2024 American Control Conference (ACC)
July 8-12, 2024. Toronto, Canada

Safe Control Synthesis for Hybrid Systems through Local Control Barrier Functions

Shuo Yang[1], Mitchell Black[2], Georgios Fainekos[2], Bardh Hoxha[2], Hideki Okamoto[2], Rahul Mangharam[1]

Abstract— **Control Barrier Functions (CBF) have provided a very versatile framework for the synthesis of safe control architectures for a wide class of nonlinear dynamical systems. Typically, CBF-based synthesis approaches apply to systems that exhibit nonlinear – but smooth – relationship in the state of the system and linear relationship in the control input. In contrast, the problem of safe control synthesis using CBF for hybrid dynamical systems, i.e., systems which have a discontinuous relationship in the system state, remains largely unexplored. In this work, we build upon the progress on CBF-based control to formulate a theory for safe control synthesis for hybrid dynamical systems. Under the assumption that local CBFs can be synthesized for each mode of operation of the hybrid system, we show how to construct CBF that can guarantee safe switching between modes. The end result is a switching CBF-based controller which provides global safety guarantees. The effectiveness of our proposed approach is demonstrated on two simulation studies.**

I. Introduction

Safety-critical control is one of the fundamental problems in autonomous systems. Among various safety control methods, control synthesis methods utilizing control barrier functions (CBF) is a recent active area of research. CBF can explicitly encode safe sets and enforce the invariance of safe sets via solving efficient online quadratic programs (QP) [1]–[4]. A special class of autonomous systems is the class of hybrid dynamical systems, which involves both continuous dynamic flow and discrete dynamical mode jumps for state evolution. Such discrete mode transitions could be needed to model physical phenomena, or high-level logical decision making. For instance, vehicle dynamics switching from dry road to wet road could be modeled by a hybrid system. Bipedal robot walking is another example of a hybrid system. Safety concerns naturally arise for safety-critical hybrid systems and many safety approaches have been proposed in the past, such as Hamilton-Jacobi reachability-based approaches [5], [6] and computing controlled invariant sets [7].

CBF-based approaches have been used to ensure safety for hybrid systems recently for its computational efficiency [8], [9], in which the authors use the notion of the global CBF (i.e., all dynamical modes share a common CBF) to ensure safety. However, such a global notion is restrictive and less necessary since it requires the CBF to satisfy the same invariance conditions for all dynamical modes, which results in the challenge of constructing a valid global CBF, and also in more conservative control behavior.

The work was primarily performed while Shuo Yang was with Toyota.
[1]University of Pennsylvania; [2] Toyota Motor North America, R&D. Correspondance to: yangs1@seas.upenn.edu

To mitigate the drawbacks of a global CBF approach, we propose to use multiple local CBFs to guarantee global safety. We assume that each dynamical mode has its own local CBF, which implies that each mode can be safe under CBF-based control without considering discrete mode switching. However, it is possible that some *unsafe behaviors can occur under discrete transitions (jumps) even if all modes can be safe independently*. To ensure global safety of hybrid systems, we first identify those safe and unsafe switching regions. Then, we refine the initial local CBFs by considering safety after mode switching. Finally, safety of hybrid systems is guaranteed under refined local CBFs.

A. Related Work

a) Safety for hybrid system: Safety is a paramount concern in hybrid systems and various methods for safety verification and control synthesis have been proposed [7], [10]–[12]. Among them, barrier function-based methods can provide provable safety guarantees [11], [13]–[15]. Furthermore, CBFs emerged as a principled control method to enforce safety for controlled (hybrid) systems [8], [9], [16]–[18]. The most relevant work to ours is [8], in which the authors define a global CBF for hybrid systems and propose a data-driven constructive method to find a valid global CBF. Our work differs from [8] since we focus on ensuring global safety using local CBFs, and we propose to refine local CBFs considering unsafe dynamical mode switching. Also, we provably show that global CBF is more conservative than local CBFs to gurantee safety for hybrid systems.

b) Stability for hybrid system: Our work is also relevant to the Lyapunov-based stability for hybrid systems [19], [20]. As demonstrated in [21], unconstrained switching might lead to global instability even if all dynamical modes are stable with corresponding local Lyapunov functions. Thus, finding switching conditions for which global Lyapunov stability is guaranteed is one of the most important and elusive problems in the hybrid systems literature, and many approaches have been proposed [22]–[26]. As a dual notion, safety is also of great importance in hybrid systems but its related research is rather limited. In this paper, we first reveal a similar result that some switching conditions might lead to lack of global safety even if all dynamical modes are safe through control. Also, we propose an algorithmic procedure to ensure global safety using multiple local CBFs. To the best of our knowledge, this is the first work addressing the safe switching problem with multiple local CBFs.

The contributions of this paper are summarized below:

979-8-3503-8265-5/$31.00 ©2024 AACC

1) we formulate the safety control problems for hybrid systems using multiple local CBFs;
2) we reveal that safety might be violated under some switching states even if all dynamical modes are safe through control;
3) we refine local CBFs by considering safe switching and provide safety guarantees which, to our knowledge, is the first work to do so using multiple local CBFs; and
4) we demonstrate the effectiveness of our approach through simulations.

II. Preliminary and Problem Formulation

In this section, we cover the background concepts relevant to this work. This includes the related definitions on Control Barrier Functions and Hybrid Systems. We denote by $\mathbb{R}$ and $\mathbb{R}^n$ the set of real numbers and real n-dimensional vectors, respectively. The set $\mathbb{N}$ denotes the natural numbers (including zero). We will be using subscripts to denote subsets of these sets, e.g., $\mathbb{R}_{>0}$ denotes the set of positive real numbers. Given a set X, $\mathcal{P}(X)$ denotes its powerset. Let α: $\mathbb{R} \to \mathbb{R}$ denote an extended class $\mathcal{K}_\infty$ function, i.e., a strictly increasing function with $\alpha(0) = 0$.

A. Control Barrier Functions

Consider a continuous-time control-affine system:

$$\dot{x} = f(x) + g(x)u, \quad x(0) = x_0, \tag{1}$$

where f and g are locally Lipschitz, $x \in D \subseteq \mathbb{R}^n$ is the state and D denotes a compact set in $\mathbb{R}^n$. Safety can be framed in the context of enforcing set invariance in the state space, i.e., the state should not exit a safe set $\mathcal{C}$. The safe set $\mathcal{C}$ is represented by the super-level set of a continuously differentiable function $h : D \to \mathbb{R}$. The algebraic expressions for the safe set $\mathcal{C}$ and its boundary $\partial\mathcal{C}$ are given by:

$$\mathcal{C} = \{x \in D \subset \mathbb{R}^n : h(x) \geq 0\}, \tag{2a}$$
$$\partial\mathcal{C} = \{x \in D \subset \mathbb{R}^n : h(x) = 0\}. \tag{2b}$$

For a locally Lipschitz continuous control law $u = k(x)$, we have that $\dot{x} = f(x) + g(x)k(x)$ is locally Lipschitz continuous. Thus, for any initial condition $x_0 \in D$, there exists a maximum time interval of existence $I(x_0) = [0, \tau_{max})$, such that $x(t)$ is the unique solution to the ordinary differential equation (1) on $I(x_0)$. We frame the safety of system (1) in terms of set invariance as shown below.

