

# Fast, Smooth and Safe: Implicit Control Barrier Functions through Reach-Avoid Differential Dynamic Programming

Athindran Ramesh Kumar, Kai-Chieh Hsu, Peter J. Ramadge, and Jaime F. Fisac

**Abstract**—Safety is a central requirement for autonomous system operation across domains. Hamilton-Jacobi (HJ) reachability analysis can be used to construct “least-restrictive” safety filters that result in infrequent, but often extreme, control overrides. In contrast, control barrier function (CBF) methods apply smooth control corrections to guard the system against an often conservative safety boundary. This paper provides an online scheme to construct an implicit CBF through HJ reach-avoid differential dynamic programming in a receding-horizon framework, enabling smooth safety filtering with infinite-time safety guarantees. Simulations with the Dubins car and 5D bicycle dynamics demonstrate the scheme’s ability to preserve safety smoothly without the conservativeness of handcrafted CBFs.

## I. INTRODUCTION

Through improved sensors, computation, learning, and decision-making techniques, autonomous robotic systems are becoming increasingly capable. However, deploying these systems in safety-critical environments requires robust fail-safe methods to ensure the avoidance of catastrophic *failure states*. To enforce safe behavior, various modern techniques rely on finding a *safe set*  $\Omega$ , that is, a subset of non-failure states such that from each state in  $\Omega$  there exists a control strategy that will keep the state in  $\Omega$ , i.e.,  $\Omega$  is controlled-invariant.

Hamilton-Jacobi (HJ) reachability analysis provides a constructive approach to obtain the *maximal* controlled-invariant safe set  $\Omega^*$  [1, 2]. Since there generally exist no closed-form solutions, a state-space discretization is often used to compute the value function and control policy using numerical dynamic programming (DP). These solvers [3] scale poorly with state space dimension  $n$  and are impractical for  $n \geq 5$ . Several approximate methods [4–8] have been proposed to overcome this issue. The safety controls obtained from reachability solvers are subsequently used in a “least-restrictive” (LR) framework where the safety control is only applied if the system attempts to cross the zero-level set of the HJ value function. While attractive in theory, this minimally invasive scheme can yield abrupt “panic” behaviors in practice, as it waits until the last chance before applying the safety control. Control barrier functions (CBF) [9, 10] offer an alternative approach to guarantee safety. Given a known safe set  $\Omega$ , these functions are designed to guarantee its invariance working in conjunction with a quadratic program (CBF-QP). The CBF-QP modifies the control well before the boundary of  $\Omega$ , and a constraint on the derivative of the CBF

prevents rapid movement towards the boundary, resulting in smoother trajectories than the LR approach. However, control barrier functions have two drawbacks. First, they are typically handcrafted for conservative invariant sets  $\Omega \subset \Omega^*$ . Second, the CBF-QP can become infeasible in the presence of tight actuation constraints, endangering safety; a backup policy is then needed to ensure persistent feasibility [11]. In contrast, the invariant sets obtained via HJ reachability analysis are maximal and guarantee feasibility, but the LR controls can be extreme.

To bridge both frameworks, [12] introduces a control-barrier value function (CBVF) derived using HJ reachability theory. The CBVF is computed offline using numerical dynamic programming. The warm-start approach of [13] alleviates the high computational burden by using a learned or handcrafted candidate CBF as initialization and iteratively refining it through numerical HJ analysis, asymptotically converging to a valid infinite-time CBF. However, numerical methods generally scale poorly to systems beyond 6 continuous state dimensions and are challenging for real-time applications.

Here, we propose a provably safe online scheme, which constructs an implicit CBF at each control cycle through reach-avoid differential dynamic programming (DDP). Our CBF-DDP scheme enables the **fast** runtime computation of control solutions that vary **smoothly** along system trajectories and are guaranteed to persistently enforce **safety** at all future times. Our race car simulations with a 5-D state space demonstrate smoother safety filtering than LR while overcoming the conservativeness of handcrafted CBFs.

## II. PRELIMINARIES

Let  $\mathcal{U} \subseteq \mathbb{R}^m$  be a non-empty compact set and consider a discrete-time system with state  $x \in \mathcal{X} \subseteq \mathbb{R}^n$  and control input  $u \in \mathcal{U}$ , evolving as  $x_{t+1} = f(x_t, u_t)$ . Let  $[t, H] := \{t, t+1, \dots, H\}$  and  $\mathbb{U}_{t:H}$  denote the set of input sequences  $\mathbf{u}_{t:H} := (u_t, \dots, u_H)$ . Let  $\mathbf{x}(\tau; x_0, \pi)$ ,  $\tau \geq 0$ , denote a state trajectory starting from  $x_0$  under control policy  $\pi: \mathcal{X} \rightarrow \mathcal{U}$  and dynamics  $f$ . For an open set  $\mathcal{F} \subseteq \mathcal{X}$  of failure states, there always exists a Lipschitz-continuous function  $g: \mathcal{X} \rightarrow \mathbb{R}$  such that  $g(x) < 0 \iff x \in \mathcal{F}$  (a signed distance function).

