Computing Probabilistic Controlled Invariant Sets

Yulong Gao, Karl Henrik Johansson, *Fellow, IEEE*, and Lihua Xie, *Fellow, IEEE*

***Abstract*—This article investigates stochastic invariance for control systems through probabilistic controlled invariant sets (PCISs). As a natural complement to robust controlled invariant sets (RCISs), we propose finite-, and infinite-horizon PCISs, and explore their relation to RICSs. We design iterative algorithms to compute the PCIS within a given set. For systems with discrete spaces, the computations of the finite-, and infinite-horizon PCISs at each iteration are based on linear programming, and mixed integer linear programming, respectively. The algorithms are computationally tractable, and terminate in a finite number of steps. For systems with continuous spaces, we show how to discretize the spaces, and prove the convergence of the approximation when computing the finite-horizon PCISs. In addition, it is shown that an infinite-horizon PCIS can be computed by the stochastic backward reachable set from the RCIS contained in it. These PCIS algorithms are applicable to practical control systems. Simulations are given to illustrate the effectiveness of the theoretical results for motion planning.**

***Index Terms*—Probabilistic controlled invariant set (PCIS), reachability analysis, stochastic control systems.**

I. Introduction

A. Motivation and Related Work

INVARIANCE is a fundamental concept in systems and control [1]–[3]. A controlled invariant set captures the region where the states can be maintained by some admissible control inputs. Robust controlled invariant sets (RCISs) are defined for control systems with bounded external disturbances and address the invariance despite any realization of the disturbances. In the past decades, there have been lots of research results on RCISs and their computations [4]–[6]. This article studies probabilistic controlled invariant sets (PCISs), which is a natural complement to RCISs suitable in many applications. A PCIS is a set within which the controller is able to keep the system state with a certain probability. Such sets not only alleviate the inherent conservatism of RCISs by allowing probabilistic violations but also enlarge the applications of RCISs by being able to address unbounded disturbances. The study of PCISs is motivated by safety-critical control [7], stochastic model predictive control (MPC) [8], [9], reliable control [10], [11], and relevant applications, e.g., air traffic management systems [12], [13] and motion planning [14].

A question at the heart of this article is

Given a set $\mathbb{Q}$ and a parameter $0 \leq \epsilon \leq 1$, how to compute a set $\tilde{\mathbb{Q}} \subseteq \mathbb{Q}$ that is invariant with probability ϵ?

To the best of authors' knowledge, this question has not been explored up to now. One essential component in iterative approaches on computing RCISs is to compute the robust backward reachable set, in which each state can be steered to the current set by an admissible input for all possible uncertainties [4]–[6]. The PCIS computation in this article follows the same idea, but the robust backward reachable set is replaced with the stochastic backward reachable sets which require different mathematical tools. Some challenges related to such an approach are highlighted as follows:

1) how to make it tractable to compute the stochastic backward reachable set, in particular for systems with continuous spaces?
2) how to mitigate the conservatism when characterizing the stochastic backward reachable set subject to the prescribed probability?
3) how to guarantee convergence of the iterations?

Controlled invariant sets have recently been extended to stochastic systems. In [18], a target set, which is similar to the PCIS of this article, is used to define stabilization in probability. In [10], a reliable control set, another similar notion to a PCIS, is used to guarantee the reliability of Markov-jump linear systems. The reliability is further studied for such systems with bounded disturbances in [11]. A definition of PCIS for nonlinear systems is provided in [15] by using reachability analysis. It is later applied to portfolio optimization [19]. Another definition of probabilistic invariance originates from stochastic MPC [16] and captures one-step invariance. In [16], an ellipsoidal approximation is given for linear systems with specific uncertainty structure. Similar invariant sets are used in [20] to construct a convex lifting function for linear stochastic control systems. A definition of a probabilistic invariant set is proposed in [17] and [21] for linear stochastic systems without control inputs.

Manuscript received May 10, 2019; revised May 12, 2019 and May 22, 2020; accepted August 6, 2020. Date of publication August 21, 2020; date of current version June 29, 2021. The work of Yulong Gao and Karl Henrik Johansson was supported in part by the Knut and Alice Wallenberg Foundation, in part by the Swedish Strategic Research Foundation, and in part by the Swedish Research Council. Recommended by Associate Editor Q.-S. Jia. *(Corresponding author: Yulong Gao.)*

Yulong Gao and Karl Henrik Johansson are with the Division of Decision, and Control Systems, KTH Royal Institute of Technology, 10044 Stockholm, Sweden (e-mail: yulongg@kth.se; kallej@kth.se).

Lihua Xie is with the School of Electrical, and Electronic Engineering, Nanyang Technological University, Singapore 639798, Singapore (e-mail: elhxie@ntu.edu.sg).

Color versions of one or more of the figures in this article are available online at https://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TAC.2020.3018438

TABLE I
COMPARISONS BETWEEN THIS ARTICLE AND OTHER WORK

	System	Invariant Set	Control	Horizon	Computation
This paper	Markov controlled process	PCIS	Yes	Finite and infinite horizons	Iteration based on stochastic backward reachable set
[15]	Nonlinear stochastic system	PCIS	Yes	Finite and infinite horizons	No
[16]	Linear stochastic system	PCIS	Yes	One step	Ellipsoidal approximation
[17]	Linear stochastic system	Probabilistic invariant set	No	Infinite horizon	Polyhedral approximation based on Chebyshevs inequality

This definition captures the probabilistic inclusion of the state at each time instant. A recent work [22] explores the correspondence between probabilistic and robust invariant sets for linear systems. In [17] and [21], polyhedral probabilistic invariant sets are approximated by using Chebyshev's inequality for linear systems with Gaussian noise. Recursive satisfaction is usually computationally intractable for general stochastic control systems.

The results of this article is built on the above-mentioned work but make significant additions and improvements. Table I summarizes the comparison between this article and the most relevant literature.

1) All the abovementioned references focus on some specific stochastic systems (e.g., linear or one-dimensional affine nonlinear systems) or on some specific class of stochastic disturbances (e.g., Gaussian or state-independent noise). In our model, we consider general Markov controlled processes, which include general system dynamics and stochastic disturbances.
2) Different from [17], [21], our invariant sets are defined based on trajectory inclusion as in [15] and, particularly, incorporate control inputs constrained by a compact set. An accompanying question is how to find an admissible control input when verifying or computing a PCIS.
3) The PCISs in this article are different from the maximal probabilistic safe sets in [23]. Every trajectory in a PCIS is required by our definition to admit the same probability level, which does not hold for the maximal probabilistic safe set.
4) The stochastic reachability analysis studied in [23] provides an important tool for maximizing the probability of staying in a set. Based on this, we compute a PCIS within a set with a prescribed probability level. This extends the results of [15], [23], and [24].

B. Main Contributions and Organization

The objective of this article is to provide a novel tool to analyze invariance in stochastic control systems. The contributions are summarized as follows.

As the first contribution, we propose two novel definitions of PCIS: N-step ϵ-PCIS and infinite-horizon ϵ-PCIS (see Definitions 3 and 4). An N-step ϵ-PCIS is a set within which the state can stay for N steps with probability ϵ under some admissible controller while an infinite-horizon ϵ-PCIS is a set within which the state can stay forever with probability ϵ under some admissible controller. These invariant sets are different from the ones proposed in [16] and [17], which address probabilistic set invariance at each time step. Our definitions are applicable for general discrete-time stochastic control systems. We provide fundamental properties of PCISs and explore their relation to RCISs. Furthermore, we propose conditions for the existence of infinite-horizon ϵ-PCIS (see Theorem 3).

The second contribution is that we design iterative algorithms to compute the largest finite- and infinite-horizon PCIS within a given set for systems with discrete and continuous spaces. The PCIS computation is based on the stochastic backward reachable set. For discrete state and control spaces, it is shown that at each iteration, the stochastic backward reachable set computation of an N-step ϵ-PCIS can be reformulated as a linear program (LP) (see Theorem 1 and Corollary 1) and an infinite-horizon ϵ-PCIS as a computationally tractable mixed-integer linear program (MILP) (see Theorem 4). Furthermore, we prove that these algorithms terminate in a finite number of steps. For continuous state and control spaces, we present a discretization procedure. Under weaker assumptions than [25], we prove the convergence of such approximations for N-step ϵ-PCISs (see Theorem 2). The approximations generalize the case in [23], which only discretizes the state space for a given discrete control space. Furthermore, in order to compute an infinite-horizon ϵ-PCIS, we propose an algorithm based on that an infinite-horizon PCIS always contains an RCIS.

The remainder of this article is organized as follows. Section II provides the system model and some preliminaries. Section III presents the definition, properties, and computation algorithms of finite-horizon PCISs. Section IV extends the results to the infinite-horizon case. Examples in Section V illustrate the effectiveness of our approach. Section VI concludes this article.

Notation: Let $\mathbb{N}$ denote the set of nonnegative integers and $\mathbb{R}$ the set of real numbers. For some $q, s \in \mathbb{N}$ and $q < s$, let $\mathbb{N}_{\geq q}$ and $\mathbb{N}_{[q,s]}$ denote the sets $\{r \in \mathbb{N} \mid r \geq q\}$ and $\{r \in \mathbb{N} \mid q \leq r \leq s\}$, respectively. For two sets $\mathbb{X}$ and $\mathbb{Y}$, $\mathbb{X} \setminus \mathbb{Y} = \{x \mid x \in \mathbb{X}, x \notin \mathbb{Y}\}$ and $\mathbb{X} \triangle \mathbb{Y} = (\mathbb{X} \setminus \mathbb{Y}) \cup (\mathbb{Y} \setminus \mathbb{X})$. When $\leq$, $\geq$, $<$, and $>$ are applied to vectors, they are interpreted element-wise. $\Pr$ denotes the probability. For a set $\mathbb{X}$, $\mathcal{B}(\mathbb{X})$, and $\mathcal{P}(\mathbb{X})$ denote the Boreal σ-algebra generated by $\mathbb{X}$ and the space of probability distributions on $\mathbb{X}$, respectively. The indicator function of a set $\mathbb{X}$ is denoted by $\mathbb{1}_{\mathbb{X}}(x)$, that is, if $x \in \mathbb{X}$, $\mathbb{1}_{\mathbb{X}}(x) = 1$ and otherwise, $\mathbb{1}_{\mathbb{X}}(x) = 0$.

II. System Description and Preliminaries

Consider a stochastic control system described by a Markov controlled process $\mathcal{S} = (\mathbb{X}, \mathbb{U}, T)$, where

1) $\mathbb{X}$ is a state-space endowed with a Borel σ-algebra $\mathcal{B}(\mathbb{X})$;
2) $\mathbb{U}$ is a compact control space endowed with a Borel σ-algebra $\mathcal{B}(\mathbb{U})$;
3) $T : \mathcal{B}(\mathbb{X}) \times \mathbb{X} \times \mathbb{U} \to \mathbb{R}$ is a Borel-measurable stochastic kernel given $\mathbb{X} \times \mathbb{U}$, which assigns to each $x \in \mathbb{X}$ and $u \in \mathbb{U}$ a probability measure on the Borel space $(\mathbb{X}, \mathcal{B}(\mathbb{X}))$: $T(\cdot|x, u)$.

Let us denote by $\mathbb{U}_x$ the set of the admissible control actions for each $x \in \mathbb{X}$. Assume that $\mathbb{U}_x$ is nonempty for each $x \in \mathbb{X}$.

Consider a finite horizon $N \in \mathbb{N}$. A policy is said to be a Markov policy if the control inputs are only dependent on the current state, i.e., $u_k = \mu_k(x_k)$.

Definition 1 (Markov Policy): A Markov policy $\boldsymbol{\mu}$ for system $\mathcal{S}$ is a sequence $\boldsymbol{\mu} = (\mu_0, \mu_1, \ldots, \mu_{N-1})$ of universally measurable maps

$$\mu_k : \mathbb{X} \to \mathbb{U} \quad \forall k \in \mathbb{N}_{[0,N-1]}.$$

Remark 1: Given a space $\mathbb{Y}$, a subset $\mathbb{A}$ in this space is universally measurable if it is measurable with respect to every complete probability measure on $\mathbb{Y}$ that measures all Borel sets in $\mathcal{B}(\mathbb{Y})$. A function $\mu : \mathbb{Y} \to \mathbb{W}$ is universally measurable if $\mu^{-1}(\mathbb{A})$ is universally measurable in $\mathbb{Y}$ for every $\mathbb{A} \in \mathcal{B}(\mathbb{W})$. As stated in [23] and [26], the condition of universal measurability is weaker than the condition of Borel measurability for showing the existence of a solution to a stochastic optimal problem. Roughly speaking, this is because the projections of measurable sets are analytic sets and analytic sets are universally measurable but not always Borel measurable [26], [27].

Remark 2: For a large class of stochastic optimal control problems, Markov policies are sufficient to characterize the optimal policy [26]. Furthermore, since a randomized Markov policy does not increase the largest probability that the states remain in a set, we focus on deterministic Markov policies in the following.

We denote the set of Markov policies as $\mathcal{M}$. Consider a set $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$. Given an initial state $x_0 \in \mathbb{X}$ and a Markov policy $\boldsymbol{\mu} \in \mathcal{M}$, an execution is a sequence of states $(x_0, x_1, \ldots, x_N)$. Introduce the probability with which the state x_k will remain within $\mathbb{Q}$ for all $k \in \mathbb{N}_{[0,N]}$

$$p^{\boldsymbol{\mu}}_{N,\mathbb{Q}}(x_0) = \Pr\{\forall k \in \mathbb{N}_{[0,N]}, x_k \in \mathbb{Q}\}.$$

Let $p^*_{N,\mathbb{Q}}(x) = \sup_{\boldsymbol{\mu} \in \mathcal{M}} p^{\boldsymbol{\mu}}_{N,\mathbb{Q}}(x), \forall x \in \mathbb{Q}$. We call $p^*_{N,\mathbb{Q}}(x)$ the N-step invariance probability at x in the set $\mathbb{Q}$. Following the dynamic program (DP) in [23], define the value function $V^*_{k,\mathbb{Q}} : \mathbb{X} \to [0, 1], k = 0, 1, \ldots, N$, by the backward recursion

$$V^*_{k,\mathbb{Q}}(x) = \sup_{u \in \mathbb{U}} \mathbb{1}_{\mathbb{Q}}(x) \int_{\mathbb{Q}} V^*_{k+1,\mathbb{Q}}(y) T(dy|x, u), x \in \mathbb{X} \quad (1)$$

with initialization $V^*_{N,\mathbb{Q}}(x) = 1, x \in \mathbb{Q}$.