Definition 1. *(Forward invariance and safety) The set $\mathcal{C}$ is* forward invariant *if for every $x_0 \in \mathcal{C}$, $x(t) \in \mathcal{C}$ holds for all $t \in I(x_0)$. If $\mathcal{C}$ is forward invariant, we say that* (1) *is safe.*

To verify invariance of $\mathcal{C}$, a control barrier function can be used as a certificate which characterizes the admissible set of control inputs that render $\mathcal{C}$ forward invariant.

Definition 2. *(Control barrier function [1]) Let $\mathcal{C} \subset D \subset \mathbb{R}^n$ be the superlevel set of a continuously differentiable function $h : D \to \mathbb{R}$, then h is a control barrier function for safe set $\mathcal{C}$ if there exists an extended class $\mathcal{K}_\infty$ function $\alpha(\cdot)$ such that for the control system (1):*

$$\sup_{u \in U} \left[\frac{\partial h(x)}{\partial x} \big(f(x) + g(x)u \big) \right] \geq -\alpha(h(x)), \tag{3}$$

for all $x \in D$.

Given the CBF $h(x)$, the set of all control values that render $\mathcal{C}$ safe is given by:

$$K_{cbf}(x) = \{u \in U : \frac{\partial h(x)}{\partial x} \big[f(x) + g(x)u \big] + \alpha(h(x)) \geq 0\} \tag{4}$$

which we denote as the safe control set. The following theorem shows that the existence of a control barrier function implies that the control system (1) is safe:

Theorem 1. *([1, Theorem 2]) Assume $h(x)$ is a CBF on $D \supset \mathcal{C}$ and $\frac{\partial h}{\partial x}(x) \neq 0$ for all $x \in \partial\mathcal{C}$. Then any Lipschitz continuous controller $u(x)$ such that $u(x) \in K_{cbf}(x)$ for all $x \in \mathcal{C}$ will render the set $\mathcal{C}$ forward invariant.*

B. Hybrid Automaton

A hybrid automaton is a model of a system with both a continuous dynamic flow and discrete dynamic jumps. The state of a hybrid automaton is a pair (q, x) where q is the discrete mode and x is the continuous state vector.

Definition 3. *A hybrid input automaton H_I is a tuple $H_I = \langle X, Q, U, U_q, F, \texttt{Guard} \rangle$:*

- *$X \subseteq \mathbb{R}^n$ is the continuous state space.*
- *Q is a finite set of modes.*
- *$U \subseteq \mathbb{R}^m$ denotes the continuous space of inputs, and $U_q \subseteq U$ is the admissible control input set for each mode $q \in Q$.*
- *$F : Q \times X \times U_q \to X$ is a vector field that describes the real-time dynamic flow of the continuous state x. For a mode $q \in Q$, we define F as a control affine system with admissible control set U_q:*

$$\dot{x} = F_q(x, u) = f_q(x) + g_q(x)u. \tag{5}$$

- *$\texttt{Guard} : Q \times Q \to \mathcal{P}(X)$ denotes the guard set that triggers mode switching.*

We only consider deterministic systems in this work, i.e., there is no uncertainty in both the dynamic flow and discrete jumps. Notice that that Def. 3 does not allow jumps in the value of the continuous state of the system. That is, we assume that the continuous state component of a hybrid system solution is continuous with respect to time.

To provide execution semantics for the hybrid automaton, we will need to define switching feedback control law.

Definition 4. *A switching feedback control law is defined as $k : Q \times X \to U$, where $k_q(x)$ is locally Lipschitz continuous w.r.t. x for any mode q.*

The composition ($\|$) of a switching feedback controller k_q with a hybrid input automaton H_I will be referred to in the following as a hybrid system. We will denote a *hybrid system* by $H = H_I \| k$.

Definition 5. *(Hybrid system solution) For a hybrid system H and a set of initial conditions $Q_0 \times X_0 \subseteq Q \times X$, a solution (trajectory) of H is a sequence $(q_i, \varphi_i, \delta_i)_{i \in N}$, where N is $\mathbb{N}$ or a bounded subset of $\mathbb{N}$, $q_i \in Q$ represents the discrete mode, $\varphi_i : X \times \mathbb{R}_{\geq 0} \to X$ represents the continuous state evolution, and $\delta_i \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ represents the duration of operating in mode i (i.e., dwell time), such that*

1) *$(q_0, x_0) \in Q_0 \times X_0$ is the initial state of the hybrid system at time $\tau_0 = 0$.*
2) *If there is some $j \in N$ such that $\tau_j = \infty$, then $j = \max N < \infty$, i.e., N is finite and j is unique.*
3) *For $i \in N$, we let*

$$\tau_{i+1} = \sum_{j=0}^{i} \delta_i.$$

If $\tau_i < \infty$ for $i > 0$, then τ_i is the switching time from mode q_{i-1} to q_i.

4) *For all $i \in N$ and for $t \in [\tau_i, \tau_{i+1}]$, $\varphi_i(x_i, t)$ is the solution of (5) for mode q_i and initial condition $x_i = \varphi_{i-1}(x_{i-1}, \tau_i)$, unless $i = 0$ since x_0 is defined. When $\tau_{i+1} = \infty$, then t ranges over $[\tau_i, \tau_{i+1})$.*
5) *For all $i \in N$ with $i > 0$, if $\tau_i < \infty$, then*

$$\varphi_{i-1}(\tau_i) \in \texttt{Guard}(q_{i-1}, q_i).$$

The above conditions require that a mode transition happens when the continuous state belongs to the Guard set. Then, the system follows the continuous flow of the new mode until the next mode transition occurs.

In the following, we will denote by $\mathcal{L}_H(Q_0, X_0)$ the set containing all solutions of H with initial conditions $Q_0 \times X_0$. If $(Q_0, X_0) = (Q, X)$, i.e., any initial condition is possible, then we just write $\mathcal{L}_H$ for the language. We will also use the notation $q \to q'$ to represent the transition (q, q') when $\texttt{Guard}(q, q') \neq \emptyset$.

