Cooptimizing Safety and Performance with a Control-Constrained Formulation

Hao Wang[1*], Adityaya Dhande[2*], and Somil Bansal[1,3], *Member, IEEE*

***Abstract*— Autonomous systems have witnessed a rapid increase in their capabilities, but it remains a challenge for them to perform tasks both effectively and safely. The fact that performance and safety can sometimes be competing objectives renders the cooptimization between them difficult. One school of thought is to treat this cooptimization as a constrained optimal control problem with a performance-oriented objective function and safety as a constraint. However, solving this constrained optimal control problem for general nonlinear systems remains challenging. In this work, we use the general framework of constrained optimal control, but given the safety state constraint, we convert it into an equivalent *control constraint*, resulting in a *state and time-dependent control-constrained optimal control problem*. This equivalent optimal control problem can readily be solved using the dynamic programming principle. We show the corresponding value function is a viscosity solution of a certain Hamilton-Jacobi-Bellman Partial Differential Equation (HJB-PDE). Furthermore, we demonstrate the effectiveness of our method with a two-dimensional case study, and the experiment shows that the controller synthesized using our method consistently outperforms the baselines, both in safety and performance. The implementation of the case study can be found on the project website [1].**

***Index Terms*— Autonomous systems, Optimal control, Robotics**

I. Introduction

PERFORMANCE and safety are two crucial factors we must consider when designing algorithms for autonomous systems. Clearly, we would like the systems to be effective in performing useful tasks. At the same time, the systems must satisfy safety requirements so that they do not inflict damage or injury. As a result, these two factors must be considered simultaneously when we are designing control algorithms.

Manuscript received September 10, 2024; revised November 16, 2024; accepted November 27, 2024. This work is supported in part by the NSF CAREER Program under award 2240163, the DARPA ANSR program, and the IUSSTF-Viterbi program.(Corresponding author: Hao Wang)

*The authors contributed equally to this work.

[1]The authors are associated with the Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California. (email: haowwang@usc.edu)

[2]The author is associated with the Department of Electrical Engineering, Indian Institute of Technology Bombay. (email: adityaya@iitb.ac.in)

[3]The author is associated with the Aeronautics and Astronautics Department, Stanford University. (email: somil@stanford.edu)

[1]https://github.com/haowwang/cooptimize_safety_performance

From the optimal control point of view, the existing methods can roughly be divided into two categories based on whether the safety requirement is posed as a constraint or objective in the optimization problem. Semantically, the latter means that safe behaviors are encouraged but not enforced. A large number of data-driven techniques [5], [18] fall into this category. One drawback of these techniques is that they do not provide any safety guarantees.

The methods that treat the safety requirement as a constraint can be subdivided into two categories based on whether the safety requirement is considered simultaneously with the performance objective. One popular family of methods is safety filtering [3], [6], [13], [19], which provides safety-preserving interventions when necessary in runtime. They generally lead to myopic and suboptimal behaviors as, by design, the safety requirement is often not considered during the performance controller synthesis.

On the other hand, one could formulate the problem as a state-constrained optimal control problem. With this framework, we can optimize the performance objective within the confines of the safety requirement and synthesize controllers that cooptimize safety and performance. However, state-constrained optimal control problems are notoriously difficult to solve using the dynamic programming principle unless certain controllability assumptions are satisfied [7], [17]. Alternatively, Model Predictive Control (MPC) techniques [11], [15] have also been used to solve this problem. However, it is difficult to achieve optimality when the underlying problem involves nonlinear dynamics and/or non-convex state constraints. Recently, the authors in [2] proposed a new framework to circumvent the controllability assumptions by characterizing the epigraph of the value function of the state-constrained optimal control problem. Though theoretically attractive, the method increases the dimensionality of the underlying optimal control problem and, in practice, is susceptible to several numerical challenges, as we demonstrate later in this paper.

In this work, we pose the problem of cooptimizing safety and performance as a state-constrained optimal control problem. Our key idea for overcoming the aforementioned challenges associated with state-constrained optimal control problems is to convert the state constraint into a control constraint using Hamilton-Jacobi reachability analysis. This results in an equivalent optimal control problem free of state constraints and can readily be solved using dynamic programming. We prove that the corresponding value function is a viscosity solution of a certain HJB-PDE, and it can be computed using existing Level Set methods and packages.

To summarize, the contribution of this letter is two-fold: 1) we propose a systematic way of converting a state-constrained optimal control problem into a *control-constrained* optimal control problem and prove that two problems are equivalent, and 2) we show that the value function of the control-constrained optimal control problem is a viscosity solution to a final-value problem for a certain HJB-PDE.

II. Problem Formulation

In this work, we are interested in synthesizing controllers that optimize performance objectives for the given system while respecting the imposed safety constraint. We consider deterministic, continuous-time, and control-affine systems, governed by the ordinary differential equation $\frac{dx}{dt} = f(x,u) = f_1(x) + f_2(x)u$, where $x \in \mathcal{X} \subseteq \mathbb{R}^{n_x}$ and $u \in \mathcal{U} \subseteq \mathbb{R}^{n_d}$ are the state and control of the system. We assume f is bounded and Lipschitz. We further assume the control space $\mathcal{U}$ is convex.

Let $r : \mathcal{X} \times \mathcal{U} \to \mathbb{R}$ and $\phi : \mathcal{X} \to \mathbb{R}$ be the running cost over finite time horizon $[0,T)$ and final cost encoding the performance objectives. We assume both $r(x,u)$ and $\phi(x)$ are bounded and Lipschitz, and we further assume that $r(x,u)$ is convex in u. Furthermore, the safety constraint is given by $l(x) \geq 0 \ \forall x \in \mathcal{X}$, where l is Lipschitz but is not required to be convex.

We formalize the problem of interest as a state-constrained optimal control problem in Prob. 1. Let us use $\xi^{\mathbf{u}}_{x,t} : [t,T] \to \mathcal{X}$ to denote the state trajectory starting from state x at time t evolved with control signal $\mathbf{u} : [t,T) \to \mathcal{U}$. With a slight abuse of the notation, we use $\xi^{\mathbf{u}}_{x,t}(\tau)$ to denote the state at time $\tau \geq t$ along the trajectory $\xi^{\mathbf{u}}_{x,t}$.

