

Characterizing Smooth Safety Filters via the Implicit Function Theorem

Max H. Cohen[✉], Pio Ong[✉], *Member, IEEE*, Gilbert Bahati, *Student Member, IEEE*,
and Aaron D. Ames[✉], *Fellow, IEEE*

Abstract—Optimization-based safety filters, such as control barrier function (CBF) based quadratic programs (QPs), have demonstrated success in controlling autonomous systems to achieve complex goals. These CBF-QPs can be shown to be continuous, but are generally not smooth, let alone continuously differentiable. In this letter, we present a general characterization of smooth safety filters – smooth controllers that guarantee safety in a minimally invasive fashion – based on the Implicit Function Theorem. This characterization leads to families of smooth universal formulas for safety-critical controllers that quantify the conservatism of the resulting safety filter, the utility of which is demonstrated through illustrative examples.

Index Terms—Constrained control, Lyapunov methods, nonlinear systems, safety-critical control.

I. INTRODUCTION

CONTROL barrier functions (CBFs) [1] have proven to be a powerful tool for designing controllers enforcing safety on nonlinear systems. The properties of CBFs naturally lead to their use as *safety filters* for nominal controllers that may not have been designed a priori to ensure safety. Most often, such safety filters are instantiated via optimization problems – typically a quadratic program (QP) – to minimize the deviation from a nominal controller while satisfying Lyapunov-like conditions that ensure forward invariance of a designated safe set [2], [3], [4]. Under certain regularity conditions, these optimization-based safety filters are locally Lipschitz functions of the system state [5], [6], allowing one to leverage set-theoretic tools [7] to conclude forward invariance of safe sets. Although such controllers are pointwise optimal, they are typically not smooth even if the problem data is.

The lack of smoothness exhibited by optimization-based safety filters has not been overly detrimental to the development of safety-critical controllers to date; however, more recent developments in the literature motivate the

consideration of smooth safety filters. For example, [8] recently proposed a barrier-backstepping methodology that enables the systematic construction of CBFs for systems in strict-feedback form. As in Lyapunov backstepping [9], such an approach requires differentiating through virtual CBF controllers at intermediate layers to construct a composite CBF for the overall system. Similar ideas are leveraged in [10], [11] to construct CBFs for robotic systems based on reduced-order models – an approach successfully used to safely control complex robotic systems, e.g., walking robots and drones. Ultimately, smoothness plays a crucial role in applying CBFs to more complex systems.

Given the similarities between control Lyapunov functions (CLFs) and CBFs, one may wonder if it is possible to adapt smooth universal formulas for CLFs [12] to CBFs. The answer is affirmative – with some slight modifications. Sontag’s Universal Formula for stabilization [12] can be applied to safety as noted in [13], [14], [15], [16]. Despite this, Sontag’s formula is scarcely used as a safety filter since, in its most common form, such a controller tends to be overly invasive, overriding inputs from the nominal controller even when not necessary to ensure safety. Alternative smooth universal formulas have been proposed in [17] based on computing weighted centroids of the set of all control values satisfying CBF and/or CLF conditions using the probability density function of a Gaussian distribution. In a different approach, the authors of [18] leverage Sontag’s formula to combine stabilization and safety objectives. Yet questions remain around the connections between smoothness and safety filters.

The main objective of this letter is to provide a general characterization of smooth safety filters – smooth controllers that guarantee safety in a minimally invasive fashion. Our characterization is motivated by the original development of Sontag’s formula in [12]. Sontag’s formula is derived by computing the roots of an algebraic equation parameterized by the Lie derivatives of a CLF/CBF. When certain regularity conditions are met, the smoothness of such roots as a function of the Lie derivatives follows directly from the Implicit Function Theorem. In this letter, we seek a deeper understanding of the properties of this equation. In particular: What properties should this equation satisfy so that one of its solutions produces a smooth minimally invasive controller that guarantees safety? We answer this question by constructing an ordinary differential equation (ODE) from a given algebraic one such that the trajectories of this ODE coincide with the

Manuscript received 15 September 2023; revised 20 November 2023; accepted 6 December 2023. Date of publication 11 December 2023; date of current version 9 January 2024. This work was supported by NSF Award under Grant 1932091. Recommended by Senior Editor J. Daafouz. (Corresponding author: Max H. Cohen.)

The authors are with the Department of Mechanical and Civil Engineering, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: maxcohen@caltech.edu; pioong@caltech.edu; gbahati@caltech.edu; ames@caltech.edu).