Notice that in Def. 5, we do not explicitly impose any conditions on the input signal u. However, Def. 5.4 requires that a solution to (5) exists. In addition, condition 4 enforces continuity of the continuous state vector at discrete mode transition times, i.e., $\varphi_{i+1}(x_{i+1}, \tau_{i+1}) = x_{i+1} = \varphi_i(x_i, \tau_{i+1})$. Therefore, in the following, we can view the solution of the hybrid system as a function of time, i.e., $x : \mathbb{R}_{\geq 0} \to X$, and ignore the jump index i, or hybrid mode q when they are not important.

C. Problem Formulation

This paper is concerned about the safety of the hybrid system in Def. 3. Before we formulate the problem statement, we first introduce an illustrative example.

Example 1. *Let us consider an adaptive cruise control system. As shown in Figure (1a), the ego car and the leading car are driving on a straight road. The ego car is required to follow the leading and maintain a safe distance. The road is partitioned into dry road and ice road with different frictions and control input bounds. So there are two modes of dynamics in this example. The hybrid system model is shown in Fig. (1b). The guard set Guard for the transition from "dry road" to "ice road" contains all states whose position p of the ego car is greater than $100m$.*

Intuitively, the ego car should avoid high speed while switching from dry road to ice road, since both the friction and control bound of the ice road dynamics are smaller so the ego might not be able to decrease the speed as fast as in dry road. Unsafe behavior can occur after switching from dry to ice even if the ego was safe on the dry road. This implies that having two CBFs, one for dry and one for ice road dynamics, and applying them as safety filters is not sufficient for global safety guarantees. Therefore, we are aiming to study the safety when considering switching dynamics.

We first define (q, q')-safety for hybrid systems, and then define global safety.

Definition 6. *((q, q')-safety for hybrid system) For a hybrid system H, a pair of modes $(q, q') \in Q^2$, and safe sets $\mathcal{C}_q, \mathcal{C}_{q'} \subset D$ for modes q and q', respectively, we say that H is (q, q')-safe w.r.t. $\mathcal{C}_q$ and $\mathcal{C}_{q'}$ if for any initial state (q_0, x_0) with $q_0 = q$ and $x_0 \in \mathcal{C}_q$, the (potentially bounded) resulting trajectory $(q_i, \varphi_i, \delta_i)_{i \in \{0,1\}}$ of H with $q_1 = q'$ satisfies:*

- *$\varphi_0(x_0, t) \in \mathcal{C}_q$ for all $t \in [\tau_0, \tau_1]$, and*
- *$\varphi_1(x_1, t) \in \mathcal{C}_{q'}$ for all $t \in [\tau_1, \tau_2]$ if $\delta_1 < \infty$ or for all $t \geq \tau_1$ otherwise.*

The above (q, q')-safety definition enforces safety requirements for any trajectory transitioning from mode q to q'. Note that the safe set $\mathcal{C}_q$ for flow mode q and $\mathcal{C}_{q'}$ for q' will be different in general. Nevertheless, for (q, q')-safety to hold (to be enforceable) it must be the case that $\mathcal{C}_q \cap \mathcal{C}_{q'} \cap \texttt{Guard}(q, q') \neq \emptyset$.

We define the set of possible mode transition pairs of a hybrid system H as

$$\mathcal{T}(H) = \{(q_i, q_{i+1})_{i \in N \setminus \sup N} \mid (q_i, \varphi_i, \delta_i)_{i \in N} \in \mathcal{L}_H\}.$$

We can now define global safety based on (q, q')-safety.

Definition 7. *(Global safety for hybrid system) For a hybrid system H, for given safe sets $\mathcal{C}_q \subset D$ for any mode $q \in Q$, we say that H is globally safe w.r.t. $\{\mathcal{C}_q\}_{q \in Q}$ if H is (q, q')-safe for any $(q, q') \in \mathcal{T}(H)$.*

Similar to (q, q')-safety, $\mathcal{C}_q \cap \mathcal{C}_{q'} \cap \texttt{Guard}(q, q') \neq \emptyset$ must also hold for any $(q, q') \in \mathcal{T}(H)$ to ensure that global safety is enforceable. In the following, we will assume that the safe sets for each mode of the hybrid automaton are provided since our goal is to synthesize a switching controller that guarantees global safety.

Assumption 1. *For any mode $q \in Q$ of H_I, we assume that we have a local control barrier function h_q for (5) and the corresponding safe set $\mathcal{C}_q = \{x \in D \subset \mathbb{R}^n : h_q(x) \geq 0\}$. For every $q \in Q$, we denote its safe control set by*

$$K_q(x) = \{u \in U_q : \frac{\partial h_q(x)}{\partial x}\left[f_q(x) + g_q(x)u\right] \geq -\alpha_q(h_q(x))\},$$

where α_q is the corresponding extended class $\mathcal{K}_\infty$ function.

The above assumption says that every mode of the hybrid automaton is equipped with its own local CBF. This is a mild

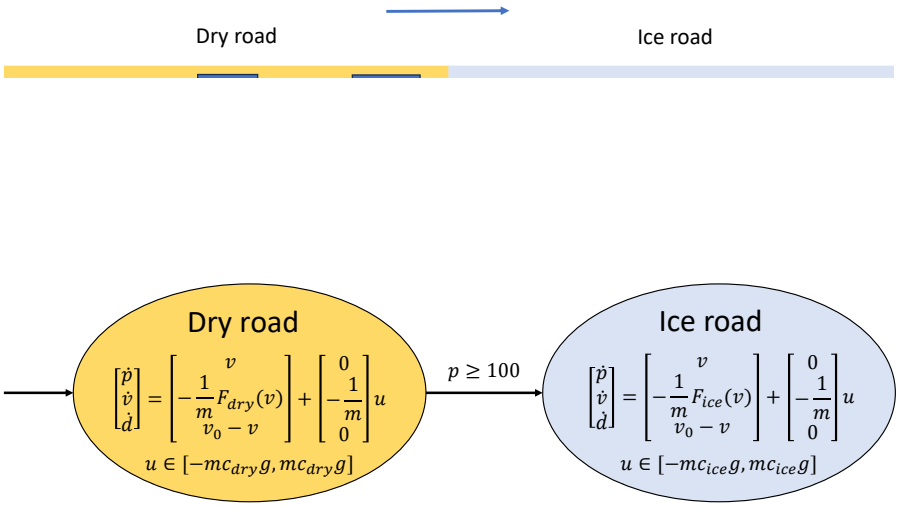

(b) Hybrid dynamics for adaptive cruise control. Please refer to Sec. IV-A for dynamics details.

