

Adaptive Safety with Control Barrier Functions

Andrew J. Taylor and Aaron D. Ames

Abstract—Adaptive Control Lyapunov Functions (aCLFs) were introduced 20 years ago, and provided a Lyapunov-based methodology for stabilizing systems with parameter uncertainty. The goal of this paper is to revisit this classic formulation in the context of safety-critical control. This will motivate a variant of aCLFs in the context of safety: *adaptive Control Barrier Functions (aCBFs)*. Our proposed approach adaptively achieves safety by keeping the system’s state within a safe set even in the presence of parametric model uncertainty. We unify aCLFs and aCBFs into a single control methodology for systems with uncertain parameters in the context of a Quadratic Program (QP) based framework. We validate the ability of this unified framework to achieve stability and safety in an Adaptive Cruise Control (ACC) simulation.

I. INTRODUCTION

In many modern control applications, safety is of critical importance. It is impossible to model the system dynamics in these applications exactly—that is, parameters of the model may not match the real system. For instance, the mass and electrical properties of robotic systems are often approximate values. Thus, to truly enforce safety, it is necessary to quantify safety in the context of unknown parameters.

The use of Control Barrier Functions (CBFs) [1], [2] for ensuring safety of nonlinear control systems has become increasingly popular [18], [24], [25]. Controllers synthesized via CBFs rely on a model, and the guarantees they achieve may fail in the presence of model uncertainty. Robust control methods can ensure safety [7], [27] or quantify how safety properties degrade [9] in the presence of model uncertainty, but may be overly conservative in restricting the behavior of the system. Data-driven methods employing machine learning [19], [5] provide probabilistic safety guarantees, but may require episodic, offline training to improve model estimates [6].

In this paper, we focus on an online, *adaptive* approach to ensuring that a system remains safe in the presence of model uncertainty. Adaptive control seeks to update a model of the system as it evolves to achieve stability or a desired level of performance [10]. In particular, we build upon the idea of *adaptive* Control Lyapunov Functions (aCLFs) [11], which have been used to stabilize nonlinear systems in the presence of parametric model uncertainty [12], [13], [15]. That is, the goal of this paper is to find conditions for adaptive safety (via Control Barrier Functions) equivalent to those derived for adaptive stability (via Control Lyapunov Functions).

One challenge in developing adaptive control methods that guarantee safety is ensuring that a nonlinear system’s state

remains within a prescribed safe set at all times. In contrast, guarantees on stability provided by aCLFs describe the behavior of the state and parameter estimation error jointly, allowing the state to grow large before stabilizing, as long as the parameter estimation error diminishes. To achieve this stricter guarantee of safety, we leverage stronger assumptions on the initial parameter estimation error. The end result are conditions for safety even under the presence of model uncertainty, i.e., that a system with unknown parameters can be rendered safe for all time.

The main contribution of this paper is a formal methodology for ensuring safety in nonlinear (control affine) systems with parameter uncertainty through the formulation of *adaptive Control Barrier Functions (aCBFs)*. Like aCLFs, aCBFs provide a framework for updating model parameter estimates online, but do so to ensure safety. Unlike aCLFs, aCBFs require a different viewpoint on adaptive control to make stronger statements on the behavior of the system’s state. To the best of our knowledge, our approach is the first that adaptively ensures safety utilizing CBFs. The definitions and results in this paper provide the first steps towards a framework for adaptive safety unifying both online and data-driven, episodic updates of model parameters.

This paper is organized as follows. Section II reviews CLFs and aCLFs and how quadratic program based controllers can be synthesized to adaptively stabilize a system. Section III discusses CBFs and how they can be used to ensure the safety of a system. Section IV provides the main result of the paper by defining aCBFs, and shows how a system can be rendered adaptively safe in the presence of model uncertainty. Section V presents simulation results for an Adaptive Cruise Control system using both a safety-critical controller and a quadratic program based controller implementing an aCLF and an aCBF simultaneously.

II. ADAPTIVE CONTROL LYAPUNOV FUNCTIONS

To develop provably correct controllers for nonlinear systems, it is typically assumed that the model is known. Yet there are many practical applications where this assumption is not adequate. A simple illustration is a mechanical system whose parameters (masses, inertias, etc) are not completely known—and one may not want to treat the unknown model parameters as a perturbation from nominal parameters since this would only guarantee stability to a region corresponding to a bound on this difference (which also may not be known). The purpose of this section, therefore, is to review the framework of adaptive Control Lyapunov Functions.

Consider a state space $\mathcal{X} \subset \mathbb{R}^n$ and a control input space $\mathcal{U} \subset \mathbb{R}^m$, where it is assumed \mathcal{X} is path-connected and $\mathbf{0} \in$

Both authors are with the Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125, USA
ajtaylor@caltech.edu, ames@caltech.edu

\mathcal{X} . Consider the affine dynamic system given by:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u} \quad (1)$$

where $\mathbf{x} \in \mathcal{X}$, $\mathbf{u} \in \mathcal{U}$, $\mathbf{f} : \mathcal{X} \rightarrow \mathbb{R}^n$ and $\mathbf{g} : \mathcal{X} \rightarrow \mathbb{R}^{n \times m}$ are smooth on \mathcal{X} . We additionally assume $\mathbf{f}(\mathbf{0}) = \mathbf{0}$. We will use the following definition, found in [8], to study the stability of (1).

Definition 1 (Class \mathcal{K} Function). A continuous function $\alpha : [0, a) \rightarrow \mathbb{R}_+$, with $a > 0$, is a *class \mathcal{K} function* ($\alpha \in \mathcal{K}$) if $\alpha(0) = 0$ and α is strictly monotonically increasing. If $a = \infty$ and $\lim_{r \rightarrow \infty} \alpha(r) = \infty$, then α is said to be a *class \mathcal{K}_∞ function* ($\alpha \in \mathcal{K}_\infty$).

Given this definition, we can define a Control Lyapunov Function as in [3], [14].

Definition 2 (Control Lyapunov Function (CLF)). A smooth function $V : \mathcal{X} \rightarrow \mathbb{R}_+$ is a *Control Lyapunov Function (CLF)* for (1) if there exists $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{K}_\infty$ such that:

$$\alpha_1(\|\mathbf{x}\|) \leq V(\mathbf{x}) \leq \alpha_2(\|\mathbf{x}\|) \quad (2)$$

$$\inf_{\mathbf{u} \in \mathcal{U}} \dot{V}(\mathbf{x}, \mathbf{u}) \leq -\alpha_3(\|\mathbf{x}\|) \quad (3)$$

for all $\mathbf{x} \in \mathcal{X}$.

