Scalable Underapproximative Verification of Stochastic LTI Systems using Convexity and Compactness*

Abraham P. Vinod†
University of New Mexico
Albuquerque, New Mexico, USA
aby.vinod@gmail.com

Meeko M. K. Oishi‡
University of New Mexico
Albuquerque, New Mexico, USA
oishi@unm.edu

ABSTRACT

We present a scalable algorithm to construct a polytopic underapproximation of the terminal hitting time stochastic reach-avoid set, for the verification of high-dimensional stochastic LTI systems with arbitrary stochastic disturbance. We prove the existence of a polytopic underapproximation by characterizing the sufficient conditions under which the stochastic reach-avoid set and the proposed open-loop underapproximation are compact and convex. We construct the polytopic underapproximation by formulating and solving a series of convex optimization problems. These set-theoretic properties also characterize circumstances under which the stochastic reach-avoid problem admits a bang-bang optimal Markov policy. We demonstrate the scalability of our algorithm on a 40D chain of integrators, the highest dimensional example demonstrated to date for stochastic reach-avoid problems, and compare its performance with existing approaches on a spacecraft rendezvous and docking problem.

CCS CONCEPTS

• **Theory of computation** → **Stochastic control and optimization**; *Convex optimization*; • **Computing methodologies** → *Control methods*; Computational control theory;

KEYWORDS

Stochastic reachability; Convex optimization; Stochastic optimal control

ACM Reference Format:
Abraham P. Vinod and Meeko M. K. Oishi. 2018. Scalable Underapproximative Verification of Stochastic LTI Systems using Convexity and Compactness. In *HSCC '18: 21st International Conference on Hybrid Systems: Computation and Control (part of CPS Week), April 11–13, 2018, Porto, Portugal*. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3178126.3178148

1 INTRODUCTION

Reachability analysis of discrete-time stochastic dynamical systems is an established verification tool that provides probabilistic guarantees of safety or performance, and has been applied a wide range of applications [1–6]. In [1], two classes of problems characterize verification over a finite horizon — first hitting time and terminal hitting time – and dynamic programming approaches are formulated to solve both (similarly to [7, 8]). We focus on the finite time horizon *terminal* hitting time stochastic reach-avoid problem (referred to here as the *terminal time problem* for convenience), that is, computing the maximum probability of hitting a target set at the terminal time, while avoiding an unsafe set during all the preceding time steps using a state-feedback policy. Specifically, we analyze the convexity and compactness properties of the corresponding *stochastic reach-avoid set*, the set of initial states where its optimal terminal time (reach-avoid) probability is greater than a given threshold. We use these properties to construct scalable polytopic representations. These properties also determine the sufficient conditions for admittance of a bang-bang optimal control.

The dynamic programming-based discretization approach (DPBDA), proposed in [8], approximately computes value functions for the terminal time problem via gridding, and hence suffers from the well-known *curse of dimensionality*. Attempts to circumvent this problem, via approximate dynamic programming [9, 10], Gaussian mixtures [9], particle filters [5, 10], convex chance-constrained optimization [5], and semi-definite programming [11] have been applied to systems that are at most 10-dimensional – far beyond the scope of what is possible with DPBDA, but not scalable to larger problems. Recently, a scalable Fourier transform-based underapproximation of the stochastic reach-avoid problem was proposed for stochastic linear systems with Gaussian perturbations, and demonstrated on a 40-dimensional system [12]. This method focused on underapproximating the stochastic reach-avoid probability when starting from a known initial condition. In this paper, we build on the method [12] to compute a polytopic underapproximation of the stochastic

[*] This material is based upon work supported by the National Science Foundation under Grant Number CMMI-1254990 (Oishi, CAREER) and CNS-1329878. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

[†] A. P. Vinod and M. Oishi are with Electrical and Computer Engineering, University of New Mexico, Albuquerque, NM 87131 USA.

[‡] Corresponding author

HSCC '18, April 11–13, 2018, Porto, Portugal
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-5642-8/18/04…$15.00
https://doi.org/10.1145/3178126.3178148

reach-avoid set. Using the proposed method, we compute the underapproximation of the stochastic reach-avoid set for a 40-dimensional system, which to our best knowledge, is the largest system that has been verified via the stochastic reach-avoid formulation.

We first characterize sufficient conditions for closedness, compactness, and convexity of stochastic reach-avoid sets and the open-loop underapproximation [12]. We then propose a scalable algorithm to construct a tight polytopic underapproximation of the latter. Our approach does not require gridding of state, input, or disturbance spaces, and has low memory requirements. We compare our approach with the underapproximation technique proposed using Lagrangian methods in [6] and the convex chance-constrained optimization technique in [5]. Using the convexity results, we also show that the *stochastic avoid problem*, computing the minimum probability of hitting a target set at the terminal time irrespective of control, admits a bang-bang solution, potentially enabling more efficient computation strategies.

Our main contributions are as follows: 1) sufficient conditions for the closedness, compactness, and convexity of the stochastic reach-avoid set and the underapproximation in [12], 2) a scalable algorithm for the computation of a polytopic underapproximation of the stochastic reach-avoid set, and 3) the sufficient conditions for the admittance of bang-bang optimal policy in stochastic reach-avoid problems.

In Section 2, we describe the terminal time problem, its open-loop approximation, and relevant properties from probability theory and real analysis. Section 3 presents sufficient conditions for compactness, closedness, and convexity, and discusses how they provide sufficient conditions for bang-bang controllers in stochastic reach-avoid problems. Section 4 presents a scalable algorithm for the computation of the polytopic underapproximation, and we discuss our recommendations for computational complexity, implementation, and the synthesis of probabilistically safe open-loop controllers. We demonstrate the proposed algorithm on several numerical examples in Section 5. We conclude and provide directions for future work in Section 6.

2 PRELIMINARIES AND PROBLEM FORMULATION

We denote the Borel σ-algebra by $\mathscr{B}(\cdot)$, a discrete-time time interval which inclusively enumerates all integers in between a and b for $a, b \in \mathbb{N}$ and $a \leq b$ by $\mathbb{N}_{[a,b]}$, random vectors with bold case, and non-random vectors with an overline. The indicator function of a non-empty set $\mathcal{E}$ is denoted by $1_{\mathcal{E}}(\overline{y})$, such that $1_{\mathcal{E}}(\overline{y}) = 1$ if $\overline{y} \in \mathcal{E}$ and is zero otherwise. We denote the p-dimensional identity matrix by I_p, and the matrix with all entries as $\overline{x} \in \mathbb{R}$ by $\overline{x}_{p\times q} \in \mathbb{R}^{p\times q}$. We denote the affine and convex hull of a set $\mathcal{E}$ by affineHull($\mathcal{E}$) and convexHull($\mathcal{E}$) respectively.

2.1 Real analysis

The relative interior of a set $\mathcal{E} \subseteq \mathbb{R}^n$ is defined as

$$\mathrm{relint}(\mathcal{E}) = \{\overline{x} \in \mathbb{R}^n : \exists r > 0, \mathrm{Ball}(\overline{x}, r) \cap \mathrm{affineHull}(\mathcal{E}) \subseteq \mathcal{E}\}$$

where $\mathrm{Ball}(\overline{x}, r)$ denotes a ball in $\mathbb{R}^n$ centered at $\overline{x}$ and of radius r with respect to any Euclidean norm [13, Sec. 2.1.3]. The relative boundary is $\partial\mathcal{E} = \mathrm{closure}(\mathcal{E}) \setminus \mathrm{relint}(\mathcal{E})$. From the Heine-Borel theorem [14, Thm 12.5.7], we know that $\mathcal{E}$ is compact if and only if it is closed and bounded.

A point $\overline{x} \in \mathcal{E}$ is an *extreme point* of the set $\mathcal{E}$ if and only if the only way to express $\overline{x}$ as a convex combination $(1-\theta)\overline{y} + \theta\overline{z}$, such that $\overline{y} \in \mathcal{E}$, $\overline{z} \in \mathcal{E}$, and $0 < \theta < 1$, is by taking $\overline{y} = \overline{z} = \overline{x}$ [15, Sec. 18]. Given $\mathcal{E}_{\mathrm{extreme}}$ as the set of all the extreme points of a compact and convex set $\mathcal{E}$, we have $\mathrm{convexHull}(\partial\mathcal{E}) = \mathrm{convexHull}(\mathcal{E}_{\mathrm{extreme}}) = \mathcal{E}$ [15, Thm. 18.4, Corr. 18.5.1]. For a compact $\mathcal{E}$, $\mathcal{E}_{\mathrm{extreme}} \subseteq \partial\mathcal{E}$ [15, Corr. 19.1.1]. Depending on $\mathcal{E}$, the set $\mathcal{E}_{\mathrm{extreme}}$ could be countable, for *e.g.* vertices of a polytope, or uncountable, for *e.g.* the boundary of an ellipsoid.