Assumption 1: The set

$$\mathbb{U}_k(x, \lambda) = \left\{ u \in \mathbb{U} \mid \int_{\mathbb{X}} V^*_{k+1,\mathbb{Q}}(y) T(dy|x, u) \geq \lambda \right\}$$

is compact for all $x \in \mathbb{Q}$, $\lambda \in \mathbb{R}$, and $k \in \mathbb{N}_{[0,N-1]}$.

Lemma 1 (see [23]): For all $x \in \mathbb{Q}$, $p^*_{N,\mathbb{Q}}(x) = V^*_{0,\mathbb{Q}}(x)$. If Assumption 1 holds, the optimal Markov policy $\boldsymbol{\mu}^*_{\mathbb{Q}} = (\mu^*_{0,\mathbb{Q}}, \mu^*_{1,\mathbb{Q}}, \ldots, \mu^*_{N-1,\mathbb{Q}})$ exists and is given by

$$\mu^*_{k,\mathbb{Q}}(x) = \arg\sup_{u \in \mathbb{U}} \mathbb{1}_{\mathbb{Q}}(x) \int_{\mathbb{Q}} V^*_{k+1,\mathbb{Q}}(y) T(dy|x, u) \quad x \in \mathbb{Q}, k \in \mathbb{N}_{[0,N-1]}.$$

Extending the finite horizon to infinite horizon, we need to introduce stationary Markov policies.

Definition 2 (Stationary Markov Policy): A Markov policy $\boldsymbol{\mu} \in \mathcal{M}$ is said to be stationary if $\boldsymbol{\mu} = (\bar{\mu}, \bar{\mu}, \ldots)$ with $\bar{\mu} : \mathbb{X} \to \mathbb{U}$ universally measurable.

Given an initial state $x_0 \in \mathbb{X}$ and a stationary Markov policy $\boldsymbol{\mu} \in \mathcal{M}$, an execution is denoted by a sequence of states $(x_0, x_1, \ldots)$. We introduce the probability with which the state x_k will remain within $\mathbb{Q}$ for all $k \in \mathbb{N}_{\geq 0}$

$$p^{\boldsymbol{\mu}}_{\infty,\mathbb{Q}}(x_0) = \Pr\{\forall k \in \mathbb{N}, x_k \in \mathbb{Q}\}.$$

Denote $p^*_{\infty,\mathbb{Q}}(x_0) = \sup_{\boldsymbol{\mu} \in \mathcal{M}} p^{\boldsymbol{\mu}}_{\infty,\mathbb{Q}}(x_0)$. We call $p^*_{\infty,\mathbb{Q}}(x)$ the infinite-horizon invariance probability at x in the set $\mathbb{Q}$. Define the value function $G^*_{k,\mathbb{Q}} : \mathbb{X} \to [0, 1], k \in \mathbb{N}_{\geq 0}$, through the forward recursion

$$G^*_{k+1,\mathbb{Q}}(x) = \sup_{u \in \mathbb{U}} \mathbb{1}_{\mathbb{Q}}(x) \int_{\mathbb{Q}} G^*_{k,\mathbb{Q}}(y) T(dy|x, u), x \in \mathbb{X} \quad (2)$$

initialized with $G^*_{0,\mathbb{Q}}(x) = 1, x \in \mathbb{Q}$.

Assumption 2: There exists a $\bar{k} \geq 0$ such that the set

$$\mathbb{U}_k(x, \lambda) = \left\{ u \in \mathbb{U} \mid \int_{\mathbb{X}} G^*_{k,\mathbb{Q}}(y) T(dy|x, u) \geq \lambda \right\}$$

is compact for all $x \in \mathbb{Q}$, $\lambda \in \mathbb{R}$, and $k \in \mathbb{N}_{\geq \bar{k}}$.

Lemma 2 (see [23]): Suppose that Assumption 2 holds. Then, for all $x \in \mathbb{Q}$, the limit $G^*_{\infty,\mathbb{Q}}(x)$ exists and satisfies

$$G^*_{\infty,\mathbb{Q}}(x) = \sup_{u \in \mathbb{U}} \mathbb{1}_{\mathbb{Q}}(x) \int_{\mathbb{Q}} G^*_{\infty,\mathbb{Q}}(y) T(dy|x, u) \quad (3)$$

and $p^*_{\infty,\mathbb{Q}}(x) = G^*_{\infty,\mathbb{Q}}(x)$. Furthermore, an optimal stationary Markov policy $\boldsymbol{\mu}^*_{\mathbb{Q}} = (\bar{\mu}^*_{\mathbb{Q}}, \bar{\mu}^*_{\mathbb{Q}}, \ldots)$ exists and is given by

$$\bar{\mu}^*_{\mathbb{Q}}(x) = \arg\sup_{u \in \mathbb{U}} \mathbb{1}_{\mathbb{Q}}(x) \int_{\mathbb{Q}} G^*_{\infty,\mathbb{Q}}(y) T(dy|x, u), x \in \mathbb{Q}.$$

In the following two sections, we explore finite- and infinite-horizon PCISs and how to compute them.

III. Finite-Horizon ϵ-PCIS

In this section, we first define finite-horizon ϵ-PCIS for the system $\mathcal{S}$ and provide the properties of this set. Then, we explore how to compute the finite-horizon ϵ-PCIS within a given set.

Definition 3 (N-step ϵ-PCIS): Consider a stochastic control system $\mathcal{S} = (\mathbb{X}, \mathbb{U}, T)$. Given a confidence level $0 \leq \epsilon \leq 1$, a set $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$ is an N-step ϵ-PCIS for $\mathcal{S}$ if for any $x \in \mathbb{Q}$, there exists at least one Markov policy $\boldsymbol{\mu} \in \mathcal{M}$ such that $p^{\boldsymbol{\mu}}_{N,\mathbb{Q}}(x) \geq \epsilon$.

We define the stochastic backward reachable set $\mathbb{S}^*_{\epsilon,N}(\mathbb{Q})$ by collecting all the states $x \in \mathbb{Q}$ at which the N-step invariance

probability $p^*_{N,\mathbb{Q}}(x) \geq \epsilon$, i.e.,

$$\begin{aligned}\mathbb{S}^*_{\epsilon,N}(\mathbb{Q}) &= \{x \in \mathbb{Q} \mid \exists \boldsymbol{\mu} \in \mathcal{M}, p^{\boldsymbol{\mu}}_{N,\mathbb{Q}}(x) \geq \epsilon\} \\ &= \{x \in \mathbb{Q} \mid \sup_{\boldsymbol{\mu}\in\mathcal{M}} p^{\boldsymbol{\mu}}_{N,\mathbb{Q}}(x) \geq \epsilon\} \\ &= \{x \in \mathbb{Q} \mid V^*_{0,\mathbb{Q}}(x) \geq \epsilon\}.\end{aligned}$$

If $\mathbb{S}^*_{\epsilon,N}(\mathbb{Q}) = \mathbb{Q}$, it yields from $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$ that $\mathbb{S}^*_{\epsilon,N}(\mathbb{Q})$ is also Borel-measurable. If $\mathbb{S}^*_{\epsilon,N}(\mathbb{Q}) \subset \mathbb{Q}$, the following lemma addresses the measurability of the set $\mathbb{S}^*_{\epsilon,N}(\mathbb{Q})$.

Lemma 3: For any $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$, the set $\mathbb{S}^*_{\epsilon,N}(\mathbb{Q}) \subseteq \mathbb{Q}$ is universally measurable.

Proof: See Appendix A. ■

Let us denote by $\mathcal{P}(\mathbb{X})$ the set of all probability measures on $\mathbb{X}$. The following proposition shows that despite of the universal measurability of $\mathbb{S}^*_{\epsilon,N}(\mathbb{Q})$, one can find another Borel-measurable set $\tilde{\mathbb{S}}^*_{\epsilon,N}(\mathbb{Q}))$ for which the difference to $\mathbb{S}^*_{\epsilon,N}(\mathbb{Q})$ is measure-zero for any probability measure on $\mathbb{X}$.

Proposition 1: For any $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$, there exists a set $\tilde{\mathbb{S}}^*_{\epsilon,N}(\mathbb{Q}) \in \mathcal{B}(\mathbb{X})$ with $\tilde{\mathbb{S}}^*_{\epsilon,N}(\mathbb{Q}) \subseteq \mathbb{Q}$ such that $p(\tilde{\mathbb{S}}^*_{\epsilon,N}(\mathbb{Q}) \triangle \mathbb{S}^*_{\epsilon,N}(\mathbb{Q})) = 0$ for any $p \in \mathcal{P}(\mathbb{X})$.

Proof: It follows from the universal measurability of $\mathbb{S}^*_{\epsilon,N}(\mathbb{Q})$ as shown in Lemma 3, the Borel measurability of $\mathbb{Q}$, $\mathbb{S}^*_{\epsilon,N}(\mathbb{Q}) \subseteq \mathbb{Q}$, and [26, Lemma 7.26]. ■

From Lemma 1 and the definition of $\mathbb{S}^*_{\epsilon,N}(\mathbb{Q})$, we can verify whether a set $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$ is an N-step ϵ-PCIS or not by checking if either $\mathbb{S}^*_{\epsilon,N}(\mathbb{Q}) = \mathbb{Q}$, or $V^*_{0,\mathbb{Q}}(x) \geq \epsilon, \forall x \in \mathbb{Q}$, where $V^*_{0,\mathbb{Q}}(x)$ is defined in (1).

Remark 3: The stochastic backward reachable set $\mathbb{S}^*_{\epsilon,N}(\mathbb{Q})$ is called the maximal probabilistic safe set in [23]. The N-step ϵ-PCIS $\mathbb{Q}$ in Definition 3 refines the maximal probabilistic safe set by requiring that for any initial state $x_0 \in \mathbb{Q}$, the N-step invariance probability $p^*_{\infty,\mathbb{Q}}(x_0)$ is no less than ϵ.

In the following, we show that finite-horizon PCISs are closed under union.

Proposition 2: Consider a collection of sets $\mathbb{Q}_i \in \mathcal{B}(\mathbb{X})$, $i = 1, \ldots, r$. If each $\mathbb{Q}_i$ is an N_i-step ϵ_i-PCIS for the same system $\mathbb{S}$, then the union $\bigcup_{i=1}^r \mathbb{Q}_i$ is an N-step ϵ-PCIS, where $N = \min_i N_i$ and $\epsilon = \min_i \epsilon_i$.

Proof: The result follows from the following two facts: (i) for any $\mathbb{Q}, \mathbb{P} \in \mathcal{B}(\mathbb{X})$ with $\mathbb{Q} \subseteq \mathbb{P}$, $\sup_{\boldsymbol{\mu}\in\mathcal{M}} p^{\boldsymbol{\mu}}_{N,\mathbb{Q}}(x) \leq \sup_{\boldsymbol{\mu}\in\mathcal{M}} p^{\boldsymbol{\mu}}_{N,\mathbb{P}}(x)$, $\forall N \in \mathbb{N}$ and $\forall x \in \mathbb{Q}$; (ii) for any $N, N' \in \mathbb{N}$ with $N \leq N'$, $\sup_{\boldsymbol{\mu}\in\mathcal{M}} p^{\boldsymbol{\mu}}_{N',\mathbb{Q}}(x) \leq \sup_{\boldsymbol{\mu}\in\mathcal{M}} p^{\boldsymbol{\mu}}_{N,\mathbb{Q}}(x), \forall Q \in \mathcal{B}(\mathbb{X})$ and $\forall x \in \mathbb{Q}$. ■

A. Finite-Horizon ϵ-PCIS Computation

This section will address the following problem.

Problem 1: Given a set $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$ and a prescribed probability $0 \leq \epsilon \leq 1$, compute an N-step ϵ-PCIS $\tilde{\mathbb{Q}} \subseteq \mathbb{Q}$.

To handle this problem, our basic idea is to iteratively compute stochastic backward reachable sets until convergence. A general procedure is presented in the following algorithm.

In Algorithm 1, we compute the stochastic backward reachable set $\mathbb{S}^*_{\epsilon,N}(\mathbb{P}_i)$ within $\mathbb{P}_i$ and update $\mathbb{P}_{i+1}$ to be the corresponding Borel-measurable set $\tilde{\mathbb{S}}^*_{\epsilon,N}(\mathbb{P}_i)$. The following theorem shows convergence of $\mathbb{P}_i$. The terminal condition guarantees

Algorithm 1: N-Step ϵ-PCIS.

1: Initialize $i = 0$ and $\mathbb{P}_i = \mathbb{Q}$.
2: Compute $V^*_{0,\mathbb{P}_i}(x), \forall x \in \mathbb{P}_i$.
3: Compute $\mathbb{S}^*_{\epsilon,N}(\mathbb{P}_i)$ and $\mathbb{P}_{i+1} = \tilde{\mathbb{S}}^*_{\epsilon,N}(\mathbb{P}_i)$, where $\tilde{\mathbb{S}}^*_{\epsilon,N}(\cdot)$ is defined in Proposition 1.
4: If $\mathbb{P}_{i+1} = \mathbb{P}_i$, stop. Else, set $i = i + 1$ and go to step 2.

that the resulting set by this algorithm is an N-step ϵ-PCIS $\tilde{\mathbb{Q}} \subseteq \mathbb{Q}$.

Theorem 1: Let Assumption 1 hold. For any $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$, Algorithm 1 converges, i.e., $\lim_{i\to\infty} \mathbb{P}_i$ exists. If $\lim_{i\to\infty} \mathbb{P}_i \neq \emptyset$, it is the largest N-step ϵ-PCIS within $\mathbb{Q}$.

