Automatica 70 (2016) 43–56

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

The stochastic reach-avoid problem and set characterization for diffusions ☆

Peyman Mohajerin Esfahani [a], Debasish Chatterjee [b], John Lygeros [a]

[a] Automatic Control Laboratory, Physikstrasse 3, ETH Zürich, 8092 Zürich, Switzerland

[b] Systems and Control Engineering, Indian Institute of Technology Bombay, Powai, Mumbai 400076, India

ARTICLE INFO

Article history:
Received 16 July 2014
Received in revised form
21 January 2016
Accepted 2 March 2016
Available online 13 April 2016

Keywords:
Reachability
Stochastic control
Dynamic programming
Partial differential equations
Discontinuous viscosity solutions

ABSTRACT

In this article we approach a class of stochastic reachability problems with state constraints from an optimal control perspective. Preceding approaches to solving these reachability problems are either confined to the deterministic setting or address almost-sure stochastic requirements. In contrast, we propose a methodology to tackle problems with less stringent requirements than almost sure. To this end, we first establish a connection between two distinct stochastic reach-avoid problems and three classes of stochastic optimal control problems involving discontinuous payoff functions. Subsequently, we focus on solutions of one of the classes of stochastic optimal control problems—the exit-time problem, which solves both the two reach-avoid problems mentioned above. We then derive a weak version of a dynamic programming principle (DPP) for the corresponding value function; in this direction our contribution compared to the existing literature is to develop techniques that admit discontinuous payoff functions. Moreover, based on our DPP, we provide an alternative characterization of the value function as a solution of a partial differential equation (PDE) in the sense of discontinuous viscosity solutions, along with boundary conditions both in Dirichlet and viscosity senses. Theoretical justifications are also discussed to pave the way for deployment of off-the-shelf PDE solvers for numerical computations. Finally, we validate the performance of the proposed framework on the stochastic Zermelo navigation problem.

1. 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 deals with the determination of the set of initial states for which one can find at least one control strategy to steer the system to a target set while avoiding certain obstacles. This problem finds applications in, for example, air traffic management (Lygeros, Tomlin, & Sastry, 2000) and security of power networks (Mohajerin Esfahani, Vrakopoulou, Margellos, Lygeros, & Andersson, 2011).

The set representing the solution of this problem is known as a capture basin (Aubin, 1991). A direct approach to compute the capture basin is formulated in the language of viability theory in Cardaliaguet (1996) and Cardaliaguet, Quincampoix, and Saint-Pierre (2002). An alternative and indirect approach to reachability problems proceeds via level set methods defined by value functions that are solutions of appropriate optimal control problems. Employing dynamic programming techniques for reachability and viability problems, one can in turn characterize these value functions by solutions of the standard Hamilton–Jacobi–Bellman (HJB) equations corresponding to these optimal control problems (Lygeros, 2004). The focus of this article is on the stochastic counterpart of this problem.

1.1. The literature in the stochastic setting

In the literature, probabilistic analogs of reachability problems have mainly been studied from an almost-sure perspective. For example, stochastic viability and controlled invariance are treated

☆ The material in this paper was partially presented at the 50th IEEE Conference on Decision and Control and European Control Conference, December 12–15, 2011, Orlando, FL, USA. This paper was recommended for publication in revised form by Associate Editor Kok Lay Teo under the direction of Editor Ian R. Petersen.

E-mail addresses: mohajerin@control.ee.ethz.ch (P. Mohajerin Esfahani), dchatter@iitb.ac.in (D. Chatterjee), lygeros@control.ee.ethz.ch (J. Lygeros).

http://dx.doi.org/10.1016/j.automatica.2016.03.016
0005-1098/

in Aubin and Da Prato (1990), Aubin, Da Prato, and Frankowska (2000) and Bardi and Jensen (2002). Methods involving stochastic contingent sets (Aubin & Da Prato, 1998; Aubin et al., 2000), viscosity solutions of second-order partial differential equations (Bardi & Goatin, 1999; Bardi & Jensen, 2002; Buckdahn, Peng, Quincampoix, & Rainere, 1998), derivatives of the distance function (Da Prato & Frankowska, 2001), and equivalence relation to certain deterministic control systems (Da Prato & Frankowska, 2004) were all developed in this context.

Geared towards similar almost-sure reachability objective, the article (Soner & Touzi, 2002a) introduced a new class of the so-called stochastic target problems, and characterized the solution via a dynamic programming approach. The differential properties of the almost-sure reachable set were also studied based on the geometrical partial differential equation which is the analogue of the HJB equation (Soner & Touzi, 2002b) in that setting.

Although almost sure versions of reachability specifications are interesting in their own right, they may be a too strict concept in some applications, particularly when a common specification is only to control the probability that undesirable events take place. In this regard, the authors of Bouchard, Elie, and Touzi (2009–2010) recently extended the stochastic target framework of Soner and Touzi (2002a) to allow for unbounded control set, which together with the martingale representation theory, addresses the aforementioned almost-sure limitation in an augmented state space; see also the recent book (Touzi, 2013). This article approaches the same question, but indirectly and from an optimal control perspective.