A function $f : \mathbb{R}^n \to \mathbb{R}$ is upper semi-continuous (u.s.c.) if its superlevel sets $\{\overline{x} \in \mathbb{R}^n : f(\overline{x}) \geq \alpha\}$ for some $\alpha \in \mathbb{R}$ are closed [16, Defs. 2.3 and 2.8]. A function $f : \mathbb{R}^n \to \mathbb{R}$ is log-concave if $f(\overline{x}) \geq 0$ for all $\overline{x}$ and $\log f$ is concave with $\log 0 \triangleq -\infty$ [13, Sec. 3.5.1]. Many standard distributions are log-concave, for example, Gaussian, uniform, and exponential [13, Eg. 3.40]. From these definitions, we also see that the indicator function of a closed set is u.s.c., and the indicator function of a convex set is log-concave (See [13, Eg. 3.1 and Sec. 3.1.7]).

2.2 Probability theory

A random vector $\boldsymbol{y}$ is a measurable transformation defined in the probability space $(\Omega, \mathscr{Y}, \mathbb{P})$ with sample space Ω, σ-algebra $\mathscr{Y}$, and probability measure $\mathbb{P}$ over $\mathscr{Y}$. We consider Borel-measurable random vectors, $\boldsymbol{y} : \mathbb{R}^p \to \mathbb{R}^p$ with $\Omega = \mathbb{R}^p$ and $\mathscr{Y} = \sigma(\boldsymbol{y}) = \mathscr{B}(\mathbb{R}^p)$. For $N \in \mathbb{N}$, a random process is a sequence of random vectors $\{\boldsymbol{y}_k\}_{k=0}^{N}$ where the random vectors $\boldsymbol{y}_k$ are defined in the probability space $(\Omega, \mathscr{Y}, \mathbb{P})$. The random vector $\boldsymbol{Y} = [\boldsymbol{y}_0\ \boldsymbol{y}_1\ \ldots\ \boldsymbol{y}_N]^\top$ is defined in the probability space $(\Omega^{N+1}, \sigma(\times_{k=0}^{N} \mathscr{Y}_k), \mathbb{P}_{\boldsymbol{Y}})$, with $\mathbb{P}_{\boldsymbol{Y}}$ induced from $\mathbb{P}$. See [17, 18] for details.

2.3 Terminal stochastic reach-avoid analysis

Consider the discrete-time stochastic LTI system,

$$\boldsymbol{x}_{k+1} = A\boldsymbol{x}_k + B\overline{u}_k + \boldsymbol{w}_k \tag{1}$$

with state $\boldsymbol{x}_k \in \mathcal{X} = \mathbb{R}^n$, input $\overline{u}_k \in \mathcal{U} \subseteq \mathbb{R}^m$, disturbance $\boldsymbol{w}_k \in \mathcal{W} \subseteq \mathbb{R}^n$, and matrices A, B assumed to be of appropriate dimensions. We assume that $\mathcal{U}$ is compact, $\boldsymbol{w}_k$ is absolutely continuous with a known probability density function (PDF) $\psi_{\boldsymbol{w}}$, and the random process $\boldsymbol{w}[\cdot]$ is independent and identically distributed (IID). Let N be a finite time horizon. For any given sequence of (non-random) inputs $\overline{u}[\cdot]$ and an initial condition $\overline{x}_0 \in \mathcal{X}$, the state $\boldsymbol{x}_k$ is a random vector for all $k \in \mathbb{N}_{[1,N]}$ via (1).

The system (1) can be equivalently described by a Markov control process with stochastic kernel that is a Borel-measurable function $Q : \mathscr{B}(\mathcal{X}) \times \mathcal{X} \times \mathcal{U} \to [0,1]$, which assigns to each $\overline{x} \in \mathcal{X}$ and $\overline{u} \in \mathcal{U}$, a probability measure on the Borel space $(\mathcal{X}, \mathscr{B}(\mathcal{X}))$. For (1),

$$Q(d\overline{y}|\overline{x},\overline{u}) = \psi_{\boldsymbol{w}}(\overline{y} - A\overline{x} - B\overline{u})d\overline{y}. \tag{2}$$

We define a *Markov policy* $\pi = (\mu_0, \mu_1, \ldots, \mu_{N-1}) \in \mathcal{M}$ as a sequence of universally measurable maps $\mu[\cdot] : \mathcal{X} \to \mathcal{U}$. The random vector $\boldsymbol{X} = [\boldsymbol{x}_1^\top\ \boldsymbol{x}_2^\top\ \ldots\ \boldsymbol{x}_N^\top]^\top$, defined in $(\mathcal{X}^N, \mathscr{B}(\mathcal{X}^N), \mathbb{P}_{\boldsymbol{X}}^{\overline{x}_0,\pi})$ [1], has probability measure $\mathbb{P}_{\boldsymbol{X}}^{\overline{x}_0,\pi}$ defined using Q [19, Prop. 7.45].

Let $\mathcal{S}, \mathcal{T} \in \mathscr{B}(\mathcal{X})$. Define the *terminal time probability*, $\hat{r}_{\overline{x}_0}^{\pi}(\mathcal{S},\mathcal{T})$, for known $\overline{x}_0$ and π, as the probability that the execution with policy π is inside the target set $\mathcal{T}$ at time N and stays within the safe set $\mathcal{S}$ for all time up to N. From [1],

$$\hat{r}_{\overline{x}_0}^{\pi}(\mathcal{S},\mathcal{T}) = \mathbb{P}_{\boldsymbol{X}}^{\overline{x}_0,\pi}\left\{\boldsymbol{x}_N \in \mathcal{T} \wedge \boldsymbol{x}_k \in \mathcal{S}\ \forall k \in \mathbb{N}_{[0,N-1]}\right\}. \tag{3}$$

From [1, Def. 10], a Markov policy π^* is a *maximal reach-avoid policy in the terminal sense* if and only if it is the optimal solution of the problem

$$\hat{r}_{\overline{x}_0}^{\pi^*}(\mathcal{S},\mathcal{T}) = \sup_{\pi \in \mathcal{M}} \hat{r}_{\overline{x}_0}^{\pi}(\mathcal{S},\mathcal{T}). \tag{4}$$

The solution of (4) is characterized via dynamic programming [1, Thm. 11]. Define $\hat{V}_k^* : \mathcal{X} \to [0,1],\ k \in \mathbb{N}_{[0,N]}$, by the backward recursion for $\overline{x} \in \mathcal{X}$,

$$\hat{V}_N^*(\overline{x}) = 1_{\mathcal{T}}(\overline{x}) \tag{5a}$$

$$\hat{V}_k^*(\overline{x}) = \sup_{\overline{u} \in \mathcal{U}} 1_{\mathcal{S}}(\overline{x}) \int_{\mathcal{X}} \hat{V}_{k+1}^*(\overline{y}) Q(d\overline{y}|\overline{x},\overline{u}). \tag{5b}$$

Then, the optimal value of (4) is $\hat{r}_{\overline{x}_0}^{\pi^*}(\mathcal{S},\mathcal{T}) = \hat{V}_0^*(\overline{x}_0)$ for every $\overline{x}_0 \in \mathcal{X}$. A sufficient condition for the existence of an optimal Markov policy was first given in [1, Thm. 11]. Lemma 1 provides another (stricter) sufficient condition that is relatively easier to ensure.

Lemma 1. *[12, Thm. 1] If $\mathcal{U}$ is compact, $\mathcal{X}$, $\mathcal{S}$, and $\mathcal{T}$ are Borel, and $Q(\cdot|\overline{x},\overline{u})$ is continuous, then an optimal Markov policy π^* exists to solve* (4).

Lemma 2. *[8, Thm. 2] If $\mathcal{S}$, $\mathcal{U}$, and $\mathcal{T}$ are compact and $Q(\cdot|\overline{x},\overline{u})$ is Lipschitz, then $\hat{V}_k^*(\cdot)$ is Lipschitz over $\mathcal{S}$ for $k \in \mathbb{Z}_{[0,N-1]}$.*

While Lemma 2 allows for the implementation of (5) by discretizing $\mathcal{X}, \mathcal{U}$, and $\mathcal{W}$ [1], this approach suffers from the curse of dimensionality.

For $\alpha \in [0,1]$, we define the stochastic reach-avoid set $\mathcal{L}^{\pi^*}(\alpha,\mathcal{S},\mathcal{T})$ as the set of states which satisfy the terminal hitting time stochastic reach-avoid objective with a probability greater than or equal to α under the optimal Markov policy π^*. Formally, $\mathcal{L}^{\pi^*}(\alpha,\mathcal{S},\mathcal{T})$ is the α-super level set of $\hat{r}_{\overline{x}_0}^{\pi^*}(\mathcal{S},\mathcal{T})$,

$$\mathcal{L}^{\pi^*}(\alpha,\mathcal{S},\mathcal{T}) = \{\overline{x}_0 \in \mathcal{X} : \hat{r}_{\overline{x}_0}^{\pi^*}(\mathcal{S},\mathcal{T}) \geq \alpha\}. \tag{6}$$

2.4 Open-loop underapproximation

In [5, 12], the authors proposed an underapproximation to the stochastic reach-avoid problem by restricting the search for the optimal control policy to open-loop control policies. An open-loop policy $\rho : \mathcal{X} \to \mathcal{U}^N$ provides an open-loop sequence of inputs $\rho(\overline{x}_0) = [\overline{u}_0^\top\ \overline{u}_1^\top\ \ldots\ \overline{u}_{N-1}^\top]^\top$ for every initial condition $\overline{x}_0$. Then the random vector describing the extended state $\boldsymbol{X}$, under the action of $\rho(\overline{x}_0)$, lies in the probability space $(\mathcal{X}^N, \mathscr{B}(\mathcal{X}^N), \mathbb{P}_{\boldsymbol{X}}^{\overline{x}_0,\rho})$, with $\mathbb{P}_{\boldsymbol{X}}^{\overline{x}_0,\rho}$ defined using Q [19, Prop. 7.45].