Proof: From Algorithm 1 and Lemma 1, we have that if the termination condition does not hold, $\mathbb{P}_{i+1} \subset \mathbb{P}_i$. It follows that the sequence $\{\mathbb{P}_i\}_{i\in\mathbb{N}}$ is nonincreasing. Then

$$\liminf_{i\to\infty} \mathbb{P}_i = \bigcup_{i\geq 1} \bigcap_{j\geq i} \mathbb{P}_j = \bigcap_{j\geq 1} \mathbb{P}_j = \bigcap_{i\geq 1} \bigcup_{j\geq i} \mathbb{P}_j = \limsup_{i\to\infty} \mathbb{P}_i$$

which suggests the existence of $\lim_{i\to\infty} \mathbb{P}_i$. Furthermore, if $\lim_{i\to\infty} \mathbb{P}_i$ is nonempty, we conclude that it is the largest N-step PCIS within $\mathbb{Q}$ based on the fixed-point theory. ■

To facilitate the practical implementation of Algorithm 1, we need to address two important properties: the computational tractability of $V^*_{0,\mathbb{P}_i}(x), \forall x \in \mathbb{P}_i$, and the finite-step convergence of Algorithm 1. In the following, we will derive these two properties for discrete and continuous spaces, respectively. It is shown that if the spaces are discrete, the properties are guaranteed and in particular at each iteration we only need to solve an LP to compute the exact value of $V^*_{0,\mathbb{P}_i}$. If the spaces are continuous, we will design a discretization algorithm with convergence guarantee, which enables us to preserve the abovementioned two properties.

1) Discrete State and Control Spaces: If the state and control spaces are discrete, i.e., they are finite sets, the stochastic kernel $T(y|x,u)$ denotes the transition probability from state $x \in \mathbb{X}$ to state $y \in \mathbb{X}$ under control action $u \in \mathbb{U}_x$, which satisfies that $\sum_{y\in\mathbb{X}} T(y|x,u) = 1, \forall x \in \mathbb{X}$ and $u \in \mathbb{U}_x$.

In this case, according to [28, Th. 1], we can exactly compute $V^*_{0,\mathbb{P}_i}(x)$ via an LP. Moreover, the existence of the optimal Markov policy can be always guaranteed.

Lemma 4: Given any set $\mathbb{P}_i \subset \mathbb{X}$, the value functions $V^*_{k,\mathbb{P}_i}$ in (1) can be obtained by solving an LP

$$\min \quad \sum_{k=0}^{N} \sum_{x\in\mathbb{P}_i} v_k(x) \tag{4a}$$

$$\text{subject to} \quad \forall x \in \mathbb{P}_i$$

$$v_k(x) \geq \sum_{y\in\mathbb{P}_i} v_{k+1}(y) T(y|x,u)$$

$$\forall u \in \mathbb{U}_x \quad \forall k \in \mathbb{N}_{[0,N-1]} \tag{4b}$$

$$v_N(x) \geq 1 \tag{4c}$$

which gives $V^*_{k,\mathbb{P}_i}(x) = v^*_k(x),\ \forall x \in \mathbb{P}_i$, and $\forall k \in \mathbb{N}_{[0,N]}$, where v^*_k is the optimal solution of (4). The optimal Markov policy $\boldsymbol{\mu}^*_{\mathbb{P}_i} = (\mu^*_{0,\mathbb{P}_i}, \mu^*_{1,\mathbb{P}_i}, \ldots, \mu^*_{N-1,\mathbb{P}_i})$ is given by $\mu^*_{k,\mathbb{P}_i}(x) = u$ where $u \in \mathbb{U}_x$ is such that

$$v^*_k(x) = \sum_{y \in \mathbb{P}_i} v^*_{k+1}(y) T(y|x,u). \tag{5}$$

Proof: See [28, Th. 1] for the proof. ■

Corollary 1: For discrete state and control spaces, Algorithm 1 converges in a finite number of iterations. Furthermore, at each iteration, the N-step invariance probability $V^*_{0,\mathbb{P}_i}(x)$, $\forall x \in \mathbb{P}_i$, can be computed via the LP (4) and the corresponding optimal policy is determined by (5).

Proof: The finite-step convergence of Algorithm 1 follows from Theorem 1 and the finite cardinality of $\mathbb{Q}$. The remaining part follows from Lemma 4. ■

Remark 4: When implementing Algorithm 1 to a system with discrete spaces, the maximal number of iterations is $|\mathbb{Q}|$. At each iteration, an LP is solved to compute the value of $V^*_{0,\mathbb{P}_i}(x)$, $\forall x \in \mathbb{P}_i$. The number of the decision values in the LP is at most $|\mathbb{Q}|(N+1)$ and the number of constraints is at most $|\mathbb{Q}|(N|\mathbb{U}|+1)$. It follows from [29] that Algorithm 1 can be implemented in $O(|\mathbb{Q}|^2(N|\mathbb{U}|+1))$ time.

2) Continous State and Control Action Spaces: In order to preserve the computational tractability of $V^*_{0,\mathbb{P}_i}$ and the finite-step convergence of Algorithm 1, if the state and control spaces are both continuous, we first discretize the spaces with convergence guarantee. Then, we adapt Algorithm 1 to compute an approximate N-step ϵ-PCIS within a given set.

Assume that $\mathbb{X} \subseteq \mathbb{R}^{n_x}$ and $\mathbb{U} \subseteq \mathbb{R}^{n_u}$ for some $n_x, n_u \in \mathbb{N}$. For simplicity, we use Euclidean metric for the spaces $\mathbb{X}$ and $\mathbb{U}$. For any $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$, we define $\phi(\mathbb{Q}) = \text{Leb}(\mathbb{Q})$ where $\text{Leb}(\cdot)$ denotes the Lebesgue measure of sets. We suppose that the stochastic kernel $T(\cdot|x,u)$ admits a density $t(y|x,u)$, which represents the probability density of y given the current state x and the control action u.

Now we consider Problem 1, where we assume that the given set $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$ is compact, which implies that $\phi(\mathbb{Q})$ is bounded. We further suppose that the density function satisfies the following assumption.

Assumption 3: For any $x, x,' y, y' \in \mathbb{Q}$, and $u, u' \in \mathbb{U}$, there exists a constant L such that $|t(y|x,u) - t(y'|x,' u')| \le L(\|y - y'\| + \|x - x'\| + \|u - u'\|)$.

Discretization: We discretize the compact set $\mathbb{Q} \subset \mathbb{X}$ into m_x pair-wise disjoint nonempty Borel sets $\mathbb{Q}_i, i \in \mathbb{N}_{[1,m_x]}$, i.e., $\mathbb{Q} = \cup_{i=1}^{m_x} \mathbb{Q}_i$. We pick a representative state from each set $\mathbb{Q}_i$, denoted by q_i. Let $\hat{\mathbb{Q}} = \{q_i, i \in \mathbb{N}_{[1,m_x]}\}$, $d_i = \sup_{x,y \in \mathbb{Q}_i} \|x - y\|$, and $D_x = \max_{i \in \mathbb{N}_{[1,m_x]}} d_i$.

Similarly, the compact control space $\mathbb{U}$ is divided into m_u pair-wise disjoint nonempty Borel sets $\mathbb{C}_i$, $i \in \mathbb{N}_{[1,m_u]}$, i.e., $\mathbb{U} = \cup_{i=1}^{m_u} \mathbb{C}_i$. We pick a representative element from the set $\mathbb{C}_i$, denoted by $\hat{u}_i$. Let $\hat{\mathbb{U}} = \{\hat{u}_i, i \in \mathbb{N}_{[1,m_u]}\}$, $l_i = \sup_{x,y \in \mathbb{C}_i} \|x - y\|$, and $D_u = \max_{i \in \mathbb{N}_{[1,m_u]}} l_i$.

Let the grid size be a constant $\delta \ge \max\{D_x, D_u\}$. For each $x \in \mathbb{Q}$, define the set of admissible discrete control actions as

$$\hat{\mathbb{U}}_x = \{\hat{u} \in \hat{\mathbb{U}} \mid \|u - \hat{u}\| \le \delta \text{ for some } u \in \mathbb{U}_{s_x}\} \tag{6}$$

where s_x is the representative state of $\mathbb{Q}_i$ to which x belongs, i.e., $s_x = q_i$ if $x \in \mathbb{Q}_i$. Following [25], the following lemma shows that each $x \in \mathbb{Q}$ has a nonempty admissible discretized control set.

Lemma 5: For each $q_i \in \hat{\mathbb{Q}}$, the set $\hat{\mathbb{U}}_{q_i}$ is nonempty and $\hat{\mathbb{U}}_x = \hat{\mathbb{U}}_{q_i}, \forall x \in \mathbb{Q}_i$.

Proof: Since the admissible control set $\mathbb{U}_{s_x}$ is nonempty, $\forall x \in \mathbb{Q}$, there exists $\hat{u} \in \hat{\mathbb{U}}$ such that $\|u - \hat{u}\| \le \delta, \forall u \in \mathbb{U}_{s_x}$. Hence, by the definition of s_x, we have that the set $\hat{\mathbb{U}}_{q_i}$ is nonempty for each $q_i \in \hat{\mathbb{Q}}$. Furthermore, from (6), it is easy to obtain that $\hat{\mathbb{U}}_x = \hat{\mathbb{U}}_{q_i}, \forall x \in \mathbb{Q}_i$. ■

As in [25], let us define the function $\hat{t} : \mathbb{Q} \times \mathbb{Q} \times \hat{\mathbb{U}} \to \mathbb{R}$

$$\hat{t}(y|x,\hat{u}) = \begin{cases} \frac{t(s_y|s_x,\hat{u})}{\int_{\mathbb{Q}} t(s_z|s_x,\hat{u})dz}, & \text{if } \int_{\mathbb{Q}} t(s_z|s_x,\hat{u})dz \ge 1 \\ t(s_y|s_x,\hat{u}), & \text{otherwise.} \end{cases} \tag{7}$$

From (7), we observe that all states $y \in \mathbb{Q}_i$ enjoy the same stochastic kernel. An approximate stochastic control system is given by a triple $\hat{\mathcal{S}}_{\mathbb{Q}} = (\hat{\mathbb{Q}}, \hat{\mathbb{U}}, \hat{T})$. Here, the transition probability $\hat{T}(q_j|q_i,\hat{u})$ is defined by $\hat{T}(q_j|q_i,\hat{u}) = \int_{\mathbb{Q}_j} \hat{t}(y|q_i,\hat{u})dy$, where $q_i, q_j \in \hat{\mathbb{Q}}$ with $q_i \in \mathbb{Q}_i$ and $q_j \in \mathbb{Q}_j$, and $\hat{u} \in \hat{\mathbb{U}}$. *Approximation of PCISs.* For the approximate system $\hat{\mathcal{S}}_{\mathbb{Q}}$, the discretized version of the DP (1) is given by

$$\begin{cases} \hat{V}^*_{N,\mathbb{Q}}(q_i) = 1 \\ \hat{V}^*_{k,\mathbb{Q}}(q_i) = \max_{\hat{u} \in \hat{\mathbb{U}}} \left(\sum_{j=1}^{m_x} \hat{V}^*_{k+1,\mathbb{Q}}(q_j) \hat{T}(q_j|q_i,\hat{u}) \right) \\ \qquad\qquad \forall k \in \mathbb{N}_{[0,N-1]}. \end{cases}$$

For each $x \in \mathbb{Q}_i$, $\hat{V}^*_{k,\mathbb{Q}}(x) = \hat{V}^*_{k,\mathbb{Q}}(q_i), \forall k \in \mathbb{N}_{[0,N]}$. We define the discretized optimal Markov policy $\hat{\boldsymbol{\mu}}^*_{\mathbb{Q}} = (\hat{\mu}^*_{0,\mathbb{Q}}, \ldots, \hat{\mu}^*_{N-1,\mathbb{Q}})$ as

$$\begin{aligned} \hat{\mu}^*_{k,\mathbb{Q}}(q_i) &= \arg\max_{\hat{u} \in \hat{\mathbb{U}}} \int_{\mathbb{Q}} \hat{V}^*_{k+1,\mathbb{Q}}(y) \hat{t}(y|q_i,\hat{u})dy \\ &= \arg\max_{\hat{u} \in \hat{\mathbb{U}}} \left(\sum_{j=1}^{m_x} \hat{V}^*_{k+1,\mathbb{Q}}(q_j) \hat{T}(q_j|q_i,\hat{u}) \right). \end{aligned}$$

For each $x \in \mathbb{Q}_i$, $\hat{\mu}^*_{k,\mathbb{Q}}(x) = \hat{\mu}^*_{k,\mathbb{Q}}(q_i),\ \forall k \in \mathbb{N}_{[0,N-1]}$.

Remark 5: Since the state and control action spaces of the approximated system $\hat{\mathcal{S}}$ are finite, the value of $\hat{V}^*_{k,\mathbb{Q}}$ can be computed via the LP (4) and the corresponding optimal policy can be determined by (5). In addition, all the states in each $\mathbb{Q}_i$ share the same approximate N-step invariance probability and optimal policy as the representative state $q_i \in \mathbb{Q}_i$.

Lemma 6: Under Assumptions 1 and 3, the functions $V^*_{k,\mathbb{Q}}(x)$ and $\hat{V}^*_{k,\mathbb{Q}}(x)$ satisfy that $\forall x \in \mathbb{Q}$

$$|V^*_{k,\mathbb{Q}}(x) - \hat{V}^*_{k,\mathbb{Q}}(x)| \le \tau_k(\mathbb{Q})\delta \tag{8}$$

where

$$\begin{cases} \tau_N(\mathbb{Q}) = 0 \\ \tau_k(\mathbb{Q}) = 4\phi(\mathbb{Q})L + \tau_{k+1}(\mathbb{Q}) \quad \forall k \in \mathbb{N}_{[0,N-1]}. \end{cases} \tag{9}$$

Proof: See Appendix B. ■

Algorithm 2: Approximate N-Step ϵ-PCIS.

1: Choose grid size $0 < \delta < \frac{1-\epsilon}{\tau_0(\mathbb{Q})}$, discretize the sets $\mathbb{Q}$ and $\mathbb{U}$, construct an approximate system $\hat{\mathcal{S}}_{\mathbb{Q}} = (\hat{\mathbb{Q}}, \hat{\mathbb{U}}, \hat{T})$.
2: Initialize $i = 0$, $\mathbb{P}_i = \mathbb{Q}$, and $\hat{\mathbb{P}}_i = \hat{\mathbb{Q}}$.
3: Compute $\hat{V}^*_{0,\mathbb{P}_i}(q_j), \forall q_j \in \hat{\mathbb{P}}_i$.
4: Compute $\tau_0(\hat{\mathbb{P}}_i)$ by (9) and $\hat{\epsilon} = \epsilon + \tau_0(\mathbb{P}_i)\delta$.
5: Compute the set $\hat{\mathbb{P}}_{i+1} = \mathbb{S}^*_{\hat{\epsilon},N}(\hat{\mathbb{P}}_i)$ for $\hat{\mathcal{S}}_{\mathbb{Q}}$ and $\mathbb{P}_i = \cup_{q_j \in \hat{\mathbb{P}}_i} \mathbb{Q}_j$
6: If $\hat{\mathbb{P}}_{i+1} = \hat{\mathbb{P}}_i$, stop. Else, set $i = i + 1$ and go to step 3.

