 2.1 and 4.1 hold. Then:

○ the lower semicontinuous function of V introduced in (5) is a viscosity supersolution of

$$-\sup_{u \in \mathbb{U}} \mathcal{L}^u V_*(t,x) \geq 0 \qquad on \quad [0,T] \times \overline{O}^c,$$

○ the upper semicontinuous function of V is a viscosity subsolution of

$$-\sup_{u \in \mathbb{U}} \mathcal{L}^u V^*(t,x) \leq 0 \qquad on \quad [0,T] \times \overline{O}^c,$$

both with boundary conditions
$$\begin{cases} V(t,x) = \ell(x) & \forall(t,x) \in [0,T] \times \overline{O} \quad \textit{(Lateral)}, \\ V(T,x) = \ell(x) & \forall x \in \mathbb{R}^n \quad \textit{(Terminal)}. \end{cases}$$

Proof: we briefly sketch the proof and refer to [30, Proposition 4.10] for more details. In view of the definition of the stopping time $\bar{\tau}$ and based on the assumption 4.1.a. and 4.1.b., one sees that $\bar{\tau}(t,x) = t$ for all (t,x) on the lateral and terminal boundaries which immediately arrives at the boundary assertion. In the *supersolution* part, for the sake of contradiction, one can assume that there exists (t_0,x_0), and a smooth function ϕ dominated by the value function V_* where $(V_* - \phi)(t_0,x_0) = 0$, such that for some $\delta > 0$, $-\sup_{u \in \mathbb{U}} \mathcal{L}^u\phi(t_0,x_0) < -2\delta$. Since ϕ is smooth, the map $(t,x) \mapsto \mathcal{L}^u\phi(t,x)$ is continuous. Therefore, there exist $u \in \mathbb{U}$ and $r > 0$ such that $\text{B}_r(t_0,x_0) \subset [0,T] \times \overline{O}^c$ and $-\mathcal{L}^u\phi(t,x) < -\delta$ for all (t,x) in $\text{B}_r(t_0,x_0)$. Let us define the stopping time $\theta(t,x) \in \mathcal{T}_{[t,T]}$ as the first exit time of trajectory $X^{t,x;u}$ from the ball $\text{B}_r(t_0,x_0)$. Note that by continuity of solutions to (1), $t < \theta(t,x) < T$ $\mathbb{P}$- a.s. for all $(t,x) \in \text{B}_r(t_0,x_0)$. Therefore, selecting $r > 0$ sufficiently small so that $\theta(t,x) < \tau_O$, and applying Itô's formula, we see that for all $(t,x) \in \text{B}_r(t_0,x_0)$, $\phi(t,x) < \mathbb{E}\left[\phi\left(\theta(t,x), X_{\theta(t,x)}^{t,x;u}\right)\right]$. Now it suffices to take a sequence $(t_n,x_n,V(t_n,x_n))_{n \in \mathbb{N}}$ converging to $(t_0,x_0,V_*(t_0,x_0))$. For sufficiently large n we have $V(t_n,x_n) < \mathbb{E}\left[V_*\left(\theta(t_n,x_n), X_{\theta(t_n,x_n)}^{t_n,x_n;u}\right)\right]$ which, in view of the fact that $\theta(t_n,x_n) < \tau_O \wedge T$, contradicts the DPP in (8). The *subsolution* property is proved in a similar fashion. ∎

V. A CONNECTION BETWEEN THE REACH-AVOID PROBLEM AND PDE CHARACTERIZATION

In this section we draw a connection between the reach-avoid problem of §II and the stochastic optimal control problems stated in §III. To this end, note that on the one hand, an assumption on the sets A and B in the reach-avoid problem (Definition 2.3) within the time interval $[0,T]$ is that they are closed. On the other hand, our solution to the stochastic optimal control problem (defined in §III and solved in §IV) relies on lower semicontinuity of the payoff function ℓ in (5), see Assumption 4.1.c.