### A. Controlled-invariant Set

A (closed) set  $\Omega \subset \mathcal{X}$  is *controlled-invariant* under policy  $\pi^b$  if for all  $x_0 \in \Omega$  and all  $\tau \geq 0$ ,  $\mathbf{x}(\tau; x_0, \pi^b) \in \Omega$ . The goal of safety analysis is to find a controlled-invariant set  $\Omega$ , and its corresponding  $\pi^b$ , such that  $\Omega \cap \mathcal{F} = \emptyset$ . In this

All authors are with the Department of Electrical and Computer Engineering, Princeton University, New Jersey, USA. email: arkumar@princeton.edu

context,  $\Omega$  is called the *safe set*. Additionally, to reduce conservativeness, we desire  $\Omega$  to be as large as possible (ideally, close to the maximal safe set). Once such a safe set is constructed, a *least-restrictive safety filter*  $\pi^{\text{sf}}$  can be applied to any performance-driven control policy  $\pi^{\text{p}}: \mathcal{X} \rightarrow \mathcal{U}$ , enforcing the safety-oriented control  $\pi^{\text{b}}(x)$  whenever using  $\pi^{\text{p}}(x)$  would lead to leaving the safe set  $\Omega$ :

$$\pi^{\text{sf}}(x) := \begin{cases} \pi^{\text{p}}(x), & f(x, \pi^{\text{p}}(x)) \in \Omega \\ \pi^{\text{b}}(x), & \text{otherwise.} \end{cases} \quad (1)$$

### B. Control Barrier Functions

Control barrier functions (CBFs) provide Lyapunov-like conditions to guarantee the forward invariance of a (safe) set and have been applied in various safety-critical control problems [9]. Agrawal and Sreenath [10] extend the CBF initially developed for the continuous-time dynamic systems to the discrete-time domain. Formally, they consider a continuously differentiable function  $B: \mathcal{X} \rightarrow \mathbb{R}$  whose zero superlevel set is the safe set  $\Omega \subseteq \mathcal{X}$ , i.e.,  $\Omega = \{x \in \mathcal{X} \mid B(x) \geq 0\}$ . A function  $B$  is a discrete-time exponential CBF for dynamics  $f$  if, for some  $\gamma \in (0, 1]$  and for each  $x \in \Omega$ , there exists a control  $u \in \mathcal{U}$  such that  $B(f(x, u)) \geq \gamma B(x)$  [10]. In other words, such control input  $u$  keeps the state trajectory inside  $\Omega$  and ensures that its boundary can only be approached asymptotically. We can incorporate this CBF constraint into an online trajectory optimization as

$$\pi^{\text{sf}}(x) = \underset{u \in \mathcal{U}}{\operatorname{argmin}} \quad \|u - \pi^{\text{p}}(x)\|_2^2 \quad (2a)$$

$$\text{s.t. } B(f(x, u)) \geq \gamma B(x). \quad (2b)$$

If dynamics are control-affine and the CBF is quadratic, (2) can be solved with a quadratically constrained quadratic program (QCQP) [10]. This enables real-time computation even for high-dimensional dynamics, making it attractive for many robotics applications. However, finding a valid CBF for general dynamic systems is difficult. Although there are successes of heuristically constructed CBFs, these do not recover the maximal safe set and may be overly conservative [12]. Further, the safety guarantee is contingent on the feasibility of the quadratic program, which could be lost in the presence of tight actuation limits or other control constraints.

### C. Hamilton-Jacobi Reachability Analysis

Hamilton-Jacobi (HJ) reachability analysis defines an optimal control problem to find the maximal safe set (or *viability kernel*)  $\Omega^*$ . The value function  $V: \mathcal{X} \times [0, H] \rightarrow \mathbb{R}$ ,  $V(x_t, t) := \max_{\pi^{\text{b}}} \min_{\tau \in [t, H]} g(\mathbf{x}(\tau - t; x_t, \pi^{\text{b}}))$  encodes the minimum safety margin that the safety policy  $\pi^{\text{b}}$  can maintain over all times up to a planning horizon  $H$ . The value function is then computed by finding the solution of the finite-horizon terminal-value problem with safety Bellman equation [6]

$$V(x_t, t) := \min \left\{ g(x_t), \max_{u_t \in \mathcal{U}} V(f(x_t, u_t), t + 1) \right\} \quad (3a)$$

$$V(x_H, H) := g(x_H). \quad (3b)$$

Note that this value function only allows finite-horizon safety. So in practice, HJ reachability analysis usually considers a sufficiently long horizon to guarantee infinite-horizon safety, i.e.,  $V^\infty(x) := \lim_{H \rightarrow \infty} \max_{\pi^{\text{b}}} \min_{\tau \in [0, H]} g(\mathbf{x}(\tau; x, \pi^{\text{b}}))$ . Numerical level-set methods have been developed to compute the HJ value function, and require discretizing the state space into a finite grid; thus, computation scales exponentially with the dimensionality of the state space, becoming practically prohibitive beyond 6 continuous state variables [14]. If this value function can be computed, we can retrieve the viability kernel by  $\Omega^* := \{x \in \mathcal{X} \mid V^\infty(x) \geq 0\}$  and the safety policy  $\pi^{\text{b}*}(x) = \operatorname{argmax}_{u \in \mathcal{U}} V^\infty(f(x, u))$ .

Additionally, wherever this value function is differentiable, it is a valid CBF; thus, HJ reachability analysis provides a principled way to construct a CBF. Due to the minimum operator in (3),  $V^\infty(f(x, \pi^{\text{b}*}(x))) \geq V^\infty(x)$ , the safety policy  $\pi^{\text{b}*}$  always prevents the safety value from decreasing. Thus, this HJ value function is often paired with the least-restrictive control law in (1), only overriding the proposed control input near the boundary of the safe set  $\Omega^*$  [2]. However, this least-restrictive scheme can result in jerky control switches and is susceptible to numerical issues near the safe set boundary. In contrast, CBF-constrained optimization (2) provides a smoother transition at the cost of additional restrictions to the control input away from the safe set boundary.