This definition can be constructed with $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{K}$, with resulting stability guarantees holding locally. The existence of a CLF for (1) implies there exists a smooth (except at $\mathbf{x} = \mathbf{0}$) state-feedback control law $\mathbf{k} : \mathcal{X} \rightarrow \mathcal{U}$, that renders the origin globally asymptotically stable [3], [21], noting that global refers to the state space \mathcal{X} . \mathbf{k} can be made continuous at $\mathbf{0}$ if V satisfies the small control property [22].

Following the classic formulation of aCLFs in [11], uncertainty in the dynamics (1) appears as:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{F}(\mathbf{x})\boldsymbol{\theta}^* + \mathbf{g}(\mathbf{x})\mathbf{u}, \quad (4)$$

where $\boldsymbol{\theta}^* \in \Theta \subset \mathbb{R}^p$ is a vector of unknown parameters and $\mathbf{F} : \mathcal{X} \rightarrow \mathbb{R}^{n \times p}$ is assumed to be smooth on \mathcal{X} with $\mathbf{F}(\mathbf{0}) = \mathbf{0}$. The impossibility of designing explicit controllers that are robust to unbounded unknown parameters suggests that we need to consider a larger class of controllers to stabilize (4). In particular, controllers that *update* an estimate of the unknown parameters. These are called adaptive controllers, and take the form:

$$\mathbf{u} = \mathbf{k}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \quad (5)$$

$$\dot{\hat{\boldsymbol{\theta}}} = \boldsymbol{\Gamma}\boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}), \quad (6)$$

where $\hat{\boldsymbol{\theta}} \in \Theta$ represents an estimate of the parameters $\boldsymbol{\theta}^*$ maintained by the controller, $\mathbf{k} : \mathcal{X} \times \Theta \rightarrow \mathcal{U}$ is a control law, $\boldsymbol{\Gamma} \in \mathbb{R}^{p \times p}$ is a matrix adaptive gain, and $\boldsymbol{\tau} : \mathcal{X} \times \Theta \rightarrow \mathbb{R}^p$ is the adaptation law. We make the following assumption on these functions:

- (A1) \mathbf{k} is locally Lipschitz continuous on $(\mathcal{X} \setminus \{\mathbf{0}\}) \times \Theta$ and $\mathbf{k}(\mathbf{0}, \hat{\boldsymbol{\theta}}) = \mathbf{0}$,
- (A2) $\boldsymbol{\tau}$ is locally Lipschitz continuous on $\mathcal{X} \times \Theta$,
- (A3) $\boldsymbol{\Gamma} \in \mathbb{R}^{p \times p}$ is symmetric and positive-definite.

Introducing this parameter update results in a composite

dynamic system:

$$\begin{bmatrix} \dot{\mathbf{x}} \\ \dot{\hat{\boldsymbol{\theta}}} \end{bmatrix} = \begin{bmatrix} \mathbf{f}(\mathbf{x}) + \mathbf{F}(\mathbf{x})\boldsymbol{\theta}^* + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \\ \boldsymbol{\Gamma}\boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \end{bmatrix} \quad (7)$$

Solutions to this system evolve in $\mathcal{X} \times \Theta$. Given this construction we introduce the following definition from [11]:

Definition 3 (Globally Adaptively Stabilizable). The system with unknown parameters (4) is *globally adaptively stabilizable* if there exists a dynamic controller of the form (5)-(6) satisfying (A1)-(A3) such that solutions $(\mathbf{x}(t), \hat{\boldsymbol{\theta}}(t))$ of (7) are globally bounded and $\lim_{t \rightarrow \infty} \mathbf{x}(t) = \mathbf{0}$.

Remark 1. Note that the requirements for global adaptive stabilizability are rather weak in the sense that $\hat{\boldsymbol{\theta}}(t)$ is not required to converge to $\boldsymbol{\theta}^*$. We will see, in fact, that convergence of $\hat{\boldsymbol{\theta}}(t)$ to $\boldsymbol{\theta}^*$ is not necessary for $\mathbf{x}(t)$ to converge to the equilibrium.

The strategy in designing adaptive controllers is to show that this problem is equivalent to a non-adaptive controller design problem. Such equivalence is shown via the notion of adaptive Control Lyapunov Functions as in [11]:

Definition 4 (Adaptive Control Lyapunov Function (aCLF)). Let $\alpha_1(\cdot, \boldsymbol{\theta})$, $\alpha_2(\cdot, \boldsymbol{\theta})$, $\alpha_3(\cdot, \boldsymbol{\theta}) \in \mathcal{K}_\infty$ for all $\boldsymbol{\theta} \in \Theta$. A smooth function $V_a : \mathcal{X} \times \Theta \rightarrow \mathbb{R}_+$, satisfying:

$$\alpha_1(\|\mathbf{x}\|, \boldsymbol{\theta}) \leq V_a(\mathbf{x}, \boldsymbol{\theta}) \leq \alpha_2(\|\mathbf{x}\|, \boldsymbol{\theta}), \quad (8)$$

is called an *adaptive Control Lyapunov Function (aCLF)* for (4) if there exists a symmetric positive-definite matrix $\boldsymbol{\Gamma} \in \mathbb{R}^{p \times p}$ such that for every $\boldsymbol{\theta} \in \Theta$, V_a is a CLF for the system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{F}(\mathbf{x})\boldsymbol{\lambda}_{clf}(\mathbf{x}, \boldsymbol{\theta}) + \mathbf{g}(\mathbf{x})\mathbf{u}, \quad (9)$$

where

$$\boldsymbol{\lambda}_{clf}(\mathbf{x}, \boldsymbol{\theta}) \triangleq \boldsymbol{\theta} + \boldsymbol{\Gamma} \left(\frac{\partial V_a}{\partial \boldsymbol{\theta}}(\mathbf{x}, \boldsymbol{\theta}) \right)^T. \quad (10)$$

That is,

$$\inf_{\mathbf{u} \in \mathcal{U}} \left[\frac{\partial V_a}{\partial \mathbf{x}}(\mathbf{f}(\mathbf{x}) + \mathbf{F}(\mathbf{x})\boldsymbol{\lambda}_{clf}(\mathbf{x}, \boldsymbol{\theta}) + \mathbf{g}(\mathbf{x})\mathbf{u}) \right] \leq -\alpha_3(\|\mathbf{x}\|, \boldsymbol{\theta}). \quad (11)$$

Adaptive Control Lyapunov Functions can be used to obtain the following result establishing the equivalence between the original adaptive controller design problem and a non-adaptive one.