To achieve a reconciliation between the two sets of hypotheses, given sets A and B satisfying Assumption 3.1, we construct a smaller measurable set $A_\epsilon \subset A^\circ$ such that $A_\epsilon := \{x \in A \mid \text{dist}(x, A^c) \geq \epsilon\}$ and A_ϵ satisfies Assumption 4.1.b. Note that this is always possible if $O := A \cup B$ satisfies Assumption 4.1.b.—indeed, simply take $\epsilon < h/2$ to see this, where h is as defined in Assumption 4.1.b. Figure 3 depicts this case.

Analytically, we define

$$V_\epsilon(t,x) := \sup_{\boldsymbol{u} \in \mathcal{U}_t} \mathbb{E}\left[\mathbb{1}_{A^\circ}\left(X_{\tau_\epsilon}^{t,x;\boldsymbol{u}}\right)\right], \qquad \tau_\epsilon := \tau_{A_\epsilon \cup B} \wedge T. \tag{9}$$

In the following Theorem, we show that the above technique affords an ϵ-conservative but precise way of characterizing the solution to the reach-avoid problem defined in Definition 2.3 in the framework of §IV.

Theorem 5.1: Consider the system (1), and suppose that Assumptions 2.1, 3.1, 4.1.a. and 4.1.b. hold. Then, for all (t,x) in $\mathbb{S}$ and $\epsilon_1 \geq \epsilon_2 > 0$, we have $V_{\epsilon_2}(t,x) \geq V_{\epsilon_1}(t,x)$, and $V(t,x) = \lim_{\epsilon \downarrow 0} V_\epsilon(t,x)$ where the functions V and V_ϵ are defined as (3a) and (9) respectively.

Proof: Here we briefly sketch the proof and refer to [30, Theorem 5.2] for more details. Since the family of the sets $(A_\epsilon)_{\epsilon > 0}$ is nested and increasing, it is straightforward to deduce that the family of functions $(V_\epsilon)_{\epsilon > 0}$ is also increasing pointwise as $\epsilon \downarrow 0$. Moreover, one can observe that the difference between the value functions V and $\lim_{\epsilon \downarrow 0} V_\epsilon$ is indeed the probability measure of the trajectories that either hit A at terminal time T or hit the boundary of the set A and do not enter A. Both of these events are in fact negligible duo to the Assumptions 4.1.b. and 4.1.c. ∎

Observe also that for the problem of reachability at the time T, defined in Definition 2.4, the above procedure is unnecessary if the set A is open, see the required conditions for Proposition 3.3.

VI. NUMERICAL EXAMPLE: ZERMELO NAVIGATION PROBLEM

To illustrate the theoretical results of the preceding sections, we apply the proposed reach-avoid formulation to the *Zermelo* navigation problem with constraints and stochastic uncertainties. In control theory, the Zermelo navigation

problem consists of a swimmer who aims to reach an island (Target) in the middle of a river while avoiding the waterfall, with the river current leading towards the waterfall. The situation is depicted in Figure 4. We say that the swimmer "succeeds" if he reaches the target before going over the waterfall, the latter forming a part of his Avoid set.

A. Mathematical modeling

The dynamics of the river current are nonlinear; we let $f(x, y)$ denote the river current at position (x, y) [12]. We assume that the current flows with constant direction towards the waterfall, with the magnitude of f decreasing with increasing distance from the middle of the river:

$$f(x,y) := \begin{bmatrix} 1 - ay^2 \\ 0 \end{bmatrix}.$$

This model may not describe the behavior of a realistic river current, so we consider some uncertainties in the river current modeled by a diffusion term as

$$\sigma(x,y) := \begin{bmatrix} \sigma_x & 0 \\ 0 & \sigma_y \end{bmatrix}.$$

We assume that the swimmer moves with constant velocity V_S, and we assume that he can change his direction α instantaneously. The complete dynamics of the swimmer in the river is given by

$$\begin{bmatrix} \mathrm{d}x_s \\ \mathrm{d}y_s \end{bmatrix} = \begin{bmatrix} 1 - ay^2 + V_S\cos(\alpha) \\ V_S\sin(\alpha) \end{bmatrix} \mathrm{d}s + \begin{bmatrix} \sigma_x & 0 \\ 0 & \sigma_y \end{bmatrix} \mathrm{d}W_s,$$
(10)

where W_s is a two-dimensional Brownian motion, and $\alpha \in\ ]-\pi, \pi]$ is the direction of the swimmer with respect to the x axis and plays the role of the controller for the swimmer.