Fig. 1: Hybrid adaptive cruise control.

assumption given the growing literature on the synthesis of CBF [2], [27]–[29]. However, local CBF cannot guarantee safety when switching between different modes. In other words, applying a control input $u \in K_q$ when the mode is q is not enough to ensure safety when the system switches to a mode q'. Thus, we formulate the following problem which is addressed in this paper.

Problem 1. *Given a hybrid automaton H_I, under Assumption 1, we are interested in solving two sub-problems:*

1) *find a switching control law k that can ensure the (q, q')-safety of H, where $q, q' \in Q$;*
2) *find a switching control law k that can ensure the global safety of H.*

To solve the above problem, we first identify safe and unsafe switching sets, and then compute the unsafe backward reachable set. Finally, we refine the initial local CBFs by avoiding the new unsafe sets. The refined CBFs can guarantee safety for the hybrid system.

III. Safe Control for Hybrid Systems

In this section, we formalize our notions of safety for hybrid dynamical systems, and we provide sufficient conditions for safe control synthesis. We start by defining what safe and unsafe switching sets are.

Definition 8. *(Safe and unsafe switching sets) For any mode jump $q \to q'$ in H, the corresponding safe switching set is defined by $\mathcal{S}_{q,q'} = \texttt{Guard}(q, q') \cap \mathcal{C}_q \cap \mathcal{C}_{q'}$ and the corresponding unsafe switching set is defined by $\mathcal{U}_{q,q'} = (\texttt{Guard}(q, q') \cap \mathcal{C}_q) \backslash \mathcal{S}_{q,q'}$.*

Safety can be preserved when the switching state is in the safe switching set $\mathcal{S}_{q,q'}$. However, safety is jeopardized when the switching state is in the unsafe switching set. This intuition is formalized below.

Proposition 1. *For a hybrid system H and a given initial condition (q_0, x_0), for any (potentially bounded) trajectory $(q_i, \varphi_i, \delta_i)_{i\in\{0,1\}}$ of H that satisfies $q_0 = q$ and $q_1 = q'$, we have that:*

- *if $\varphi_0(x_0, t) \in \mathcal{C}_q$ for all $t \in [\tau_0, \tau_1]$, and $\varphi_0(x_0, \tau_1) \in \mathcal{S}_{q,q'}$, then $\varphi_1(x_1, t) \in \mathcal{C}_{q'}$ under $k_{q'}(\varphi_1(x_1, t)) \in K_{q'}(\varphi_1(x_1, t))$ for all $t \in [\tau_1, \tau_2]$, i.e., H is (q, q')-safe;*
- *if $\varphi_0(x_0, t) \in \mathcal{C}_q$ for all $t \in [\tau_0, \tau_1]$ and $\varphi_0(x_0, \tau_1) \in \mathcal{U}_{q,q'}$, then H is not (q, q')-safe.*

Proof. First, if the switching state $\varphi_0(x_0, \delta_0) \in \mathcal{S}_{q,q'} = \texttt{Guard}(q, q') \cap \mathcal{C}_q \cap \mathcal{C}_{q'}$, then we have $\varphi_0(x_0, \delta_0) \in \mathcal{C}_{q'}$. Since we already know that $\mathcal{C}_{q'}$ is forward invariant under the safe control set $K_{q'}$, then we know that $\varphi_1(\varphi_0(x_0, \delta_0), t) \in \mathcal{C}_{q'}$ can hold for all $t \in [\tau_1, \tau_2]$. Second, however, if $\varphi_0(x_0, \tau_1) \in \mathcal{U}_{q,q'}$, this means that the switching state is not in the safe set $\mathcal{C}_{q'}$, which directly implies that H is not (q, q')-safe. □

The above proposition states that the continuous state must be in the safe switching set to ensure the (q, q')-safety. It also implies that the safety of a hybrid system can still be violated even if each mode is safe under the corresponding CBF safety filter. Next, Algorithm 1 is provided to guide the system state to reach the safe switching set and avoid the unsafe switching set when the system is must switching.

Step 1 of Alg. 1 is typically easy to compute depending on the computational representation of the sets, i.e., safe sets and guard sets. A polyhedral set representation is closed under intersection, union and complementation and it is typically used when modeling guards in hybrid systems [30]. However, the results on the synthesis of CBF-based controllers with polyhedral safety sets are limited [31]. Therefore, set under-approximations using ellipsoidal sets will typically need to be computed [32].

Next, we discuss the computation of the backward reachable sets of unsafe guard conditions in Step 2. Let $C(U)$ be

Algorithm 1: Procedure for (q, q')-safety synthesis

1 Identify the safe switching set $\mathcal{S}_{q,q'}$ and the unsafe switching set $\mathcal{U}_{q,q'}$ for each $q \to q'$;
2 Compute the backward reachable set $\texttt{BackUnsafe}_{q,q'}$ for the unsafe switching set $\mathcal{U}_{q,q'}$;
3 Obtain the new CBF $h_{q,q'}$ for $q \to q'$ by refining the initial CBF of mode q (i.e., h_q) via considering the backward unsafe set using dynamic programming;
4 Control the system with $h_{q,q'}$ when system has mode q, and with $h_{q'}$ when system has mode q';

the set of all functions from positive reals to some set U, i.e., $C(U) = U^{\mathbb{R}_{\geq 0}}$.

Definition 9. *(Unsafe backward set) For any jump $q \to q'$ in H_I, the corresponding unsafe backward set is defined by*

$$\texttt{BackUnsafe}_{q,q'} = \{x_0 \in C_q \mid \forall u \in C(U_q), \exists T \in \mathbb{R}_{\geq 0}, \text{ s.t. } \varphi_q(x_0, T) \in \mathcal{U}_{q,q'}\}. \quad (6)$$

$\texttt{BackUnsafe}_{q,q'}$ contains all states that will inevitably enter the unsafe switching set $\mathcal{U}_{q,q'}$ in finite time no matter what control signal is applied. Therefore, we need to control the system to avoid the unsafe backward set, otherwise it will definitely enter $\mathcal{U}_{q,q'}$ and invalidate safety. There are some approaches to compute the backwards reachable set of the given target set, see, e.g., [6], [33]–[35]. We use Hamilton-Jacobi reachability in this paper to compute $\texttt{BackUnsafe}_{q,q'}$. The readers are refered to [6] for technical details of Hamilton-Jacobi reachability.

Remark 1. *In general, the set $\mathcal{C}_q \backslash \texttt{BackUnsafe}_{q,q'}$ is not a controlled invariant set*[1]*. This motivates us to find a new CBF $h_{q,q'}$ such that for the new safe set*

$$\mathcal{C}_{q,q'} = \{x \in D \subset \mathbb{R}^n \mid h_{q,q'}(x) \geq 0\},$$

we have that $\mathcal{C}_{q,q'} \subseteq \mathcal{C}_q \backslash \texttt{BackUnsafe}_{q,q'}$. Here, we find the new CBF $h_{q,q'}$ by refining the initial local CBF h_q.