Similarly to the value function $\hat{V}_0^*(\cdot)$, we define $\hat{W}_0^*(\overline{x}_0) : \mathcal{X} \to [0,1]$ as the maximum terminal time reach-avoid probability attained by evolving (1) from $\overline{y}$, when restricted to open-loop controllers. Denoting the optimal open-loop controller as ρ^*, we define $\hat{W}_0^*(\overline{x}_0)$ for every $\overline{x}_0 \in \mathcal{X}$ through the following optimization problem defined using $\mathbb{P}_{\boldsymbol{X}}^{\overline{x}_0,\rho}$,

$$\hat{W}_0^*(\overline{x}_0) \triangleq \hat{r}_{\overline{x}_0}^{\rho^*}(\mathcal{S},\mathcal{T}) = \sup_{\rho(\overline{x}_0) \in \mathcal{U}^N} \hat{r}_{\overline{x}_0}^{\rho}(\mathcal{S},\mathcal{T}), \tag{7}$$

$$\hat{r}_{\overline{x}_0}^{\rho}(\mathcal{S},\mathcal{T}) \triangleq \mathbb{P}_{\boldsymbol{X}}^{\overline{x}_0,\rho}\left\{\boldsymbol{x}_N \in \mathcal{T} \wedge \boldsymbol{x}_k \in \mathcal{S}\ \forall k \in \mathbb{N}_{[0,N-1]}\right\}. \tag{8}$$

The probability measure $\mathbb{P}_{\boldsymbol{X}}^{\overline{x}_0,\rho}$ is linked to the *forward stochastic reach probability measure* [2] and can be computed for arbitrary disturbances using Fourier transforms [12]. The difference between (3) and (8) is the input policy considered. Equation (8) can also be expressed as

$$\hat{r}_{\overline{x}_0}^{\rho}(\mathcal{S},\mathcal{T}) = \int_{\mathcal{S}^{N-1} \times \mathcal{T}} \psi_{\boldsymbol{X}}(\bar{X}; \bar{x}_0, \rho) d\bar{X} \tag{9}$$

where $\bar{X} = [\bar{x}_1^\top\ \bar{x}_2^\top\ \ldots\ \bar{x}_N^\top]^\top \in \mathcal{X}^N$, $\bar{x}_k \in \mathcal{X}\ \forall k \in \mathbb{N}_{[1,N]}$, $d\bar{X}$ is short for $d\bar{x}_1 d\bar{x}_2 \ldots d\bar{x}_N$, and $\psi_{\boldsymbol{X}}(\cdot; \bar{x}_0, \rho)$ is the PDF induced from $\mathbb{P}_{\boldsymbol{X}}^{\overline{x}_0,\rho}$. From (9), $\hat{r}_{\overline{x}_0}^{\rho}(\mathcal{S},\mathcal{T})$ is a nN-dimensional integral over the set $\mathcal{S}^{N-1} \times \mathcal{T} \subseteq \mathcal{X}^N$ of the joint PDF of the extended state $\boldsymbol{X}$. Similarly to (6), we define the α-superlevel set of $\hat{r}_{\overline{x}_0}^{\rho^*}(\mathcal{S},\mathcal{T})$ as $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$,

$$\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T}) = \{\overline{x}_0 \in \mathcal{X} : \hat{r}_{\overline{x}_0}^{\rho^*}(\mathcal{S},\mathcal{T}) \geq \alpha\}. \tag{10}$$

The set $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$ is an underapproximation of the stochastic reach-avoid set since (7) underapproximates (4) [12, Thm. 2]. Consequently, we have Lemma 3.

Lemma 3. *$\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T}) \subseteq \mathcal{L}^{\pi^*}(\alpha,\mathcal{S},\mathcal{T})$ for all $\alpha \in [0,1]$.*

For a Gaussian $\boldsymbol{w}$ and polytopic $\mathcal{S}$ and $\mathcal{T}$, the objective $\hat{r}_{\overline{x}_0}^{\rho}(\mathcal{S},\mathcal{T})$ is the integration of a Gaussian random vector over a polytope. We have closed-form expressions for $\psi_{\boldsymbol{X}}$ using Fourier transforms. Capitalizing on the log-concavity of (7) [12, Prop. 2], existing scalable algorithms to compute the integration of multivariate Gaussian distributions over polytopes [20], and Lemma 3, an underapproximative method was proposed in [12] to verify high-dimensional stochastic linear systems with affine Gaussian disturbance. While restricting the search for optimal policies to only open-loop controllers in (7) does introduce conservativeness as described in Lemma 3, we will demonstrate in this paper that the computational effort, especially the memory requirements, for computing a tight underapproximation of $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$ is significantly lower than $\mathcal{L}^{\pi^*}(\alpha,\mathcal{S},\mathcal{T})$.

2.5 Problem statement

We use properties of the stochastic reach-avoid set $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$ and its open-loop underapproximation $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$, as well as the resulting optimal policies, π^* and ρ^*, to develop a scalable algorithm to compute a tight polytopic underapproximation of $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$ and therefore of $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$.

Problem 1. *Characterize the sufficient conditions under which the stochastic reach-avoid set $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$ and its open-loop underapproximation $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$ are 1) compact and convex, and 2) closed and convex.*

Problem 1.a. *Characterize the sufficient conditions for the admittance of bang-bang control policies as the optimal control policies π^* and ρ^* that solve* (4) *and* (7) *respectively.*

Problem 2. *Construct a scalable algorithm to compute a tight polytopic underapproximation of the open-loop underapproximation of the stochastic reach-avoid set $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$* (10).

3 COMPACTNESS AND CONVEXITY OF STOCHASTIC REACH-AVOID SETS

From [19, Defn. 7.12], [21, App. B], the stochastic kernel $Q(\cdot|\bar{x}, \bar{u})$ is said to be continuous if for every sequence $(\bar{x}_i, \bar{u}_i) \xrightarrow{i\to\infty} (\bar{x}, \bar{u}) \in \mathcal{X} \times \mathcal{U}$ and every bounded Borel-measurable function $f(\cdot)$ over $\mathcal{X}$,

$$\lim_{i\to\infty} \int_{\mathcal{X}} f(\bar{y}) Q(d\bar{y}|(\bar{x}_i, \bar{u}_i)) = \int_{\mathcal{X}} f(\bar{y}) Q(d\bar{y}|(\bar{x}, \bar{u})).$$

Continuous $\psi_{\boldsymbol{w}}$ yield continuous $Q(\cdot|\overline{x}, \overline{u})$ [12, Lem. 2]. Since affine transformations preserve log-concavity [13, Sec. 3.2.2], we have Lemma 4.

Lemma 4. *If $\psi_{\boldsymbol{w}}$ is a log-concave PDF, then $Q(\cdot|\overline{x}, \overline{u})$ defined in* (2) *is log-concave over $\mathcal{X} \times \mathcal{X} \times \mathcal{U}$.*

To exploit the set-theoretic properties of $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$, we use an alternative representation of (9) derived from the Markov process description (1) and the stochastic kernel (2),

$$\hat{r}^{\rho}_{\bar{x}_0}(\mathcal{S}, \mathcal{T}) = \int_{\mathcal{S}^{N-1}\times\mathcal{T}} \prod_{k=0}^{N-1} Q(\bar{x}_{k+1}; \bar{x}_k, \bar{u}_k) d\bar{X}. \tag{11}$$

The proofs for Section 3.1 and 3.2 are in the appendix.

3.1 Sufficient conditions for compactness

Proposition 1. *$\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T}) \subseteq \mathcal{S}$ for all $\alpha \in (0, 1]$ if $\mathcal{S}, \mathcal{T} \subseteq \mathcal{X}$ and $N \geq 1$.*

By Proposition 1, bounded $\mathcal{S}$ implies bounded $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$ for $\alpha \in (0, 1]$. Moreover, Lipschitz continuity implies u.s.c., and thereby, closedness of the superlevel sets. Therefore, Lemma 2 provides sufficient conditions to guarantee compact $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$. However, we demonstrate compactness for $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$ under weaker requirements on $\mathcal{T}$ and $Q(\cdot|\overline{x}, \overline{u})$.