B. Reach-Avoid formulation

Obviously, the probability of the swimmer's "success" starting from some initial position in the navigation region depends on starting point (x, y). As shown in §III, this probability can be characterized as the level set of a value function, and by Theorem 4.8 this value function is the discontinuous viscosity solution of a certain differential equation on the navigation region with particular lateral and terminal boundary conditions. The differential operator $\mathcal{L}$ in Theorem 4.8 can be analytically calculated in this case as follows:

$$\sup_{u\in\mathbb{U}}\mathcal{L}^u\Phi(t,x,y) = \sup_{\alpha\in[-\pi,\pi]} \Big(\partial_t\Phi(t,x,y)$$
$$+ \big(1 - ay^2 + V_S\cos(\alpha)\big)\partial_x\Phi(t,x,y)$$
$$+ V_S\sin(\alpha)\partial_y\Phi(t,x,y)$$
$$+ \frac{1}{2}\sigma_x^2\partial_x^2\Phi(t,x,y) + \frac{1}{2}\sigma_y^2\partial_y^2\Phi(t,x,y)\Big).$$

It can be shown that the controller value maximizing the above Dynkin operator is

$$\alpha^*(t,x,y) := \arg\max_{\alpha\in[-\pi,\pi]} \Big(\cos(\alpha)\partial_x\Phi(t,x,y)$$
$$+ \sin(\alpha)\partial_y\Phi(t,x,y)\Big)$$
$$= \arctan\Big(\frac{\partial_y\Phi}{\partial_x\Phi}\Big)(t,x,y).$$

Therefore, the differential operator can be simplified to

$$\sup_{u\in\mathbb{U}}\mathcal{L}^u\Phi(t,x,y) = \partial_t\Phi(t,x,y) + (1 - ay^2)\partial_x\Phi(t,x,y)$$
$$+ \frac{1}{2}\sigma_x^2\partial_x^2\Phi(t,x,y) + \frac{1}{2}\sigma_y^2\partial_y^2\Phi(t,x,y)$$
$$+ V_S\|\nabla\Phi(t,x,y)\|,$$

where $\nabla\Phi(t,x,y) := \begin{bmatrix} \partial_x\Phi(t,x,y) & \partial_y\Phi(t,x,y) \end{bmatrix}$.

C. Simulation results

For the following numerical simulations we fix the diffusion coefficients $\sigma_x = 0.5$ and $\sigma_y = 0.2$. We investigate three different scenarios: First, we assume that the river current is uniform, i.e., $a = 0\ \mathrm{m^{-1}s^{-1}}$ in (10). Moreover, we consider the case that the swimmer velocity is less than the current flow, e.g., $V_S = 0.6\ \mathrm{ms^{-1}}$. Based on the above calculations, Figure 5(a) depicts the value function which is the numerical solution of the differential operator equation in Theorem 4.8 with the corresponding terminal and lateral conditions. As expected, since the swimmer's speed is less than the river current, if he starts from the beyond the target he has less chance of reach the island. This scenario is also captured by the value function shown in Figure 5(a).

(a) (b) (c)
The first scenario: the swimmer's speed is slower than the river current momentum, assumed uniform.

The third scenario: the swimmer's speed is slower than the river current.

Fig. 5. The three different scenarios considered in the Zermelo navigation problem.

Second, we assume that the river current is non-uniform and decreases with respect to the distance from the middle of

(a) (b)
The The
first third
sce- sce-
nario: nario:
the the
swim- swim-
mer's mer's
speed speed
is with
slower slower
than than
the the
river max-
cur- imum
rent. thrust
ing river
as- cur-
sumed rent.
uni-
form.

Fig. 6. The level sets of the value functions for the different scenarios considered in the Zermelo navigation problem.

the river. This means that the swimmer, even in the case that his speed is less than the current, has a non-zero probability of success if he initially swims to the sides of the river partially against its direction, followed by swimming in the direction of the current to reaches the target. This scenario is depicted in Figure 5(b), where a non-uniform river current $a = 0.04 \text{ m}^{-1}\text{s}^{-1}$ in (10) is considered.