Problem 1 (State-Constrained Optimal Control Problem):

$$\inf_{\mathbf{u}} \quad J(x,t,\mathbf{u}) = \int_t^T r(\xi^{\mathbf{u}}_{x,t}(\tau), \mathbf{u}(\tau))d\tau + \phi(\xi^{\mathbf{u}}_{x,t}(T)) \tag{1a}$$

$$s.t. \quad \frac{d}{d\tau}\xi^{\mathbf{u}}_{x,t}(\tau) = f(\xi^{\mathbf{u}}_{x,t}(\tau), \mathbf{u}(\tau)) \ \forall \tau \in [t,T) \tag{1b}$$

$$l(\xi^{\mathbf{u}}_{x,t}(\tau)) \geq 0 \ \forall \tau \in [t,T] \tag{1c}$$

Our goal in this work is finding the state-feedback controller $\pi^* : \mathcal{X} \times [t,T) \to \mathcal{U}$ that solves Prob. 1 at each state $x \in \mathcal{X}$ and time $t \in [0,T]$. Since solving Prob. 1 is challenging, we will present an equivalent optimal control problem whose solution is π^*.

III. Background: Hamilton-Jacobi Reachability Analysis

In this section, we provide a brief overview of Hamilton-Jacobi (HJ) reachability analysis, an approach we use to convert the state constraint (1c) into a control constraint. Given a state constraint (1c), we use HJ reachability to determine the *safe set* $\mathcal{S}$, the set of state x and time t starting from which the system can satisfy (1c) over time horizon $[t,T]$. The construction of $\mathcal{S}$ is formulated as a minimum cost optimal control problem [9], [14] with the cost functional $J_s(x,t,\mathbf{u}) = \min_{\tau \in [t,T]} l(\xi^{\mathbf{u}}_{x,t}(\tau))$. The *safety value function* at state x and time t is defined as

$$V_s(x,t) = \sup_{\mathbf{u}} J_s(x,t,\mathbf{u}) = \sup_{\mathbf{u}} \min_{\tau \in [t,T]} l(\xi^{\mathbf{u}}_{x,t}(\tau)) \tag{2}$$

Then, the safe set $\mathcal{S}$ can be characterized using $V_s(x,t)$ as $\mathcal{S} = \{(x,t) \in \mathcal{X} \times [0,T] | V_s(x,t) \geq 0\}$.

HJ reachability analysis provides a tractable means to compute the safety value function $V_s(x,t)$. It has been shown that $V_s(x,t)$ is the viscosity solution of the Hamilton-Jacobi-Bellman Variational Inequality (HJB-VI) [9], [14]:

$$\min\left\{\frac{\partial V_s}{\partial t} + \max_{u \in \mathcal{U}}\{\frac{\partial V_s}{\partial x}^\top f(x,u)\}, l(x) - V_s(x,t)\right\} = 0$$
$$\forall x \in \mathcal{X} \text{ and } \forall t \in [0,T), \ V_s(x,T) = l(x) \ \forall x \in \mathcal{X} \tag{3}$$

IV. Method

At its core, our method converts the state-constrained optimal control problem (Prob. 1) into a *state and time-dependent control-constrained* optimal control problem, by explicitly characterizing the set of controls that leads the system to satisfy the state constraint, referred to as the *set of safe controls*, at each state x and time t. We first formalize the notion of set of safe controls and use it to formulate the state and time-dependent control-constrained optimal control problem (Prob. 2). We then show Prob. 2 is equivalent to Prob. 1. Subsequently, we show the value function of Prob. 2 is a viscosity solution of a final-value problem for a certain HJB-PDE. Finally, we show one specific way of constructing the set of safe controls using HJ reachability analysis.

A. State and Time-Dependent Control-Constrained Optimal Control Problem

We first provide the definition of the set of safe controls, inspired by a similar notion in [6].

Definition 1 (Set of Safe Controls): The set of safe controls at state x and time t, denoted by $\mathcal{U}_s(x,t)$, is the set of controls that can instantaneously keep the system within the safe set $\mathcal{S}$. More precisely,

$$\mathcal{U}_s(x,t) = \{u \in \mathcal{U} | \lim_{\epsilon \to 0} V_s(\xi^u_{x,t}(t+\epsilon), t+\epsilon) \geq 0\} \tag{4}$$

A set of safe controls is maximal if it contains all other sets of safe control, and we denote the maximal set of safe control by $\mathcal{U}^*_s$.

Note 1: $\mathcal{U}_s(x,t)$ and $\mathcal{U}^*_s(x,t)$ can both be seen as set-value maps from $\mathcal{X} \times [0,T)$ to $\mathcal{U}$.

The state and time-dependent control-constrained optimal control problem is presented below in Prob. 2. It is worthwhile to note that Prob. 2 is identical to Prob. 1, only with the state constraint (1c) replaced by the control constraint (5c) in Prob. 2. We will also show that the optimal value of Prob. 2 is identical to that of Prob. 1 for any state $x \in \mathcal{X}$ and time $t \in [0,T]$, in Theorem. 1.

Problem 2 (Control-Constrained Optimal Control Problem):

$$\inf_{\mathbf{u}} \quad J(x,t,\mathbf{u}) = \int_t^T r(\xi^{\mathbf{u}}_{x,t}(\tau), \mathbf{u}(\tau))d\tau + \phi(\xi^{\mathbf{u}}_{x,t}(T)) \tag{5a}$$

$$s.t. \quad \frac{d}{d\tau}\xi^{\mathbf{u}}_{x,t}(\tau) = f(\xi^{\mathbf{u}}_{x,t}(\tau), \mathbf{u}(\tau)) \ \forall \tau \in [t,T] \tag{5b}$$

$$\mathbf{u}(\tau) \in \mathcal{U}^*_s(\xi^{\mathbf{u}}_{x,t}(\tau), \tau) \ \forall \tau \in [t,T] \tag{5c}$$

Theorem 1: Let us denote the optimal value of Prob. 1 and Prob. 2, at state $x \in \mathcal{X}$ and time $t \in [0,T]$, by $V_1(x,t)$ and $V(x,t)$. Then $V_1(x,t) = V(x,t)\ \forall x \in \mathcal{X}$ and $\forall t \in [0,T]$.