Next, we introduce how to find the new control barrier function $h_{q,q'}$ using dynamic programming (Step 3 of Alg. 1). By leveraging techniques from [36], we update CBF h_q recursively using Hamilton-Jacobi reachability. When the process terminates, we obtain a valid CBF $h_{q,q'}$ on $\mathcal{C}_q \backslash \texttt{BackUnsafe}_{q,q'}$. The new safe set $\mathcal{C}_{q,q'}$ is the superlevel set of $h_{q,q'}$. The validity of $h_{q,q'}$ is established in the following results.

Lemma 1. *(Adapting Theorem 1 from [36]) The refined CBF $h_{q,q'}$ is valid upon convergence, i.e., there exists an extended class $\mathcal{K}_\infty$ function $\alpha_{q,q'}(\cdot)$ such that:*

$$\sup_{u \in U_q} \frac{\partial h_{q,q'}(x)}{\partial x}[f_q(x) + g_q(x)u] \geq -\alpha_{q,q'}(h_{q,q'}(x)), \quad (7)$$

for all $x \in \mathcal{C}_{q,q'}$.

[1]A set C is called a controlled invariant set if any trajectory starting within C can always be controlled to remain inside C. For example, the superlevel set of a CBF is a controlled invariant set.

Our initial h_q can be considered as a good warmstarting for the CBF refinement, which can accelerate the convergence of the recursive update. However, note that this dynamic programming (DP)-based CBF refinement is generally limited to low-dimensional systems [36]. This is typically the case since spatially discretized DP recursion results in exponential computational complexity w.r.t. the system dimensionality.

Remark 2. *From [36], [37], we know that the converged CBF recovers a valid control barrier-value function (CBVF), and that CBVF recovers the largest controlled invariant set. Hence, our refined CBF-based method is not conservative under the given multiple local CBFs.*

Then, we can obtain the (q, q')-safety guarantees. Using notation similar to (4), $K_{q,q'}(x)$ denotes the safe control set defined by $h_{q,q'}$.

Theorem 2. *For any initial state $x_0 \in \mathcal{C}_{q,q'}$ at mode q, a hybrid system H is (q, q')-safe under any switching feedback controller k s.t. $k_q(x) \in K_{q,q'}(x)$ and $k_{q'}(x) \in K_{q'}(x)$.*

Proof. Consider any trajectory $(q_i, \varphi_i, \delta_i)_{i \in \{0,1\}}$ of H satisfying $q_0 = q$ and $q_1 = q'$. For any $x \in \mathcal{C}_{q,q'}$, any control input $k_q(x) \in K_{q,q'}(x)$ maintains the forward invariance of the safe set $\mathcal{C}_{q,q'}$. Thus, $\varphi_0(x_0, t) \in \mathcal{C}_{q,q'} \subseteq \mathcal{C}_q$ for all $t \in [0, \delta_0]$. Now, according to Prop. 1, we only need to prove that the switching state is in the safe switching set, i.e., $\varphi_0(x_0, \delta_0) \in \mathcal{S}_{q,q'}$. Since $\varphi_0(x_0, \delta_0) \in \mathcal{C}_{q,q'}$ and $\mathcal{C}_{q,q'} \subseteq \mathcal{C}_q \backslash \texttt{BackUnsafe}_{q,q'}$, then $\varphi_0(x_0, \delta_0) \notin \texttt{BackUnsafe}_{q,q'}$. Also, since $\mathcal{U}_{q,q'} \subseteq \texttt{BackUnsafe}_{q,q'}$, then we have that $\varphi_0(x_0, \delta_0) \notin \mathcal{U}_{q,q'}$. Note that $\varphi_0(x_0, \delta_0) \in \texttt{Guard}_{q,q'} \cap \mathcal{C}_{q,q'} \subseteq \texttt{Guard}_{q,q'} \cap \mathcal{C}_q = \mathcal{S}_{q,q'} \cup \mathcal{U}_{q,q'}$. Hence, we finally obtain $\varphi_0(x_0, \delta_0) \in \mathcal{S}_{q,q'}$. □

Remark 3. *Note that when $\mathcal{C}_{q,q'} \cap \mathcal{S}_{q,q'} = \emptyset$, then there is no safe switching from mode q to mode q'. In this case, the safe control input set $K_{q,q'}$ will prevent the continuous state from entering the guard set and switching mode. Also, when the initial state is in $\mathcal{C}_q \backslash \mathcal{C}_{q,q'}$, then the system cannot be (q, q')-safe, since the 0-superlevel set of $h(q, q')$ has already recovered the largest controlled invariant set [36], [37].*

After deriving the safety control method for (q, q')-safety, we can now present our main result on global safety. Given a mode q, we define a new safe set $\mathcal{C}_q^* = \bigcap_{(q,q') \in \mathcal{T}(H)} \mathcal{C}_{q,q'}$ which is the intersection of all safe sets for the safe jump to any next possible mode. The resulting safe control set $\mathcal{K}_q^*(x) = \bigcap_{(q,q') \in \mathcal{T}(H)} K_{q,q'}(x)$ is the intersection of all safe control sets for safe transition to any next possible mode.

Theorem 3. *Under the assumptions that*

1) *For the given set of initial conditions $Q_0 \times X_0 \subseteq Q \times X$, we have $x_0 \in \mathcal{C}_{q_0}^* \neq \emptyset$ for any initial state $(q_0, x_0) \in Q_0 \times X_0$; and*
2) *$\forall (q, q') \in \mathcal{T}(H) : \mathcal{C}_q^* \cap \mathcal{C}_{q'}^* \cap \texttt{Guard}(q, q') \neq \emptyset$ or $\mathcal{C}_q^* \cap \texttt{Guard}(q, q') = \emptyset$; and*
3) *$\forall q \in Q, \forall x \in \mathcal{C}_q^* : \mathcal{K}_q^*(x) \neq \emptyset$ if $\mathcal{C}_q^* \neq \emptyset$,*

then for any initial state $(q_0, x_0) \in Q_0 \times X_0$, the hybrid system H is globally safe under any switching k control which satisfies $k_q(x(t)) \in \mathcal{K}_q^*(x(t))$.