### D. Model Predictive Control

Model predictive control (MPC) [15] verifies the system's safety on the fly instead of relying on a pre-computed value function or CBF. In the context of safety filters, MPC can be used to optimize a control sequence such that the first control is the closest to the task-oriented control but satisfies all the constraints below

$$\min_{u_{k|t}, x_{k|t}} \|\pi^{\text{p}}(x_t) - u_{0|t}\| \quad (4a)$$

$$\text{s.t. } \forall k \in [0, H - 1]:$$

$$x_{0|t} = x_t, \quad x_{k+1|t} = f(x_{k|t}, u_{k|t}), \quad (4b)$$

$$x_{k|t} \notin \mathcal{F}, \quad u_{k|t} \in \mathcal{U}, \quad (4c)$$

$$x_{H|t} \in \mathcal{T}. \quad (4d)$$

In addition to the dynamics constraint (4b) and state and control constraints (4c), MPC adds a terminal constraint (4d) confining the final state to a target  $\mathcal{T} \subset \mathcal{X}$ , rendered controlled-invariant by a known  $\pi^{\mathcal{T}}$ . After the optimization, only the first control will be executed, i.e.,  $\pi^{\text{sf}}(x_t) = u_{0|t}$ . At the next time step  $t + 1$ , (4) is solved again with a shifted time horizon. Importantly, constraint (4d) ensures recursive feasibility, as we can repeat the control sequence  $(u_{1:H-1|t}, \pi^{\mathcal{T}}(x_{H|t}))$  at  $t + 1$ . This recursive feasibility in turn guarantees all-time safety.

### E. Reach-Avoid Analysis

MPC requires the system's state to be inside a known controlled-invariant set exactly at the end of the planning

horizon. However, we can loosen this requirement by considering a reach-avoid problem, which finds a control policy to guide the system's state into a (closed) target set at any time within the planning horizon without breaching the constraints at prior times. Specifically, the *reach-avoid* set for horizon  $H$  is

$$\mathcal{RA}(\mathcal{T}, \mathcal{F}, H) := \{x \in \mathcal{X} \mid \exists \pi^b, t \in [0, H], \\ \mathbf{x}(t; x, \pi^b) \in \mathcal{T} \wedge \forall \tau \in [0, t], \mathbf{x}(\tau; x, \pi^b) \notin \mathcal{F}\}. \quad (5)$$

Introducing an additional Lipschitz-continuous target margin function  $\ell: \mathcal{X} \rightarrow \mathbb{R}$  such that  $\ell(x) \geq 0 \iff x \in \mathcal{T}$ , we can follow an analogous formulation to II-C, defining the reach-avoid objective along horizon  $[t, H]$  for arbitrary  $t \in [0, H]$ :

$$J(x_t, \pi^b, t) := \max_{\tau \in [t, H]} \min \left\{ \ell(x_\tau), \min_{s \in [t, \tau]} g(x_s) \right\} \quad (6)$$

with  $x_\tau = \mathbf{x}(\tau - t; x_t, \pi^b)$ . The associated value function  $V(x_t, t) := \max_{\pi^b} J(x_t, \pi^b, t)$  satisfies the reach-avoid Bellman equation [5, Appendix A]:

$$V(x_t, t) = \min \left\{ g(x_t), \max \left\{ \ell(x_t), \max_{u_t \in \mathcal{U}} V(f(x_t, u_t), t + 1) \right\} \right\}, \quad (7)$$

with terminal condition  $V(x_H, H) = \min \{g(x_H), \ell(x_H)\}$ . The reach-avoid set (5) is characterized by the sign of the value function at  $t = 0$ , as  $V(x, 0) \geq 0 \iff x \in \mathcal{RA}(\mathcal{T}, \mathcal{F}, H)$ .

Recent work addresses the *curse of dimensionality* (the exponential scaling of computation) by approximate methods such as state-space decomposition [4], deep reinforcement learning [5, 6], sum-of-squares (SOS) optimization [7] or differential dynamic programming (DDP) [16], in particular, using iterative linear-quadratic (ILQ) optimization [8, 17].

### III. REACH-AVOID CONTROL BARRIER FUNCTIONS

While least-restrictive safety frameworks provide maximum flexibility to task-driven control policies, they may result in abrupt control switches. On the other hand, hand-crafted CBFs are often overly conservative and difficult to design for high-dimensional systems. All of the above methods rely on offline computation and suffer under unforeseen scenarios, such as obstacles emerging into the system's field of view.

This section introduces a novel scheme to systematically construct control barrier functions at runtime, by computing a reach-avoid value function through differential dynamic programming. Specifically, we solve a *receding-horizon* ILQ optimization problem until convergence and use the resulting value function to construct a quadratically constrained quadratic program (QCQP), allowing for a smooth transition between task-oriented and safety-oriented control. We prove that the proposed CBF-DDP ensures infinite-horizon safety through recursively feasible finite-horizon planning by reaching a controlled-invariant target set.

#### A. Barrier Safety Filter with Reach-Avoid Value Function

We begin by laying out the theoretical foundations of our approach. In order to provide infinite-horizon safety guarantees, we assume the existence of a known target set with the following two properties: (1) the target set does not intersect with the failure set,  $\mathcal{T} \cap \mathcal{F} = \emptyset$ , and (2) the target set is controlled-invariant under a known policy  $\pi^T$ .