Third, we consider the case that the swimmer can swim faster than river current. In this case we expect the swimmer to succeed with some probability even if he starts from beyond the target. This scenario is captured in Figure 5(c), where the reachable set (of course in probabilistic fashion) covers the entire navigation region of the river except the region near the waterfall.

In the following we show the level sets of the aforementioned value functions for $p = 0.9$. To wit, as defined in §III (and in particular in Proposition 3.2), these level sets, roughly speaking, correspond to the reachable sets with probability $p = 90\%$ in certain time horizons while the swimmer is avoiding the waterfall. By definition, as shown by the following figures, these sets are nested with respect to the time horizon.

All simulations were obtained using the Level Set Method Toolbox [29] (version 1.1), with a grid 101×101 in the region of simulation.

VII. Concluding Remarks and Future Direction

In this article we presented a new method to address a class of stochastic reachability problems with state constraints. The proposed framework provides a set characterization of the stochastic reach-avoid set based on discontinuous viscosity solutions of a second order PDE. In contrast to earlier approaches, this methodology is not restricted to almost-sure notions and one can compute the desired set with any non-zero probability. In future works, we aim to extend our framework to stochastic differential games.

References

[1] J. Aubin, *Viability Theory*, Systems & Control: Foundations & Applications, Birkhäuser Boston Inc., Boston, MA, 1991.

[2] J. Aubin, J. Lygeros, M. Quincampoix, S. Sastry, and N. Seube, *Impulse differential inclusions: a viability approach to hybrid systems*, Institute of Electrical and Electronics Engineers. Transactions on Automatic Control, 47 (2002), pp. 2–20.

[3] J. Aubin and G. D. Prato, *The viability theorem for stochastic differential inclusions*, Stochastic Analysis and Applications, 16 (1998), pp. 1–15.

[4] J. Aubin, G. D. Prato, and H. Frankowska, *Stochastic invariance for differential inclusions*, Set-Valued Analysis. An International Journal Devoted to the Theory of Multifunctions and its Applications, 8 (2000), pp. 181–201.

[5] M. Bardi and P. Goatin, *Invariant sets for controlled degenerate diffusions: a viscosity solutions approach*, in Stochastic analysis, control, optimization and applications, Systems Control Found. Appl., Birkhäuser Boston, Boston, MA, 1999, pp. 191–208.

[6] M. Bardi and R. Jensen, *A geometric characterization of viable sets for controlled degenerate diffusions*, Set-Valued Analysis, 10 (2002), pp. 129–141. Calculus of variations, nonsmooth analysis and related topics.

[7] O. Bokanowski, N. Forcadel, and H. Zidani, *Reachability and minimal times for state constrained nonlinear problems without any controllability assumption*, SIAM J. Control and Optimization, 48(7) (2010), pp. 4292–4316.

[8] V. S. Borkar, *Controlled diffusion processes*, Probability Surveys, 2 (2005), pp. 213–244 (electronic).

[9] B. Bouchard and N. Touzi, *Weak dynamic programming principle for viscosity solutions*, to appear in SIAM Journal on Control and Optimization, (2009).

[10] R. Buckdahn, S. Peng, M. Quincampoix, and C. Rainere, *Existence of stochastic control under state constraints*, Comptes Rendus de l'Académie des Sciences. Série I. Mathématique, 327 (1998), pp. 17–22.

[11] P. Cardaliaguet, *A differential game with two players and one target*, SIAM Journal on Control and Optimization, 34 (1996), pp. 1441–1460.

[12] P. Cardaliaguet, M. Quincampoix, and P. Saint-Pierre, *Optimal times for constrained nonlinear control problems without local controllability*, Applied Mathematics and Optimization, 36 (1997), pp. 21–42.

[13] ———, *Differential Games with State-Constraints*, in ISDG2002, Vol. I, II (St. Petersburg), St. Petersburg State Univ. Inst. Chem., St. Petersburg, 2002, pp. 179–182.