Theorem 1. [11] *System (4) is globally adaptively stabilizable iff there exists an aCLF for (4).*

It is useful to give a sketch of the proof for the sufficiency portion of this result, as it will inform the proof of the analogous result in the context of aCBFs.

Sketch. Assume that we have an aCLF V_a for (4). As V_a is a CLF for (9) with $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$, we can construct a Lipschitz continuous (away from $\mathbf{x} = \mathbf{0}$) control law $\mathbf{k} : \mathcal{X} \times \Theta \rightarrow \mathcal{U}$ stabilizing (9) (a specific example of such a controller will be given after the proof), i.e., we can construct a control law

$\mathbf{u} = \mathbf{k}(\mathbf{x}, \hat{\boldsymbol{\theta}})$ such that:

$$L_{\tilde{\mathbf{f}}_{clf}} V_a(\mathbf{x}, \hat{\boldsymbol{\theta}}) + L_{\mathbf{g}} V_a(\mathbf{x}, \hat{\boldsymbol{\theta}}) \mathbf{k}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \leq -\alpha_3(\|\mathbf{x}\|, \hat{\boldsymbol{\theta}}), \quad (12)$$

where $\tilde{\mathbf{f}}_{clf}$ is given by:

$$\tilde{\mathbf{f}}_{clf}(\mathbf{x}, \hat{\boldsymbol{\theta}}) = \mathbf{f}(\mathbf{x}) + \mathbf{F}(\mathbf{x}) \boldsymbol{\lambda}_{clf}(\mathbf{x}, \hat{\boldsymbol{\theta}}). \quad (13)$$

We note that this controller only depends on the current estimate of the parameters $\hat{\boldsymbol{\theta}}$, and does not depend on the actual parameters $\boldsymbol{\theta}^*$. Define the parameter error:

$$\tilde{\boldsymbol{\theta}} = \boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}} \quad (14)$$

to be the difference between the actual and estimated parameters. Consider now the candidate composite Lyapunov function:

$$V(\mathbf{x}, \hat{\boldsymbol{\theta}}) = V_a(\mathbf{x}, \hat{\boldsymbol{\theta}}) + \frac{1}{2} \tilde{\boldsymbol{\theta}}^T \boldsymbol{\Gamma}^{-1} \tilde{\boldsymbol{\theta}}. \quad (15)$$

Computing its derivative we obtain:

$$\begin{aligned} \dot{V} &= \dot{V}_a - \tilde{\boldsymbol{\theta}}^T \boldsymbol{\Gamma}^{-1} \dot{\hat{\boldsymbol{\theta}}} \\ &\leq -\alpha_3(\|\mathbf{x}\|, \hat{\boldsymbol{\theta}}) + \tilde{\boldsymbol{\theta}}^T \mathbf{a}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \\ &\quad - \left(\frac{\partial V_a}{\partial \boldsymbol{\theta}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \right) \boldsymbol{\Gamma} \mathbf{a}(\mathbf{x}, \hat{\boldsymbol{\theta}}). \end{aligned}$$

where:

$$\mathbf{a}(\mathbf{x}, \hat{\boldsymbol{\theta}}) = \left(\left(\frac{\partial V_a}{\partial \mathbf{x}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \mathbf{F}(\mathbf{x}) \right)^T - \boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \right). \quad (16)$$

It is now easy to see that using the update law

$$\boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}) = \left(\frac{\partial V_a}{\partial \mathbf{x}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \mathbf{F}(\mathbf{x}) \right)^T \quad (17)$$

implies

$$\dot{V} \leq -\alpha_3(\|\mathbf{x}\|, \hat{\boldsymbol{\theta}}), \quad (18)$$

from which we conclude that the equilibrium point $(\mathbf{0}, \boldsymbol{\theta}^*)$ of (7) is stable. In particular, we see that $(\mathbf{x}(t), \hat{\boldsymbol{\theta}}(t))$ is globally bounded. It now follows from the LaSalle invariance principle that $\mathbf{x}(t)$ converges to the largest invariant subset of the collection of points $\mathbf{x} \in \mathcal{X}$ satisfying $\alpha_3(\|\mathbf{x}\|, \hat{\boldsymbol{\theta}}) = 0$ which is the singleton $\mathbf{x} = \mathbf{0}$. \square

As noted in the preceding proof, given an aCLF V_a , we can correspondingly synthesize a Lipschitz continuous (away from $\mathbf{x} = \mathbf{0}$) control law $\mathbf{k} : \mathcal{X} \times \Theta \rightarrow \mathcal{U}$. This can be achieved in a point-wise optimal fashion by considering an optimization based control framework. In particular, since the aCLF condition (11) is satisfied, we can consider the following quadratic program:

$$\begin{aligned} \mathbf{k}(\mathbf{x}, \hat{\boldsymbol{\theta}}) &= \underset{\mathbf{u} \in \mathcal{U}}{\operatorname{argmin}} \quad \frac{1}{2} \|\mathbf{u}\|^2 && \text{(aCLF-QP)} \\ \text{s.t.} \quad &\frac{\partial V_a}{\partial \mathbf{x}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \left(\tilde{\mathbf{f}}_{clf}(\mathbf{x}, \hat{\boldsymbol{\theta}}) + \mathbf{g}(\mathbf{x}) \mathbf{u} \right) \leq -\alpha_3(\|\mathbf{x}\|, \hat{\boldsymbol{\theta}}) \end{aligned}$$

This QP based controller will be guaranteed to have a solution, again because (11) is satisfied, and is Lipschitz continuous [16] (away from $\mathbf{0}$). Moreover, a closed form solution to this optimization problem, termed the *min-norm*

controller, can be obtained via the KKT conditions [4]. To see this, define:

$$\begin{aligned} \phi_0(\mathbf{x}, \hat{\boldsymbol{\theta}}) &\triangleq \frac{\partial V_a}{\partial \mathbf{x}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \tilde{\mathbf{f}}_{clf}(\mathbf{x}, \hat{\boldsymbol{\theta}}) + \alpha_3(\|\mathbf{x}\|, \hat{\boldsymbol{\theta}}) \\ \phi_1^T(\mathbf{x}, \hat{\boldsymbol{\theta}}) &\triangleq \frac{\partial V_a}{\partial \mathbf{x}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \mathbf{g}(\mathbf{x}) \end{aligned}$$

wherein the solution to (aCLF-QP) follows from :

$$\mathbf{k}(\mathbf{x}, \hat{\boldsymbol{\theta}}) = \begin{cases} -\frac{\phi_0(\mathbf{x}, \hat{\boldsymbol{\theta}}) \phi_1(\mathbf{x}, \hat{\boldsymbol{\theta}})}{\phi_1^T(\mathbf{x}, \hat{\boldsymbol{\theta}}) \phi_1(\mathbf{x}, \hat{\boldsymbol{\theta}})} & \text{if } \phi_0(\mathbf{x}, \hat{\boldsymbol{\theta}}) > 0 \\ 0 & \text{if } \phi_0(\mathbf{x}, \hat{\boldsymbol{\theta}}) \leq 0 \end{cases}$$