We note that the above are mild assumptions, satisfied by broad classes of systems of practical interest: for example, many robots and vehicles are able to come to a stop, or alternatively transition into a safe loiter orbit, after which the system can remain safe indefinitely under benign engineering assumptions. While more sophisticated methods may be used to determine the controlled-invariant regions of attraction around these states (e.g. Lyapunov analysis), here we illustrate a simpler approach. We construct such a target set using a conservative stopping policy  $\pi^{\text{stop}}$ , which brings the robot to a complete halt and maintains the robot at rest. The target margin function  $\ell$  can be rapidly computed from each state as the signed distance to the failure set along the robot's hypothetical stopping path from that state. Being in the target set, then, is equivalent to being able to safely come to a halt using  $\pi^{\text{stop}}$ . Specifically,  $\ell$  is defined as:

$$\ell(x_t) = \min g(\mathbf{x}(\tau - t; x_t, \pi^{\text{stop}})).$$

At each state  $x_t$ ,  $\tau \geq t$  are provided with an arbitrary task-oriented control  $u^p$  proposed by the task policy  $\pi^p$ . For the corresponding control cycle  $t$ , we formulate a constrained program, which seeks to minimize deviation  $\delta u$  from  $u^p$  while satisfying a control barrier constraint as shown below:

$$\min_{\delta u: \delta u + u^p \in \mathcal{U}} \frac{1}{2} \|\delta u\|_2^2 \quad (8a)$$

$$\text{s.t.} \quad V(x_{1|t}, 0) \geq \gamma V(x_t, 0), \quad (8b)$$

where  $x_{1|t} = f(x_t, u^p + \delta u)$ ,  $V(x_t, 0)$  is the reach-avoid value (7) at the current state and  $V(x_{1|t}, 0)$  is the reach-avoid value at the state reached after applying control  $u^p + \delta u$ . Generally, computing these two values is computationally prohibitive for real-time decision-making. Instead, we use DDP to compute the quadratic approximation of the reach-avoid value function and the (locally) optimal safe control.

We use the iterative linear-quadratic approach outlined in Alg. 2 to compute these value functions. Additionally,  $V(x_{1|t}, 0)$  in (8b) is approximated by

$$V(x_{1|t}, 0) \approx V(\bar{x}_{1|t}, 0) + \delta x_{1|t}^\top z + \frac{1}{2} \delta x_{1|t}^\top Z \delta x_{1|t}, \quad (9)$$

where  $\bar{x}_{1|t} = f(x_t, u^p)$  and  $\delta x_{1|t} = x_{1|t} - \bar{x}_{1|t} \approx f_u \delta u$  is obtained from dynamics Jacobian  $f_u := \frac{\partial f}{\partial u}|_{x_t, u^p}$ . Then, (8b) is approximated as a quadratic constraint:  $\delta u^\top P \delta u + p^\top \delta u + c \geq 0$ , where  $P = \frac{1}{2} f_u^\top Z f_u$  is symmetric,  $p = f_u^\top z$ , and  $c = V(\bar{x}_{1|t}, 0) - \gamma V(x_t, 0)$ . Thus, (8) is reduced to a quadratically constrained quadratic program (QCQP) as shown below:

$$\min_{\delta u: \delta u + u^p \in \mathcal{U}} \|\delta u\|_2^2 \quad (10a)$$

$$\text{s.t.} \quad \delta u^\top P \delta u + p^\top \delta u + c \geq 0. \quad (10b)$$

The final executed control is  $u_t = \pi^{\text{sf}}(x_t) := u^p + \delta u^*$ , where  $\delta u^*$  is an optimal solution to (8). Henceforth, we represent the parameters of the QCQP in (10) as  $\Phi = (P, p, c)$ .

*Lemma 1:* For the QCQP with parameters  $\Phi := (P, p, c)$  in (10), if  $P$  is negative definite ( $P \prec 0$ ), then optimal solution  $\delta u^*$  is unique and continuous in  $\Phi$ . Additionally, if  $P = 0$ , then  $\delta u^*$  is unique and Lipschitz continuous in  $\Phi$ .

*Proof:* If  $P = 0$ , the CBF-QP has a known closed-form solution. Thm. 3 in [9] proves Lipschitz continuity. If  $P \prec 0$ , the QCQP is a strictly convex QP with a strictly convex constraint. Thus,  $\delta u^*$  is unique. The level sets of the constraint are bounded ellipsoids and are hence compact sets. The objective and quadratic constraint of (10) are continuous in  $\delta u, \Phi$ . From Berge's maximum theorem [18], the set of optimizers is compact and upper hemi-continuous (UHC). Since  $\delta u^*$  is unique, UHC is equivalent to continuity. ■

*Lemma 2:* Given  $\forall x, u$ , almost everywhere (a.e.),  $V \in \mathcal{C}^2$  with respect to (w.r.t.)  $x$  and  $f \in \mathcal{C}^2$  w.r.t.  $(x, u)$ ,  $\pi^p$  is continuous in  $x$  and  $P \preceq 0$ , then  $\exists$  unique  $\delta u^*$  continuous in  $x$  a.e..

*Proof:* Given  $V, f$  is  $\mathcal{C}^2$  a.e.,  $V, f$  is  $\mathcal{C}^2$  at  $f(x, \pi^p(x))$  almost surely. Hence,  $\Phi := (P, p, c)$  is continuous in  $x$  a.e.. Using Lemma 1 and the continuity of composition of continuous functions,  $\delta u^*$  is unique, continuous in  $x$  a.e.. ■ A sufficient condition for continuous  $\delta u^*$  is if  $V, f$  is  $\mathcal{C}^2$ . We choose  $\ell, g, f$  to be piecewise  $\mathcal{C}^2$  to aid this requirement. To use Lemma 1, we check whether  $P \prec 0$  and switch to solving with a linear constraint if this condition is violated. In practice,  $\ell$  and  $g$  are chosen as the maximum of multiple safety constraints that are independently  $\mathcal{C}^2$ . There can be isolated points of discontinuity in  $u$ . Compared to the least-restrictive baseline, these discontinuities appears away from the boundary of  $\Omega^*$  and does not result in chatter along the boundary.