Proof: Take an initial state $x \in \mathcal{X}$ and initial time $t \in [0,T]$. Let us denote the solutions to Prob. 1 and Prob. 2, from x and t, by $\mathbf{u}_1^*$ and $\mathbf{u}^*$, respectively.

The system never violates the state constraint (1c) if $\mathbf{u}^*$ is applied over $[t,T)$, because $\mathbf{u}^*(\tau)$ keeps the system within the safe set $\mathcal{S}$ instantaneously for all time $\tau \in [t,T]$ by definition of the set of safe controls. With this fact established, we can now compare $V_1(x,t)$ and $V(x,t)$.

Case 1: $t = T$. In this case, $V_1(x,T) = V(x,T) = \phi(x)$.

Case 2: $t \in [0,T)$. By definition of the state-constrained optimal control problem, we have $V_1(x,t) \leq V(x,t)\ \forall x \in \mathcal{X}$ and $\forall t \in [0,T)$.

Now we would like to prove $V(x,t) \leq V_1(x,t)\ \forall x \in \mathcal{X}$ and $\forall t \in [0,T)$. Before proceeding, we will establish the fact that $\mathbf{u}_1^*$ satisfies the control constraint (5c) for all $\tau \in [t,T)$. Suppose that is not the case. Then $\exists \tau \in [t,T)$ such that $\mathbf{u}_1^*(\tau) \notin \mathcal{U}_s(\xi_{x,t}^{\mathbf{u}_1^*}(\tau),\tau)$. As a result, $\lim_{\epsilon \to 0} V_s(\xi_{x,t}^{\mathbf{u}_1^*}(\tau+\epsilon),\tau+\epsilon) < 0$. By definition of the safety value function $V_s(x,t)$ (2), the system would certainly violate the state constraint at some point over the time horizon $[\tau+\epsilon, T]$. However, $\mathbf{u}_1^*$ is a solution of Prob. 1 and hence will not lead the system to violate the state constraint over the time horizon $[t,T]$. We have reached a contradiction, and therefore $\mathbf{u}_1^*(\tau)$ satisfies the control constraint (5c) for all $\tau \in [t,T)$.

Take solution $\mathbf{u}_1^*$ of Prob. 1 at initial state x and time t. Since $\mathbf{u}_1^*(\tau)$ satisfies the control constraint (5c) for all $\tau \in [t,T)$, $\mathbf{u}_1^*$ is feasible for Prob. 2. Therefore, $V(x,t) \leq V_1(x,t)$.

Hence, we have shown that $V_1(x,t) = V(x,t)$ for any state $x \in \mathcal{X}$ and time $t \in [0,T]$. ∎

B. Solving Control-Constrained Optimal Control Problem

For the remainder of this letter, we use $V(x,t)$ to denote the value function of Prob. 2. We introduce the following result regarding $V(x,t)$, and the proof of this result is heavily inspired by the proof of Theorem 10.2 in [8].

Theorem 2: Assume the set-valued map $\mathcal{U}_s^* : \mathcal{X} \times [0,T] \rightrightarrows \mathcal{U}$ is lower hemicontinuous. The value function $V(x,t)$ is a viscosity solution of the following final-value problem for the HJB-PDE

$$\frac{\partial V}{\partial t} + \min_{u \in \mathcal{U}_s^*(x,t)} \{f(x,u)^\top \frac{\partial V}{\partial x} + r(x,u)\} = 0 \quad (6)$$
$$\forall x \in \mathcal{X} \text{ and } \forall t \in [0,T), V(x,T) = \phi(x)\ \forall x \in \mathcal{X}$$

Proof: The continuity of $V(x,t)$ can be established using similar arguments presented in Lemma 10.3.3 in [8]. For brevity, we will not show $V(x,t)$ is continuous in this proof. We will first show that $V(x,t)$ is a viscosity supersolution. Take test function $\psi \in C^1(\mathcal{X} \times [0,T])$ and assume that $V - \psi$ has a local maximum at (x_0,t_0). We must show that $\frac{\partial \psi}{\partial t}|_{(x_0,t_0)} + \min_{u \in \mathcal{U}_s^*(x_0,t_0)} \{\frac{\partial \psi}{\partial x}|_{(x_0,t_0)}^\top f(x_0,u_0) + r(x_0,u)\} \geq 0$. Suppose that is not the case. Then $\exists u_0 \in \mathcal{U}_s^*(x_0,t_0)$ and $\exists \theta > 0$ such that

$$\frac{\partial \psi}{\partial t}|_{(x_0,t_0)} + \frac{\partial \psi}{\partial x}|_{(x_0,t_0)}^\top f(x_0,u_0) + r(x_0,u_0) \leq -\theta < 0 \quad (7)$$

Since f and r are continuous in x and u, for (x,u,t) that is sufficiently close to (x_0,u_0,t_0), or equivalently $||x-x_0||_2 + ||u-u_0||_2 + |t-t_0| < \delta$ for some $\delta > 0$, condition (7) holds. We denote the neighborhoods $||x-x_0||_2 + |t-t_0| < \frac{\delta}{2}$ and $||u-u_0||_2 < \frac{\delta}{2}$, by $\hat{\mathcal{N}}$ and $\hat{\mathcal{U}}$, respectively.