Remark 6: Lemma 6 guarantees convergence as the grid size tends to zero and generalizes the case considered in [23], which only discretizes the state space for a given finite control space. To prove Lemma 6, we need to show that (i) the value functions in (1) are Lipschitz continuous (Lemma 8), which is similar to [23, Th. 8], and (ii) the difference between the approximate density function and the original density function is bounded (Lemma 9), which is different from that in [23].

Theorem 2: Let Assumptions 1 and 3 hold. Consider a compact set $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$ and a corresponding discretized set $\hat{\mathbb{Q}}$ of $\mathbb{Q}$. If $\hat{\mathbb{Q}}$ is an N-step $\hat{\epsilon}$-PCIS for the approximate system $\hat{\mathcal{S}}_{\mathbb{Q}} = (\hat{\mathbb{Q}}, \hat{\mathbb{U}}, \hat{T})$, and $\hat{\epsilon} \geq \tau_0(\mathbb{Q})\delta$, the set $\mathbb{Q}$ is an N-step ϵ-PCIS for the system $\mathcal{S}$, where $\epsilon = \hat{\epsilon} - \tau_0(\mathbb{Q})\delta$.

Proof: According to the construction of the discretized system $\hat{\mathcal{S}}_{\mathbb{Q}}$, we have that $\forall k \in \mathbb{N}_{[0,N]}$, $\forall i \in \mathbb{N}_{[1,m_x]}$, and $\forall x \in \mathbb{Q}_i$, $\hat{V}^*_{k,\mathbb{Q}}(x) = \hat{V}^*_{k,\mathbb{Q}}(q_i)$. Since $\hat{\mathbb{Q}}$ is an N-step $\hat{\epsilon}$-PCIS, it follows that $\forall x \in \mathbb{Q}$, $\hat{V}^*_{0,\mathbb{Q}}(x) \geq \hat{\epsilon}$. By Lemma 6 and triangular inequality, we have

$$V^*_{0,\mathbb{Q}}(x) \geq \hat{V}^*_{0,\mathbb{Q}}(x) - \tau_0(\mathbb{Q})\delta \geq \hat{\epsilon} - \tau_0(\mathbb{Q})\delta \quad \forall x \in \mathbb{Q}.$$

Then, when $\hat{\epsilon} \geq \tau_0(\mathbb{Q})\delta$, we conclude that the set $\mathbb{Q}$ is an N-step ϵ-PCIS where $0 \leq \epsilon = \hat{\epsilon} - \tau_0(\mathbb{Q})\delta$. ■

Remark 7: From Theorem 2, if $0 \leq \epsilon < 1$, by choosing a suitable grid size $0 < \delta \leq \frac{1-\epsilon}{\tau_0(\mathbb{Q})}$, the problem of computing an N-step ϵ-PCIS within $\mathbb{Q}$ for $\mathcal{S}$ can be transformed into that of computing an approximate N-step $\hat{\epsilon}$-PCIS with probability $\hat{\epsilon} \geq \epsilon + \tau_0(\mathbb{Q})\delta$ for $\hat{\mathcal{S}}_{\mathbb{Q}}$.

Computation algorithm: Assume that a probability level $0 \leq \epsilon < 1$ is given. After discretizing the set $\mathbb{Q}$ and the control space $\mathbb{U}$, we modify Algorithm 1 to compute an N-step ϵ-PCIS $\tilde{\mathbb{Q}} \subseteq \mathbb{Q}$, as shown in the following.

In Algorithm 2, we first construct an approximate system $\hat{\mathcal{S}}_{\mathbb{Q}} = (\hat{\mathbb{Q}}, \hat{\mathbb{U}}, \hat{T})$ with grid size $0 < \delta < \frac{1-\epsilon}{\tau_0(\mathbb{Q})}$. Then, following similar steps as in Algorithm 1, we compute the stochastic backward reachable set iteratively for the system $\hat{\mathcal{S}}_{\mathbb{Q}}$. At each iteration, an LP is solved to obtain the N-step invariance probability. One difference is that the stochastic backward reachable set is computed with respect to $\hat{\epsilon} = \epsilon + \tau_0(\mathbb{P}_i)\delta$ and the updated set for the system $\mathcal{S}$ is the union of the subsets of $\mathbb{Q}$ corresponding to the stochastic backward reachable set. By Theorem 2, the resulting set by Algorithm 2 is an N-step ϵ-PCIS.

Corollary 2: Let Assumptions 1 and 3 hold. For continuous state and control spaces, Algorithm 2 converges in a finite number of iterations and generates an N-step ϵ-PCIS. Furthermore, at each iteration, the N-step invariance probability $\hat{V}^*_{0,\mathbb{P}_i}(q_j)$, $\forall q_j \in \hat{\mathbb{P}}_i$, can be computed via the LP (4) and the corresponding optimal policy is determined by (5).

Proof: By Theorem 2 and the Borel measurability of the subsets $\mathbb{Q}_i, \forall i \in \mathbb{N}_{[1,m_x]}$, it follows that the set generated by Algorithm 2 is an N-step ϵ-PCIS. The remaining part is similar to the proof of Corollary 1. ■

Remark 8: When implementing Algorithm 2 to a system with continuous spaces, it follows from [29] that Algorithm 2 can be implemented in $O(m_x^2(Nm_u + 1))$ time, cf. Remark 4.

IV. Extension to Infinite-Horizon ϵ-PCIS

Now let us extend finite-horizon ϵ-PCISs to infinite-horizon ϵ-PCISs. In this section, we define the infinite-horizon ϵ-PCIS and explore the conditions of its existence. Furthermore, we provide algorithms to compute an infinite-horizon ϵ-PCIS within a given set.

Definition 4 (Infinite-horizon PCIS): Consider a stochastic control system $\mathcal{S} = (\mathbb{X}, \mathbb{U}, T)$. Given a confidence level $0 \leq \epsilon \leq 1$, a set $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$ is an infinite-horizon ϵ-PCIS for $\mathcal{S}$ if for any $x \in \mathbb{Q}$, there exists at least one stationary Markov policy $\boldsymbol{\mu} \in \mathcal{M}$ such that $p^{\boldsymbol{\mu}}_{\infty,\mathbb{Q}}(x) \geq \epsilon$.

We define the stochastic backward reachable set $\mathbb{S}^*_{\epsilon,\infty}(\mathbb{Q})$ by collecting all the states $x \in \mathbb{Q}$ at which the infinite-horizon invariance probability $p^*_{\infty,\mathbb{Q}}(x) \geq \epsilon$, i.e.,

$$\begin{aligned}\mathbb{S}^*_{\epsilon,\infty}(\mathbb{Q}) &= \{x \in \mathbb{Q} \mid \exists \boldsymbol{\mu} \in \mathcal{M}, p^{\boldsymbol{\mu}}_{\infty,\mathbb{Q}}(x) \geq \epsilon\} \\ &= \{x \in \mathbb{Q} \mid \sup_{\boldsymbol{\mu} \in \mathcal{M}} p^{\boldsymbol{\mu}}_{\infty,\mathbb{Q}}(x) \geq \epsilon\} \\ &= \{x \in \mathbb{Q} \mid G^*_{\infty,\mathbb{Q}}(x) \geq \epsilon\}.\end{aligned}$$

For the infinite-horizon case, Lemma 3 and Proposition 1 still hold. That is, the set $\mathbb{S}^*_{\epsilon,\infty}(\mathbb{Q})$ is universally measurable and there exists another Borel-measurable set $\tilde{\mathbb{S}}^*_{\epsilon,\infty}(\mathbb{Q}) \subseteq \mathbb{Q}$ such that $p(\tilde{\mathbb{S}}^*_{\epsilon,\infty}(\mathbb{Q}) \triangle \mathbb{S}^*_{\epsilon,\infty}(\mathbb{Q})) = 0$ for any $p \in \mathcal{P}(\mathbb{X})$.

Under Assumption 2, by Lemma 2 and the definition of $\mathbb{S}^*_{\epsilon,\infty}(\mathbb{Q})$, we can verify whether a set $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$ is an infinite-horizon ϵ-PCIS or not by checking if either $\mathbb{S}^*_{\epsilon,\infty}(\mathbb{Q}) = \mathbb{Q}$, or $G^*_{\infty,\mathbb{Q}}(x) \geq \epsilon$, $\forall x \in \mathbb{Q}$, where $G^*_{\infty,\mathbb{Q}}(x)$ is defined by (2) and (3).

Definition 5: Consider a stochastic control system $\mathcal{S} = (\mathbb{X}, \mathbb{U}, T)$. An RCIS $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$ for $\mathcal{S}$ is an N-step ϵ-PCIS with $N = 1$ and $\epsilon = 1$.

Remark 9: Another interpretation of RCIS in Definition 5 is that a set $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$ is an RCIS if for any $x \in \mathbb{Q}$, there exists at least one control input $u \in \mathbb{U}$ such that $T(\mathbb{Q}|x, u) = 1$. It is easy to verify that an RCIS is also an infinite-horizon ϵ-PCIS with $\epsilon = 1$. It is called an absorbing set in [30] where there is no control input. In the following, we show that the RCIS plays an important role in the existence of infinite-horizon PCIS and provide how to design an algorithm to compute such PCIS based on RCIS.

Remark 10: Note that infinite-horizon ϵ-PCISs are also closed under union, as shown in Proposition 2 when N is replaced by ∞.

A. Existence of Infinite-Horizon PCIS

Intuitively, the monotone decrease of $G^*_{\infty,\mathbb{Q}}(x)$ may imply that the value of $G^*_{\infty,\mathbb{Q}}(x)$ is one or zero. However, it is possible to get $0 < G^*_{\infty,\mathbb{Q}}(x) < 1$ in some cases (see Examples 1 and 2 in Section V). The following theorem provides necessary conditions and sufficient conditions for the existence of infinite-horizon ϵ-PCIS with $\epsilon > 0$.

Theorem 3: Suppose that Assumption 2 holds and let $0 < \epsilon \leq 1$ be fixed. Given a nonempty set $\mathbb{Q}$, let u_x be the control input such that (3) holds for each $x \in \mathbb{Q}$. The set $\mathbb{Q}$ is an infinite-horizon ϵ-PCIS

i) *only if* there exists an RCIS $\mathbb{Q}_f \subseteq \mathbb{Q}$ such that $\forall x \in \mathbb{Q} \setminus \mathbb{Q}_f$

$$T(\mathbb{Q}_f|x,u_x) + \int_{\mathbb{Q}\setminus\mathbb{Q}_f} T(\mathbb{Q}_f|y,u_y)T(dy|x,u_x) + \frac{\rho^2}{1-\rho} \geq \epsilon \quad (10)$$

where $\rho = \sup_{x\in\mathbb{Q}\setminus\mathbb{Q}_f} \int_{\mathbb{Q}\setminus\mathbb{Q}_f} T(dy|x,u_x)$;

ii) *if* there exists an RCIS $\mathbb{Q}_f \subseteq \mathbb{Q}$ such that $\forall x \in \mathbb{Q} \setminus \mathbb{Q}_f$

$$T(\mathbb{Q}_f|x,u_x) + \int_{\mathbb{Q}\setminus\mathbb{Q}_f} T(\mathbb{Q}_f|y,u_y)T(dy|x,u_x) \geq \epsilon. \quad (11)$$

Proof: See Appendix C. ■

Remark 11: The value of ρ is the largest probability that the next state y remains outside the RCIS $\mathbb{Q}_f$ from any $x \in \mathbb{Q} \setminus \mathbb{Q}_f$ under the optimal stationary Markov policy in Lemma 2. Note that $\frac{\rho^2}{1-\rho}$ is the gap between the necessary condition and the sufficient condition. In addition, the second item in (10) and (11) denotes the probability that the state is steered into the RCIS $\mathbb{Q}_f$ by two transitions from $x \in \mathbb{Q} \setminus \mathbb{Q}_f$ with an intermediate state y outside $\mathbb{Q}_f$.

Corollary 3: Suppose that Assumption 2 holds and let $0 < \epsilon \leq 1$ be fixed. A nonempty set $\mathbb{Q}$ is an infinite-horizon ϵ-PCIS

i) *only if* there exists an RCIS $\mathbb{Q}_f \subseteq \mathbb{Q}$ such that $\forall x \in \mathbb{Q} \setminus \mathbb{Q}_f$, $T(\mathbb{Q}|x,u) \geq \epsilon$ for some $u \in \mathbb{U}$;

ii) *if* there exists an RCIS $\mathbb{Q}_f \subseteq \mathbb{Q}$ such that $\forall x \in \mathbb{Q} \setminus \mathbb{Q}_f$, $T(\mathbb{Q}_f|x,u) + \epsilon T(\mathbb{Q} \setminus \mathbb{Q}_f|x,u) \geq \epsilon$ for some $u \in \mathbb{U}$.

Proof: See Appendix D. ■

Remark 12: A nonempty set $\mathbb{Q}$ is an infinite-horizon ϵ-PCIS if there exists an RCIS $\mathbb{Q}_f \subseteq \mathbb{Q}$ such that $\forall x \in \mathbb{Q} \setminus \mathbb{Q}_f$, $T(\mathbb{Q}_f|x,u) \geq \epsilon$ for some $u \in \mathbb{U}$. This implication will facilitate the design of an algorithm for an infinite-horizon ϵ-PCIS, see Algorithm 4.

Remark 13: Considering the similarity between the reliability defined in [11] and the infinite-horizon invariance probability in this article, we can extend the results on infinite-horizon PICSs, including the existence condition above and the computational algorithms in the following, to the reliable control set in [10] to general stochastic systems.

B. Infinite-Horizon ϵ-PCIS Computation

This section will address the following problem.

Problem 2: Given a set $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$ and a prescribed probability $0 \leq \epsilon \leq 1$, compute an infinite-horizon ϵ-PCIS $\tilde{\mathbb{Q}} \subseteq \mathbb{Q}$.