*Lemma 3:* Suppose  $x_t \in \mathcal{RA}(\mathcal{T}, \mathcal{F}, H)$  and  $x_t \notin \mathcal{T}$ , then constraint (8b) is feasible. A feasible solution is  $\delta u = u^b - u^p$  with  $u^b = \arg \max_{u \in \mathcal{U}} V(f(x_t, u), 1)$ .

*Proof:* Since  $x_t \in \mathcal{RA}(\mathcal{T}, \mathcal{F}, H)$ , we have  $V(x_t, 0) \geq 0$ . Also, since  $x_t \notin \mathcal{T}$ ,  $\ell(x_t) < 0$ . Consequently, from (7),

$$V(x_t, 0) = \min \left\{ g(x_t), \max_{u \in \mathcal{U}} V(f(x_t, u), 1) \right\} \leq \max_{u \in \mathcal{U}} V(f(x_t, u), 1).$$

Now, using (6) and (7), for any  $x_0$ ,

$$\begin{aligned} V(x_0, 0) &= \max_{\pi^b} \max_{\tau \in [0, H]} \min \left\{ \ell(x_\tau), \min_{s \in [0, \tau]} g(x_s) \right\} \\ &\geq \max_{\pi^b} \max_{\tau \in [0, H-1]} \min \left\{ \ell(x_\tau), \min_{s \in [0, \tau]} g(x_s) \right\} = V(x_0, 1), \end{aligned}$$

where  $x_\tau = \mathbf{x}(\tau; x_0, \pi^b)$ . From above observations, we note:

$$\max_{u \in \mathcal{U}} V(f(x_t, u), 0) \geq \max_{u \in \mathcal{U}} V(f(x_t, u), 1) > \gamma V(x_t, 0).$$

for any  $\gamma \in [0, 1)$ . This ensures that  $\exists \delta u: \delta u + u^p \in \mathcal{U}$  which guarantees the feasibility of (8b), e.g.,  $\delta u = u^b - u^p$ . ■

With the above Lemmas, we can prove our results of interest.

*Theorem 1 (Recursive Safety):* Suppose that  $\mathcal{T} \cap \mathcal{F} = \emptyset$  and  $\mathcal{T}$  is controlled-invariant. If  $x_t \in \mathcal{RA}(\mathcal{T}, \mathcal{F}, H)$ , then  $\exists u_t$  such that  $x_{t+1} = f(x_t, u_t) \in \mathcal{RA}(\mathcal{T}, \mathcal{F}, H)$ . Furthermore,  $u^b$  is a feasible solution for  $x_{t+1} \in \mathcal{RA}(\mathcal{T}, \mathcal{F}, H)$ .

*Proof:* If  $x_t \in \mathcal{T}$ , the theorem is proven by choice of  $\mathcal{T}$  being controlled-invariant with a known policy  $\pi^{\mathcal{T}}$ . By definition,  $u^b$  achieves  $V(f(x_t, u^b), 1) \geq V(f(x_t, \pi^{\mathcal{T}}(x_t)), 1) \geq 0$ . If  $x_t \notin \mathcal{T}$ , feasibility of (8b) from Lemma 3 ensures  $V(x_{t+1}, 0) \geq \gamma V(x_t, 0) \geq 0$  and therefore  $x_{t+1} \in \mathcal{RA}(\mathcal{T}, \mathcal{F}, H)$ . Choosing  $\pi^{\text{sf}}(x) = u^b$  is sufficient for feasibility of (8b). ■

*Theorem 2 (Infinite-Horizon Safety):* Given  $x_0 \in \mathcal{T}$ , Alg. 1 ensures  $x_t \in \mathcal{RA}(\mathcal{T}, \mathcal{F}, H), \forall t$ . Further, Alg. 1 maintains a fallback policy  $\pi^b$  ensuring safety for infinite-horizon.

*Proof:* Suppose  $x_t \in \mathcal{RA}(\mathcal{T}, \mathcal{F}, H)$ , then from Thm. 1, we get  $x_{t+1} \in \mathcal{RA}(\mathcal{T}, \mathcal{F}, H)$ . Given  $x_0 \in \mathcal{T}$ , we see  $x_t \in \mathcal{RA}(\mathcal{T}, \mathcal{F}, H), \forall t$ . Also, given,  $x_t \in \mathcal{RA}(\mathcal{T}, \mathcal{F}, H)$ ,  $\exists \pi^b$  that ensures  $\mathbf{x}(\tau - t; x_t, \pi^b) \in \mathcal{T}$  for some  $\tau \in [t, t + H]$ . Alg. 1 maintains a sequence of fallback policies  $\pi^b$  constructed by Alg. 3 capable of guiding the robot inside the target set. The construction of  $\pi^b$  is valid by choice of using closed-loop policies from DDP until entering  $\mathcal{T}$  and  $\pi^{\mathcal{T}}$  thereafter. ■

*Theorem 3 (Continuity of controls):* Given  $\forall x, u$ , a.e.,  $V \in \mathcal{C}^2$  w.r.t.  $x, f \in \mathcal{C}^2$  w.r.t.  $(x, u)$  and  $\pi^p$  is continuous in  $x$ , then Alg. 1 provides  $u_t$  continuous in  $x_t$  a.e.