Theorem 1. *If $\mathcal{U}$ is compact, $\mathcal{S}$ and $\mathcal{T}$ are closed, and $Q(\cdot|\overline{x}, \overline{u})$ is continuous, then 1) $\hat{V}^*_k(\cdot)$ and $\hat{W}^*_k(\cdot)$ are u.s.c. over $\mathcal{S}$ for $k \in \mathbb{N}_{[0,N-1]}$, and 2) $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$ and $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$ are closed for $\alpha \in (0, 1]$.*

The sufficient conditions in Theorem 1 also satisfy the requirements from Lemma 1 for the existence of an optimal Markov policy.

Theorem 2. *If $\mathcal{U}$ and $\mathcal{S}$ are compact, $\mathcal{T}$ is closed, $Q(\cdot|\overline{x}, \overline{u})$ is continuous, and $N \geq 1$, then $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$ and $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$ are compact for $\alpha \in (0, 1]$.*

Unlike Lemma 2, Theorem 2 guarantees that $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$ is compact, even when $\mathcal{T}$ is unbounded (and therefore is not compact), *e.g.*, when the target set is specified as a half-space. The boundedness of $\mathcal{S}$ is the only additional requirement in Theorem 2, as compared to Theorem 1. We show that this can not be weakened by a counterexample. Consider a two-dimensional point mass system,

$$\begin{bmatrix} x_{k+1} \\ y_{k+1} \end{bmatrix} = \begin{bmatrix} x_k \\ y_k \end{bmatrix} + \overline{u} + \boldsymbol{w} \tag{12}$$

with the state $[x_k\ y_k]^\top \in \mathcal{X} = \mathbb{R}^2$, input $\overline{u} \in [-1, 1]^2$, and a stochastic disturbance $\boldsymbol{w} \sim \mathcal{N}(\overline{0}_{2\times1}, I_2)$. If $\mathcal{S} = \mathcal{X}$ (closed but unbounded) and $\mathcal{T} = [-1, 1] \times \mathbb{R}$, then the line $x = 0$ (the y-axis) is contained in $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$ whenever $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T}) \neq \emptyset$. This follows from the observation that the terminal time problem is independent of the state y_k. Since $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$ is not bounded, $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$ is not compact.

3.2 Sufficient conditions for convexity

Proposition 2. *If $\mathcal{S}$ and $\mathcal{T}$ are convex, and $Q(\cdot|\overline{x}, \overline{u})$ is log-concave, then $\hat{r}^{\rho}_{\overline{x}_0}(\mathcal{S}, \mathcal{T})$ is log-concave over $\mathcal{X} \times \mathcal{U}^N$.*

Note that while closed sets (and therefore, compact sets) are Borel-measurable, convex sets need not be Borel, e.g., the union of an open ball in $\mathbb{R}^n$ with any non-measurable subset of its boundary. Hence, the sufficient conditions must satisfy conditions for convexity as well as the requirements in Lemma 1, to ensure the existence of an optimal Markov control policy for (4).

Theorem 3. *If $\mathcal{U}$ is compact and convex, $\mathcal{S}$ and $\mathcal{T}$ are Borel and convex, and $Q(\cdot|\overline{x}, \overline{u})$ is continuous and log-concave, then 1) $\hat{V}^*_k(\cdot)$ and $\hat{W}^*_k(\cdot)$ are log-concave over $\mathcal{X}$ for $k \in \mathbb{N}_{[0,N]}$, and 2) $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$ and $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$ are convex for $\alpha \in (0, 1]$.*

We combine Theorems 1, 2, and 3 in Theorem 4, which solves Problem 1.

Theorem 4. *If $\mathcal{U}$ is compact and convex, $\mathcal{S}$ and $\mathcal{T}$ are closed and convex, and $\mathcal{Q}(\cdot|\overline{x}, \overline{u})$ is continuous and log-concave, then $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$ and $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$ are closed and convex for $\alpha \in$* $(0, 1]$. *In addition, if $\mathcal{S}$ is bounded (making $\mathcal{S}$ compact) and $N \geq 1$, then $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$ and $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$ are compact and convex for $\alpha \in (0, 1]$.*

3.3 Existence of bang-bang optimal control

We now address Problem 1.a. While it is established that the differential game formulation of the reach-avoid problem for uncertain systems leads to a bang-bang controller [22, Prop. 1], no formal results similarly exist for the stochastic reach-avoid problem. We use sufficient conditions for the log-concavity of the value functions (Theorem 3) to characterize the sufficient conditions for bang-bang optimal Markov control policies in the stochastic reach-avoid problem.

Proposition 3. *Let $\mathcal{U}$ be a polytope with a finite set of vertices $\mathcal{U}_{\mathrm{vertices}} \subseteq \mathbb{R}^m$, such that $\mathcal{U} = \mathrm{convexHull}(\mathcal{U}_{\mathrm{vertices}})$. If $\mathcal{T}$ is convex and Borel, and $Q(\cdot|\overline{x},\overline{u})$ is continuous and log-concave, then the stochastic reach-avoid problem* (4)*, with target set $\mathcal{X}\backslash\mathcal{T}$ and safe set as $\mathcal{X}$, and its underapproximation* (7) *admit bang-bang solutions, π^*_{bang} and ρ^*_{bang}, respectively.*

Proof: We first prove that (4) admits a bang-bang optimal Markov policy. Lemma 1 ensures the existence of a Markov policy for (4) under the given conditions. From (3) and (4), for a known $\overline{x}_0 \in \mathcal{X}$,

$$\begin{aligned}\sup_{\pi\in\mathcal{M}} \hat{r}^{\pi}_{\overline{x}_0}(\mathcal{X},\mathcal{X}\setminus\mathcal{T}) &= \sup_{\pi\in\mathcal{M}} \mathbb{P}^{\overline{x}_0,\pi}_{\boldsymbol{X}}\left\{\boldsymbol{x}_N \in \mathcal{X}\setminus\mathcal{T}\right\}\\ &= \sup_{\pi\in\mathcal{M}} \left(1 - \mathbb{P}^{\overline{x}_0,\pi}_{\boldsymbol{X}}\left\{\boldsymbol{x}_N \in \mathcal{T}\right\}\right)\\ &= 1 - \inf_{\pi\in\mathcal{M}} \mathbb{P}^{\overline{x}_0,\pi}_{\boldsymbol{X}}\left\{\boldsymbol{x}_N \in \mathcal{T}\right\}\\ &= 1 - \inf_{\pi\in\mathcal{M}} \hat{r}^{\pi}(\mathcal{X},\mathcal{T}). \end{aligned} \tag{13}$$

To solve the optimization problem in (13), we formulate the recursion of the value functions $\hat{V}^*_{k,\mathrm{bang}}(\cdot)$ for $k \in \mathbb{N}_{[0,N]}$ that solve (13) (similarly to (5b)),

$$\hat{V}^*_{N,\mathrm{bang}}(\overline{x}) = 1_{\mathcal{T}}(\overline{x}) \tag{14}$$

$$\hat{V}^*_{k,\mathrm{bang}}(\overline{x}) = \inf_{\overline{u}\in\mathcal{U}} \int_{\mathcal{X}} \hat{V}^*_{k+1,\mathrm{bang}}(\overline{y}) Q(d\overline{y}|\overline{x},\overline{u}). \tag{15}$$

As in Theorem 3, we can show that $\hat{V}^*_{k,\mathrm{bang}}(\overline{x})$ is log-concave for $k \in \mathbb{N}_{[0,N]}$. From [15, Thm. 32.2], we see that (15) admits an optimal $\bar{u}^* \in \mathcal{U}_{\mathrm{vertices}}$ for π^*_{bang} to (13).

For the open-loop underapproximation (7), we use similar arguments and Proposition 2 for ρ^*_{bang}. ■

For the terminal time problem described in Proposition 3, the admittance of bang-bang optimal control policies implies that discretization of the input space for DPBDA that includes all the extreme points of $\mathcal{U}$ will not introduce any approximation. Note that if instead of a polytopic $\mathcal{U}$, we had a generic compact and convex input set, then π^*_{avoid} would still be bang-bang with the optimal inputs lying in the corresponding (possibly uncountable) set of extreme points.

For the stochastic avoid problem in e.g., obstacle avoidance scenarios [3, 4], a proof can be constructed similarly to that in Proposition 3, to characterize sufficient conditions that admit bang-bang optimal solutions.

4 POLYTOPIC UNDERAPPROXIMATION OF COMPACT, CONVEX $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$

Since the convex hull of the relative boundary points of a compact and convex set is equal to the set itself, we can obtain an arbitrarily tight polytopic underapproximation of $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$ by using a finite subset of the relative boundary points of $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$ as the vertices of the polytopic underapproximation. We denote this polytope by $\underline{\mathcal{K}}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T},\mathcal{D})$. Note that when these vertices include all of the extreme points of $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$ (possible for a polytopic $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$), we have $\underline{\mathcal{K}}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T},\mathcal{D}) = \mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$ [23, Thm. 2.6.16]. We propose an algorithm to compute $\underline{\mathcal{K}}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T},\mathcal{D})$ in three steps: 1) check whether the set is $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$ is non-empty, 2) compute a finite subset of relative boundary points of the set $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$, and 3) compute the convex hull of the computed relative boundary points to obtain $\underline{\mathcal{K}}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T},\mathcal{D})$.