To handle this problem, the key point is to compute the infinite-horizon invariance probability $G^*_{\infty,\mathbb{Q}}$. For discrete spaces, it is shown that computationally tractable MILP can be used to compute the exact value of $G^*_{\infty,\mathbb{Q}}$. In this case, we can compute the largest infinite-horizon ϵ-PCIS by computing iteratively the stochastic backward reachable sets until convergence. For continuous spaces, it is in general computationally intractable to compute $G^*_{\infty,\mathbb{Q}}$ and the discretization method fails to work since the approximation error in (8) increases with the horizon. In this case, we design another computational algorithm based on the sufficient conditions in Remark 12.

1) Discrete State and Control Spaces: If the state and control spaces are discrete, we adopt the same assumptions as in Section III-A1. We will first show how to compute the exact value of $G^*_{\infty,\mathbb{Q}}$ in (2) and (3) through an MILP. Then, we will adapt Algorithm 1 to compute the largest infinite-horizon ϵ-PCIS within a given set.

MILP reformulation: Since 0 is a trivial solution of (3), we cannot directly reformulate (2) and (3) as an LP, which is the traditional way to deal with infinite-horizon stochastic optimal control problems [31].

The following lemma provides a computationally tractable MILP reformulation when computing $G^*_{\infty,\mathbb{Q}}$.

Lemma 7: Given any set $\mathbb{Q} \subseteq \mathbb{X}$, the value of $G^*_{\infty,\mathbb{Q}}$ in (3) can be obtained by solving the MILP

$$\max_{g(x),\kappa(x,u)} \sum_{x\in\mathbb{Q}} g(x) \quad (12a)$$

subject to $\forall x \in \mathbb{Q}$

$$g(x) \geq \sum_{y\in\mathbb{Q}} g(y)T(y|x,u) \quad \forall u \in \mathbb{U}_x \quad (12b)$$

$$g(x) \leq \sum_{y\in\mathbb{Q}} g(y)T(y|x,u) + (1-\kappa(x,u))\Delta \quad \forall u \in \mathbb{U}_x \quad (12c)$$

$$\sum_{u\in\mathbb{U}_x} \kappa(x,u) \geq 1 \quad (12d)$$

$$0 \leq g(x) \leq 1, \kappa(x,u) \in \{0,1\} \quad \forall u \in \mathbb{U}_x \quad (12e)$$

where Δ is a constant greater than one. That is, $G^*_{\infty,\mathbb{Q}}(x) = g^*(x)$, $\forall x \in \mathbb{Q}$, where g^* is the optimal solution of the MILP (12). The optimal stationary Markov policy is $\bar{\mu}^*_{\mathbb{Q}}(x) = u$ where $u \in \mathbb{U}_x$ such that $\kappa^*(x,u) = 1$ and κ^* is the optimal solution of the MILP (12).

Proof: See Appendix E. ■

Computational algorithm: As an adaption of Algorithm 1, the following algorithm provides a way to compute the largest infinite-horizon ϵ-PCIS within $\mathbb{Q}$.

The difference between Algorithms 1 and 3 is that the value of $G^*_{\infty,\mathbb{P}_i}(x)$, instead of $V^*_{0,\mathbb{P}_i}(x)$, $\forall x \in \mathbb{P}_i$, is computed by (12)

Algorithm 3: Infinite-Horizon ϵ-PCIS.

1: Initialize $i = 0$ and $\mathbb{P}_i = \mathbb{Q}$.
2: Compute $G^*_{\infty,\mathbb{P}_i}(x)$ for all $x \in \mathbb{P}_i$.
3: Compute the set $\mathbb{P}_{i+1} = \mathbb{S}^*_{\epsilon,\infty}(\mathbb{P}_i)$.
4: If $\mathbb{P}_{i+1} = \mathbb{P}_i$, stop. Else, set $i = i + 1$ and go to step 2.

Algorithm 4: Infinite-Horizon ϵ-PCIS.

1: Compute the RCIS within $\mathbb{Q}$, denoted by $\mathbb{Q}_f$.
2: Compute the stochastic backward reachable set from $\mathbb{Q}_f$, i.e.,
$\tilde{\mathbb{Q}} = \{x \in \mathbb{Q} \mid \exists u \in \mathbb{U}, \int_{\mathbb{Q}_f} T(dy|x,u) \geq \epsilon\}$.

(replacing $\mathbb{Q}$ with $\mathbb{P}_i$). Furthermore, the updated set $\mathbb{P}_{i+1} = \mathbb{S}^*_{\epsilon,\infty}(\mathbb{P}_i)$, which is a stochastic backward reachable set within $\mathbb{P}_i$ with respect to infinite horizon and a probability level ϵ. The following theorem provides the convergence of $\mathbb{P}_i$ and shows that the resulting set $\tilde{\mathbb{Q}}$ by this algorithm is an infinite-horizon ϵ-PCIS.

Theorem 4: For discrete state and control spaces, Algorithm 3 converges in a finite number of iterations and generates the largest infinite-horizon ϵ-PCIS within $\mathbb{Q}$. Furthermore, at each iteration, the infinite-horizon invariance probability $G^*_{\infty,\mathbb{P}_i}(x)$, $\forall x \in \mathbb{P}_i$, can be computed via the MILP (12).

Proof: The finite-step convergence of Algorithm 3 follows from the finite cardinality of the set $\mathbb{Q}$. Similar to Theorem 1, the generated infinite-horizon ϵ-PCIS is the largest one within $\mathbb{Q}$. The MILP reformulation refers to Lemma 7. ■

Remark 14: When implementing Algorithm 3 to a system with discrete spaces, the maximal iteration number is $|\mathbb{Q}|$. An MILP is used to compute the value of $G^*_{\infty,\mathbb{P}_i}(x), \forall x \in \mathbb{P}_i$, at each iteration. The number of real-valued decision values is at most $|\mathbb{Q}|$, the number of binary decision values is at most $|\mathbb{Q}||\mathbb{U}|$, and the number of constraints is at most $|\mathbb{Q}|(2|\mathbb{U}| + 3)$. In general, MILPs are NP-hard and can be solved by cutting plane algorithm or branch-and-bound algorithm [32]. Some advanced softwares have been developed to solve large MILPs efficiently [33], [34].

2) Continuous State and Control Spaces: If the state and control spaces are continuous, it is computationally intractable to compute the exact value of infinite-horizon invarinace probability $G^*_{\infty,\mathbb{Q}}(x)$. Based on Remark 12, this section provides another way to compute an infinite-horizon ϵ-PCIS within a given set $\mathbb{Q}$.

Different from Algorithm 3, which computes iteratively the stochastic backward reachable sets, the following algorithm generates an infinite-horizon ϵ-PCIS by computing a backward stochastic reachable set from the RCIS $\mathbb{Q}_f$ contained in $\mathbb{Q}$.

The first step in Algorithm 4 is the computation of RCIS within a given set, which is a well-studied topic in the literature [4]–[6]. Then, based on RCIS $\mathbb{Q}_f$ within $\mathbb{Q}$, the stochastic backward reachable set

$$\tilde{\mathbb{Q}} = \left\{ x \in \mathbb{Q} \mid \exists u \in \mathbb{U}, \int_{\mathbb{Q}_f} T(dy|x,u) \geq \epsilon \right\}$$

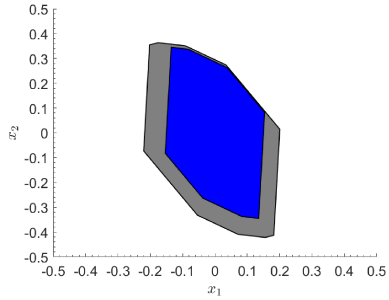

Fig. 1. Computations of the largest RCIS (blue) and an infinite-horizon ϵ-PCIS with $\epsilon = 0.80$ (gray) by Algorithm 4 for Example 1.

is an infinite-horizon ϵ-PCIS within $\mathbb{Q}$. In comparison with Algorithms 1–3, the iteration is avoided in Algorithm 4, which only needs two steps.

Remark 15: Note that the resulting set from Algorithm 4 is in general not the largest infinite-horizon ϵ-PCIS within the given set $\mathbb{Q}$. It is possible to obtain a larger infinite-horizon ϵ-PCIS if we can reformulate the existence conditions in Theorem 3 and Corollary 3 in a recursive form and thereby modify Algorithm 4 to be a recursive algorithm.

Remark 16: The complexity of Algorithm 4 depends on the computation of the RCIS [3]–[6], and the computation of the backward stochastic reachable set. The later can be reformulated as a chance-constrained problem and then approximately solved. Some results on computation of the backward stochastic reachable set have been reported in [35]. The first example in Section V will show how to compute the backward stochastic reachable set.

V. Examples

In this section, two examples are provided to illustrate the effectiveness of the proposed theoretical results. The first one is concerned with comparison between PCIS and RCIS. Then, we consider an application to motion planning of a mobile robot in a partitioned space with obstacles.

A. Example 1: Comparison Between PCIS and RCIS

Consider the following example from [36]:

$$x_{k+1} = Ax_k + Bu_k + w_k$$

where $A = \begin{bmatrix} 1.6 & 1.1 \\ -0.7 & 1.2 \end{bmatrix}$ and $B = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$. The control input is constrained by $|u_k| \leq 0.25$. We consider w_k to be either non-stochastic or stochastic when computing RCIS and PCIS, respectively. The region of interest is $\mathbb{Q} = \{x \in \mathbb{R}^2 \mid \|x\|_\infty \leq 0.5\}$. We will compare the largest RCIS and PCIS within $\mathbb{Q}$.

To derive an RCIS for this system, we assume the disturbance belongs to the compact set $\mathbb{W} = \{w \in \mathbb{R}^2 \mid \|w\|_\infty \leq 0.05\}$. By using the methods in [1] and [6], we obtain the largest RCIS, which is the blue region shown in Fig. 1. The gray region is an infinite-horizon ϵ-PCIS described in the end of this example.

When computing a finite-horzion PCIS, assume that elements of w_k are independent identically distributed (i.i.d.) Gaussian

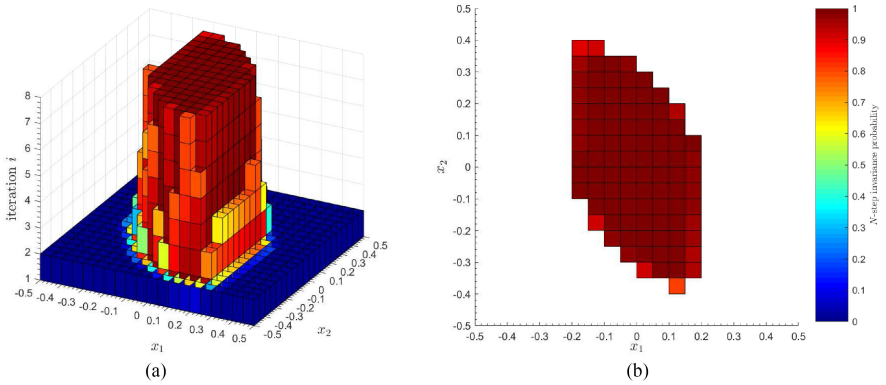

Fig. 2. Computation of N-step ϵ-PCIS with $N = 5$ and $\epsilon = 0.80$ for Example 1. (a) Sets $\mathbb{P}_i$ and the corresponding N-step invariance probability in Algorithm 2. (b) N-step ϵ-PCIS $\tilde{\mathbb{Q}}$.

random variables with zero mean and variance $\sigma^2 = 1/30^2$. This system can be represented as a triple $\mathcal{S} = \{\mathbb{X}, \mathbb{U}, T\}$

$$\begin{cases} \mathbb{X} = \mathbb{R}^2 \\ \mathbb{U} = \{u \in \mathbb{R} \mid |u| \leq 0.1\} \\ t(x_{k+1}|x_k, u_k) = \psi(\Lambda^{-1}(x_{k+1} - Ax_k - Bu_k)) \end{cases}$$

where $\psi(\cdot)$ is the density function of the standard normal distribution and $\Lambda = \text{diag}\{\sigma, \sigma\}$. In this case, since the Lipschitz constant L in Assumption 3 is small, we ignore the approximation error τ_0 in (9). We discretize the continuous spaces and implement Algorithm 2 to compute the N-step ϵ-PCIS $\tilde{\mathbb{Q}}$. First consider $N = 5$ and $\epsilon = 0.80$. Fig. 2(a) shows the evolution of the set $\mathbb{P}_i$ in Algorithm 2. The color indicates the corresponding N-step invariance probability $p^*_{N,\mathbb{P}_i}(x)$ and the z-axes the iteration index i. The algorithm converges in 8 steps. Fig. 2(b) shows $\mathbb{P}_8$, which corresponds to the N-step ϵ-PCIS $\tilde{\mathbb{Q}}$ for $N = 5$ and $\epsilon = 0.80$.

When computing an infinite-horizon PCIS, we choose the same bound on the disturbance as for the RCIS. The elements of w_k are truncated i.i.d. Gaussian random variables with zero mean and variance $\sigma^2 = 1/30^2$. Denote the largest RCIS computed above by $\mathbb{Q}_f = \{x \in \mathbb{R}^2 \mid Hx \leq h\}$, where the matrix H and the vector h are with appropriate dimensions. As stated in Algorithm 4, the one-step stochastic backward reachable set from the RCIS associated with probability 0.80 is an infinite-horizon ϵ-PCIS with $\epsilon = 0.80$, i.e.,

$$\tilde{\mathbb{Q}} = \{x \in \mathbb{Q} \mid \exists u \in \mathbb{U}, \Pr\{H(Ax + Bu + w) \leq h\} \geq 0.80\}.$$

This set can be represented as

$$\tilde{\mathbb{Q}} = \{x \in \mathbb{Q} \mid \exists u \in \mathbb{U}, H(Ax + Bu) + h' \leq h\}$$

where h' is the optimal solution of the chance constrained program

$$\min \sum_j h'_j$$

$$\text{subject to } \Pr\{Hw \leq h'\} = 0.8.$$

This program can be numerically solved by using the methods in [37] and [38]. The resulting infinite-horizon ϵ-PCIS with $\epsilon = 0.80$ is the gray region shown in Fig. 1. This region is obviously a superset of the RCIS in blue.