Proof. Consider any trajectory $(q_i, \varphi_i, \delta_i)_{i \in N}$ of H. For any $i \in N \backslash \sup N$ and any feedback switching control $k_{q_i}(\varphi_i(x_i,t)) \in \mathcal{K}_{q_i}^*(\varphi_i(x_i,t))$, since $\mathcal{K}_{q_i}^*(\varphi_i(x_i,t)) \subseteq K_{q_i,q_{i+1}}(\varphi_i(x_i,t))$, so $k_{q_i}(\varphi_i(x_i,t)) \in K_{q_i,q_{i+1}}(\varphi_i(x_i,t))$; also, $k_{q_{i+1}}(\varphi_{i+1}(x_{i+1},t)) \in \mathcal{K}_{q_{i+1}}^*(\varphi_{i+1}(x_{i+1},t))$ ensures that $\varphi_{i+1}(x_{i+1},t) \in \mathcal{C}_{q_{i+1}}^* \subseteq \mathcal{C}_{q_{i+1}}$. Thus, (q_i, q_{i+1})-safety is guaranteed according to Theorem 2 and H is globally safe.

The first assumption says the initial state should be safe. For the second assumption, it is necessary because the case $\mathcal{C}_q^* \cap \mathcal{C}_{q'}^* \cap \texttt{Guard}(q,q') \neq \emptyset$ means that safe switching set (see Def. 8) is not empty, which make the safe switching feasible (see Prop. 1). On the other hand, if the safe switching set is empty, we can also ensure safety as long as $\mathcal{C}_q^* \cap \texttt{Guard}(q,q') = \emptyset$, i.e., the switching cannot happen, which is the latter case of the second assumption.

For the third assumption, we do not assume $\mathcal{K}_q^* \neq \emptyset$ if $\mathcal{C}_q^* = \emptyset$. We illustrate its reason by proving that mode q will not appear in any trajectory if $\mathcal{C}_q^* = \emptyset$, i.e., there exists no $i \in N \backslash \sup N$ such that $q_{i+1} = q$. Suppose there exists such $i \in N \backslash \sup N$ with $q_{i+1} = q$, then $\mathcal{C}_{q_i}^* \cap \mathcal{C}_{q_{i+1}}^* \cap \texttt{Guard}(q_i, q_{i+1}) = \emptyset$ since $\mathcal{C}_{q_{i+1}}^* = \emptyset$. Also, $\mathcal{C}_{q_i}^* \cap \texttt{Guard}(q_i, q_{i+1}) \neq \emptyset$ holds, otherwise the transition from mode q_i to mode q_{i+1} is impossible under switching control $k_{q_i}(\varphi_i(x_i,t)) \in \mathcal{K}_{q_i}^*(\varphi_i(x_i,t))$. Therefore, the second assumption cannot hold, which is a contradiction. Thus, mode q will not be in any trajectory if $\mathcal{K}_q^* \neq \emptyset$, so we do not need to pose any requirement on $\mathcal{K}_q^*$. □

Note that the above theorem provides *sufficient but not necessary* conditions for safe feedback switching controller design. There might exist a safe control law even if the conditions of Theorem 3 are not satisfied. For example, consider the hybrid automaton in Fig. 2, where the transitions

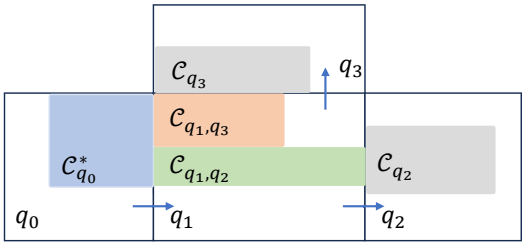

Fig. 2: One hybrid automaton example where safe switching controller exists but Theorem 3 is not applicable.

Remark 4. *It is possible that the above multiple CBF constraints sometimes lead to infeasibility, i.e., $\mathcal{K}_q^*(x(t)) = \emptyset$. This problem can be addressed by considering a subset of feasible transitions, or even by staying in a specific mode if this mode can guarantee safety without any jumps.*

In many cases, e.g., in automated driving systems, we may only need to consider one next dynamics mode transition, e.g., based on path planning or prediction, so one CBF constraint is sufficient to ensure safety for each mode transition (i.e., global safety reduces to sequential (q,q')-safety). Theoretically, improving feasibility under multiple CBFs is an important problem and has been recenlty addressed using different approaches [38]–[40].

One may ask about the relationship between global hybrid system safety (Def. 7) through multiple local CBFs and safety induced through a global CBF. In the following proposition, we show that the local CBFs-based method is generally less conservative than global safety guaranteed by a global CBF-based method. We first formally define global CBF for hybrid systems.

Definition 10. *$h_g(x)$ is a global CBF for the hybrid automaton H_I if there exists an extended class $\mathcal{K}_\infty$ function $\alpha(\cdot)$ such that for H_I:*

$$\sup_{u \in U_q} \frac{\partial h_g(x)}{\partial x}[f_q(x) + g_q(x)u] \geq -\alpha(h_g(x)) \quad (8)$$

holds for all $(q,x) \in Q \times X$.

Proposition 2. *If there exists a global CBF h_g for the hybrid automaton H_I for its corresponding safe set $\mathcal{C}_g$, then there exist a local CBF h_q for each mode $q \in Q$ such that the state can always stay inside $\mathcal{C}_g$.*

Proof. We can let $h_q = h_g$ for each mode $q \in Q$. □

The above proposition establishes that a global CBF can be viewed as a special case of the local CBFs method. To further highlight the differences, we make several observations:

- Global CBF asks for the same safe state set for each flow mode, but the local CBFs-based method can support different safe sets in each mode. Therefore, control synthesis for global safety through local CBF can support a wider range of safe control applications.
- The single global CBF should satisfy constraints for all flow modes, but any local CBF is only responsible for its own specific mode. This implies that global CBF are harder to synthesize than local CBFs.
- Global CBF impose more restrictive constraint conditions than local CBFs to enforce safety because of the previous observation. This means that local CBFs can lead to better system performance while ensuring safety.

IV. Case Studies

In this section, we present two case studies to illustrate the effectiveness of our multiple local CBFs-based safe control method. We compare our approach with a baseline approach where each initial local CBF is applied for each dynamical mode and it is unaware of the safety effects of modes

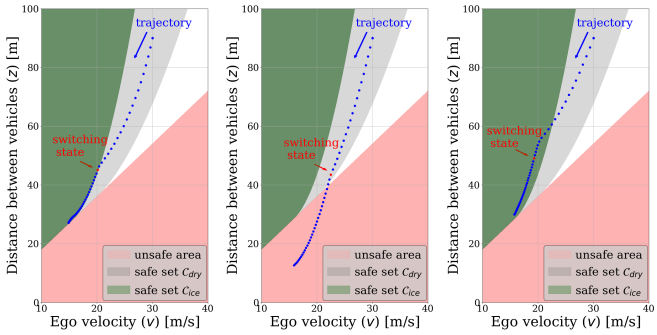

(a) Switch-aware CBF (b) Switch-unaware CBF (c) Global CBF

Fig. 3: Trajectories of our approach, switch-unaware CBF approach, and global CBF approach. The unsafe area (red) is defined by the safety constraint function $c(x)$.