4.1 Computation of $\underline{\mathcal{K}}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T},\mathcal{D})$

4.1.1 Checking if $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$ is non-empty. We first compute a point guaranteed to lie in $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$. Recall that $\hat{W}^*_0 : \mathcal{X} \to [0,1]$ is the optimal value function of (7), and it maps initial states to an optimal reach-avoid probability that can be attained using open-loop controllers. Let $\overline{x}_{\max} \in \mathcal{X}$ be the initial condition that has the highest reach-avoid probability among all the initial states. For every $\alpha \in [0,1]$, it then follows from (10) that either $\overline{x}_{\max} \in \mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$ or $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T}) = \emptyset$, as in Figure 1. We solve

$$\begin{array}{ll} \underset{\overline{x}_0,\rho(\overline{x}_0)}{\mathrm{maximize}} & \hat{r}^{\rho}_{\overline{x}_0}(\mathcal{S},\mathcal{T}) \\ \text{subject to} & \left\{\begin{array}{rcl} \hat{r}^{\rho}_{\overline{x}_0}(\mathcal{S},\mathcal{T}) & \geq & \alpha \\ \rho(\overline{x}_0) & \in & \mathcal{U}^N \\ \overline{x}_0 & \in & \mathcal{X} \end{array}\right. \end{array} \tag{16}$$

to obtain $\overline{x}_{\max}$ with the guarantee that $\hat{W}^*_0(\overline{x}_{\max}) \geq \alpha$. In contrast to (7), in which $\overline{x}_0$ is known, (16) treats $\hat{r}^{\rho}_{\overline{x}_0}(\mathcal{S},\mathcal{T})$ as a function of both $\overline{x}_0$ and $\rho(\overline{x}_0)$. By Proposition 2, (16) is a log-concave optimization problem. Note that $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$ can be empty even if $\mathcal{L}^{\pi^*}(\alpha,\mathcal{S},\mathcal{T})$ is non-empty, by Lemma 3.

4.1.2 Computing a relative boundary point of $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$. From here on, we assume $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T}) \neq \emptyset$ for a given value of $\alpha \in [0,1]$, and $\overline{x}_{\max}$ exists. We define $g : [0,\infty) \to [0,1]$ to evaluate $\hat{W}^*_0(\cdot)$ along a vector $\bar{d} \in \mathcal{X}$ (see Figure 2),

$$g(\theta;\overline{x}_{\max},\bar{d},\hat{W}^*_0) = \hat{W}^*_0(\overline{x}_{\max} + \theta\bar{d}). \tag{17}$$

For any $\theta \in [0,\infty)$, we can compute (17) by solving (7) with $\overline{x}_0 = \overline{x}_{\max} + \theta\bar{d}$. Consider the following optimization problem,

$$\begin{array}{ll} \underset{\theta}{\mathrm{maximize}} & \theta \\ \text{subject to} & \left\{\begin{array}{r} g(\theta;\overline{x}_{\max},\bar{d},\hat{W}^*_0) \geq \alpha \\ \theta \geq 0 \end{array}\right. \end{array} \tag{18}$$

Let the optimal solution to (18) be θ^*.

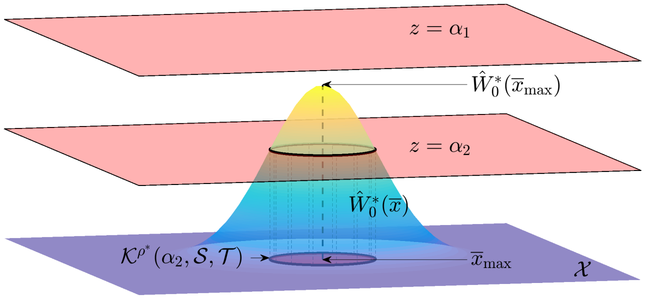
$z = \alpha_1$
$\hat{W}_0^*(\overline{x}_{\max})$
$z = \alpha_2$
$\hat{W}_0^*(\overline{x})$
$\mathcal{K}^{\rho^*}(\alpha_2, \mathcal{S}, \mathcal{T})$
$\overline{x}_{\max}$
$\mathcal{X}$

Figure 1: Feasibility of (16) guarantees non-empty $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$. For $0 \leq \alpha_2 < \hat{W}_0^*(\overline{x}_{\max}) < \alpha_1 \leq 1$, $\mathcal{K}^{\rho^*}(\alpha_1, \mathcal{S}, \mathcal{T}) = \emptyset$ and $\mathcal{K}^{\rho^*}(\alpha_2, \mathcal{S}, \mathcal{T}) \neq \emptyset$.

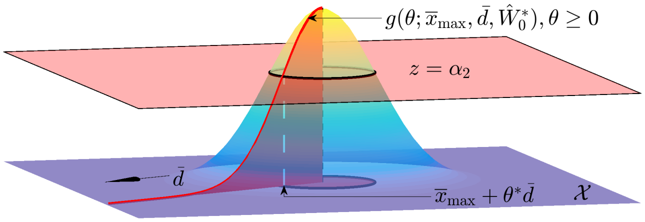
$g(\theta; \overline{x}_{\max}, \bar{d}, \hat{W}_0^*), \theta \geq 0$
$z = \alpha_2$
$\bar{d}$
$\overline{x}_{\max} + \theta^* \bar{d}$
$\mathcal{X}$

Figure 2: Evaluation of $\hat{W}_0^*(\overline{x})$ along a line $\bar{d}$ and the optimization problem (18).

Proposition 4. *For a compact and convex $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$, 1) $g(\theta; \overline{x}_{\max}, \bar{d}, \hat{W}_0^*)$ is a quasi-concave non-increasing function in θ, and 2) $\overline{x}_{\max} + \theta^* \bar{d}$ is a relative boundary point with $\theta^* < \infty$.*

Proof: From (10), $\hat{W}_0^*(\cdot)$ is quasi-concave since $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$ is convex. Further, restricting a quasi-concave function on $\mathbb{R}^n$ to a line intersecting its domain is also quasi-concave [13, Sec. 3.4.2]. Hence, from (17), $g(\theta; \overline{x}_{\max}, \bar{d}, \hat{W}_0^*)$ is a quasi-concave function. Recall the quasi-concave functions over $\mathbb{R}$ are either 1) nondecreasing, 2) non-increasing, or 3) non-increasing up to a point $\theta_0 \in [0, \infty)$ and then non-decreasing [13, Sec. 3.4.2]. From (17) and definition of $\overline{x}_{\max}$, it follows that $g(\theta; \overline{x}_{\max}, \bar{d}, \hat{W}_0^*)$ is non-increasing in θ.

The feasible set of (18) is given by

$$\begin{aligned}\mathcal{E}_{\text{feas}} &= \{\theta \in \mathbb{R} : \theta \geq 0,\ g(\theta; \overline{x}_{\max}, \bar{d}, \hat{W}_0^*) \geq \alpha\}. \\ &= \{\theta : \theta \geq 0\} \cap \{\theta : \overline{x}_{\max} + \theta \bar{d} \in \mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})\}.\end{aligned}$$

Since $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$ is compact and convex, $\mathcal{E}_{\text{feas}}$ is a compact and convex set. Hence, $\mathcal{E}_{\text{feas}} = [0, \phi]$ for some $\phi \in [0, \infty)$ and $\theta^* \leq \phi$. Further, we know that $\overline{x}_{\max} + \phi \bar{d} \in \partial \mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$ since for every $r > 0$, there is some $\nu > 0$ such that $\text{Ball}(\overline{x}_{\max} + \phi \bar{d}, r) \cap \text{affineHull}(\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T}))$ contains a relative interior point $(\overline{x}_{\max} + (\phi - \nu)\bar{d})$ and a point outside $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$ $(\overline{x}_{\max} + (\phi + \nu)\bar{d})$. Finally, assume for contradiction, $\phi \neq \theta^*$. Then $\theta^* < \phi$ which contradicts the optimality assumption of θ^*. ■

From Proposition 4, $g(\theta^*; \overline{x}_{\max}, \bar{d}, \hat{W}_0^*) = \alpha$. Hence, for any $\theta_1, \theta_2 \in [0, \infty)$ with $\theta_1 < \theta^* < \theta_2$, we have

$$g(\theta_1; \overline{x}_{\max}, \bar{d}, \hat{W}_0^*) \geq \alpha \geq g(\theta_2; \overline{x}_{\max}, \bar{d}, \hat{W}_0^*). \tag{19}$$

From Lemma 3 and Proposition 1, we see that $\theta_{\max}$, defined such that $\overline{x}_{\max} + \theta_{\max}\bar{d}$ lies on the relative boundary of $\mathcal{S}$, bounds θ^* from above. We solve (18) to a desired tolerance $\epsilon > 0$ using a bisection algorithm, and (19) on the interval $[0, \theta_{\max}]$ (See [13, Algo. 4.1]). At every iteration, we solve (7) for the given $\overline{x}_0 = \overline{x}_{\max} + \theta \bar{d}$ to obtain $\rho^*(\overline{x}_0)$ and $\hat{W}_0^*(\overline{x}_0)$.