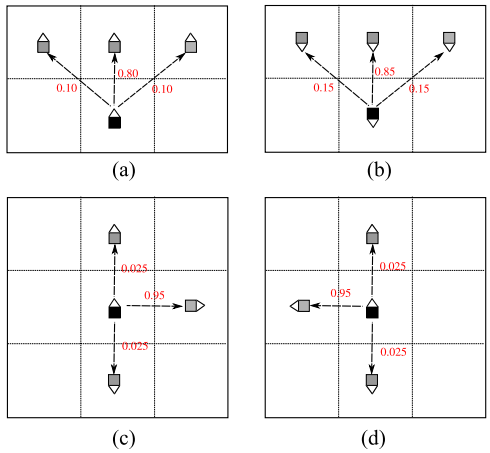

Fig. 3. Transition probability under actions for Example 2.

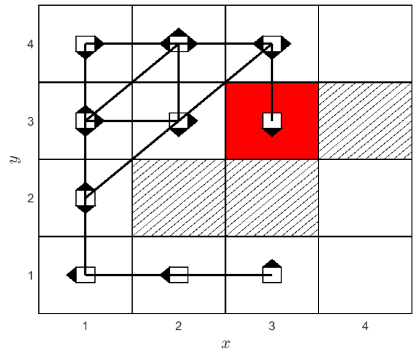

Fig. 4. One simulated state trajectory with indication of the robot orientation starting from $(3, 1, \mathcal{N})$ and ending at $(3, 4, \mathcal{S})$ in Example 2.

B. Example 2: Motion Planning

The motion planning example in [39] is adapted to seek an infinite-horizon PCIS within the workspace for a mobile

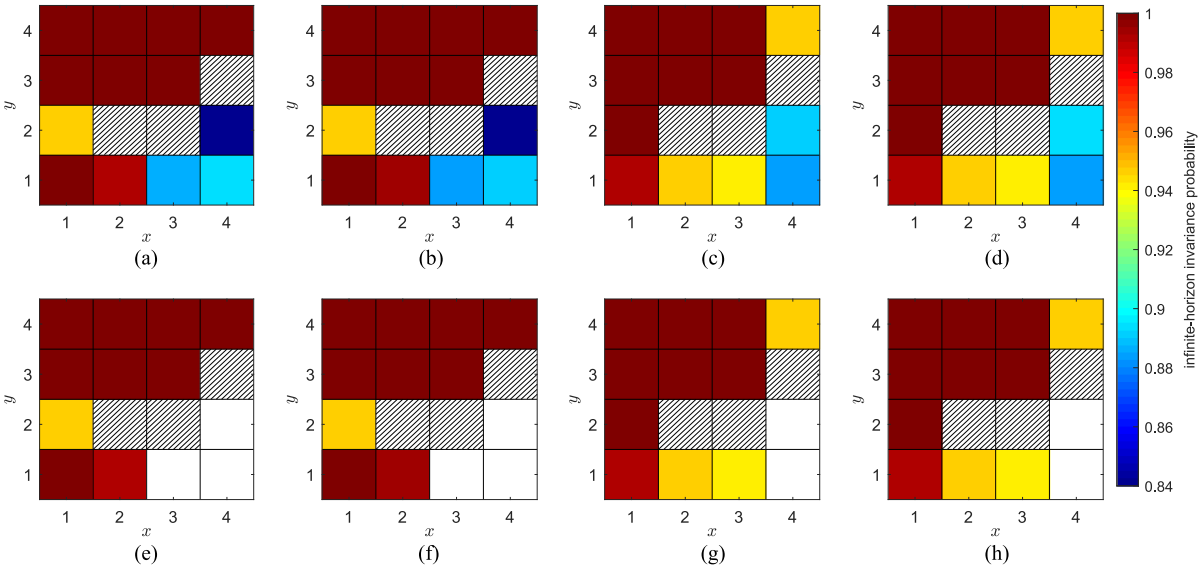

Fig. 5. Sets $\mathbb{P}_i$ and the corresponding infinite-horizon invariance probability in Example 2 when computing the largest infinite-horizon ϵ-PCIS with $\epsilon = 0.90$ by Algorithm 3.

robot. The state of the robot is abstracted by its cell coordinate, i.e., $(p_x, p_y) \in \{1,2,3,4\}^2$, and its four possible orientations $\{\mathcal{E}, \mathcal{W}, \mathcal{S}, \mathcal{N}\}$. Due to the actuation noise and drifting, the robot motion is stochastic. Here, we restrict the action space to be $\{\mathrm{FR}, \mathrm{BK}, \mathrm{TRFR}, \mathrm{TLFR}\}$, under which the possible transitions are shown in Fig. 3. Specifically, action "FR" means driving forward for 1 unit. As illustrated in the figure, the probability for that is 0.80. The probability of drifting forward to the left or the right by 1 unit is 0.10. Action "BK" can be similarly defined. Action "TRFR" means turning right $\pi/2$ and driving forward for 1 unit, of which the probability is 0.95. The probability of driving forward for 1 unit without turning right is 0.025 and the probability of turning right for π and driving forward for 1 unit is 0.025. Similarly, we can define the action "TLFR."

Consider the partitioned workspace shown in Fig. 4, where the shadowed cells are occupied by obstacles and the red cell is an absorbing region, i.e., when the robot enters in this region it will stay there forever. We construct an MDP with 64 states and 4 actions. The transition relation and probability can be defined based on the abovementioned description. We compute the largest infinite-horizon ϵ-PCIS with $\epsilon = 0.90$ within the safe state space, i.e., the remaining of the state space by excluding the states associated with the obstacles.

By implementing Algorithm 3, the computed sets $\mathbb{P}_i$ and the corresponding infinite-horizon invariance probability $p^*_{\infty,\mathbb{P}_i}(x)$ are shown in Fig. 5, of which each subfigure corresponds to one orientation in $\{\mathcal{E}, \mathcal{W}, \mathcal{S}, \mathcal{N}\}$. The first row of Fig. 5 shows the results after the first iteration, where we can see that the infinite-horizon invariance probability $p^*_{\infty,\mathbb{P}_i}(x)$ at $x = (4,2,\mathcal{E})$ and $x = (4,2,\mathcal{W})$ is less than $\epsilon = 0.90$. Algorithm 3 converges in 2 steps and generates the largest infinite-horizon ϵ-PCIS $\tilde{\mathbb{Q}}$ with $\epsilon = 0.90$ shown in Fig. 5(e)–(h). This invariant set provides a region where the admissible action can drive the robot without colliding with the obstacles with probability 0.90. By implementing the optimal policy obtained in Lemma 7, we run a state trajectory starting from $(3,1,\mathcal{N})$ as shown in Fig. 4. We can see that this trajectory is collision-free and finally ends at the absorbing region $(3,3,\mathcal{S})$.

VI. Conclusion

We investigated the extension of set invariance in a stochastic sense for control systems. We proposed finite- and infinite-horizon ϵ-PCISs, and provided some fundamental properties. We designed iterative algorithms to compute the PCIS within a given set. For systems with discrete state and control spaces, finite- and infinite-horizon ϵ-PCISs can be computed by solving an LP and an MILP at each iteration, respectively. We proved that the iterative algorithms were computationally tractable and can be terminated in a finite number of steps. For systems with continuous state and control spaces, we established the approximation of stochastic control systems and proved its convergence when computing finite-horizon ϵ-PCIS. In addition, thanks to the sufficient conditions for the existence of infinite-horizon ϵ-PCIS, we can compute an infinite-horizon ϵ-PCIS by the stochastic backward reachable set from the RCIS contained in it. Numerical examples were given to illustrate the theoretical results.

One future direction is to apply the PCISs to safety-critical control and stochastic predictive control. In particular, how to characterize stability using PCISs is an important problem to consider. Another interesting future extension of PCISs is to study reliability and mean-time-to-failure for general stochastic systems.

Appendix A
Proof of Lemma 3

Define the functions $J^*_{k,\mathbb{Q}} : \mathbb{X} \to \mathbb{R}$, $k \in \mathbb{N}_{[0,N]}$, as

$$J^*_{k,\mathbb{Q}}(x) = -V^*_{N-k,\mathbb{Q}}(x) \quad \forall x \in \mathbb{X}.$$

As shown in [23], the function $J^*_{N,\mathbb{Q}}$ is lower semianalytic for any $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$. From [26, Definitions 7.20 and 7.21], we have that the function $J^*_{N,\mathbb{Q}}$ is also analytically measurable and thus is universally measurable for any $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$. According to the definition of universal measurability, the set $J^{*,-1}_{N,\mathbb{Q}}(\mathbb{B}) = \{x \in \mathbb{X} \mid J^*_{k,\mathbb{Q}}(x) \in \mathbb{B}\}$ for $\mathbb{B} \in \mathcal{B}(\mathbb{R})$ is universally measurable.

Recall the definition of the stochastic backward reachable set $\mathbb{S}^*_{\epsilon,N}(\mathbb{Q})$, we have that

$$\begin{aligned}\mathbb{S}^*_{\epsilon,N}(\mathbb{Q}) &= \{x \in \mathbb{Q} \mid V^*_{0,\mathbb{Q}}(x) \geq \epsilon\} \\ &= \{x \in \mathbb{Q} \mid -1 \leq J^*_{N,\mathbb{Q}}(x) \leq -\epsilon\} \\ &= J^{*,-1}_{N,\mathbb{Q}}(\mathbb{B})\end{aligned}$$

where $\mathbb{B} = [-1, -\epsilon] \in \mathcal{B}(\mathbb{R})$. Thus, the set $\mathbb{S}^*_{\epsilon,N}(\mathbb{Q})$ is universally measurable for any $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$.

APPENDIX B
PROOF OF LEMMA 6

Before proving Lemma 6, we need two auxiliary lemmas. Lemma 8 shows that the value functions in (1) are Lipschitz continuous. It is adapted from [23, Th. 8]. Lemma 9 shows that the difference between the approximate density function and the original density function is bounded.

Lemma 8: Under Assumptions 1 and 3, for any $x, x' \in \mathbb{Q}$, the value functions $V^*_{k,\mathbb{Q}}$ in (1) satisfy

$$|V^*_{k,\mathbb{Q}}(x) - V^*_{k,\mathbb{Q}}(x')| \leq \phi(\mathbb{Q})L\|x - x'\| \quad \forall k \in \mathbb{N}_{[0,N]}. \tag{13}$$

Proof: Similar to [23, Th. 8]. ■

Lemma 9: Under Assumptions 3, for all $y \in \mathbb{Q}$ and $q_i \in \hat{\mathbb{Q}}$

$$\int_{\mathbb{Q}} |\hat{t}(y|q_i, \hat{u}) - t(y|q_i, \hat{u})|dy \leq 2\phi(\mathbb{Q})L\delta \quad \forall \hat{u} \in \hat{\mathbb{U}}.$$

Proof: If $\int_{\mathbb{Q}} t(s_z|s_x, \hat{u})dz < 1$, it follows from Assumption 3 that

$$\int_{\mathbb{Q}} |\hat{t}(y|q_i, \hat{u}) - t(y|q_i, \hat{u})|dy \leq \phi(\mathbb{Q})L\delta.$$

And if $\int_{\mathbb{Q}} t(s_z|s_x, \hat{u})dz \geq 1$, we first have

$$\begin{aligned}0 &\leq \int_{\mathbb{Q}} t(s_y|q_i, \hat{u})dy - 1 \\ &\leq \int_{\mathbb{Q}} t(s_y|q_i, \hat{u})dy - \int_{\mathbb{Q}} t(y|q_i, \hat{u})dy \\ &\leq \int_{\mathbb{Q}} |t(s_y|q_i, \hat{u}) - t(y|q_i, \hat{u})|dy \\ &\leq \phi(\mathbb{Q})L\delta.\end{aligned}$$

Furthermore, we have

$$\begin{aligned}&\int_{\mathbb{Q}} |\hat{t}(y|q_i, \hat{u}) - t(y|q_i, \hat{u})|dy \\ &= \int_{\mathbb{Q}} \frac{|t(s_y|q_i, \hat{u}) - t(y|q_i, \hat{u}) \int_{\mathbb{Q}} t(s_z|s_x, \hat{u})dz|}{\int_{\mathbb{Q}} t(s_z|s_x, \hat{u})dz} dy \\ &\leq \int_{\mathbb{Q}} |t(s_y|q_i, \hat{u}) - t(y|q_i, \hat{u}) \int_{\mathbb{Q}} t(s_z|s_x, \hat{u})dz|dy \\ &\leq \int_{\mathbb{Q}} |t(s_y|q_i, \hat{u}) - t(y|q_i, \hat{u})|\, dy \\ &+ \Big| \int_{\mathbb{Q}} t(s_z|s_x, \hat{u})dz - 1 \Big| \int_{\mathbb{Q}} |t(y|q_i, \hat{u})|dy \\ &\leq 2\phi(\mathbb{Q})L\delta.\end{aligned}$$