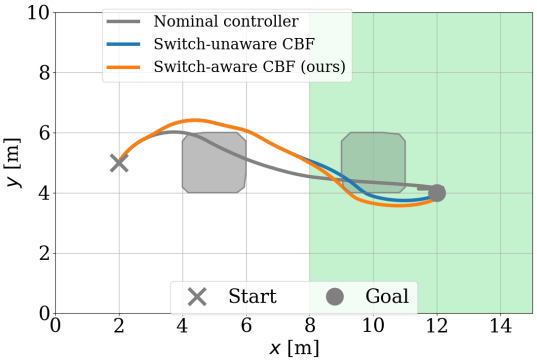

Fig. 4: Dubins car is reaching a goal and avoiding two obstacles. The white and green regions are dry and wet surfaces, respectively. Gray boxes are obstacles. Nominal controller provides reference trajectory for CBF-based approaches. Global CBF method is not applicable in this case.

switching (we name this approach briefly as switch-unaware CBF). We also compare ours with global CBF method.

A. Adaptive cruise control

We first conduct simulation on the aforementioned adaptive cruise control example. As described in Example 1, the road has two different surfaces (dry road and ice road) with different dynamical modes. Hence, two different local CBFs are required (and assumed) to ensure safety for each mode. The system state is $x = [p\ v\ d]^T$, where p is the ego car position, v is the ego car velocity, and d is the distance between the two cars. The control input u is the acceleration. The switching scenario is presented in Fig. 1a and the hybrid model is shown in Fig. 1b, where m is the mass of the vehicle, g is the gravitational constant, v_0 is the velocity of leading car, c_{dry} (c_{ice}) is the maximum g-force that can be applied on dry (ice) road, and $Fr_{dry}(v) = f_{0,d}v^2 + f_{1,d}v + f_{2,d}$ is the friction of dry road (correspondingly, $Fr_{ice}(v)$ is the friction of ice road but with different parameters).

In our simulation, the ego car is expected to drive with a desired speed v_d (where $v_d > v_0$) while maintaining a safe distance with the leading car. The safety specification is defined by the constraint function $c(x) = d - T_h \cdot v$. The CBF for dry road is $h_{dry}(x) = d - T_h v - \frac{(v_0 - v)^2}{2c_{dry}g}$, and $h_{ice}(x)$ is defined similarly by replacing c_{dry} with c_{ice} for ice road.

We compare our proposed approach with the switch-unaware method, in which h_{dry} is applied as the safety filter while on dry road and h_{ice} on the ice road, and also compare with global CBF method, where $h_{global} = h_{ice}$ is applied for both modes. For our approach, we first refine h_{dry} to consider safe switching from dry road to ice road, and obtain $h_{dry,ice}$, which is applied as the safety filter while on dry road. The CBF is switched to h_{ice} after the dynamical mode has switched. Our result is shown in Fig. 3a, in which we can notice that the switching state (red point) is in the safe set of the new dynamics ($\mathcal{C}_{ice}$ in this example). Thus, the safety will not be violated after switching to h_{ice}. However, as it can be observed in Fig. 3b, in the switch-unaware controller, the switching point is not in $\mathcal{C}_{ice}$ and safety is violated after switching. Also, global CBF method (Fig. 3c) is safe but more conservative than our approach since it decreases ego velocity more and will finally reach the destination later.

B. Dubins car collision avoidance

Consider an extended Dubins car model below, where x, y are position, θ is heading, and v is speed.

$$\dot{x} = vcos(\theta), \dot{y} = vsin(\theta), \dot{v} = a, \dot{\theta} = \omega. \tag{9}$$

The control input includes the acceleration a and angular velocity ω. The objective of the Dubins car is to navigate to reach a goal while avoiding obstacles, as shown in Fig. 4. However, there are two different road surfaces in this task: dry road (white region on the left side of Fig. 4) and wet road (green region on the right side of Fig. 4). We have two local CBFs h_{dry} and h_{wet} for them, respectively. For different surfaces, the Dubins car has different control bounds.

As a comparison, we show trajectories of 3 different approaches in Fig. 4. To demonstrate the performance of the CBF-based safety filters, a reference trajectory (gray) is generated without considering the obstacles, so it is unsafe. With the switch-unaware CBF approach, i.e., applying the original local CBF for each mode as safety filters, the trajectory (blue) turns to be unsafe after switching to the ice road. Finally, our approach refines h_{dry} and obtains $h_{dry,wet}$ to ensure safe switching from dry surface to wet surface, and $h_{dry,wet}$ and h_{wet} are applied as safety filters before and after switching respectively. Our trajectory (orange) is safe. Global CBF approach is not applicable in this example because dry road mode and wet mode have different safe regions.

V. Conclusion and Discussion

In this work, we formulate and solve the safety control problem for hybrid systems using multiple local CBFs. We discover that the safety of hybrid systems can be jeopardized even if all flow modes can be independently safe through control. Thus, we propose to refine the initial local CBFs to avoid any unsafe switching regions. Finally, we can obtain safety guarantees under refined multiple local CBFs.

In the future, we plan to apply our method to real-world challenging scenarios such as multi-frictions autonomous vehicles racing [41]–[43].

VI. Acknowledgment

The authors thank Hardik Parwana for initial discussions.

References

[1] 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):3861–3876, 2016.

[2] A. Robey, H. Hu, L. Lindemann, H. Zhang, D. V. Dimarogonas, S. Tu, and N. Matni. Learning control barrier functions from expert demonstrations. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 3717–3724. IEEE, 2020.

[3] S. Yang, S. Chen, V. M. Preciado, and R. Mangharam. Differentiable safe controller design through control barrier functions. *IEEE Control Systems Letters*, 7:1207–1212, 2022.

[4] L. Berducci, S. Yang, R. Mangharam, and R. Grosu. Learning adaptive safety for multi-agent systems. *arXiv preprint arXiv:2309.10657*, 2023.

[5] C. J. Tomlin, J. Lygeros, and S. S. Sastry. A game theoretic approach to controller design for hybrid systems. *Proceedings of the IEEE*, 88(7):949–970, 2000.

[6] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin. Hamilton-jacobi reachability: A brief overview and recent advances. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 2242–2253. IEEE, 2017.

[7] B. Legat, P. Tabuada, and R Jungers. Computing controlled invariant sets for hybrid systems with applications to model-predictive control. *IFAC-PapersOnLine*, 51(16):193–198, 2018.