4.1.3 Construction of the polytopic underapproximation. We denote a given finite set of distinct direction vectors $\bar{d}_i$ as $\mathcal{D}$, and let θ_i^* for $i \in \mathbb{N}_{[1, |\mathcal{D}|]}$ denote the corresponding optimal solution to (18). We define, $\underline{\mathcal{K}}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T}, \mathcal{D})$, the polytopic underapproximation of $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$, in (20).

Algorithm 1 solves Problem 2.

Algorithm 1 Computation of $\underline{\mathcal{K}}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T}, \mathcal{D})$

```
Input: Open-loop cost function $\hat{r}^{\rho}_{\overline{x}_0}(\mathcal{S}, \mathcal{T})$, input space $\mathcal{U}$,
    set of direction vectors $\mathcal{D}$, tolerance $\epsilon > 0$
Output: Polytope $\underline{\mathcal{K}}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T}, \mathcal{D})$
Solve (16) to compute $\overline{x}_{\max}$
if (16) is feasible then                                ▷ $\overline{x}_{\max}$ exists
    for $\bar{d}_i \in \mathcal{D}$ do
        Compute $\theta_i^*$ via bisection to solve (18)
    end for
    Compute $\underline{\mathcal{K}}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T}, \mathcal{D})$ using (20)
else
    $\underline{\mathcal{K}}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T}, \mathcal{D}) \leftarrow \emptyset$ by (20)
end if
```

4.2 Discussion

4.2.1 Time and memory complexity. The key reason for Algorithm 1's scalability to high-dimensional systems is that its memory complexity is $\mathcal{O}(|\mathcal{D}|)$, in contrast to the exponential dependence of the memory requirement of DPBDA on n due to gridding [8]. This drastic reduction is the direct consequence of Algorithm 1's ability to exploit the convexity and compactness properties presented in Section 3.

We will characterize the time complexity in terms of the number of evaluations of $\hat{r}^{\rho}_{\overline{x}_0}(\mathcal{S}, \mathcal{T})$ required, the most expensive computation operation in Algorithm 1. Denoting γ_{xmax} and γ_{line} as upper bounds on the number of evaluations of $\hat{r}^{\rho}_{\overline{x}_0}(\mathcal{S}, \mathcal{T})$ in (16) and (18) respectively, and t_{intg} as the computational time to evaluate $\hat{r}^{\rho}_{\overline{x}_0}(\mathcal{S}, \mathcal{T})$, we have the worst-case computational complexity of Algorithm 1 as $\mathcal{O}\left(t_{\text{intg}}\left(\gamma_{\text{xmax}} + \gamma_{\text{line}}|\mathcal{D}|\lceil \log_2(\theta_{\max}) - \log_2(\epsilon)\rceil\right)\right)$. Clearly, γ_{xmax} and γ_{line} depend on the optimization algorithm used, the dimensions of the state space n, the input space m, and the disturbance space p, the time horizon N, the tolerances required, and other problem characteristics. In our experiments, we found γ_{xmax} and γ_{line} to be around 300–800 for the 2D system analyzed in Section 5.1.1. As expected, t_{intg}

$$\underline{\mathcal{K}}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T},\mathcal{D}) = \begin{cases} \text{convexHull}(\{\overline{x}_{\max} + \theta_i^* \bar{d}_i : i \in \mathbb{N}_{[1,|\mathcal{D}|]}\}) & \overline{x}_{\max} \text{ exists} \\ \emptyset & \text{otherwise} \end{cases}. \tag{20}$$

is heavily influenced by the problem, and in our experiments, we found it in the range of 0.02–1 second.

4.2.2 Implementation of Algorithm 1. The choice of $\mathcal{D}$ influences the performance of Algorithm 1 significantly. Choosing the vectors in $\mathcal{D}$ to be far apart improves the quality of underapproximations (in terms of volume). Increasing $|\mathcal{D}|$ will also yield better quality underapproximations, at the cost of increased computational time.

The computation of $\hat{r}^{\rho}_{\overline{x}_0}(\mathcal{S},\mathcal{T})$ requires a numerical integration of dimension nN. We use Genz's algorithm [20, 24] to evaluate $\hat{r}^{\rho}_{\overline{x}_0}(\mathcal{S},\mathcal{T})$ for a Gaussian $\boldsymbol{w}$. However, due to the quasi-Monte Carlo simulations used in Genz's algorithm, the evaluation of $\hat{r}^{\rho}_{\overline{x}_0}(\mathcal{S},\mathcal{T})$ will be noisy. We set the number of particles to ensure an accuracy of 0.01 for every evaluation of $\hat{r}^{\rho}_{\overline{x}_0}(\mathcal{S},\mathcal{T})$. As in [12], direct optimization methods [25] are recommended instead of traditional gradient-based methods. Even when the objective function is known to be convex (Theorem 3), gradient-based methods (such as MATLAB's *fmincon*) may converge to a suboptimal point due to the perturbations introduced by the noisy evaluation of $\hat{r}^{\rho}_{\overline{x}_0}(\mathcal{S},\mathcal{T})$. We use MATLAB's *patternsearch* here. The trade-off for the increased accuracy in the solution is the larger number of function evaluations required by *patternsearch* as compared to *fmincon*. However, the solutions will be sensitive to the initial guesses to solve (7) and (16).

The most significant computational challenge in Algorithm 1 stems from the noisy evaluation of $\hat{r}^{\rho}_{\overline{x}_0}(\mathcal{S},\mathcal{T})$. While our current implementation uses direct optimization methods, which tend to be slower, stochastic optimization techniques such as the stochastic branch-and-bound method, may improve computational speed and accuracy. Algorithm 1 is highly parallelizable since the computations along each of the direction $\bar{d}_i$ are independent. Additionally, γ_{line} and γ_{xmax} may be reduced by storing the evaluations of $\hat{r}^{\rho}_{\overline{x}_0}(\mathcal{S},\mathcal{T})$ in memory.

4.2.3 Open-loop controller synthesis. Algorithm 1 provides probabilistically safe open-loop controllers only for the extreme points of $\underline{\mathcal{K}}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T},\mathcal{D})$ as identified by $\mathcal{D}$. By definition, any non-extreme point in $\underline{\mathcal{K}}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T},\mathcal{D})$ can be written as a convex combination of these extreme points. An easily accessible open-loop controller for a non-extreme initial state of interest is the corresponding convex combination of the optimal open-loop controllers for the extreme states, similarly to [26]. While the feasibility of this controller follows easily from the convexity of the input space and the linearity of the system, this controller, however, is not necessarily optimal for the terminal time problem. We must solve (7) at the given initial state of interest to identify the optimal open-loop controller.

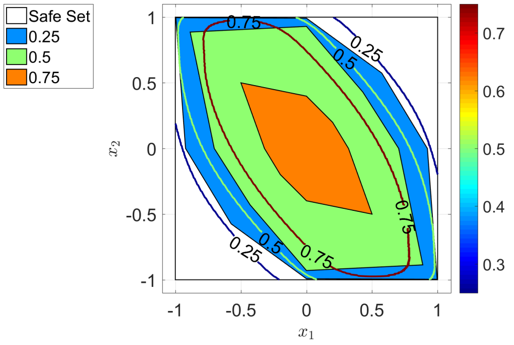

Figure 3: Stochastic reach-avoid sets (contours) and their polytopic underapproximations for a double integrator with a target set $[-0.5, 0.5]^2$.

	$n = 2$ (Fig. 3)			$n = 40$ (Fig. 4)	
$\alpha \in (0, 1]$	0.25	0.5	0.75	0.25	0.98
Algorithm 1	5.62	6.14	4.74	30.57	15.48
DPBDA	33.77	33.77	33.77	–	–

Table 1: Computation times in minutes for the chain of integrators for various α. The dynamic programming-based discretization approach (DPBDA) cannot solve the 40D problem.

5 APPLICATIONS

All computations were performed using MATLAB on an Intel Xeon CPU with 3.4GHz clock rate and 32 GB RAM. The MATLAB code for this work is available at https://github.com/unm-hscl/hscc2018. We used MPT3 [27] for the polytopic constructions and plotting.

5.1 Chain of integrators

We consider a chain of n integrators with $|\mathcal{D}| = 6$. See [12, Sec. V-B] details on the dynamics and the numerical values.

5.1.1 2D system. Consider the terminal time problem with safe set as $\mathcal{S} = [-1, 1]^2$ and the target set as $\mathcal{T} = [-0.5, 0.5]^2$. We compare the set $\mathcal{K}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T})$ computed via Algorithm 1 and the stochastic reach-avoid set $\mathcal{L}^{\pi^*}(\alpha,\mathcal{S},\mathcal{T})$ obtained via DPBDA. The grid spacing for the DPBDA was chosen to be 0.005, as in [12, Table II].