III. CONTROL BARRIER FUNCTIONS

The goal of this work is to provably enforce safety, even in the context of uncertain models. As a result, we will leverage the framework of Control Barrier Functions (CBFs) [1], [2], [27]. This section, therefore, will review the basic concepts related to these functions and corresponding controller synthesis.

In the context of safety, we consider a set S defined as the 0-superlevel set of a continuously differentiable function $h : \mathcal{X} \rightarrow \mathbb{R}$, yielding:

$$S \triangleq \{\mathbf{x} \in \mathcal{X} \mid h(\mathbf{x}) \geq 0\}, \quad (19)$$

$$\partial S \triangleq \{\mathbf{x} \in \mathcal{X} \mid h(\mathbf{x}) = 0\}, \quad (20)$$

$$\operatorname{int}(S) \triangleq \{\mathbf{x} \in \mathcal{X} \mid h(\mathbf{x}) > 0\}, \quad (21)$$

We refer to S as the *safe set*.

Consider again the known dynamics (1). A feedback control law $\mathbf{u} = \mathbf{k}(\mathbf{x})$ induces the closed loop dynamics:

$$\dot{\mathbf{x}} = \mathbf{f}_{cl}(\mathbf{x}) \triangleq \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x}) \mathbf{k}(\mathbf{x}) \quad (22)$$

which is assumed to be locally Lipschitz continuous. This assumption implies that for any initial condition $\mathbf{x}_0 \in \mathcal{X}$ there exists a maximum interval of existence $I(\mathbf{x}_0) = [0, \tau_{max})$ such that $\mathbf{x}(t)$ is the unique solution to (22) on $I(\mathbf{x}_0)$; in the case when \mathbf{f}_{cl} is forward complete, $\tau_{max} = \infty$. This notation allows us to define forward invariance and safety:

Definition 5 (Forward Invariant). The set S is *forward invariant* if for every $\mathbf{x}_0 \in S$, $\mathbf{x}(t) \in S$ for $\mathbf{x}(0) = \mathbf{x}_0$ and all $t \in I(\mathbf{x}_0)$.

Definition 6 (Safety). The system (22) is *safe* with respect to the set S if the set S is forward invariant.

It is desirable to achieve safety without the need to specify a specific controller as was done in (22). This leads to the notion of Control Barrier Functions. Before defining these, we require the following definition as in [2]:

Definition 7 (Extended Class \mathcal{K} Function). A continuous function $\alpha : (-b, a) \rightarrow \mathbb{R}$, with $a, b > 0$, is an *extended class \mathcal{K} function* ($\alpha \in \mathcal{K}_e$) if $\alpha(0) = 0$ and α is strictly monotonically increasing. If $a, b = \infty$, $\lim_{r \rightarrow \infty} \alpha(r) = \infty$, $\lim_{r \rightarrow -\infty} \alpha(r) = -\infty$. then α is said to be an *extended class \mathcal{K}_∞ function* ($\alpha \in \mathcal{K}_{\infty, e}$).

This enables the following definition as in [2]:

Definition 8 (*Control Barrier Function (CBF)*). Let $S \subset \mathcal{X}$ be the 0-superlevel set of a continuously differentiable function $h : \mathcal{X} \rightarrow \mathbb{R}$. h is a *Control Barrier Function* (CBF) for S if there exists an extended class \mathcal{K}_∞ function α such that for the system (1):

$$\sup_{\mathbf{u} \in \mathcal{U}} [L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} + \alpha(h(\mathbf{x}))] \geq 0 \quad (23)$$

for all $\mathbf{x} \in S$.

We can consider the point-wise set consisting of all control values that render S safe:

$$K_{cbf}(\mathbf{x}) = \{\mathbf{u} \in \mathcal{U} : L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} + \alpha(h(\mathbf{x})) \geq 0\}. \quad (24)$$

The main results of [1], [27] is that the existence of a CBF for S implies the system (1) can be rendered safe with respect to S :

Theorem 2. *Given a set $S \subset \mathcal{X}$ defined as the 0-superlevel set of continuously differentiable function $h : \mathcal{X} \rightarrow \mathbb{R}$, if h is a CBF on S , then any Lipschitz continuous controller \mathbf{k} such that $\mathbf{k}(\mathbf{x}) \in K_{cbf}(\mathbf{x})$ for all $\mathbf{x} \in S$ renders the system (1) safe with respect to the set S .*

In addition, if $\mathbf{k}(\mathbf{x}) \in K_{cbf}(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}$, then the set S is asymptotically stable in \mathcal{X} .

IV. ADAPTIVE CONTROL BARRIER FUNCTIONS

Motivated by the construction of adaptive Control Lyapunov Functions, we now explore the notion of an adaptive Control Barrier Function.

We again assume the control system has the form given in (4), wherein θ^* is a set of unknown parameters, and extend the previous construction of the safe set S to be parameter dependent. In this case, we construct a family of safe sets parameterized by θ and defined as the 0-superlevel sets of a continuously differentiable function $h_a : \mathcal{X} \times \Theta \rightarrow \mathbb{R}$:

$$S_\theta \triangleq \{\mathbf{x} \in \mathcal{X} \mid h_a(\mathbf{x}, \theta) \geq 0\}, \quad (25)$$

$$\partial S_\theta \triangleq \{\mathbf{x} \in \mathcal{X} \mid h_a(\mathbf{x}, \theta) = 0\}, \quad (26)$$

$$\text{int}(S_\theta) \triangleq \{\mathbf{x} \in \mathcal{X} \mid h_a(\mathbf{x}, \theta) > 0\}, \quad (27)$$

In particular, we will see this construction allows the states in the state space that are considered safe to change according to the current estimate of the parameters. If the set in the state space to be kept safe is independent of the parameters, the preceding construction is identical to that in (19)-(21).