[14] M. G. Crandall, H. Ishii, and P. L. Lions, *User's guide to viscosity solutions of second order partial differential equations*, American Mathematical Society, 27 (1992), pp. 1–67.

[15] G. Da Prato and H. Frankowska, *Stochastic viability for compact sets in terms of the distance function*, Dynamic Systems and Applications, 10 (2001), pp. 177–184.

[16] ———, *Invariance of stochastic control systems with deterministic arguments*, Journal of Differential Equations, 200 (2004), pp. 18–52.

[17] C. Dellacherie and P.-A. Meyer, *Probabilities and Potential*, vol. 29 of North-Holland Mathematics Studies, North-Holland Publishing Co., Amsterdam, 1978.

[18] P. M. Esfahani, M. Vrakopoulou, K. Margellos, J. Lygeros, and G. Andersson, *Cyber attack in a two-area power system: Impact identification using reachability*, in American Control Conference, 2010, pp. 962–967.

[19] S. Ethier and T. Kurtz, *Markov Processes: Characterization and Convergence*, Wiley Series in Probability and Mathematical Statistics, John Wiley & Sons, Ltd., New York, 1986.

[20] I. J. Fialho and T. Georgiou, *Worst case analysis of nonlinear systems*, IEEE transactions on Automatic Control, 44(6) (1999), pp. 4292–4316.

[21] W. Fleming and H. Soner, *Controlled Markov Processes and Viscosity Solution*, Springer-Verlag, 3 ed., 2006.

[22] Y. Gao, J. Lygeros, and M. Quincampoix, *The Reachability Problem for Uncertain Hybrid Systems Revisited: a Viability Theory Perspective*, in Hybrid systems: computation and control, vol. 3927 of Lecture Notes in Comput. Sci., Springer, Berlin, 2006, pp. 242–256.

[23] I. KARATZAS AND S. SHREVE, *Brownian Motion and Stochastic Calculus*, vol. 113 of Graduate Texts in Mathematics, Springer-Verlag, New York, 2 ed., 1991.

[24] N. KRYLOV, *Controlled Diffusion Processes*, vol. 14 of Stochastic Modelling and Applied Probability, Springer-Verlag, Berlin Heidelberg, 2009. Reprint of the 1980 Edition.

[25] J. LYGEROS, *On reachability and minimum cost optimal control*, Automatica. A Journal of IFAC, the International Federation of Automatic Control, 40 (2004), pp. 917–927 (2005).

[26] J. LYGEROS, C. TOMLIN, AND S. SASTRY, *A game theorretic approach to controller design for hybrid systems*, Proceedings of IEEE, 88 (2000), pp. 949–969.

[27] K. MARGELLOS AND J. LYGEROS, *Hamilton-Jacobi formulation for reach-avoid differential games*. http://arxiv.org/abs/0911.4625, 2009.

[28] I. MITCHELL, A. BAYEN, AND C. J. TOMLIN, *Validating a Hamilton-Jacobi approximation to hybrid system reachable sets*, in Hybrid systems: computation and control, vol. 2034 of Lecture Notes in Comput. Sci., Springer, New York, 2001, pp. 418–432.

[29] I. MITCHELL AND C. J. TOMLIN, *Level set methods for computation in hybrid systems*, in Hybrid systems: computation and control, vol. 1790 of Lecture Notes in Comput. Sci., Springer-Verlag, New York, 2002, pp. 310–323.

[30] P. MOHAJERIN, D. CHATTERJEE, AND J. LYGEROS, *On A Stochastic Reach-Avoid Problem and Set Characterization*. http://control.ee.ethz.ch/index.cgi?page=publications;action=details;id=3772, Mar. 2011.

[31] S. OSHER AND J. SETHIAN, *Fronts propagating with curvature-dependent speed: algorithms based on hamilton-jacobi formulations*, Journal of Computational Physics, 79 (1988), pp. 12–49.

[32] J. SETHIAN, *Level set methods and fast marching methods*, vol. 3 of Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, Cambridge, second ed., 1999. Evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science.

[33] H. SONER AND N. TOUZI, *Dynamic programming for stochastic target problems and geometric flows*, Journal of the European Mathematical Society (JEMS), 4 (2002), pp. 201–236.