*Proof:* Alg. 1 solves (10) iteratively until (8b) is satisfied or the computational budget is exhausted. In each QCQP iteration,  $\delta u^*$  is continuous in  $x$  a.e. from Lemma 2. If the QCQP is not feasible, the reduction to QP is guaranteed to be feasible with  $\delta u^*$  continuous in  $x_t$  a.e.. Given  $x_t$ ,  $u_{1|t} = \pi^p(x_t) + \delta u^*$  is a continuous function of  $x_t$  a.e. The final  $\delta u^*$  after multiple QCQP iterations is continuous if at each QCQP iteration,  $V, f \in \mathcal{C}^2$  at  $f(x_t, u_{1|t})$ . Alg. 1 almost surely does not encounter a discontinuity during the QCQP iterations. Thus, Alg. 1 provides  $u_t$  continuous in  $x_t$  a.e. ■

## B. Implementation Details

The flow of our CBF-DDP is outlined in Alg. 1<sup>1</sup>. Each time, the core subprocedure of the algorithm solves the QCQP in (10). Since this QCQP only solves a local quadratic approximation of the non-linear constraint in (8b), it necessitates repeatedly updating the initial control and re-solving the QCQP until (8b) is satisfied. To reduce the number of QCQPs we solve, we use a scaling parameter  $\lambda \in \mathbb{R}$  on  $c$ . With  $c < 0$ , by scaling  $c$  higher than needed, we get closer to satisfying (10) faster. This inner loop is allowed to find a feasible solution within five attempts or is terminated. If QCQP cannot find a feasible solution and  $V(x_{1|t}, 0) < 0$ , we use  $\pi^b$ , thus guaranteeing fully safe operation. This fallback never occurs in the simulations as long as  $\gamma$  is chosen large enough that  $V$  does not decrease rapidly to zero.

<sup>1</sup>The sequence of fallback policies can be initialized with  $\pi_{0:H-1}^b \leftarrow \pi^{\mathcal{T}}$ .

---

**Algorithm 1** CBF-DDP
 

---

**Require:**  $x_0 \in \mathcal{T}$ , fallback policies (from  $x_0$ )  $\pi^b$ , policy for target set  $\pi^T$ , task policy  $\pi^P$ , scaling  $\lambda \geq 1$ , CBF discount  $\gamma \in (0, 1)$

- 1: **for**  $t = 0, \dots, T$  **do**
- 2:  $u^P \leftarrow \pi^P(x_t), \bar{x}_{1|t} \leftarrow f(x_t, u^P)$
- 3:  $V(\bar{x}_{1|t}, 0), f_u, Z, z, \bar{x}, \bar{u}, \mathbf{K}, \mathbf{k} \leftarrow \text{ReachAvoidILQ}(\bar{x}_{1|t})$
- 4:  $P \leftarrow f_u^\top Z f_u, p \leftarrow f_u^\top z, c = V(\bar{x}_{1|t}, 0) - \gamma V(x_t, 0)$
- 5: **while** Equation (8b) violated **or not** timeout **do**
- 6:  $u^P \leftarrow \text{BarrierFilter}(P, p, c), \bar{x}_{1|t} \leftarrow f(x_t, u^P)$   $\triangleright$  Solve QCQP (8)
- 7:  $V(\bar{x}_{1|t}, 0), f_u, Z, z, \bar{x}, \bar{u}, \mathbf{K}, \mathbf{k} \leftarrow \text{ReachAvoidILQ}(\bar{x}_{1|t})$
- 8:  $P \leftarrow f_u^\top Z f_u, p \leftarrow f_u^\top z, c = V(\bar{x}_{1|t}, 0) - \gamma V(x_t, 0)$
- 9:  $c \leftarrow \lambda c$
- 10: **if**  $V(\bar{x}_{1|t}, 0) < 0$  **then**
- 11:  $u^b \leftarrow \pi_0^b(x_t), \pi_{0:H-2}^b \leftarrow \pi_{1:H-1}^b, \pi_{H-1}^b \leftarrow \pi^T$ .
- 12: **Apply**  $u^b$   $\triangleright$  Fall back to full safety filter if  $u^P$  is unsafe
- 13: **else**
- 14:  $\pi^b \leftarrow \text{Fallback}(\bar{x}, \bar{u}, \mathbf{K}, \mathbf{k}, \pi^T, H)$   $\triangleright$  Update fallback policies
- 15: **Apply**  $u^P$   $\triangleright$  Apply control from CBF-DDP

---



---

**Algorithm 2** Reach-Avoid ILQ
 

---

**Require:**  $x_0$

- 1: Initialize control sequence  $\bar{u} \leftarrow 0$  (or warm-start with previous iteration), feedback gains  $\mathbf{K} \leftarrow 0$ , open-loop gains  $\mathbf{k} \leftarrow 0$ , and roll out for  $\bar{x}$
- 2: **while** not converged **do**
- 3:  $V, f_u, \bar{x}, \bar{u}, \{\ell_t, g_t\} \leftarrow \text{Forward}(\bar{x}, \bar{u}, \mathbf{K}, \mathbf{k})$
- 4:  $\{K_t, k_t\}, Z_0, z_0 \leftarrow \text{Backward}(\bar{x}, \bar{u}, \{\ell_t, g_t\})$
- 5: **return**  $V, f_u, Z_0, z_0, \bar{x}, \bar{u}, \mathbf{K}, \mathbf{k}$

---

## IV. SIMULATION RESULTS

In many applications, such as autonomous driving, it is desirable to create a smooth safety filter that maintains safety desiderata. The least-restrictive safety filter causes jerky behavior, which affects the comfort of the rider in a car. We show our CBF-DDP method to result in a smoother and more desirable behavior. Our CBF-DDP algorithm uses  $\lambda \in [1.2, 1.3]$  in all our runs. In all simulations, we use the Runge–Kutta method (RK4) for integration with sampling time  $\delta t = 0.05$ .