Given this construction, we can define adaptively safe in a similar fashion to the definition of global adaptively stabilizable given in Definition 3 (note that in this case we opt for a local rather than global definition).

Definition 9 (*Adaptively Safe*). The system with unknown parameters (4) can be rendered *adaptively safe* with respect to a family of sets $S_{\hat{\theta}}$ if there exists a dynamic controller of the form (5)-(6) satisfying (A1)-(A3) such that solutions $(\mathbf{x}(t), \hat{\theta}(t))$ of (7) controlled by (5)-(6) satisfy $\mathbf{x}(t) \in S_{\hat{\theta}(t)}$ for all $t \in I(\mathbf{x}(0), \hat{\theta}(0))$.

This definition implies that the state of the system must remain within a potentially time-varying set, $S_{\hat{\theta}(t)}$, even in the presence of uncertainty in the dynamics. It is not necessary that the parameters converge, or even that they remain bounded, as in the adaptively stabilizable formulation. As will be seen, this is inherently connected to the fact that safety does not force the system to converge to an equilibrium point, but only requires it remains within a set.

Before defining aCBFs, we also specify that a set of adaptive gains G is defined such that:

$$\Gamma \in G \implies \Gamma \text{ satisfies A(3)}. \quad (28)$$

We note that G need not be all values of Γ satisfying A(3). We can now define aCBFs as an extension of Definitions 4 and 8.

Definition 10 (*Adaptive Control Barrier Function (aCBF)*). Let $S_\theta \subset \mathcal{X}$ be a family of 0-superlevel sets of a continuously differentiable function $h_a : \mathcal{X} \times \Theta \rightarrow \mathbb{R}$, with $\frac{\partial h_a}{\partial \mathbf{x}}$ Lipschitz continuous. Then h_a is an *adaptive Control Barrier Function* (aCBF) on the family of sets S_θ over adaptive gains G for (4) if for any $\theta \in \Theta$ and $\Gamma \in G$:

$$\sup_{\mathbf{u} \in \mathcal{U}} \left[\frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}, \theta) (\mathbf{f}(\mathbf{x}) + \mathbf{F}(\mathbf{x})\boldsymbol{\lambda}_{cbf}(\mathbf{x}, \theta) + \mathbf{g}(\mathbf{x})\mathbf{u}) \right] \geq 0. \quad (29)$$

with

$$\boldsymbol{\lambda}_{cbf}(\mathbf{x}, \theta) \triangleq \theta - \Gamma \left(\frac{\partial h_a}{\partial \theta}(\mathbf{x}, \theta) \right)^T. \quad (30)$$

Let us make a few observations of this definition:

Remark 2. As will be seen in the proof that an aCBF can ensure a system is adaptively safe, there is a requirement on the smallest eigenvalue of Γ . As not every value of Γ satisfying A(3) will satisfy this requirement, we must consider a restricted set of values for Γ , given by G . This leads to the incorporation of the set G in the definition of an aCBF. If the family of sets S_θ does not depend on θ , such that:

$$\frac{\partial h_a}{\partial \theta}(\mathbf{x}, \theta) \equiv \mathbf{0}, \quad (31)$$

then Γ will not appear in (29). This implies h_a being an aCBF for (4) will not depend on G .

Remark 3. The constraint in (29) differs from (23) in that the term $\alpha(h_a(\mathbf{x}, \theta))$ does not appear. Rather, this closely resembles early definitions of *barrier certificates* and *Lyapunov barrier functions* [20], [26], [23], which did not allow the state to approach the boundary of the safe sets, enforcing forward invariance of level sets of h_a .

We note that a QP-based Lipschitz continuous controller attaining safety can be constructed similarly to the (aCLF-QP) given an aCBF. We now have the necessary framework in which to present the main result of this paper—that the existence of an aCBF implies safety of the family of sets $S_{\hat{\theta}}$ even under parameter uncertainty.

Theorem 3. *Let $h_a : \mathcal{X} \rightarrow \mathbb{R}$ be an adaptive Control Barrier Function on the family of sets $S_{\hat{\theta}}$ over G . Assume that $\tilde{\theta}_0 =$*

$\tilde{\boldsymbol{\theta}}(0)$ with $\|\tilde{\boldsymbol{\theta}}_0\|_2 \leq c$ for $c > 0$ and $\mathbf{x}_0 = \mathbf{x}(0) \in \text{int}(S_{\hat{\boldsymbol{\theta}}_0})$. If there exists a positive definite gain matrix, $\boldsymbol{\Gamma} \in G$, such that:

$$\lambda_{\min}(\boldsymbol{\Gamma}) \geq \frac{c^2}{2h_a(\mathbf{x}_0, \hat{\boldsymbol{\theta}}_0)}, \quad (32)$$

then there exists a Lipschitz continuous function $\boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}})$ such that for the update law:

$$\dot{\hat{\boldsymbol{\theta}}} = \boldsymbol{\Gamma}\boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}), \quad (33)$$

the family of sets $S_{\hat{\boldsymbol{\theta}}}$ is forward invariant.

The main idea is to approach the proof much in the same way as the proof of Theorem 1. Yet the construction of a composite CBF as was done in (15) in the case of aCLFs requires more care. Adding the parameter error term would result in the composite safety function 0-superlevel set properly containing the 0-superlevel set of the aCBF, adding additional states to the set that can be rendered safe. This extension of the safe set can be quantified if the parameter estimates (and thus the parameter error) remains bounded, as in the case of aCLFs, but this is not guaranteed given the necessary form of $\boldsymbol{\tau}$.

Proof. Define the following composite candidate CBF for the extended system dynamics (7):

$$h(\mathbf{x}, \hat{\boldsymbol{\theta}}) = h_a(\mathbf{x}, \hat{\boldsymbol{\theta}}) - \frac{1}{2}\tilde{\boldsymbol{\theta}}^T \boldsymbol{\Gamma}^{-1}\tilde{\boldsymbol{\theta}} \quad (34)$$

By assumption, $\mathbf{x}_0 \in \text{int}(S_{\hat{\boldsymbol{\theta}}_0})$, implying that $h_a(\mathbf{x}_0, \hat{\boldsymbol{\theta}}_0) > 0$. Further, our assumption that $\|\tilde{\boldsymbol{\theta}}_0\|_2 \leq c$ implies that:

$$\frac{1}{2}\tilde{\boldsymbol{\theta}}_0^T \boldsymbol{\Gamma}^{-1}\tilde{\boldsymbol{\theta}}_0 \leq \frac{1}{2\lambda_{\min}(\boldsymbol{\Gamma})}\|\tilde{\boldsymbol{\theta}}_0\|_2^2 \leq \frac{c^2}{2\lambda_{\min}(\boldsymbol{\Gamma})} \quad (35)$$