[8] L. Lindemann, H. Hu, A. Robey, H. Zhang, D. Dimarogonas, S. Tu, and N. Matni. Learning hybrid control barrier functions from data. In *Conference on Robot Learning*, pages 1351–1370. PMLR, 2021.

[9] A. Robey, L. Lindemann, S. Tu, and N. Matni. Learning robust hybrid control barrier functions for uncertain systems. *IFAC-PapersOnLine*, 54(5):1–6, 2021.

[10] N. Athanasopoulos and R. Jungers. Combinatorial methods for invariance and safety of hybrid systems. *Automatica*, 98:130–140, 2018.

[11] S. Prajna and A. Jadbabaie. Safety verification of hybrid systems using barrier certificates. In *International Workshop on Hybrid Systems: Computation and Control*, pages 477–492. Springer, 2004.

[12] 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.

[13] P. Glotfelter, I. Buckley, and M. Egerstedt. Hybrid nonsmooth barrier functions with applications to provably safe and composable collision avoidance for robotic systems. *IEEE Robotics and Automation Letters*, 4(2):1303–1310, 2019.

[14] M. Maghenem and R. G. Sanfelice. Characterizations of safety in hybrid inclusions via barrier functions. In *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control*, pages 109–118, 2019.

[15] A. Bisoffi and D. V. Dimarogonas. A hybrid barrier certificate approach to satisfy linear temporal logic specifications. In *2018 Annual American Control Conference (ACC)*, pages 634–639. IEEE, 2018.

[16] A. Lavaei, S. Soudjani, and E. Frazzoli. Safety barrier certificates for stochastic hybrid systems. In *2022 American Control Conference (ACC)*, pages 880–885. IEEE, 2022.

[17] S. Yang, G. J. Pappas, R. Mangharam, and L. Lindemann. Safe perception-based control under stochastic sensor uncertainty using conformal prediction. *arXiv preprint arXiv:2304.00194*, 2023.

[18] M. Maghenem, A. J. Taylor, A. D. Ames, and R. G. Sanfelice. Adaptive safety using control barrier functions and hybrid adaptation. In *2021 American Control Conference (ACC)*, pages 2418–2423. IEEE, 2021.

[19] R. A. DeCarlo, M. S. Branicky, S. Pettersson, and B. Lennartson. Perspectives and results on the stability and stabilizability of hybrid systems. *Proceedings of the IEEE*, 88(7):1069–1082, 2000.

[20] S. Chen, M. Fazlyab, M. Morari, G. J. Pappas, and V. M. Preciado. Learning lyapunov functions for hybrid systems. In *Proceedings of the 24th International Conference on Hybrid Systems: Computation and Control*, pages 1–11, 2021.

[21] D. Liberzon. *Switching in systems and control*, volume 190. Springer, 2003.

[22] R. Goebel, R. G. Sanfelice, and A. R. Teel. Hybrid dynamical systems. *IEEE control systems magazine*, 29(2):28–93, 2009.

[23] M. S. Branicky. Multiple lyapunov functions and other analysis tools for switched and hybrid systems. *IEEE Transactions on automatic control*, 43(4):475–482, 1998.

[24] J. P. Hespanha. Uniform stability of switched linear systems: Extensions of lasalle's invariance principle. *IEEE transactions on Automatic Control*, 49(4):470–482, 2004.

[25] H. Lin and P. J. Antsaklis. Stability and stabilizability of switched linear systems: a survey of recent results. *IEEE Transactions on Automatic control*, 54(2):308–322, 2009.

[26] G. Zhai, X. Xu, H. Lin, and A. N. Michel. Analysis and design of switched normal systems. *Nonlinear Analysis: Theory, Methods & Applications*, 65(12):2248–2259, 2006.

[27] W. Xiao, C. G. Cassandras, and C. Belta. *Safe Autonomy with Control Barrier Functions: Theory and Applications*. Springer Nature, 2023.

[28] A. Clark. Verification and synthesis of control barrier functions. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 6105–6112. IEEE, 2021.

[29] L. Wang, D. Han, and M. Egerstedt. Permissive barrier certificates for safe stabilization using sum-of-squares. In *2018 Annual American Control Conference (ACC)*, pages 585–590. IEEE, 2018.

[30] M. Althoff, G. Frehse, and A. Girard. Set propagation techniques for reachability analysis. *Annual Review of Control, Robotics, and Autonomous Systems*, 4(1):369–395, 2021.

[31] A. Thirugnanam, J. Zeng, and K. Sreenath. Safety-critical control and planning for obstacle avoidance between polytopes with control barrier functions. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 286–292. IEEE, 2022.

[32] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[33] A. B. Kurzhanski and P. Varaiya. Ellipsoidal techniques for reachability analysis: internal approximation. *Systems & control letters*, 41(3):201–211, 2000.

[34] T. Dreossi, T. Dang, and C. Piazza. Parallelotope bundles for polynomial reachability. In *Proceedings of the 19th International Conference on Hybrid Systems: Computation and Control*, pages 297–306, 2016.

[35] A. Majumdar, R. Vasudevan, M. M. Tobenkin, and R. Tedrake. Convex optimization of nonlinear feedback controllers via occupation measures. *The International Journal of Robotics Research*, 33(9):1209–1230, 2014.

[36] S. Tonkens and S. Herbert. Refining control barrier functions through hamilton-jacobi reachability. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 13355–13362. IEEE, 2022.

[37] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. Herbert. Robust control barrier–value functions for safety-critical control. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 6814–6821. IEEE, 2021.

[38] X. Tan and D. V. Dimarogonas. Compatibility checking of multiple control barrier functions for input constrained systems. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 939–944. IEEE, 2022.

[39] J. Breeden and D. Panagou. Compositions of multiple control barrier functions under input constraints. In *2023 American Control Conference (ACC)*, pages 3688–3695. IEEE, 2023.

[40] M. Black and D. Panagou. Consolidated control barrier functions: Synthesis and online verification via adaptation under input constraints. *arXiv preprint arXiv:2304.01815*, 2023.

[41] X. Sun, M. Zhou, Z. Zhuang, S. Yang, J. Betz, and R. Mangharam. A benchmark comparison of imitation learning-based control policies for autonomous racing. In *2023 IEEE Intelligent Vehicles Symposium (IV)*, pages 1–5. IEEE, 2023.

[42] X. Sun, S. Yang, and R. Mangharam. Mega-dagger: Imitation learning with multiple imperfect experts. *arXiv preprint arXiv:2303.00638*, 2023.

[43] S. Yang, Y. Chen, X. Yin, and R. Mangharam. Learning local control barrier functions for safety control of hybrid systems. *arXiv preprint arXiv:2401.14907*, 2024.