Take $\hat{\mathcal{U}}$. $\hat{\mathcal{U}} \cap \mathcal{U}_s^*(x_0,t_0) \neq \emptyset$ because $u_0 \in \hat{\mathcal{U}}$ and $u_0 \in \mathcal{U}_s^*(x_0,t_0)$. Since by assumption $\mathcal{U}_s^*(x,t)$ is lower hemicontinuous, there exists a neighborhood $\mathcal{N}$ of (x_0,t_0) $s.t.$ $\forall (x,t) \in \mathcal{N}, \mathcal{U}_s^*(x,t) \cap \hat{\mathcal{U}} \neq \emptyset$. It follows immediately that $\forall (x,t) \in \hat{\mathcal{N}} \cap \mathcal{N}, \mathcal{U}_s^*(x,t) \cap \hat{\mathcal{U}} \neq \emptyset$. Then by continuity of f in u, there exists $t_e > t_0$, over which we can construct a control signal $\mathbf{u}^* : [t_0,t_e) \to \mathcal{U}$, along with the resulting state trajectory $\xi_{x_0,t_0}^{\mathbf{u}^*} : [t_0,t_e] \to \mathcal{X}$, such that $(\xi_{x_0,t_0}^{\mathbf{u}^*}(\tau),\tau) \in \hat{\mathcal{N}} \cap \mathcal{N}\ \forall \tau \in [t_0,t_e]$ and concurrently $\mathbf{u}^*(\tau) \in \mathcal{U}_s^*(\xi_{x_0,t_0}^{\mathbf{u}^*}(\tau),\tau) \cap \hat{\mathcal{U}}\ \forall \tau \in [t_0,t_e)$. By construction, $(\xi_{x_0,t_0}^{\mathbf{u}^*}(\tau),\tau) \in \hat{\mathcal{N}}$ and $\mathbf{u}^*(\tau) \in \hat{\mathcal{U}}\ \forall \tau \in (t_0,t_e)$, and hence $(\xi_{x_0,t_0}^{\mathbf{u}^*}(\tau), \mathbf{u}^*(\tau), \tau)\ \forall \tau \in (t_0,t_e)$ satisfies condition (7).

By assumption $V - \psi$ has a local maximum at (x_0,t_0), we have $V(x,t) - \psi(x,t) \leq V(x_0,t_0) - \psi(x_0,t_0)\ \forall (x,t) \in \hat{\mathcal{N}}$. Note that from the dynamics programming principle we have $V(\xi_{x_0,t_0}^{\mathbf{u}}(t_0),t_0) \leq \int_{t_0}^{t_e} r(\xi_{x_0,t_0}^{\mathbf{u}}(t),\mathbf{u}(t))dt + V(\xi_{x_0,t_0}^{\mathbf{u}}(t_e),t_e)$ for any control signal $\mathbf{u}$ that satisfies the control constraint (5c) over the time horizon $[t_0,T)$. Making use of this fact and rearranging the equation we arrive at the following:

$$\begin{aligned}
0 &\leq \psi(\xi_{x_0,t_0}^{\mathbf{u}^*}(t_e),t_e) - \psi(\xi_{x_0,t_0}^{\mathbf{u}^*}(t_0),t_0) \\
&\quad - V(\xi_{x_0,t_0}^{\mathbf{u}^*}(t_e),t_e) + V(\xi_{x_0,t_0}^{\mathbf{u}^*}(t_0),t_0) \\
&\leq \psi(\xi_{x_0,t_0}^{\mathbf{u}^*}(t_e),t_e) - \psi(\xi_{x_0,t_0}^{\mathbf{u}^*}(t_0),t_0) \cancel{- V(\xi_{x_0,t_0}^{\mathbf{u}^*}(t_e),t_e)} \\
&\quad + \int_{t_0}^{t_e} r(\xi_{x_0,t_0}^{\mathbf{u}^*}(t),\mathbf{u}^*(t))dt \cancel{+ V(\xi_{x_0,t_0}^{\mathbf{u}^*}(t_e),t_e)} \\
&= \int_{t_0}^{t_e} \Big[\frac{\partial}{\partial t}\psi(\xi_{x_0,t_0}^{\mathbf{u}^*}(t),t) + \frac{\partial}{\partial x}\psi(\xi_{x_0,t_0}^{\mathbf{u}^*}(t),t)^\top \\
&\qquad f(\xi_{x_0,t_0}^{\mathbf{u}^*}(t),\mathbf{u}^*(t)) + r(\xi_{x_0,t_0}^{\mathbf{u}^*}(t),\mathbf{u}^*(t)) \Big] dt \\
&\leq -\theta(t_e - t_0)
\end{aligned}$$

Since $\theta > 0$ and $(t_e - t_0) > 0$, we have reached a contradiction. Therefore, we have $\frac{\partial \psi}{\partial t}|_{(x_0,t_0)} + \min_{u \in \mathcal{U}_s^*(x_0,t_0)} \{\frac{\partial \psi}{\partial x}|_{(x_0,t_0)}^\top f(x_0,u) + r(x_0,u)\} \geq 0$, and V is a viscosity supersolution.

The proof for the viscosity subsolution follows a similar argument, and we omit the proof for brevity. Because V is both a viscosity supersolution and subsolution, it is a viscosity solution of the final-value problem for the HJB-PDE (6) as we intended to show. ∎

C. Characterizing the Set of Safe Controls

We now introduce a method to characterize the set of safe controls using HJ reachability analysis. Given a state constraint (1c), we first obtain the safety value function V_s (2) by solving the HJB-VI (3). Then, for $0 < \gamma << 1$, we characterize the set of safe control at state x and time t in (8), and we show that when interpreted as a set-value map, (8) is lower hemicontinuous.

Definition 2 (Set of Safe Controls Using HJ Reachability):

$$\mathcal{U}_s(x,t) = \begin{cases} \mathcal{U} \text{ if } V_s(x,t) > 0 \\ \{u \in \mathcal{U} | -\gamma \leq \frac{\partial V_s}{\partial t} + \frac{\partial V_s}{\partial x}^\top f(x,u) \leq 0\} \\ \text{if } V_s(x,t) \leq 0 \end{cases} \tag{8}$$

Note 2: (8) is only a set of safe control per Definition 1 when $\gamma = 0$, since when $\gamma > 0$, $\mathcal{U}_s(x,t)$ contains controls that would decrease the safety value and lead the system out of the safe set. In practice we do use $\gamma = 0$, because the construct of γ is only necessary to render $\mathcal{U}_s(x,t)$ lower hemicontinuous. There are other ways of constructing $\mathcal{U}_s(x,t)$ beyond (8), but it is challenging to construct a $\mathcal{U}_s(x,t)$ satisfying all the following criteria: 1) safety-preserving, 2) maximal, and 3) lower-hemicontinuous. We consider (8) as the best option given the fact that it is maximal, safety-preserving in practice ($\gamma = 0$), and easy to explain and implement.