Therefore, choosing $\boldsymbol{\Gamma}$ such that

$$\lambda_{\min}(\boldsymbol{\Gamma}) \geq \frac{c^2}{2h_a(\mathbf{x}_0, \hat{\boldsymbol{\theta}}_0)} \quad (36)$$

leads to:

$$h(\mathbf{x}_0, \hat{\boldsymbol{\theta}}_0) \geq 0 \quad (37)$$

Now consider the time derivative of h as given in Table I. The second equality follows the addition and subtraction of the term:

$$\frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}, \hat{\boldsymbol{\theta}})\mathbf{F}(\mathbf{x}) \left(\hat{\boldsymbol{\theta}} - \boldsymbol{\Gamma} \left(\frac{\partial h_a}{\partial \boldsymbol{\theta}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \right)^T \right) \quad (38)$$

The third equality is a rearrangement revealing the form of the aCBF time derivative as given in (29)-(30). In particular, condition (29) permits the choice of an input \mathbf{u} such that the first inequality is satisfied. Choosing the update law $\boldsymbol{\tau}$ as:

$$\boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}) = - \left(\frac{\partial h_a}{\partial \boldsymbol{\theta}}(\mathbf{x}, \hat{\boldsymbol{\theta}})\mathbf{F}(\mathbf{x}) \right)^T \quad (39)$$

results in the last inequality. This inequality, in conjunction with (37) and the comparison lemma in [8] imply that

$$h(\mathbf{x}(t), \hat{\boldsymbol{\theta}}(t)) \geq 0 \quad (40)$$

for all $t \geq 0$. Given the construction of h in (34), it follows that:

$$h_a(\mathbf{x}(t), \hat{\boldsymbol{\theta}}(t)) \geq \frac{1}{2}\tilde{\boldsymbol{\theta}}(t)^T \boldsymbol{\Gamma}^{-1}\tilde{\boldsymbol{\theta}}(t) \geq 0. \quad (41)$$

Lastly, we conclude that $\mathbf{x}(t) \in S_{\hat{\boldsymbol{\theta}}(t)}$ for $t \geq 0$. \square

The proof reveals that superlevel sets of h are forward invariant. As h can not be computed without knowing the true parameters $\boldsymbol{\theta}^*$, it is not possible to set $\dot{h} \geq -\alpha(h)$ as is typical with CBFs. Furthermore, we have that $h_a \geq h$, implying that $-\alpha(h_a) \leq -\alpha(h)$. Thus setting $\dot{h} \geq -\alpha(h_a)$ does not yield the desired lower bound on \dot{h} . One may note that setting $\dot{h} \geq -\alpha(h_a)$ leads to $\dot{h} \geq 0$ when $h_a = 0$, or when the state is on the boundary of the safe set. This fact is concurrent with the common forward invariance proof technique utilizing Nagumo's theorem [17]. Despite this, it is in fact possible to construct simple examples (in \mathbb{R}^2) such that the state must leave the safe set defined by h_a for any choice of differentiable α and $\boldsymbol{\Gamma}$.

Remark 4. The assumption on $\tilde{\boldsymbol{\theta}}$ implies that the initial parameter error must be bounded, unlike the aCLF formulation. This is due to the fact that we seek to keep a particular set forward invariant. In contrast, the only set kept provably forward invariant in the aCLF formulation is the sublevel set of the composite Lyapunov function V corresponding to the initial conditions $(\mathbf{x}(0), \hat{\boldsymbol{\theta}}(0))$. Evaluating that set would too require assumptions on the boundedness of $\tilde{\boldsymbol{\theta}}(0)$. Additionally, while this may seem restrictive, we note that the input for the system will not be chosen to be robust to all uncertainties in this initial uncertainty set. Rather, the uncertainty will be handled by adapting parameter estimates.

Remark 5. The lower bound on the adaptive gain allows us to ensure that the system can adapt quickly enough to ensure safety from the given initial condition. Initial distance from the safety set boundary and smaller possible initial parameter error allow the adaptive gain to be made smaller.

A quadratic program based controller similar to (aCLF-QP) can be constructed using an aCBF. To this end, we adopt the safety-critical control formulation in [24], [7] that filters a desired but potentially unsafe control law $\mathbf{k}_d : \mathcal{X} \times \Theta \rightarrow \mathcal{U}$ to find the nearest safe control action:

$$\begin{aligned} \mathbf{k}(\mathbf{x}, \hat{\boldsymbol{\theta}}) &= \underset{\mathbf{u} \in \mathcal{U}}{\text{argmin}} \frac{1}{2}\|\mathbf{u} - \mathbf{k}_d(\mathbf{x}, \hat{\boldsymbol{\theta}})\|^2 & (\text{aCBF-QP}) \\ \text{s.t.} \quad & \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}, \hat{\boldsymbol{\theta}})(\tilde{\mathbf{f}}_{cbf}(\mathbf{x}, \hat{\boldsymbol{\theta}}) + \mathbf{g}(\mathbf{x})\mathbf{u}) \geq 0 \end{aligned}$$

where $\tilde{\mathbf{f}}_{cbf}$ is defined like $\tilde{\mathbf{f}}_{clf}$ in (13). As with (aCLF-QP), this quadratic program has a closed form solution.

V. ADAPTIVE CRUISE CONTROL

To demonstrate how an aCBF can be used to render a system adaptively safe, we consider the problem of Adaptive Cruise Control (ACC) as posed in [1]. The dynamics of the system are given by:

$$\frac{d}{dt} \begin{bmatrix} v \\ D \end{bmatrix} = \begin{bmatrix} 0 \\ v_0 - v \end{bmatrix} - \frac{1}{m} \begin{bmatrix} 1 & v & v^2 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} f_0 \\ f_1 \\ f_2 \end{bmatrix} + \begin{bmatrix} \frac{1}{m} \\ 0 \end{bmatrix} u \quad (42)$$