■

Proof of Lemma 6: First of all, let us prove inequality (8). It is easy to check it for $k = N$ since $V^*_{N,\mathbb{Q}}(x) = \hat{V}^*_{k,\mathbb{Q}}(x) = 1, \forall x \in \mathbb{Q}$. By induction, we assume that $|V^*_{k+1,\mathbb{Q}}(x) - \hat{V}^*_{k+1,\mathbb{Q}}(x)| \leq \tau_{k+1}(\mathbb{Q})\delta,\ x \in \mathbb{Q}$. For any $q_i \in \mathbb{Q}_i,\ i \in \mathbb{N}_{[1,m_x]}$, we define $\mu^*_k = \arg\sup_{u \in \mathbb{U}} \int_{\mathbb{Q}} V^*_{k+1,\mathbb{Q}}(y)t(y|q_i, u)dy$ and $\hat{\mu}^*_k = \arg\max_{\hat{u} \in \hat{\mathbb{U}}} \int_{\mathbb{Q}} \hat{V}^*_{k+1,\mathbb{Q}}(y)\hat{t}(y|q_i, \hat{u})dy$. According to the discretization procedure of the control space, we can choose some $\hat{\nu}_k \in \hat{\mathbb{U}}$ such that $\|\mu^*_k - \hat{\nu}_k\| \leq \delta$. Then, we have that

$$\begin{aligned}&V^*_{k,\mathbb{Q}}(q_i) - \hat{V}^*_{k,\mathbb{Q}}(q_i) \\ &= \int_{\mathbb{Q}} V^*_{k+1,\mathbb{Q}}(y)t(y|q_i, \mu^*_k)dy - \int_{\mathbb{Q}} \hat{V}^*_{k+1,\mathbb{Q}}(y)\hat{t}(y|q_i, \hat{\mu}^*_k)dy \\ &\leq \int_{\mathbb{Q}} V^*_{k+1,\mathbb{Q}}(y)t(y|q_i, \mu^*_k)dy - \int_{\mathbb{Q}} \hat{V}^*_{k+1,\mathbb{Q}}(y)\hat{t}(y|q_i, \hat{\nu}_k)dy \\ &\leq |\int_{\mathbb{Q}} V^*_{k+1,\mathbb{Q}}(y)t(y|q_i, \mu^*_k)dy - \int_{\mathbb{Q}} V^*_{k+1,\mathbb{Q}}(y)t(y|q_i, \hat{\nu}_k)dy| \\ &\quad + |\int_{\mathbb{Q}} V^*_{k+1,\mathbb{Q}}(y)t(y|q_i, \hat{\nu}_k)dy - \int_{\mathbb{Q}} V^*_{k+1,\mathbb{Q}}(y)\hat{t}(y|q_i, \hat{\nu}_k)dy| \\ &\quad + |\int_{\mathbb{Q}} V^*_{k+1,\mathbb{Q}}(y)\hat{t}(y|q_i, \hat{\nu}_k)dy - \int_{\mathbb{Q}} \hat{V}^*_{k+1,\mathbb{Q}}(y)\hat{t}(y|q_i, \hat{\nu}_k)dy| \\ &\leq \phi(\mathbb{Q})L\delta + 2\phi(\mathbb{Q})L\delta + \tau_{k+1}(\mathbb{Q})\delta \\ &= (3\phi(\mathbb{Q})L + \tau_{k+1}(\mathbb{Q}))\delta\end{aligned}$$

and

$$\begin{aligned}&\hat{V}^*_{k,\mathbb{Q}}(q_i) - V^*_{k,\mathbb{Q}}(q_i) \\ &\leq \int_{\mathbb{Q}} \hat{V}^*_{k+1,\mathbb{Q}}(y)\hat{t}(y|q_i, \hat{\mu}^*_k)dy - \int_{\mathbb{Q}} V^*_{k+1,\mathbb{Q}}(y)t(y|q_i, \hat{\mu}^*_k)dy \\ &\leq |\int_{\mathbb{Q}} \hat{V}^*_{k+1,\mathbb{Q}}(y)\hat{t}(y|q_i, \hat{\mu}^*_k)dy - \int_{\mathbb{Q}} \hat{V}^*_{k+1,\mathbb{Q}}(y)t(y|q_i, \hat{\mu}^*_k)dy| \\ &\quad + |\int_{\mathbb{Q}} \hat{V}^*_{k+1,\mathbb{Q}}(y)t(y|q_i, \hat{\mu}^*_k)dy - \int_{\mathbb{Q}} V^*_{k+1,\mathbb{Q}}(y)t(y|q_i, \hat{\mu}^*_k)dy| \\ &\leq (2\phi(\mathbb{Q})L + \tau_{k+1}(\mathbb{Q}))\delta.\end{aligned}$$

Thus, we have

$$|V^*_{k,\mathbb{Q}}(q_i) - \hat{V}^*_{k,\mathbb{Q}}(q_i)| \leq (3\phi(\mathbb{Q})L + \tau_{k+1}(\mathbb{Q}))\delta.$$

For any $x \in \mathbb{Q}_i,\ i \in \mathbb{N}_{[1,m_x]}$, it follows that

$$\begin{aligned}&|V^*_{k,\mathbb{Q}}(x) - \hat{V}^*_{k,\mathbb{Q}}(x)| \\ &= |V^*_{k,\mathbb{Q}}(x) - \hat{V}^*_{k,\mathbb{Q}}(q_i)|\end{aligned}$$

$$\leq |V^*_{k,\mathbb{Q}}(x) - V^*_{k,\mathbb{Q}}(q_i)| + |V^*_{k,\mathbb{Q}}(q_i) - \hat{V}^*_{k,\mathbb{Q}}(q_i)|$$
$$\leq (4\phi(\mathbb{Q})L + \tau_{k+1}(\mathbb{Q}))\delta = \tau_k(\mathbb{Q})\delta$$

which completes the proof of inequality (8).

Appendix C
Proof of Theorem 3

Let u_x be the control input such that (3) holds for any $x \in \mathbb{Q}$.

Only-if-part: Under Assumption 2, the fact that the set $\mathbb{Q} \in \mathcal{B}(\mathbb{X})$ is an infinite-horizon ϵ-PCIS is equivalent to $G^*_{\infty,\mathbb{Q}}(x) \geq \epsilon, \forall x \in \mathbb{Q}$. Let $\theta = \sup_{x\in\mathbb{Q}} G^*_{\infty,\mathbb{Q}}(x)$. Under Assumption 2, $G^*_{\infty,\mathbb{Q}}(x)$ exists for all $x \in \mathbb{Q}$. The set $\tilde{\mathbb{Q}}_f = \{x \in \mathbb{Q} \mid G^*_{\infty,\mathbb{Q}}(x) = \theta\}$ collects all the states for which the value of $G^*_{\infty,\mathbb{Q}}$ is maximal over the set $\mathbb{Q}$. Extending Lemma 3 to infinite-horizon case, we have that the set $\tilde{\mathbb{Q}}_f$ is universally measurable. By [26, Lemma 7.16], we have that there exists a Borel-measurable set $\mathbb{Q}_f \subseteq \mathbb{Q}$ such that $p(\mathbb{Q}_f \triangle \tilde{\mathbb{Q}}_f) = 0$ for any $p \in \mathcal{P}(\mathbb{X})$.

Next we will show that the set $\mathbb{Q}_f$ is an RCIS. It follows from Assumption 2 and Lemma 2 that $\forall x \in \mathbb{Q}_f$

$$\begin{aligned}
&G^*_{\infty,\mathbb{Q}}(x)\\
&= \int_{\mathbb{Q}_f} G^*_{\infty,\mathbb{Q}}(y)T(dy|x,u_x) + \int_{\mathbb{Q}\setminus\mathbb{Q}_f} G^*_{\infty,\mathbb{Q}}(y)T(dy|x,u_x)\\
&= G^*_{\infty,\mathbb{Q}}(x)\int_{\mathbb{Q}_f} T(dy|x,u_x)\\
&\quad + \int_{\mathbb{Q}\setminus\mathbb{Q}_f} G^*_{\infty,\mathbb{Q}}(y)T(dy|x,u_x) \qquad (14)\\
&\leq G^*_{\infty,\mathbb{Q}}(x)T(\mathbb{Q}_f|x,u_x) + G^*_{\infty,\mathbb{Q}}(x)T(\mathbb{Q}\setminus\mathbb{Q}_f|x,u_x)\\
&= G^*_{\infty,\mathbb{Q}}(x)(T(\mathbb{Q}_f|x,u_x) + T(\mathbb{Q}\setminus\mathbb{Q}_f|x,u_x)) \qquad (15)
\end{aligned}$$

where (14) follows from $G^*_{\infty,\mathbb{Q}}(x) = G^*_{\infty,\mathbb{Q}}(y), \forall x, y \in \mathbb{Q}_f$ and (15) follows from that $G^*_{\infty,\mathbb{Q}}(x) > G^*_{\infty,\mathbb{Q}}(y), \forall x \in Q_f, \forall y \in \mathbb{Q}\setminus\mathbb{Q}_f$. Furthermore, since $G^*_{\infty,\mathbb{Q}}(x) \geq \epsilon > 0, \forall x \in \mathbb{Q}$, and $0 \leq T(\mathbb{Q}|x,u_x) \leq 1$, the equality in (15) holds if and only if $T(\mathbb{Q}_f|x,u_x) = 1$ and thereby $T(\mathbb{Q}\setminus\mathbb{Q}_f|x,u_x)) = 0$. Based on the recursion in (2), we have $G^*_{\infty,\mathbb{Q}}(x) = 1, \forall x \in \mathbb{Q}_f$. Hence, the set $\mathbb{Q}_f \subseteq \mathbb{Q}$ is an RCIS.

Next let us prove that $\forall x \in \mathbb{Q}\setminus\mathbb{Q}_f$, (10) holds. That is to prove that

$$G^*_{\infty,\mathbb{Q}}(x) \leq T(\mathbb{Q}_f|x,u_x) + \int_{\mathbb{Q}\setminus\mathbb{Q}_f} T(\mathbb{Q}_f|y,u_y)T(dy|x,u_x) + \frac{\rho^2}{1-\rho}. \qquad (16)$$

By [23, Th. 7], the control input u_x is also optimal to the recursion (2). For all $k \in \mathbb{N}$, we have $\forall x \in \mathbb{Q}_f$, $G^*_{k,\mathbb{Q}}(x) = 1$ and $\forall x \in \mathbb{Q}\setminus\mathbb{Q}_f$

$$G^*_{k+1,\mathbb{Q}}(x) = T(\mathbb{Q}_f|x,u_x) + \int_{\mathbb{Q}\setminus\mathbb{Q}_f} G^*_{k,\mathbb{Q}}(y)T(dy|x,u_x).$$

Let $\rho = \sup_{x\in\mathbb{Q}\setminus\mathbb{Q}_f} \int_{\mathbb{Q}\setminus\mathbb{Q}_f} T(dy|x,u_x)$. Note that $0 \leq \rho < 1$. Then, $\forall x \in \mathbb{Q}\setminus\mathbb{Q}_f$, we can follow the induction rule to prove that

$$G^*_{k,\mathbb{Q}}(x) \leq T(\mathbb{Q}_f|x,u_x) + \int_{\mathbb{Q}\setminus\mathbb{Q}_f} T(\mathbb{Q}_f|y,u_y)T(dy|x,u_x) + \frac{\rho^2-\rho^k}{1-\rho}$$

which by taking limitation yields that (16) holds.

If-part: The proof for the existence of an RCIS $\mathbb{Q}_f \subseteq \mathbb{Q}$ is the same as that of the only if part. As shown above, the condition $T(\mathbb{Q}_f|x,u_x) = 1$ is equivalent to $G^*_{\infty,\mathbb{Q}}(x) = 1, \forall x \in \mathbb{Q}_f$. We can use induction to prove that $\forall x \in \mathbb{Q}\setminus\mathbb{Q}_f$

$$G^*_{k,\mathbb{Q}}(x) \geq T(\mathbb{Q}_f|x,u_x) + \int_{\mathbb{Q}\setminus\mathbb{Q}_f} T(\mathbb{Q}_f|y,u_y)T(dy|x,u_x)$$

which further implies that $G^*_{\infty,\mathbb{Q}}(x) \geq T(\mathbb{Q}_f|x,u_x) + \int_{\mathbb{Q}\setminus\mathbb{Q}_f} T(\mathbb{Q}_f|y,u_y)T(dy|x,u_x)$. One sufficient condition to guarantee $G^*_{\infty,\mathbb{Q}}(x) \geq \epsilon$ is (11), i.e., $T(\mathbb{Q}_f|x,u_x) + \int_{\mathbb{Q}\setminus\mathbb{Q}_f} T(\mathbb{Q}_f|y,u_y)T(dy|x,u_x) \geq \epsilon$.

Appendix D
Proof of Corollary 3

By Lemma 2 and Theorem 3, the necessary condition in Corollary 3 can be proven by showing that $\forall x \in \mathbb{Q}\setminus\mathbb{Q}_f$, there exists a $u \in \mathbb{U}$ such that

$$\begin{aligned}
\epsilon &\leq G^*_{\infty,\mathbb{Q}}(x)\\
&= \int_{\mathbb{Q}_f} G^*_{\infty,\mathbb{Q}}(y)T(dy|x,u)\\
&\quad + \int_{\mathbb{Q}\setminus\mathbb{Q}_f} G^*_{\infty,\mathbb{Q}}(y)T(dy|x,u)\\
&\leq T(\mathbb{Q}_f|x,u) + T(\mathbb{Q}\setminus\mathbb{Q}_f|x,u)\\
&= T(\mathbb{Q}|x,u) \qquad (17)
\end{aligned}$$

where (17) follows from $0 < G^*_{\infty,\mathbb{Q}}(x) \leq 1, \forall x \in \mathbb{Q}$.

The sufficient condition in Corollary 3 can be proven by showing that $\forall x \in \mathbb{Q}\setminus\mathbb{Q}_f$, there exists a $u \in \mathbb{U}$ such that

$$\begin{aligned}
&G^*_{\infty,\mathbb{Q}}(x)\\
&= \int_{\mathbb{Q}_f} G^*_{\infty,\mathbb{Q}}(y)T(dy|x,u) + \int_{\mathbb{Q}\setminus\mathbb{Q}_f} G^*_{\infty,\mathbb{Q}}(y)T(dy|x,u)\\
&\geq T(\mathbb{Q}_f|x,u) + \epsilon T(\mathbb{Q}\setminus\mathbb{Q}_f|x,u) \qquad (18)
\end{aligned}$$

where (18) follows from $G^*_{\infty,\mathbb{Q}}(x) \geq \epsilon > 0, \forall x \in \mathbb{Q}$. One sufficient condition to guarantee $G^*_{\infty,\mathbb{Q}}(x) \geq \epsilon$ is $T(\mathbb{Q}_f|x,u) + \epsilon T(\mathbb{Q}\setminus\mathbb{Q}_f|x,u) \geq \epsilon$.

Appendix E
Proof of Lemma 7

Before proving Lemma 7, we need the following two lemmas to show that $G^*_{0,\mathbb{Q}}$ is the unique maximal fixed point satisfying (3). As shown in (2) and (3), $G^*_{\infty,\mathbb{Q}}(x)$ is the limitation of $G^*_{k,\mathbb{Q}}$ as $k \to \infty$. For notational convenience, we use $G^*_{k,\mathbb{Q}}$ to

denote the vector form of $G^*_{k,\mathbb{Q}}(x)$, $x \in \mathbb{Q}$. And the optimization problems $\max_{u\in\mathbb{U}_x} \sum_{y\in\mathbb{Q}} G^*_{k,\mathbb{Q}}(y)T(y|x,u), x \in \mathbb{Q}$ are rewritten as $\max_{\boldsymbol{\mu}\in\mathcal{M}} T^{\boldsymbol{\mu}} G^*_{k,\mathbb{Q}}$. The following lemma provides the uniqueness of $G^*_{\infty,\mathbb{Q}}$.

Lemma 10: The sequence $(G^*_{0,\mathbb{Q}}, G^*_{1,\mathbb{Q}}, \ldots)$ converges to a unique fixed point satisfying (3).