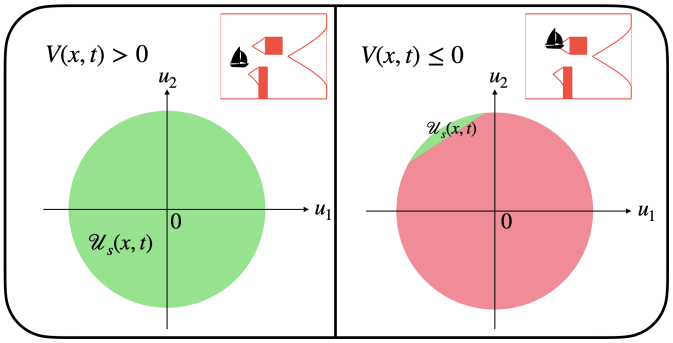

Fig. 1. An illustration of the set of safe controls using Definition 1 of the boat system used in the case study. The control set $\boldsymbol{\mathcal{U}}$ is given by $\boldsymbol{\{[u_1, u_2] \in \mathbb{R}^2 | \, ||[u_1, u_2]||_2 \leq 1\}}$. The green-shaded region is the set of safe controls at the current state $\boldsymbol{x}$ and time $\boldsymbol{t}$, and the pink-shaded region is the set of controls that would lead the system into eventual violation of the state constraint. The location of the system relative to the boundary of the safe set, denoted using solid red line, is also illustrated in the top right subfigure. (Left) When the system is in the interior of the safe set, any admissible controls are permitted/safe. (Right) When the system is on boundary or outside the safe set, only admissible controls that lead to small decrease in the safety value are permitted.

Intuitively, when the system is not at risk of exiting the safe set $\mathcal{S}$, the system is allowed to take any admissible control, and the system will remain within $\mathcal{S}$. Since $\frac{\partial V_s}{\partial t} + \frac{\partial V_s}{\partial x}^\top f(x,u)$ is the total derivative of V_s with respect to t along a state trajectory resulting from the applied control u, we can see that as we take $\gamma \to 0$, $\mathcal{U}_s(x,t)$ consists of only controls that instantaneously keep the safety value constant, when the system is on the boundary of $\mathcal{S}$. Though there is no admissible control that can render the system safe as soon as it exits $\mathcal{S}$, we define $\mathcal{U}_s(x,t)$ to be identical to the previous case. By doing so, we limit the degree to which the state constraint is violated and potential consequences, when the system finds itself outside of the safe set $\mathcal{S}$. We show an illustration of the set of safe controls based on Definition. 1 of the system used in the case study in Fig. 1. Recall that in order for the value function to be a viscosity solution of the HJB-PDE (6), the set-value map $\mathcal{U}_s(x,t)$ is required to be lower hemicontinuous. We now show (8) is lower hemicontinuous.

Proposition 1: Suppose $V_s(x,t)$ is continuously differentiable in x and t. The set-value map defined in (8) is lower hemicontinuous in x and t.

Proof:

Case 1: $V_s(x,t) > 0$. In this case, (x,t) is in the interior of the safe set $\mathcal{S}$, which we denote using $\mathcal{S}^o = \{(x,t) \in \mathcal{X} \times [0,T] | V_s(x,t) > 0\}$. Take open set $\mathcal{A} \subset \mathcal{U}$ such that $\mathcal{A} \cap \mathcal{U}_s(x,t) \neq \emptyset$. Since $\mathcal{S}^o$ is open, $\exists \epsilon > 0$ such that $\forall (x',t') \in \mathcal{B}((x,t),\epsilon)$, we have $(x',t') \in \mathcal{S}^o$. Then it follows that $\mathcal{U}_s(x',t') = \mathcal{U}$, and $\mathcal{U}_s(x',t') \cap \mathcal{A} = \mathcal{A} \neq \emptyset$. Therefore, $\mathcal{U}_s(x,t)$ is lower hemicontinuous $\forall (x,t) \in \mathcal{S}^o$.

Case 2: $V_s(x,t) < 0$. In this case, (x,t) is in the complement of the safe set $\mathcal{S}$, which we denote using $\overline{\mathcal{S}} = \{(x,t) \in \mathcal{X} \times [0,T] | V_s(x,t) < 0\}$. Take open set $\mathcal{A} \subset \mathcal{U}$ such that $\mathcal{A} \cap \mathcal{U}_s(x,t) \neq \emptyset$. Since $\mathcal{A}$ is open and $\mathcal{A} \cap \mathcal{U}_s(x,t) \neq \emptyset$, $\exists u_0 \in \mathcal{A} \cap \mathcal{U}_s(x,t)$ such that $-\gamma < \frac{\partial V_s}{\partial t} + \frac{\partial V_s}{\partial x}^\top f(x,u_0) < 0$. Let $\mathcal{N} = \mathcal{B}((x,t),\delta) \subset \overline{\mathcal{S}}$ be an open ball centered at (x,t). Take $(x',t') \in \mathcal{N}$. Then for $\delta_1 \in \mathbb{R}, \delta_2 \in \mathbb{R}^{n_x}, \delta_3 \in \mathbb{R}^{n_x}$, and $\delta_4 \in \mathbb{R}^{n_x \times n_u}$, we have the following

$$\begin{aligned} &\frac{\partial V_s}{\partial t}|_{(x',t')} + \frac{\partial V_s}{\partial x}|^\top_{(x',t')} f(x', u_0) \\ &= \frac{\partial V_s}{\partial t}|_{(x',t')} + \frac{\partial V_s}{\partial x}|^\top_{(x',t')} \big[f_1(x') + f_2(x') u_0\big] \\ &= \big[\frac{\partial V_s}{\partial t}|_{(x,t)} + \delta_1\big] + \big[\frac{\partial V_s}{\partial x}|_{(x,t)} + \delta_2\big]^\top \Big[f_1(x) \\ &\qquad + \delta_3 + \big[f_2(x) + \delta_4\big] u_0\Big] \\ &= \big[\frac{\partial V_s}{\partial t}|_{(x,t)} + \frac{\partial V_s}{\partial x}|^\top_{(x,t)} f(x, u_0)\big] \\ &+ \delta_1 + \frac{\partial V_s}{\partial x}|^\top_{(x,t)} \delta_3 + \frac{\partial V_s}{\partial x}|^\top_{(x,t)} \delta_4 u_0 + \delta_2^\top f_2(x) u_0 \\ &+ \delta_2^\top \delta_4 u_0 + \delta_2^\top f_1(x) + \delta_2^\top \delta_3 \end{aligned} \tag{9}$$