$$\begin{aligned}
\dot{h}(\mathbf{x}, \hat{\boldsymbol{\theta}}, \mathbf{u}) &= \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) (\mathbf{f}(\mathbf{x}) + \mathbf{F}(\mathbf{x})\boldsymbol{\theta}^* + \mathbf{g}(\mathbf{x})\mathbf{u}) + \frac{\partial h_a}{\partial \boldsymbol{\theta}}(\mathbf{x}, \hat{\boldsymbol{\theta}})\dot{\hat{\boldsymbol{\theta}}} + \tilde{\boldsymbol{\theta}}^\top \boldsymbol{\Gamma}^{-1}\dot{\hat{\boldsymbol{\theta}}} \\
&= \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) (\mathbf{f}(\mathbf{x}) + \mathbf{F}(\mathbf{x})\boldsymbol{\theta}^* + \mathbf{g}(\mathbf{x})\mathbf{u}) + \frac{\partial h_a}{\partial \boldsymbol{\theta}}(\mathbf{x}, \hat{\boldsymbol{\theta}})\boldsymbol{\Gamma}\boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}) + \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}, \hat{\boldsymbol{\theta}})\mathbf{F}(\mathbf{x}) \left(\hat{\boldsymbol{\theta}} - \boldsymbol{\Gamma} \left(\frac{\partial h_a}{\partial \boldsymbol{\theta}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \right)^\top \right) \\
&\quad - \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}, \hat{\boldsymbol{\theta}})\mathbf{F}(\mathbf{x}) \left(\hat{\boldsymbol{\theta}} - \boldsymbol{\Gamma} \left(\frac{\partial h_a}{\partial \boldsymbol{\theta}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \right)^\top \right) + \tilde{\boldsymbol{\theta}}^\top \boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \\
&= \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \left(\mathbf{f}(\mathbf{x}) + \mathbf{F}(\mathbf{x}) \left(\hat{\boldsymbol{\theta}} - \boldsymbol{\Gamma} \left(\frac{\partial h_a}{\partial \boldsymbol{\theta}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \right)^\top \right) + \mathbf{g}(\mathbf{x})\mathbf{u} \right) \\
&\quad + \frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}, \hat{\boldsymbol{\theta}})\mathbf{F}(\mathbf{x}) \left(\tilde{\boldsymbol{\theta}} + \boldsymbol{\Gamma} \left(\frac{\partial h_a}{\partial \boldsymbol{\theta}}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \right)^\top \right) + \frac{\partial h_a}{\partial \boldsymbol{\theta}}(\mathbf{x}, \hat{\boldsymbol{\theta}})\boldsymbol{\Gamma}\boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}) + \tilde{\boldsymbol{\theta}}^\top \boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \\
&\geq \left(\frac{\partial h_a}{\partial \boldsymbol{\theta}}(\mathbf{x}, \hat{\boldsymbol{\theta}})\boldsymbol{\Gamma} + \tilde{\boldsymbol{\theta}}^\top \right) \left(\left(\frac{\partial h_a}{\partial \mathbf{x}}(\mathbf{x}, \hat{\boldsymbol{\theta}})\mathbf{F}(\mathbf{x}) \right)^\top + \boldsymbol{\tau}(\mathbf{x}, \hat{\boldsymbol{\theta}}) \right) \\
&\geq 0
\end{aligned}$$

TABLE I. Calculation of \dot{h} as used in the proof of the main result.

with v the velocity of the vehicle, D the distance between the vehicle and a leading vehicle traveling at a fixed velocity v_0 , m the vehicle's mass, and f_0 , f_1 , and f_2 unknown parameters associated with rolling frictional force. Define the state as $\mathbf{x} = (v, D)$. In this problem, we seek to drive the velocity to a desired velocity, v_d , while simultaneously ensuring the distance between the vehicles satisfies a safety constraint given by:

$$D \geq 1.8v. \quad (43)$$

The parameters f_0 , f_1 , and f_2 are often determined empirically, and if they are not accurate, the desired velocity may not be accurately tracked. Furthermore, if the parameters do not exactly match the true parameters, it may not be possible to certify that the system will satisfy the safety constraint.

The control objective of tracking a desired velocity can be achieved with a hand-designed controller \mathbf{k}_d or encoded using a CLF, and the safety constraint can be encoded using a CBF. Additional constraints on the maximum acceleration and deceleration can be enforced to maintain passenger comfort. To handle uncertainty in the parameters, we utilize the tool of aCBFs to maintain and update estimates of these parameters. An aCBF that yields desirable results is defined as the following continuously differentiable function:

$$h_a(v, D) = \begin{cases} \alpha^2 & \text{if } D - 1.8v \geq \alpha \\ \alpha^2 - (D - 1.8v - \alpha)^2 & \text{if } D - 1.8v < \alpha \end{cases}$$

for $\alpha > 0$. This particular construction of h_a is constant away from the safety boundary and diminishes to 0 (quadratically to preserve differentiability) as the boundary is approached. In practice, this is to handle the fact that superlevel sets of the composite safety function h are forward invariant. In regions where h_a is constant, $\frac{\partial h_a}{\partial \mathbf{x}}$, and thus the update law in (39), is 0, thus making $\dot{h} = 0$ as in the first equality in Table I.

aCBF-QP Controller: A simple proportional controller on

tracking error $v - v_d$ can be implemented and achieve good tracking performance, but is not necessarily safe. A CBF alone would not ensure the safety of this controller with model uncertainty, but treating the proportional controller as \mathbf{k}_d in aCBF-QP with an aCBF, safety can be achieved.

aCLF-aCBF-QP Controller: Additionally, we can unify aCLF and aCBFs in a quadratic program based controller to receive the benefits of optimal and adaptive tracking while remaining safe. Separate estimates of the parameters are maintained for the aCLF and the aCBF, as the form of the update laws in (17) and (39) may not be simultaneously satisfiable for only one estimate of the parameters. The CLF in [1] on the velocity tracking error $v - v_d$, given by $V_a = (v - v_d)^2$, also satisfies the aCLF condition (11). Letting $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\psi}}$ be parameter estimates associated with the aCLF and aCBF, respectively, we formulate a QP-based controller:

$$\begin{aligned}
k(\mathbf{x}, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\psi}}) &= \underset{u \in \mathcal{U}}{\operatorname{argmin}} J(u) + c_V(\mathbf{x})\delta_V + c_p(\mathbf{x})\delta_p \\
\text{s.t. } &L_{\hat{\mathbf{f}}_{clf}} V_a(\mathbf{x}, \hat{\boldsymbol{\theta}}) + L_{\mathbf{g}} V_a(\mathbf{x}, \hat{\boldsymbol{\theta}})u \leq -\alpha_3(\|\mathbf{x}\|, \hat{\boldsymbol{\theta}}) + \delta_V \\
&L_{\hat{\mathbf{f}}_{cbf}} h_a(\mathbf{x}, \hat{\boldsymbol{\psi}}) + L_{\mathbf{g}} h_a(\mathbf{x}, \hat{\boldsymbol{\psi}})u \geq 0 \\
&u \leq u_{max} + \delta_p \\
&u \geq -u_{max} - \delta_p \\
&\delta_V, \delta_p \geq 0
\end{aligned}$$

with parameter updates for $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\psi}}$ as in (17) and (39), respectively. δ_V and δ_p are relaxations to the optimization problem to ensure its feasibility, while safety is ensured. The functions c_V and c_p are Lipschitz continuous and are used to achieve smoothness. With initial parameter estimates $[\hat{f}_0 \ \hat{f}_1 \ \hat{f}_2] = 10 [f_0^* \ f_1^* \ f_2^*]$ (less friction than modeled), the results of this controller appear in Figure 1.