Proof: By contradiction, assume that the sequence $(G^*_{0,\mathbb{Q}}, G^*_{1,\mathbb{Q}}, \ldots)$ could converge to two different fixed points satisfying (3), denoted by $G^{1,*}_{\infty,\mathbb{Q}}$ and $G^{2,*}_{\infty,\mathbb{Q}}$. Then, from Lemma 2, we have

$$\begin{aligned} 0 < \|G^{1,*}_{\infty,\mathbb{Q}} - G^{2,*}_{\infty,\mathbb{Q}}\| &\le \|\max_{\boldsymbol{\mu}\in\mathcal{M}} T^{\boldsymbol{\mu}} G^{1,*}_{\infty,\mathbb{Q}} - \max_{\boldsymbol{\mu}\in\mathcal{M}} T^{\boldsymbol{\mu}} G^{2,*}_{\infty,\mathbb{Q}}\| \\ &\le \max_{\boldsymbol{\mu}\in\mathcal{M}} \|T^{\boldsymbol{\mu}}(G^{1,*}_{\infty,\mathbb{Q}} - G^{2,*}_{\infty,\mathbb{Q}})\| \\ &\le \|G^{1,*}_{\infty,\mathbb{Q}} - G^{2,*}_{\infty,\mathbb{Q}}\|. \end{aligned} \tag{19}$$

In (19), the equality holds if and only if for each $x \in \mathbb{Q}$, there exists $u \in \mathbb{U}_x$ such that $\sum_{y\in\mathbb{Q}} T(y|x,u) = 1$. In this case, it is easy to check that $G^*_{\infty,\mathbb{Q}}(x) = G^*_{0,\mathbb{Q}}(x) = 1$ for each $x \in \mathbb{Q}$ so $G^*_{\infty,\mathbb{Q}}$ is unique. For other cases, we have a contradiction. Hence, the sequence $(G^*_{0,\mathbb{Q}}, G^*_{1,\mathbb{Q}}, \ldots)$ converges to a unique fixed point satisfying (3). ■

Lemma 11: The convergence point $G^*_{\infty,\mathbb{Q}}$ of the sequence $(G^*_{0,\mathbb{Q}}, G^*_{1,\mathbb{Q}}, \ldots)$ is the maximum fixed point satisfying (3).

Proof: The monotone decrease of the sequence $(G^*_{0,\mathbb{Q}}, G^*_{1,\mathbb{Q}}, \ldots)$ and the unique convergence point imply that $G^*_{\infty,\mathbb{Q}}$ is the maximum fixed point satisfying (3). ■

Proof of Lemma 7: From Lemmas 10 and 11, $G^*_{\infty,\mathbb{Q}}$ is the maximum fixed point satisfying (3). Hence, the equivalent form of $G^*_{\infty,\mathbb{Q}}$ can be written as MILP (12), where the constraints (12b)–(12d) guarantee that there exists $u \in \mathbb{U}_x$ such that the equality in (3) holds.

Acknowledgment

The authors would like to thank Prof. A. Abate for helpful discussions and feedback and to anonymous reviewers for their constructive comments.

References

[1] D. Bertsekas, "Infinite time reachability of state-space regions by using feedback control," *IEEE Trans. Autom. Control*, vol. 17, no. 5, pp. 604–613, Oct. 1972.
[2] F. Blanchini, "Set invariance in control," *Automatica*, vol. 35, no. 11, pp. 1747–1767, 1999.
[3] F. Blanchini and S. Miani, *Set-Theoretic Methods in Control*. Berlin, Germany: Springer, 2007.
[4] S. V. Raković, E. C. Kerrigan, K. I. Kouramas, and D. Q. Mayne, "Invariant approximations of the minimal robust positively invariant set," *IEEE Trans. Autom. Control*, vol. 50, no. 3, pp. 406–410, Mar. 2005.
[5] M. Rungger and P. Tabuada, "Computing robust controlled invariant sets of linear systems," *IEEE Trans. Autom. Control*, vol. 62, no. 7, pp. 3665–3670, Jul. 2017.
[6] E. Gilbert and K. T. Tan, "Linear systems with state and control constraints: The theory and practice of maximal admissible sets," *IEEE Trans. Autom. Control*, vol. 36, no. 9, pp. 1008–1020, Sep. 1991.
[7] I. M. Mitchell, S. Kaynama, M. Chen, and M. Oishi, "Safety preserving control synthesis for sampled data systems," *Nonlinear Anal.: Hybrid Syst.*, vol. 10, pp. 63–82, 2013.
[8] A. Mesbah, "Stochastic model predictive control: An overview and perspectives for future research," *IEEE Control Syst.*, vol. 36, no. 6, pp. 30–44, Dec. 2016.
[9] M. Cannon, B. Kouvaritakis, S. Raković, and Q. Cheng, "Stochastic tubes in model predictive control with probabilistic constraints," *IEEE Trans. Autom. Control*, vol. 56, no. 1, pp. 194–200, Jan. 2011.
[10] M. A. Hernández-Mejías, A. Sala, C. Ariño, and A. Querol, "Reliable controllable sets for constrained Markov-jump linear systems," *Int. J. Robust Nonlinear Control*, vol. 10, no. 26, pp. 2075–2089, 2016.
[11] M. A. Hernández-Mejías and A. Sala, "Reliability and timetofailure bounds for discretetime constrained Markov jump linear systems," *Int. J. Robust Nonlinear Control*, vol. 10, no. 27, pp. 1773–1791, 2017.
[12] J. Hu, M. Prandini, and S. Sastry, "Aircraft conflict prediction in the presence of a spatially correlated wind field," *IEEE Trans. Intell. Transp. Syst.*, vol. 6, no. 3, pp. 326–340, Sep. 2005.
[13] J. Ding, M. Kamgarpour, S. Summers, A. Abate, J. Lygeros, and C. Tomlin, "A stochastic games framework for verification and control of discrete time stochastic hybrid systems," *Automatica*, vol. 49, no. 9, pp. 2665–2674, 2013.
[14] J. Burlet, O. Aycard, and T. Fraichard, "Robust motion planning using Markov decision processes and quadtree decomposition," in *Proc. IEEE Conf. Robot. Autom.*, 2004, pp. 2820–2825.
[15] G. Pola, J. Lygeros, and M. D. Di Benedetto, "Invariance in stochastic dynamical control systems," in *Proc. 17th Int. Symp. Math. Theory Netw. Syst.*, 2006.
[16] M. Cannon, B. Kouvaritakis, and X. Wu, "Probabilistic constrained MPC for multiplicative and additive stochastic uncertainty," *IEEE Trans. Autom. Control*, vol. 54, no. 7, pp. 1626–1632, Jul. 2009.
[17] E. Kofman, J. A. De Doná, and M. M. Seron, "Probabilistic set invariance and ultimate boundedness," *Automatica*, vol. 48, no. 10, pp. 2670–2676, 2012.
[18] S. Battilotti and A. De Santis, "Stabilization in probability of nonlinear stochastic systems with guaranteed region of attraction and target set," *IEEE Trans. Autom. Control*, vol. 48, no. 9, pp. 1585–1599, Sep. 2003.
[19] G. Pola and G. Pola, "A stochastic reachability approach to portfolio construction in finance industry," *IEEE Trans. Control Syst. Technol.*, vol. 20, no. 1, pp. 189–195, Jan. 2012.
[20] N. A. Nguyen, "Stochastic output feedback control: convex lifting approach," *Automatica*, vol. 20, no. 1, pp. 212–220, 2018.
[21] E. Kofman, J. A. De Doná, M. M. Seron, and N. Pizzi, "Continuous-time probabilistic ultimate bounds and invariant sets: Computation and assignment," *Automatica*, vol. 71, pp. 98–105, 2016.
[22] L. Hewing, A. Carron, K. Wabersich, and M. Zeilinger, "On a correspondence between probabilistic and robust invariant sets for linear systems," in *Proc. Eur. Control Conf.*, 2018, pp. 876–881.
[23] A. Abate, "Probabilistic reachability for stochastic hybrid systems: Theory, computations, and applications," Ph.D. dissertation, Dept. Elect. Eng. Comput. Sci., University of California, Berkeley, Berkeley, CA, USA, 2007.
[24] S. Amin, A. Abate, M. Prandini, S. Sastry, and J. Lygeros, "Reachability analysis for controlled discrete time stochastic hybrid systems," in *Proc. Int. Workshop Hybrid Syst.: Comput. Control*, 2014, pp. 49–63.
[25] C. S. Chow and J. N. Tsitsiklis, "An optimal one-way multigrid algorithm for discrete-time stochastic control," *IEEE Trans. Autom. Control*, vol. 36, no. 8, pp. 898–914, Aug. 1991.
[26] D. P. Bertsekas and S. Shreve, *Stochastic Optimal Control: The Discrete-Time Case*. Belmont, MA, USA: Athena Scientific, 2004.
[27] M. B. Stinchcombe and H. White, "Some measurability results for extrema of random functions over random sets," *Rev. Econ. Stud.*, vol. 59, no. 3, pp. 495–514, 1992.
[28] A. Bhattacharya and J. P. Kharoufeh, "Linear programming formulation for non-stationary, finite-horizon Markov decision process models," *Oper. Res. Lett.*, vol. 45, no. 6, pp. 570–574, 2017.
[29] N. Megiddo, "Linear programming in linear time when the dimension is fixed," *J. ACM*, vol. 31, no. 1, pp. 114–127, 1984.
[30] I. Tkachev and Abate,"On infinite-horizon probabilistic properties and stochastic bisimulation functions," in *Proc. 50th IEEE Conf. Decision Control Eur. Control Conf.*, 2011, pp. 526–531.
[31] D. Bertsekas, *Dynamic Programming and Optimal Control: Vol II*. Belmont, MA, USA: Athena Scientific, 2012.
[32] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*. Chelmsford, MA, USA: Courier Corporation, 1998.
[33] J. T. Linderoth and A. Lodi, "MILP software, ",*Wiley Encyclopedia of Operations Research and Management Science*. Hoboken, NJ, USA: Wiley, 2010.
[34] P. Bonami, M. Kilinç, and J. Linderoth, "Algorithms and software for convex mixed integer nonlinear programs," in *Mixed Integer Nonlinear Programming*, J. Lee and S. Leyffer, Eds. Berlin, Germany: Springer, 2012, pp. 1–39.

[35] M. Prandini and J. Hu, "A stochastic approximation method for reachability computations," in *Stochastic Hybrid Systems*, H. A. Blom and J. Lygeros, Eds. Berlin, Germany: Springer, 2006, pp. 107–139.
[36] B. Kouvaritakis, M. Cannon, S. Raković, and Q. Cheng, "Explicit use of probabilistic distributions in linear predictive control," *Automatica*, vol. 46, no. 10, pp. 1719–1724, 2010.
[37] M. Lorenzen, F. Dabbene, R. Tempo, and F. Allgöwer, "Constraint-tightening and stability in stochastic model predictive control," *IEEE Trans. Autom. Control*, vol. 62, no. 7, pp. 3165–3177, Jul. 2017.
[38] A. Prékopa, *Stochastic Programming*. Berlin, Germany: Springer, 2013.
[39] M. Guo and M. M. Zavlanos, "Probabilistic motion planning under temporal tasks and soft constraints," *IEEE Trans. Autom. Control*, vol. 63, no. 12, pp. 4051–4066, Dec. 2018.

Yulong Gao received the B.E. degree in automation and the M.E. degree in control science and engineering from Beijing Institute of Technology, Beijing, China, in 2013 and 2016, respectively. He is currently working toward the Ph.D. degree under KTH-NTU joint Ph.D. program with the Division of Decision and Control Systems, KTH Royal Institute of Technology, Stockholm, Sweden.

He was a Visiting Student with the Department of Computer Science, University of Oxford, Oxford, U.K., in 2019. His research interests include automatic verification, stochastic control and model predictive control with application to safety-critical systems.

Karl Henrik Johansson (Fellow, IEEE) received the M.Sc. degree in electrical engineering and Ph.D. degree in automatic control from Lund University, Lund, Sweden, in 1992 and 1997, respectively.

He is currently a Professor with the School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, Stockholm, Sweden, and the Director of Digital Futures. He has held visiting positions with UC Berkeley, Caltech, NTU, HKUST Institute of Advanced Studies, and NTNU.

His research interests are in networked control systems and cyber-physical systems with applications in transportation, energy, and automation networks.

Dr. Johansson is a member of the Swedish Research Council's Scientific Council for Natural Sciences and Engineering Sciences. He has served on the IEEE Control Systems Society Board of Governors, the IFAC Executive Board, and is currently Vice-President of the European Control Association. He was the recipient of several best paper awards and other distinctions from IEEE, IFAC, and ACM. He has been awarded Distinguished Professor with the Swedish Research Council and Wallenberg Scholar with the Knut and Alice Wallenberg Foundation and the Future Research Leader Award from the Swedish Foundation for Strategic Research and the triennial Young Author Prize from IFAC. He is Fellow of the Royal Swedish Academy of Engineering Sciences, and he is IEEE Control Systems Society Distinguished Lecturer.

Lihua Xie (Fellow, IEEE) received the B.E. and M.E. degrees in electrical engineering from the Nanjing University of Science and Technology, Nanjing, China, in 1983 and 1986, respectively, and the Ph.D. degree in electrical engineering from the University of Newcastle, Callaghan, NSW, Australia, in 1992.

Since 1992, he has been with the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore, where he is currently a Professor and the Director, Delta-NTU Corporate Laboratory for Cyber-Physical Systems. He served as the Head of Division of Control and Instrumentation from 2011 to 2014. He held teaching appointments with the Department of Automatic Control, Nanjing University of Science and Technology from 1986 to 1989. His research interests include robust control and estimation, networked control systems, multiagent networks, localization and unmanned systems.

Dr Xie is an Editor-in-Chief for *unmanned systems* and an Associate Editor for IEEE TRANSACTIONS ON NETWORK CONTROL SYSTEMS. He has served as an Editor of IET book series in control and an Associate Editor for a number of journals including IEEE TRANSACTIONS ON AUTOMATIC CONTROL, *Automatica*, IEEE TRANSACTIONS ON CONTROL SYSTEMS TECHNOLOGY, and IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS-II. He was an IEEE Distinguished Lecturer (2012– 2014) and an Elected Member of Board of Governors, IEEE Control System Society (2016–2018). He is a Fellow of IFAC, and a Fellow of Academy of Engineering Singapore.