As shown in Figure 3, Algorithm 1 provides a very good underapproximation of the true stochastic reach-avoid set, particularly for low values of α. We believe that this is because the optimal closed-loop controller π^* and the closed-loop controller ρ^* have the same values for the states along

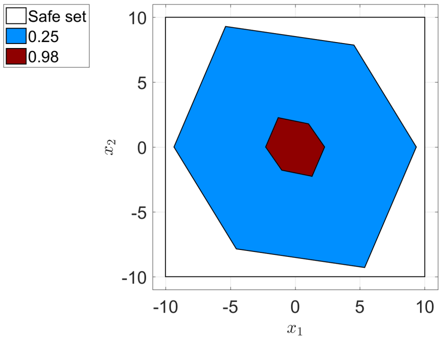

Figure 4: Slice of the polytopic underapproximation $\underline{\mathcal{K}}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T},\mathcal{D})$ for a 40D chain of integrators computed via Algorithm 1 for a target set $[-8,8]^{40}$.

the boundary. The stochastic reach-avoid probability computed via the open-loop control is lower because it does not incorporate feedback [12, Thm. 2]. Note that the maximum values for V and W are 0.86 and 0.75, respectively, and that $\underline{\mathcal{K}}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T},\mathcal{D}) = \emptyset$ for $\alpha \in [0.76, 0.86]$, while $\mathcal{L}^{\pi^*}(\alpha,\mathcal{S},\mathcal{T}) \neq \emptyset$.

5.1.2 40D System. To demonstrate scalability, we compute $\underline{\mathcal{K}}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T},\mathcal{D})$ for $n = 40$, $\alpha \in \{0.25, 0.98\}$ (with $\hat{W}_0^*(\overline{x}_{\max}) = 0.98$), $\mathcal{S} = [-10,10]^{40}$, $\Sigma_{\boldsymbol{w}} = I_{40}$, and $\mathcal{T} = [-8,8]^{40}$. Figure 4 shows a slice of $\underline{\mathcal{K}}^{\rho^*}(\alpha,\mathcal{S},\mathcal{T},\mathcal{D})$ that verifies $\overline{x}_0$ of the form $[x_1\ x_2\ 0\ \ldots\ 0]^\top \in \mathbb{R}^{40}$. To the best of our knowledge, this is the largest stochastic LTI system verified to date through a stochastic reach-avoid formulation.

Computational times in Table 1 show that as expected, Algorithm 1 scales well with dimension. Note that the approach in [12] also handles a 40D system, but would require gridding the state space to compute the stochastic reach-avoid set, making the problem intractable. The grid-independent approach of Algorithm 1 also allows it to analyze verification problems involving large safe and target sets as seen in Figure 4.

5.2 Spacecraft Rendezvous and Docking

We consider two spacecraft in the same elliptical orbit. One spacecraft, referred to as the deputy, must approach and dock with another spacecraft, referred to as the chief, while remaining in a line-of-sight cone, in which accurate sensing of the other vehicle is possible. The relative dynamics are described by the Clohessy-Wiltshire-Hill (CWH) equations [28] with additive stochastic noise.

$$\ddot{x} - 3\omega x - 2\omega\dot{y} = m_d^{-1}F_x, \qquad \ddot{y} + 2\omega\dot{x} = m_d^{-1}F_y. \tag{21}$$

The chief is located at the origin, the position of the deputy is $x, y \in \mathbb{R}$, $\omega = \sqrt{\mu/R_0^3}$ is the orbital frequency, μ is the gravitational constant, and R_0 is the orbital radius of the spacecraft. See [5, 6] for further details and numerical values.

We define the state as $z = [x, y, \dot{x}, \dot{y}] \in \mathbb{R}^4$ and input as $u = [F_x, F_y] \in \mathcal{U} \subseteq \mathbb{R}^2$. We discretize the dynamics (21) in time to obtain the discrete-time LTI system,

$$z_{k+1} = Az_k + Bu_k + w_k \tag{22}$$

with $w_k \in \mathbb{R}^4$ a Gaussian i.i.d. disturbance, with $\mathbb{E}[w_k] = 0$, $\Sigma = \mathbb{E}[w_k w_k^\top] = 10^{-4} \times \mathrm{diag}(1, 1, 5\times 10^{-4}, 5\times 10^{-4})$.

We define the target set and the constraint set as in [5]

$$\mathcal{T} = \{z \in \mathbb{R}^4 : |z_1| \leq 0.1, -0.1 \leq z_2 \leq 0, \\ |z_3| \leq 0.01, |z_4| \leq 0.01\} \tag{23}$$

$$\mathcal{K} = \{z \in \mathbb{R}^4 : |z_1| \leq z_2, |z_3| \leq 0.05, |z_4| \leq 0.05\} \tag{24}$$

with a horizon of $N = 5$. We consider two verification problems: 1) with initial velocity $\dot{x} = \dot{y} = 0$ km/s and $\mathcal{U} = [-0.1, 0.1]^2$ as in [6], and 2) with initial velocity $\dot{x} = \dot{y} = 0.01$ km/s and $\mathcal{U} = [-0.01, 0.01]^2$ as in [5].

Method	Algorithm 1	Chance constrained [5]	Lagrangian [6]
Figure 5	6.52	106.53	0.24
Figure 6	9.88	13.12	–

Table 2: Computation times of various methods in minutes for the CWH dynamics.

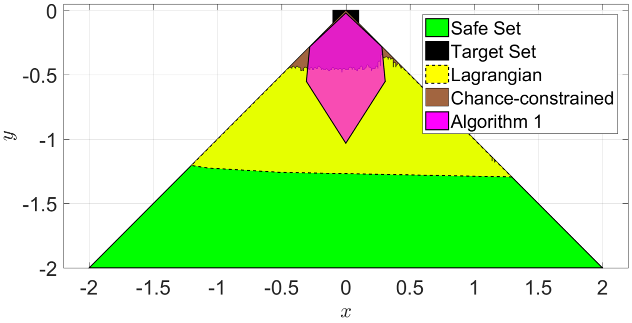

Figure 5: Comparison of Algorithm 1 with the Lagrangian [6, Fig. 4] and the chance-constrained [5] approaches for initial velocity $\dot{x} = \dot{y} = 0$ km/s.

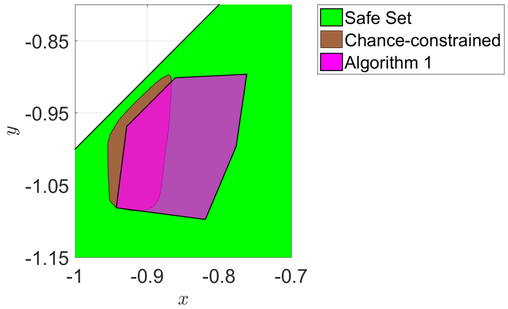

Figure 6: Comparison of Algorithm 1 with the chance-constrained approach [5, Fig. 2] for initial velocity $\dot{x} = \dot{y} = 0.01$ km/s. The Lagrangian method fails to provide a solution in this case.

The safety problem can be posed as a terminal time problem which, while intractable for DPBDA, may be solved using approximately using Algorithm 1, Lagrangian methods [6], or the chance-constrained approach [5]. Figures 5 and 6 show a slice of the stochastic reach-avoid underapproximations for both the verification problems with $|\mathcal{D}| = 10$ and $\alpha = 0.8$. The chance-constrained approach uses a grid over $\mathcal{S}$ in Figure 5 and over the visible portion in Figure 6. Computational times are summarized in Table 2.

The Lagrangian method is significantly faster than Algorithm 1 in this case, and since it provides closed-loop controllers, the resulting underapproximation of the stochastic reach-avoid set is larger in volume. However, because the Lagrangian method suffers from the well known vertex-facet enumeration problem, particularly for large time horizons or small target or safe sets, it fails to solve an underapproximation for the initial velocities in Figure 6.

The chance-constrained approach results in a smaller underapproximation than Algorithm 1. This is because it further underapproximates the cost function $\hat{r}^{\rho}_{\overline{x}_0}(\mathcal{S}, \mathcal{T})$ using chance-constraints. Additionally, since the chance-constrained approach relies on gridding for computing the stochastic reach-avoid set, it does not scale well with dimension or the size of the safe and target sets (see Table 2).

While Algorithm 1 has clear advantages in speed and scalability, the implementation of line 4 (bisection algorithm) can be problematic because of the noisy nature of the objective function evaluation. In particular, the bisection algorithm may terminate prematurely and hence not extend the polytopic underapproximation as far in the given direction as it could go. This is apparent in Figures 5 and 6, where the result via Algorithm 1 does not fully contain the result via the chance-constrained method, and the number of vertices for $\underline{\mathcal{K}}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T}, \mathcal{D})$ is smaller than $|\mathcal{D}| = 10$ in Figure 5. Stochastic optimization techniques, like the stochastic branch-and-bound method, may improve accuracy in this case. This is an area of current work.