## A. Dubins Car

We illustrate CBF-DDP in a Dubins car with velocity fixed at 0.7. The equations of motion are  $\dot{x} = v \cos(\theta)$ ,  $\dot{y} = v \sin(\theta)$ ,  $\dot{\theta} = u_1$ . A circular footprint with a radius of 0.15 is chosen for the ego-vehicle. The steering control limits are  $[-1.0, 1.0]$ . We use a test case where the Dubins car goes around an obstacle. The task-oriented policy is to turn in the direction facing the goal. In this setup, we don't use a reach-avoid objective function.

---

**Algorithm 3** Construct Fallback Policy
 

---

**Require:**  $\bar{x}, \bar{u}, \mathbf{K}, \mathbf{k}, \pi^T, H$ .

- 1: **for**  $t \leftarrow 0$  to  $H - 1$  **do**
- 2: **if**  $\bar{x}_t \in \mathcal{T}$  **then**  $\pi_{t:H-1}^b \leftarrow \pi^T$  and break the for-loop
- 3: **else**  $\pi_t^b(x_t) \leftarrow \bar{u}_t + K_t(x_t - \bar{x}_t) + k_t$
- 4: **return**  $\pi^b$

---




Fig. 1: Dubins car: Comparison of safety filters CBF-DDP, LR-DDP and a manual CBF. (Left) Trajectories of the car. (Right) Steer control deviation from  $\pi^P$ .

We use  $g$  in a *reachability* [8] formulation where the car is constrained to be in a set from which it avoids all obstacles.

The comparison of CBF-DDP with LR-DDP is illustrated in Fig. 1 with  $H = 40$ . We see that LR-DDP filter frequently switches between the task and safety control near the unsafe boundary, resulting in jerky movements. In contrast, the CBF-DDP filter activates earlier and smoothly controls the car. The number of QCQP iterations is limited to 2. We also compare it to a manual CBF that is known to be sub-optimal. The manual CBF provides smooth controls but is more conservative because of the sub-optimality.

## B. Kinematic Bicycle Dynamics

We consequently test on dynamics with two controls, i.e., acceleration and steering. The equations of motion are  $\dot{x} = v \cos(\theta)$ ,  $\dot{y} = v \sin(\theta)$ ,  $\dot{v} = u_1$ ,  $\dot{\theta} = \frac{v}{L} \tan \delta$ ,  $\dot{\delta} = u_2$ . A circular footprint with a radius of 0.1 is chosen for ego-vehicle. The acceleration control limits  $u_1$  are  $[-0.5, 0.5]$  and steering control limits  $u_2$  are  $[-0.8, 0.8]$ . Here, we use a reach-avoid value function with the target margin function computed based on the stopping path criterion. For  $\pi^P$ , we use a naive heuristic where the steering is pursuing a look-ahead point near the center of the road. This policy has a switch when the car is already moving in a desired direction. Further, this  $\pi^P$  feeds acceleration to maintain velocity at 0.9 using linear feedback. We run our experiments on a 3.6 GHz CPU with 64 GB RAM.

The CBF-DDP is robust and provides smooth controls (Fig. 2) for these dynamics, but LR-DDP makes a U-turn or comes to a safe halt resulting in task failure. LR-DDP can be improved by constraining the horizontal yaw deviation  $\delta\theta$ . The  $\ell, g$  are then defined using the worst-case margin of all constraints. The CBF-DDP provides smoother trajectories and fast task completion (Fig. 2) even with the constraint, but the controls get jerkier with tighter limits on  $\delta\theta$ . Discontinuities in derivatives of  $V$  from interacting constraints between two obstacles or between obstacle and yaw can induce discontinuities in controls from CBF-DDP. Future practical improvements can incorporate soft versions of minimum operators to ease some of these discontinuities. Table 1 compares LR-DDP and CBF-DDP with  $H = 45$  and a constraint on  $\delta\theta$ . LR-DDP completes the task, but the trajectories are quantitatively jerkier than CBF-DDP. From Table 1, the total integral deviation ( $\|\pi^P - \pi^{\text{sf}}\|_1$ ) from  $\pi^P$  is less for CBF-DDP, indicating that it is not overly restrictive. With stronger task policies, the yaw constraint is not needed giving an advantage to CBF-DDP.

| Constraints |                    | Task Success |     | Acceleration Jerk (mean and std) |              | Steer Jerk (mean and std) |             | Total Deviation from $\pi^P$ |       | Control cycle time (s) |             |
|-------------|--------------------|--------------|-----|----------------------------------|--------------|---------------------------|-------------|------------------------------|-------|------------------------|-------------|
| Road limit  | yaw $\delta\theta$ | LR           | CBF | LR                               | CBF          | LR                        | CBF         | LR                           | CBF   | LR                     | CBF         |
| 1.2         | None               | ×            | ✓   | ×                                | 0.009 ± 0.02 | ×                         | 0.05 ± 0.06 | ×                            | 52.81 | ×                      | 0.10 ± 0.10 |
| 1.2         | 0.5 $\pi$          | ✓            | ✓   | 0.03 ± 0.14                      | 0.01 ± 0.02  | 0.06 ± 0.19               | 0.05 ± 0.08 | 93.23                        | 50.46 | 0.21 ± 0.22            | 0.27 ± 0.39 |
| 1.2         | 0.4 $\pi$          | ✓            | ✓   | 0.03 ± 0.15                      | 0.01 ± 0.04  | 0.10 ± 0.27               | 0.08 ± 0.13 | 59.50                        | 53.79 | 0.22 ± 0.20            | 0.38 ± 0.52 |
| 1.0         | None               | ×            | ✓   | ×                                | 0.009 ± 0.02 | ×                         | 0.04 ± 0.06 | ×                            | 52.81 | ×                      | 0.10 ± 0.09 |
| 1.0         | 0.5 $\pi$          | ×            | ✓   | ×                                | 0.01 ± 0.03  | ×                         | 0.06 ± 0.11 | ×                            | 49.47 | ×                      | 0.28 ± 0.39 |
| 1.0         | 0.4 $\pi$          | ✓            | ✓   | 0.05 ± 0.17                      | 0.01 ± 0.03  | 0.11 ± 0.32               | 0.07 ± 0.11 | 78.15                        | 51.93 | 0.20 ± 0.19            | 0.37 ± 0.52 |