1.2. Our methodology and contributions

The stochastic "reach-avoid" problems studied in this article are as follows:

RA: *Given an initial state $x \in \mathbb{R}^n$, a horizon $T > 0$, a number $p \in [0, 1]$, and two disjoint sets $A, B \subset \mathbb{R}^n$, determine whether there exists a control policy such that the process reaches A prior to entering B within the interval $[0, T]$ with probability at least p.*

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 of 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 addition, we address the following slightly different reach-avoid problem compared to RA above, that requires the process to be in the set A at time T:

$\widetilde{\text{RA}}$: *Given an initial state $x \in \mathbb{R}^n$, a horizon $T > 0$, a number $p \in [0, 1]$, and two disjoint sets $A, B \subset \mathbb{R}^n$, determine whether there exists a policy such that with probability at least p the controlled process resides in A at time T while avoiding B on the interval $[0, T]$.*

Our methodology and contributions toward the above problems are summarized below:

(i) We establish a link from the problems RA and $\widetilde{\text{RA}}$ to three different classes of stochastic optimal control problems involving *discontinuous* payoff functions in Section 3;
(ii) focusing on the class of exit-time problems that addressed both the reach-avoid problems alluded above, we propose a *weak* dynamic programming principle (DPP) leading to a (*discontinuous*) PDE characterization along with appropriate boundary conditions;
(iii) finally, in Section 5 we provide theoretical justification that pave the analytical ground to deploy existing (*continuous*) off-the-shelf PDE solvers for our numerical purposes.

More specifically, we first show that the desired set of initial conditions for the reach-avoid problems RA and $\widetilde{\text{RA}}$ can be translated as super level sets of particular functions described in the context of stochastic optimal control problems (Propositions 3.3 and 3.4). Different classes of optimal control problems are suggested for each of the two reach-avoid problems, and it turns out that the class of exit-time problems with discontinuous payoff functions can adequately address both the reach-avoid problems. This connection is relatively straightforward and does not require any assumption on the underlying dynamics. We, however, are not aware of any results in the literature reflecting this connection.

The exit-time problem with a continuous payoff function is a classical stochastic optimal control problem whose alternative PDE characterizations have been established in the literature; see for instance (Fleming & Soner, 2006, Section IV.7). However, these results are not directly applicable to our reach-avoid problems due to the discontinuity of the payoff function. We address this technical issue by developing a DPP in a weak sense in the spirit of Bouchard and Touzi (2011) (Theorem 4.4). We emphasize that the results of Bouchard and Touzi (2011) were developed in the framework of fixed time horizon and the optimal stopping time. Neither of these settings is applicable to the exit-time problem. To that end, it turns out that we require some technical continuity properties which are essential for the proposed weak DPP (Proposition 4.2) as well as the respective boundary conditions (Proposition 4.8). To the best of our knowledge, these continuity results are also new in the literature. It is also worth noting that this weak formulation avoids delicate issues related to a measurable selection in the context of optimal control problems.

Based on the proposed DPP, we characterize the value function as the (discontinuous) viscosity solution of a PDE (Theorem 4.7) along with boundary conditions in both viscosity and Dirichlet (pointwise) senses (Theorem 4.9). We remark that due to the discontinuity of the payoff function, the viscosity boundary conditions involves a non-trivial regularity condition which is a stronger version of the requirement for the proposed DPP (see Proposition 4.8). These technical details are required to rigorously settle the PDE characterization for a stochastic exit-problem problem and we cannot find them elsewhere in the existing literature.

Finally, we provide theoretical justifications (Theorem 5.1) so that the Reach-Avoid problem is amenable to numerical solutions by means of off-the-shelf PDE solvers, which have been mainly developed for continuous solutions. Preliminary results of this study were reported in Mohajerin Esfahani, Chatterjee, and Lygeros (2011) without covering the technical details and mathematical proofs.

Organization of the article: In Section 2 we formally introduce the stochastic reach-avoid problems RA and $\widetilde{\text{RA}}$ above. In Section 3 we characterize the set of initial conditions that solve the reach-avoid problems in terms of super level sets of three different value functions. Focusing on the class of exit-time problems, in Section 4 we establish a DPP and characterize it as the solution of a PDE along with some boundary conditions. Finally, Section 5 presents results connecting those in Sections 3 and 4 and justifies the deployment of the existing PDE solvers for numerical purposes. To illustrate the performance of our technique, the theoretical results developed in preceding sections are applied to solve the stochastic Zermelo navigation problem in Section 6. We conclude with some remarks and directions for future work in Section 7. For better readability, some of the technical proofs are given in appendices.

Notation. Given $a, b \in \mathbb{R}$, we define $a \wedge b := \min\{a, b\}$ and $a \vee b := \max\{a, b\}$. We denote by A^c (resp. A°) the complement