We see that the proportional controller fails to keep the vehicle safe, but filtering it with the aCBF-QP keeps it safe (with $D \geq 1.8v$ for all time) even with model uncertainty.




Fig. 1. Comparison of different adaptive and non-adaptive control methodologies. The aCBF-QP is able to enforce safety of the proportional controller (left). An aCLF controller is able to track the desired velocity with zero steady state error (center). Both aCBF controllers are able to keep the vehicle within the safe region for all time (right).

A CLF-CBF controller with no adaptive elements fails to either track the desired velocity (with steady state error) or keep the vehicle safe. The CLF-aCBF controller keeps the vehicle safe but has steady state tracking error, while an aCLF-aCBF controller accurately tracks the desired velocity with no steady state error, and keeps the vehicle safe.

VI. CONCLUSION

We presented a novel approach for ensuring the safety of a system under a form of parametric uncertainty. This approach builds off the structure established with adaptive Control Lyapunov Functions, and highlights the differences that must be considered when ensuring the forward invariance of a specific set. Future work includes considering this framework within a batched-data framework, in which initial parametric uncertainty can be iteratively and episodically reduced to permit less conservative safe sets.

REFERENCES

- [1] A. Ames, J. Grizzle, and P. Tabuada. Control barrier function based quadratic programs with application to adaptive cruise control. In *Conference on Decision & Control (CDC)*, pages 6271–6278. IEEE, 2014.
- [2] 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, 2017.
- [3] Z. Artstein. Stabilization with relaxed controls. *Nonlinear Analysis: Theory, Methods & Applications*, 7(11):1163–1173, 1983.
- [4] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [5] R. Cheng, G. Orosz, R. M. Murray, and J. W. Burdick. End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3387–3395, 2019.
- [6] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin. A general safety framework for learning-based control in uncertain robotic systems. *IEEE Transactions on Automatic Control*, 2018.
- [7] T. Gurriet, A. Singletary, J. Reher, L. Ciarletta, E. Feron, and A. Ames. Towards a framework for realizable safety critical control through active set invariance. In *International Conference on Cyber-Physical Systems*, pages 98–106. IEEE Press, 2018.
- [8] H.K. Khalil. *Nonlinear Systems - 3rd Edition*. PH, Upper Saddle River, NJ, 2002.
- [9] S. Kolathaya and A. D. Ames. Input-to-state safety with control barrier functions. *IEEE Control Systems Letters*, 3(1):108–113, 2018.
- [10] M. Krstić, I. Kanellakopoulos, and P. V. Kokotović. *Nonlinear and adaptive control design*. Wiley New York, 1995.
- [11] M. Krstić and P. V. Kokotović. Control lyapunov functions for adaptive nonlinear stabilization. *Systems & Control Letters*, 26(1):17–23, 1995.
- [12] M. Krstić and P. V. Kokotović. Modular approach to adaptive nonlinear stabilization. *Automatica*, 32(4):625–629, 1996.
- [13] Z. Li and M. Krstić. Optimal design of adaptive tracking controllers for non-linear systems. *Automatica*, 33(8):1459–1473, 1997.
- [14] Y. Lin and E. D. Sontag. A universal formula for stabilization with bounded controls. *Systems & Control Letters*, 16(6):393–397, 1991.
- [15] J. Moore and R. Tedrake. Adaptive control design for underactuated systems using sums-of-squares optimization. In *American Control Conference (ACC)*, pages 721–728. IEEE, 2014.
- [16] B. J. Morris, M. J. Powell, and A. D. Ames. Continuity and smoothness properties of nonlinear optimization-based feedback controllers. In *Conference on Decision and Control (CDC)*, pages 151–158. IEEE, 2015.
- [17] M. Nagumo. Über die lage der integralkurven gewöhnlicher differentialgleichungen. *Proceedings of the Physico-Mathematical Society of Japan. 3rd Series*, 24:551–559, 1942.
- [18] Q. Nguyen and K. Sreenath. Exponential control barrier functions for enforcing high relative-degree safety-critical constraints. In *American Control Conference (ACC)*, pages 322–328. IEEE, 2016.
- [19] Motoya Ohnishi, Li Wang, Gennaro Notomista, and Magnus Egerstedt. Barrier-certified adaptive reinforcement learning with applications to brushbot navigation. *IEEE Transactions on Robotics*, 35(5):1186–1205, 2019.
- [20] S. Prajna. *Optimization-based methods for nonlinear and hybrid systems verification*. PhD thesis, Caltech, 2005.
- [21] E. D. Sontag. Smooth stabilization implies coprime factorization. *IEEE Transactions on Automatic Control*, 34(4):435–443, 1989.
- [22] E. D. Sontag. A universal construction of artstein’s theorem on nonlinear stabilization. *Systems & Control Letters*, 13(2):117–123, 1989.
- [23] K. P. Tee, S. S. Ge, and E. H. Tay. Barrier lyapunov functions for the control of output-constrained nonlinear systems. *Automatica*, 45(4):918–927, 2009.
- [24] L. Wang, A. D. Ames, and M. Egerstedt. Safe certificate-based maneuvers for teams of quadrotors using differential flatness. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3293–3298. IEEE, 2017.
- [25] L. Wang, E. A. Theodorou, and M. Egerstedt. Safe learning of quadrotor dynamics using barrier certificates. In *International Conference on Robotics and Automation (ICRA)*, pages 2460–2465. IEEE, 2018.
- [26] P. Wieland and F. Allgöwer. Constructive safety using control barrier functions. *IFAC Proceedings Volumes*, 40(12):462–467, 2007.
- [27] X. Xu, P. Tabuada, J. W. Grizzle, and A. D. Ames. Robustness of control barrier functions for safety critical control. *IFAC-PapersOnLine*, 48(27):54–61, 2015.