Using triangle inequality and definition of dot product, we have the following, where $||\cdot||$ denotes the Euclidean norm for a vector and the spectral norm for a matrix, and $|\cdot|$ denotes the absolute value of a real number.

$$\begin{aligned} &\delta_1 + \frac{\partial V_s}{\partial x}|^\top_{(x,t)} \delta_3 + \frac{\partial V_s}{\partial x}|^\top_{(x,t)} \delta_4 u_0 + \delta_2^\top f_2(x) u_0 \\ &\qquad + \delta_2^\top \delta_4 u_0 + \delta_2^\top f_1(x) + \delta_2^\top \delta_3 \end{aligned} \tag{10a}$$

$$\begin{aligned} &\leq |\delta_1| + ||\frac{\partial V_s}{\partial x}|_{(x,t)}|| \cdot ||\delta_3|| + ||\frac{\partial V_s}{\partial x}|_{(x,t)}|| \cdot ||\delta_4|| \\ &\cdot ||u_0|| + ||\delta_2|| \cdot ||f_2(x)|| \cdot ||u_0|| + ||\delta_2|| \cdot ||\delta_4|| \\ &\cdot ||u_0|| + ||\delta_2|| \cdot ||f_1(x)|| + ||\delta_2|| \cdot ||\delta_3|| \end{aligned} \tag{10b}$$

Since $V_s(x,t)$ is continuously differentiable in x and t, and $f_1(x)$ as well as $f_2(x)$ are continuous in x, we can choose δ such that $\forall (x',t') \in \mathcal{N} = \mathcal{B}((x,t),\delta)$ we have (10b) $\leq \min\{\big|\frac{\partial V_s}{\partial t}|_{(x,t)} + \frac{\partial V_s}{\partial x}|^\top_{(x,t)} f(x,u_0)\big|, \big|-\gamma - \frac{\partial V_s}{\partial t}|_{(x,t)} - \frac{\partial V_s}{\partial x}|^\top_{(x,t)} f(x,u_0)\big|\}$, or equivalently $-\gamma \leq \frac{\partial V_s}{\partial t}|_{(x',t')} + \frac{\partial V_s}{\partial x}|^\top_{(x',t')} f(x',u_0) \leq 0$. We have shown that $\forall (x',t') \in \mathcal{N}, u_0 \in \mathcal{U}_s(x',t')$, and as a result $\mathcal{U}_s(x',t') \cap \mathcal{A} \neq \emptyset$. Therefore, $\mathcal{U}_s(x,t)$ is lower hemicontinuous $\forall (x,t) \in \overline{\mathcal{S}}$.

Case 3: $V_s(x,t) = 0$. In this case, (x,t) is on the boundary of the safe set $\mathcal{S}$, which we denote using $\partial\mathcal{S} = \{(x,t) \in \mathcal{X} \times [0,T] | V_s(x,t) = 0\}$. Take open set $\mathcal{A} \subset \mathcal{U}$ such that $\mathcal{A} \cap$

$\mathcal{U}_s(x,t) \neq \emptyset$. We select $\epsilon > 0$ and construct $\epsilon-$neighborhood around (x,t), $\mathcal{N} = \mathcal{B}((x,t),\epsilon)$ such that $\forall (x',t') \in \mathcal{N} \cap \overline{\mathcal{S}^o}$, we have $\mathcal{U}_s(x',t') \cap \mathcal{A} \neq \emptyset$, using the argument presented above in Case 2. Note that $\forall (x',t') \in \mathcal{N} \cap \mathcal{S}^o$, we have $\mathcal{U}_s(x',t') = \mathcal{U}$ and hence $\mathcal{U}_s(x',t') \cap \mathcal{A} = \mathcal{A} \neq \emptyset$. Since $(\mathcal{N} \cap \overline{\mathcal{S}^o}) \cup (\mathcal{N} \cap \mathcal{S}^o) = \mathcal{N}$, we have show that $\forall (x',t') \in \mathcal{N}, \mathcal{U}_s(x',t') \cap \mathcal{A} \neq \emptyset$, and $\mathcal{U}_s(x,t)$ is lower hemicontinuous $\forall (x,t) \in \partial\mathcal{S}$.

We have exhausted all the cases, and therefore $\mathcal{U}_s(x,t)$ is lower hemicontinuous in x and t. ■

D. Synthesizing the Cooptimization Controller

After obtaining the value function $V(x,t)$, we synthesize the closed-loop controller as follows

$$\pi^*(x,t) = \underset{u \in \mathcal{U}_s^*(x,t)}{\operatorname{argmin}} \{ \frac{\partial V}{\partial x}^\top f(x,u) + r(x,u) \} \qquad (11)$$

It is important to note that the controller synthesis problem (11) is convex under the assumption that the running cost $r(x,u)$ is convex in u and the dynamics $f(x,u)$ are control-affine. Then it is clear that $\frac{\partial V}{\partial x}^\top f(x,u) + r(x,u)$ is convex in u. Using the set of safe controls proposed in (8), for (x,t) such that $V_s(x,t) > 0$, $\mathcal{U}_s(x,t)$ is the entire control space $\mathcal{U}$, which is a convex set. On the other hand, for (x,t) such that $V_s(x,t) \leq 0, \mathcal{U}_s(x,t) = \{u \in \mathcal{U} | \frac{\partial V_s}{\partial t} + \frac{\partial V_s}{\partial x}^\top f(x,u) = 0\}$, the intersection of a hyperplane and a convex set, is also convex. Therefore, (11) is an optimization problem with a convex objective and a convex constraint for any state $x \in \mathcal{X}$ and time $t \in [0,T)$. Very often $r(x,u)$ depends quadratically on u (e.g., to minimize the control energy), and for common choices of the control space $\mathcal{U}$, such as hypercubes or Euclidean norm balls, (11) is a quadratic program (QP) or a quadratically-constrained quadratic program (QCQP), both of which can be solved efficiently and reliably online.