TABLE 1: Bicycle: Comparison of LR-DDP vs CBF-DDP in obstacle avoidance with naive  $\pi^P$ . The maximum allowable yaw deviation from the horizontal  $\delta\theta$  is varied. CBF-DDP is robust to road boundaries and yaw constraints. The discontinuity of controls in CBF-DDP increases as the yaw constraint is tightened. On the other hand, LR-DDP tends to halt or U-turn safely without a yaw constraint.




Fig. 2: (Left) Comparison of CBF-DDP and LR-DDP with varying  $\delta\theta$ . Trajectories from CBF-DDP are much more well-behaved. Videos available at <https://tinyurl.com/2pr8dy6j>. (Right) Comparison of controls from CBF-DDP and LR-DDP with favorable  $\delta\theta$  for each. CBF-DDP gives smoother controls in the blue-shaded regions where the filter is active.

## V. CONCLUSIONS

We propose a principled approach to construct a CBF online with reach-avoid DDP. We use a receding horizon to provide infinite-horizon recursive feasibility despite finite-horizon planning. The construction of our reach-avoid value functions guarantees safety. Simulations on a bicycle model show that CBF-DDP improves the smoothness over least-restrictive control and avoids the conservativeness of handcrafted CBFs.

## REFERENCES

- [1] J. F. Fisac, M. Chen, C. J. Tomlin, and S. S. Sastry. “Reach-Avoid Problems with Time-Varying Dynamics, Targets and Constraints”. *Proceedings of the International Conference on Hybrid Systems: Computation and Control*. Seattle, Washington, 2015, pp. 11–20.
- [2] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin. “Hamilton-Jacobi reachability: A brief overview and recent advances”. *Proceedings of the IEEE Annual Conference on Decision and Control (CDC)*. 2017, pp. 2242–2253.
- [3] I. M. Mitchell. “The flexible, extensible and efficient toolbox of level set methods”. *Journal of Scientific Computing* 35.2 (2008), pp. 300–329.
- [4] M. Chen, S. L. Herbert, M. S. Vashishtha, et al. “Decomposition of Reachable Sets and Tubes for a Class of Nonlinear Systems”. *IEEE Transactions on Automatic Control* 63.11 (2018), pp. 3675–3688.
- [5] K.-C. Hsu, V. Rubies-Royo, C. J. Tomlin, and J. F. Fisac. “Safety and Liveness Guarantees through Reach-Avoid Reinforcement Learning”. *Proceedings of Robotics: Science and Systems*. Virtual, July 2021.
- [6] J. F. Fisac, N. F. Lugovoy, V. Rubies-Royo, et al. “Bridging Hamilton-Jacobi Safety Analysis and Reinforcement Learning”. *Proceedings of the International Conference on Robotics and Automation (ICRA)*. 2019, pp. 8550–8556.
- [7] B. Landry, M. Chen, S. Hemley, and M. Pavone. “Reach-Avoid Problems via Sum-or-Squares Optimization and Dynamic Programming”. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2018, pp. 4325–4332.
- [8] D. Fridovich-Keil and C. J. Tomlin. “Approximate Solutions to a Class of Reachability Games”. *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 2021, pp. 12610–12617.
- [9] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada. “Control Barrier Function Based Quadratic Programs for Safety Critical Systems”. *IEEE Transactions on Automatic Control* 62.8 (2017), pp. 3861–3876.
- [10] A. Agrawal and K. Sreenath. “Discrete Control Barrier Functions for Safety-Critical Control of Discrete Systems with Application to Bipedal Robot Navigation”. *Proceedings of Robotics: Science and Systems*. Cambridge, Massachusetts, July 2017.
- [11] Y. Chen, M. Jankovic, M. Santillo, and A. D. Ames. “Backup Control Barrier Functions: Formulation and Comparative Study”. *Proceedings of the 60th IEEE Conference on Decision and Control (CDC)*. 2021, pp. 6835–6841.
- [12] J. J. Choi, D. Lee, K. Sreenath, et al. “Robust Control Barrier-Value Functions for Safety-Critical Control”. *Proceedings of the 60th IEEE Conference on Decision and Control (CDC)*. 2021, pp. 6814–6821.
- [13] S. Tonkens and S. Herbert. “Refining Control Barrier Functions through Hamilton-Jacobi Reachability”. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2022, pp. 13355–13362.
- [14] M. Bui, G. Giovanis, M. Chen, and A. Shriraman. “OptimizedDP: An Efficient, User-friendly Library For Optimal Control and Dynamic Programming”. 2022.
- [15] D. Mayne, J. Rawlings, C. Rao, and P. Scokaert. “Constrained model predictive control: Stability and optimality”. *Automatica* 36.6 (2000), pp. 789–814.
- [16] J. F. A. D. O. PANTOJA. “Differential dynamic programming and Newton’s method”. *International Journal of Control* 47.5 (1988), pp. 1539–1553.
- [17] D. R. Anthony, D. P. Nguyen, D. Fridovich-Keil, and J. F. Fisac. “Back to the Future: Efficient, Time-Consistent Solutions in Reach-Avoid Games”. arXiv, 2021.
- [18] C. Berge. “Topological Spaces: including a treatment of multi-valued functions, vector spaces, and convexity”. Courier Corporation, 1997.