6 CONCLUSION

We proposed a scalable technique to compute a polytopic underapproximation of compact and convex stochastic reach-avoid sets using an open-loop control. We characterized the sufficient conditions under which stochastic reach-avoid sets and their open-loop underapproximation are compact and convex, then posed two convex optimization problems to tightly underapproximate the open-loop underapproximation. We also presented sufficient conditions for the admittance of bang-bang optimal Markov controllers. Our approach is highly scalable, does not suffer from the vertex-facet enumeration problem, and is amenable to parallelization for higher fidelity calculations. We presented the first demonstration of safety verification through stochastic reach-avoid sets on a high dimensional (40D) system. Future work will exploit the bang-bang solution in alternative implementations of Algorithm 1 and explore the use of stochastic optimization techniques to further improve the quality of underapproximation.

REFERENCES

[1] S. Summers and J. Lygeros. Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem. *Automatica*, 46(12):1951–1961, 2010.

[2] A. P. Vinod, B. HomChaudhuri, and M. M. K. Oishi. Forward stochastic reachability analysis for uncontrolled linear systems using Fourier transforms. In *Proc. Hybrid Syst.: Comput. and Ctrl.*, pages 35–44, 2017.

[3] B. HomChaudhuri, A. P. Vinod, and M. M. K. Oishi. Computation of forward stochastic reach sets: Application to stochastic, dynamic obstacle avoidance. In *Proc. American Ctrl. Conf.*, Seattle, WA, 2017.

[4] N. Malone, K. Lesser, M. M. K. Oishi, and L. Tapia. Stochastic reachability based motion planning for multiple moving obstacle avoidance. In *Proc. Hybrid Syst.: Comput. and Ctrl.*, pages 51–60, 2014.

[5] K. Lesser, M. M. K. Oishi, and R. Erwin. Stochastic reachability for control of spacecraft relative motion. In *IEEE Conf. Dec. Ctrl.*, 2013.

[6] J. Gleason, A. Vinod, and M. M. K. Oishi. Underapproximation of reach-avoid sets for discrete-time stochastic systems via Lagrangian methods. In *IEEE Conf. Dec. Ctrl.*, 2017.

[7] A. Abate, M. Prandini, J. Lygeros, and S. Sastry. Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 44(11):2724–2734, 2008.

[8] A. Abate, S. Amin, M. Prandini, J. Lygeros, and S. Sastry. Computational approaches to reachability analysis of stochastic hybrid systems. In *Proc. Hybrid Syst.: Comput. and Ctrl.*, pages 4–17, 2007.

[9] N. Kariotoglou, K. Margellos, and J. Lygeros. On the computational complexity and generalization properties of multi-stage and stage-wise coupled scenario programs. *Sys. & Ctr. Lett.*, 94:63–69, 2016.

[10] G. Manganini, M. Pirotta, M. Restelli, L. Piroddi, and M. Prandini. Policy search for the optimal control of Markov Decision Processes: A novel particle-based iterative scheme. *IEEE Trans. Cybern.*, 2015.

[11] D. Drzajic, N. Kariotoglou, M. Kamgarpour, and J. Lygeros. A semidefinite programming approach to control synthesis for stochastic reach-avoid problems. In *Int'l Wrk. on App. Verif. for Cts. and Hybrid Syst.*, pages 134–143, 2016.

[12] A. P. Vinod and M. M. K. Oishi. Scalable Underapproximation for the Stochastic Reach-Avoid Problem for High-Dimensional LTI Systems Using Fourier Transforms. *IEEE Control Syst. Lett.*, 1(2):316–321, 2017.

[13] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge Univ. Press, 2004.

[14] T. Tao. *Analysis II*. Hindustan Book Agency, 2 edition, 2009.

[15] R. T. Rockafellar. *Convex analysis*. Princeton Univ. Press, 1970.

[16] W. Rudin. *Real and complex analysis*. Tata McGraw-Hill Ed., 1987.

[17] J. A. Gubner. *Probability and random processes for electrical and computer engineers*. Cambridge Univ. Press, 2006.

[18] Y.S. Chow and H. Teicher. *Probability Theory: Independence, Interchangeability, Martingales*. Springer New York, 1997.

[19] D. Bertsekas and S. Shreve. *Stochastic optimal control: The discrete time case*. Academic Press, 1978.

[20] A. Genz. QSCMVNV. [Online]. Available: http://www.math.wsu.edu/faculty/genz/software/matlab/qscmvnv.m.

[21] M. L. Putterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2005.

[22] O. Shakernia, G. J. Pappas, and S. Sastry. Decidable controller synthesis for classes of linear systems. In *Proc. Hybrid Syst.: Comput. and Ctrl.*, pages 407–420, 2000.

[23] Roger Webster. *Convexity*. Oxford University Press, 1994.

[24] A. Genz. Numerical computation of multivariate normal probabilities. *J. of Comp. and Graph. Stat.*, 1(2):141–149, 1992.

[25] T. G. Kolda, R. M. Lewis, and V. Torczon. Optimization by direct search: New perspectives on some classical and modern methods. *SIAM review*, 45(3):385–482, 2003.

[26] B. Schurmann and M. Althoff. Convex Interpolation Control with Formal Guarantees for Disturbed and Constrained Nonlinear

Systems. In *Proc. Hybrid Syst.: Comput. and Ctrl.*, pages 121–130, 2017.

[27] M. Herceg, M. Kvasnica, C.N. Jones, and M. Morari. Multi-Parametric Toolbox 3.0. In *Proc. of the European Control Conference*, pages 502–510, Zürich, Switzerland, July 17–19 2013. http://control.ee.ethz.ch/~mpt.

[28] W. Wiesel. *Spaceflight Dynamics*. McGraw-Hill, New York, 1989.

A PROOFS IN SECTIONS 3.1 AND 3.2

Fact 1: Upper semi-continuity is preserved under multiplication, integration using a continuous stochastic kernel, and partial supremum over compact sets [21, Props. B.1 and B.4 and Thm. B.5].

Fact 2: Log-concavity is preserved under multiplication, partial integration, and partial supremum over convex sets [13, Secs. 3.2.5 and 3.5.2].

Proof of Proposition 1. From (5b), the range of $\hat{V}_0^*(\overline{x})$, and $N \geq 1$, $\hat{V}_0^*(\overline{x}) \leq 1_{\mathcal{S}}(\overline{x}) \; \forall \; \overline{x} \in \mathcal{X}$. From (6), $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T}) \subseteq \{\overline{x} \in \mathcal{X} : 1_{\mathcal{S}}(\overline{x}) \geq \alpha\}$. For $\alpha \in (0, 1]$, $\{\overline{x} \in \mathcal{X} : 1_{\mathcal{S}}(\overline{x}) \geq \alpha\} = \{\overline{x} \in \mathcal{X} : 1_{\mathcal{S}}(\overline{x}) = 1\} = \mathcal{S}$ implying $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T}) \subseteq \mathcal{S}$.

Proof of Theorem 1. We prove the u.s.c property of $\hat{V}_k^*(\cdot)$ via induction using (5). Since $\mathcal{T}$ is closed, $\hat{V}_N^*(\cdot)$ is u.s.c by (5a). We know $\hat{V}_{N-1}^*(\cdot)$ is u.s.c by (5b) and Fact 1, which proves the base case of the induction. Assume for induction that $\hat{V}_{k+1}^*(\cdot)$ is u.s.c. We then conclude that $\hat{V}_k^*(\cdot)$ is u.s.c. by similar arguments, completing the proof.

Similarly, $\hat{W}_0^*(\cdot)$ is u.s.c using (11) and Fact 1.

Since $\hat{V}_0^*(\cdot)$ and $\hat{W}_0^*(\cdot)$ are u.s.c, we conclude that the sets $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$ and $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$ are closed for $\alpha \in (0, 1]$.

Proof of Theorem 2. The addition of the boundedness requirement on $\mathcal{S}$ guarantees boundedness of $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$ and $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$ by Lemma 3. The Heine-Borel theorem and Theorem 1 completes the proof.

Proof of Proposition 2. Cartesian products preserve convexity [13, Sec. 2.3.2]. The proof follows from Fact 2 and (11).

Proof of Theorem 3. The proof of the log-concavity of $\hat{V}_k^*(\cdot)$ is similar to Theorem 1 and follows from Fact 2. We need $\mathcal{S}$ and $\mathcal{T}$ to be convex for their respective indicator functions to be log-concave. The additional restrictions of compactness of $\mathcal{U}$, continuity of $Q(\cdot|\overline{x}, \overline{u})$, and Borel $\mathcal{S}, \mathcal{T}$ guarantee the existence of a Markov policy [12, Thm. 1].

Proposition 2 and Fact 2 ensures $\hat{W}_k^*(\cdot)$ is log-concave.

Log-concave functions are quasi-concave implying the convexity of the sets $\mathcal{L}^{\pi^*}(\alpha, \mathcal{S}, \mathcal{T})$ and $\mathcal{K}^{\rho^*}(\alpha, \mathcal{S}, \mathcal{T})$.