V. Case Study

Since our method is ultimately solving a state-constrained optimal control problem using dynamic programming, it is most similar to [2]. To better compare against [2], we implement the numerical example from the paper with some minor modifications. The 2D system has the following dynamics $[\dot{x_1}, \dot{x_2}]^\top = [u_1 + 2 - 0.5x_2^2, u_2]$, with the control space $\mathcal{U} = \{[u_1,u_2] \in \mathbb{R}^2 | \; ||[u_1,u_2]||_2 \leq 1\}$.

The rectangular arena is given by $[-3,2] \times [-2,2]$, and the states outside of this arena are considered unsafe. There are two additional obstacles situated within the arena. The arena and the obstacle configuration are shown in Fig 2.

The objective of the system is to minimize its distance to the goal location $[1.5,0]^\top$, giving rise to the cost functional $J(x,t,\mathbf{u}) = \int_0^2 \sqrt{(x_1-1.5)^2 + x_2^2}\, d\tau$, while maintaining safety (i.e. not running into the obstacles), over a time horizon of two seconds. We obtain the safety value function $V_s(x,t)$ and the value function $V(x,t)$ using LevelSetToolbox [16] and HelperOC [1] using a grid size of 70×70.

We evaluate our method and the baselines by synthesizing closed-loop control signals from 100 random initial states, and we focus primarily on 1) rollout success rate: the percent of trajectories that are safe over the entire time horizon, 2) rollout cost: the cost functional J evaluated with the resulting control signals, and 3) offline and online computation time.

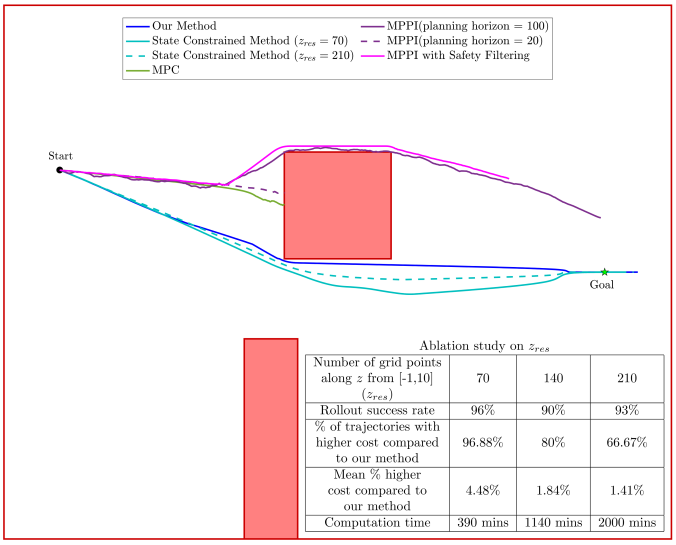

Ablation study on z_{res}

Number of grid points along z from [-1,10] (z_{res})	70	140	210
Rollout success rate	96%	90%	93%
% of trajectories with higher cost compared to our method	96.88%	80%	66.67%
Mean % higher cost compared to our method	4.48%	1.84%	1.41%
Computation time	390 mins	1140 mins	2000 mins

Fig. 2. Trajectories from initial state $[-\mathbf{2.58, 0.77}]^\top$. Costs incurred are (i) Our method: 3.51, (ii) State constrained method, $(\boldsymbol{z_{res} = 70})$: 3.62, (iii) State constrained method, $(\boldsymbol{z_{res} = 210})$: 3.52, (iv) MPPI (horizon = 20): Violates safety constraint (v) MPPI (horizon = 100): 4.55, (vi) MPPI with safety filtering: 4.95, (vii) MPC: Violates safety constraint

The baselines we considered are Baseline 1) solving the state-constrained optimal control problem directly [2], Baseline 2) converting the state constraint (1c) into an obstacle penalty and solving the problem using Model Predictive Path Integral Control (MPPI) [20], Baseline 3) performing safety filtering [6] on the output of Baseline 2, and Baseline 4) solving the state-constrained optimal control problem in a receding horizon fashion (MPC). The results are compiled in TABLE . I.

TABLE I
Comparison of metrics for our method and the baselines

Method	Our Method	State-Constrained Method [2] (Baseline 1)	MPPI (Baseline 2)	MPPI+ Filtering (Baseline 3)	MPC (Baseline 4)
Rollout success rate	100%	96%	19%	100%	11%
% of trajectories with higher cost compared to our method	-	96.88%	84.21%	92%	72.73%
Mean % higher cost compared to our method	-	4.48%	14.96%	23.17%	2.13%
Offline computation time	21 mins	390 mins	-	3s	-
Online computation time	0.0015s	0.0015s	0.1s	0.1s	0.6s

We first analyze the rollout success rate, the metric indicative of the methods' ability to satisfy the safety requirement. In theory, Baseline 1 guarantees the satisfaction of the state constraint over the entire time horizon. However, Baseline 1 fails to achieved 100% rollout success rate due to numerical inaccuracies that arise from the discretization of the state space. Baseline 2 and 4 performs poorly in this metric primarily due to the highly non-convex state constraint (disjoint obstacles in this case). On the other hand, our method and Baseline 3 are able to achieve 100% rollout success rate.

In terms of the rollout cost, our method consistently outperforms Baseline 2 and 3 mostly due the fact that MPPI, in finite data regime, is only able to find locally optimal solution.

Similarly, our method outperforms Baseline 4, because the non-convex optimization used in Baseline 4 is not solved to global optimum. Perhaps surprisingly, our method consistently outperforms Baseline 1. Though Baseline 1 and our method are computed using the same numerical tool, Baseline 1 is more severely affected by the discretization of the state space. Note that Baseline 1 augments its state space with an auxiliary state z that is used to determine the actual value of the state. The discretization of the auxiliary state z has a significant effect on the quality of the synthesized control signals, and we will demonstrate the effect using an ablation study on the number of grid points z_{res} used in z's dimension. The result of this ablation study is compiled in the bottom right table of Fig. 2. The performance of Baseline 1, in terms of trajectory cost, improves as the number of grid points in z increases. However, the improvement of performance comes with a negative consequence of significant increase in offline computation time.

We now examine the computation time. Compared to other methods, Baseline 1 and our method require the most offline computation, given the fact that the value functions are computed using dynamic programming on a grid [16]. It is worthwhile to point out that our method requires significantly less offline computation, primarily due to the fact that our method does not require an auxiliary state as in [2], and the reduction in computation scales with the discretization of the auxiliary state. Though our method requires solving two optimal control problems, it is typically much faster than [2] as the number of grid points used for the auxiliary state is much larger than 2. Baseline 3 requires some minimal offline computation for the safety value function. On the other hand, online methods Baseline 2 and Baseline 4 do not require any offline computation. For online computation time, Baseline 1 and our method outperform the rest of the baselines, as both methods solve quadratic programs, for which we use fast and reliable solver Gurobi [12], for control synthesis online.

We demonstrate the qualitative behaviors of the methods by showing the state trajectories, obtained using the synthesized closed-loop control signals over the entire time horizon, starting from a particular initial state in Fig. 2. The trajectory from our method is quite similar to that of Baseline 1, though the trajectory from Baseline 1 is slightly suboptimal for the aforementioned reasons. Baseline 2 and 3 unsurprisingly enter into a local minimum early on and are never able to recover. Baseline 4 fails to be safe as the corresponding optimization problem does not return the optimal solution satisfying the state constraint.

VI. Conclusion

In this work, we proposed a method to synthesize controllers that cooptimize safety and performance for autonomous systems by formulating the problem as a control-constrained optimal control problem. We also show that the value function of the optimal control problem is a viscosity solution to a certain HJB-PDE. Although our method is shown to provide safety guarantee for the system and outperform other methods in terms of performance, our method has several drawbacks. First, while the theory is general, our method does not scale to high-dimensional systems. In the future, we will look into computing the value function using deep learning techniques [4], [10]. Furthermore, to synthesize controllers, our method assumes that the safety value function $V_s(x,t)$ and the value function $V(x,t)$ are differentiable everywhere, which is typically not the case. We will explore overcoming this challenge using a smooth overapproximation of the value functions [6].

Acknowledgement

We would like to thank Sanat Mulay for his insights and help in the proof of Proposition 1.

References

[1] helperOC Library, 2019. `https://github.com/HJReachability/helperOC`.

[2] Altarovici, Albert, Bokanowski, Olivier, and Zidani, Hasnaa. A general hamilton-jacobi framework for non-linear state-constrained control problems. *ESAIM: COCV*, 19(2):337–357, 2013.

[3] Aaron D. Ames, Xiangru Xu, Jessy W. Grizzle, and Paulo Tabuada. Control barrier function based quadratic programs for safety critical systems. *IEEE Transactions on Automatic Control*, 62(8):3861–3876, 2017.

[4] Somil Bansal and Claire J. Tomlin. Deepreach: A deep learning approach to high-dimensional reachability. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1817–1824, 2021.

[5] Homanga Bharadhwaj, Aviral Kumar, Nicholas Rhinehart, Sergey Levine, Florian Shkurti, and Animesh Garg. Conservative safety critics for exploration. *arXiv preprint arXiv:2010.14497*, 2020.

[6] Javier Borquez, Kaustav Chakraborty, Hao Wang, and Somil Bansal. On safety and liveness filtering using hamilton-jacobi reachability analysis. *IEEE Transactions on Robotics*, pages 1–16, 2024.

[7] Italo Capuzzo-Dolcetta and P-L Lions. Hamilton-jacobi equations with state constraints. *Transactions of the American mathematical society*, 318(2):643–683, 1990.

[8] Lawrence C Evans. *Partial Differential Equations*. Graduate studies in mathematics. American Mathematical Society, 2010.

[9] I.J. Fialho and T.T. Georgiou. Worst case analysis of nonlinear systems. *IEEE Transactions on Automatic Control*, 44(6):1180–1196, 1999.

[10] Jaime F. Fisac, Neil F. Lugovoy, Vicenç Rubies-Royo, Shromona Ghosh, and Claire J. Tomlin. Bridging hamilton-jacobi safety analysis and reinforcement learning. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8550–8556, 2019.

[11] Carlos E Garcia, David M Prett, and Manfred Morari. Model predictive control: Theory and practice—a survey. *Automatica*, 25(3):335–348, 1989.

[12] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2024.

[13] Kai-Chieh Hsu, Haimin Hu, and Jaime F Fisac. The safety filter: A unified view of safety-critical control in autonomous systems. *Annual Review of Control, Robotics, and Autonomous Systems*, 7, 2023.

[14] John Lygeros. On reachability and minimum cost optimal control. *Automatica*, 40(6):917–927, 2004.

[15] D.Q. Mayne, J.B. Rawlings, C.V. Rao, and P.O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.

[16] Ian M Mitchell et al. A toolbox of level set methods. *UBC Department of Computer Science Technical Report TR-2007-11*, page 31, 2007.

[17] Halil Mete Soner. Optimal control with state-space constraint i. *SIAM Journal on Control and Optimization*, 24(3):552–561, 1986.

[18] Krishnan Srinivasan, Benjamin Eysenbach, Sehoon Ha, Jie Tan, and Chelsea Finn. Learning to be safe: Deep rl with a safety critic. *arXiv preprint arXiv:2010.14603*, 2020.

[19] Kim P. Wabersich, Andrew J. Taylor, Jason J. Choi, Koushil Sreenath, Claire J. Tomlin, Aaron D. Ames, and Melanie N. Zeilinger. Data-driven safety filters: Hamilton-jacobi reachability, control barrier functions, and predictive methods for uncertain systems. *IEEE Control Systems Magazine*, 43(5):137–177, 2023.

[20] Grady Williams, Paul Drews, Brian Goldfain, James M. Rehg, and Evangelos A. Theodorou. Information-theoretic model predictive control: Theory and applications to autonomous driving. *IEEE Transactions on Robotics*, 34(6):1603–1622, 2018.