Digital Object Identifier 10.1109/LCSYS.2023.3341345

2475-1456 © 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See <https://www.ieee.org/publications/rights/index.html> for more information.

solutions of the algebraic equation. Leveraging invariance-like tools, we introduce sufficient conditions for this ODE so that its trajectories produce a smooth safety filter. This characterization leads to various smooth universal formulas for safety-critical control that allow one to assess the conservatism of the resulting safety filter, the utilities of which we illustrate through their application to safety-critical control based on reduced order models [10], [11].

II. PRELIMINARIES AND PROBLEM FORMULATION

Consider the nonlinear control affine system:

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

where $\mathbf{x} \in \mathbb{R}^n$ is the system state, $\mathbf{u} \in \mathbb{R}^m$ is the control input, $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is the drift vector field, and $\mathbf{g} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ captures the control directions. Throughout this letter, we assume that $\mathbf{f}$ and $\mathbf{g}$ are smooth¹ functions of the state. Applying a smooth feedback controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ to (1) produces the smooth closed-loop system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x}), \quad (2)$$

which, for each initial condition $\mathbf{x}_0 \in \mathbb{R}^n$, generates a unique smooth solution $\mathbf{x} : I(\mathbf{x}_0) \rightarrow \mathbb{R}^n$ satisfying (2) on some maximal interval of existence $I(\mathbf{x}_0) \subseteq \mathbb{R}_{\geq 0}$.

A. CLFs and Sontag's Universal Formula

Before discussing smooth safety filters, we recount the main ideas behind CLFs as presented in [12]. Recall that a smooth, proper, positive definite function $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ is a CLF for (1) if for all $\mathbf{x} \in \mathbb{R}^n \setminus \{0\}$

$$\inf_{\mathbf{u} \in \mathbb{R}^m} \{ \underbrace{\nabla V(\mathbf{x}) \cdot \mathbf{f}(\mathbf{x})}_{L_f V(\mathbf{x})} + \underbrace{\nabla V(\mathbf{x}) \cdot \mathbf{g}(\mathbf{x})\mathbf{u}}_{L_g V(\mathbf{x})} \} < 0.$$

The existence of a CLF implies that for each $\mathbf{x} \in \mathbb{R}^n$ there exists a $\mathbf{u} \in \mathbb{R}^m$ that enforces V to decrease, and allows for constructing a feedback controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ that renders the origin asymptotically stable by ensuring that:

$$\forall \mathbf{x} \in \mathbb{R}^n \setminus \{0\} : L_f V(\mathbf{x}) + L_g V(\mathbf{x})\mathbf{k}(\mathbf{x}) < 0. \quad (3)$$

In this letter, we are concerned with designing *smooth* feedback controllers satisfying a general class of affine inequalities, such as the one in (3). In [12] Sontag provides one example of such a controller, now known as Sontag's Universal Formula for stabilization, which is given by:

$$\mathbf{k}(\mathbf{x}) = \lambda_{\text{CLF}}(L_f V(\mathbf{x}), \|L_g V(\mathbf{x})\|^2) L_g V(\mathbf{x})^\top, \quad (4)$$

$$\lambda_{\text{CLF}}(a, b) := \begin{cases} 0 & b = 0 \\ -\frac{a - \sqrt{a^2 + q(b)b}}{b} & b \neq 0, \end{cases} \quad (5)$$

where $q : \mathbb{R} \rightarrow \mathbb{R}$ is smooth and satisfies $q(0) = 0$ and $q(b) > 0$ for all $b \neq 0$. In [12], the smoothness of Sontag's formula (5) is proven using the Implicit Function Theorem [19, Th. 11.2]. Specifically, the following result [19, Th. 11.1] is useful for establishing smoothness of (5).

Theorem 1 [19]: Let $(a, b, p) \mapsto F(a, b, p)$ be a smooth function. If a continuous function $\lambda : \mathcal{S} \rightarrow \mathbb{R}$ satisfies:

$$F(a, b, \lambda(a, b)) = 0, \quad (6)$$

¹By smooth, we mean differentiable as many times as necessary.

$$\frac{\partial F}{\partial p}(a, b, \lambda(a, b)) \neq 0, \quad (7)$$

for all $(a, b) \in \mathcal{S}$, then λ is smooth for all $(a, b) \in \mathcal{S}$ and its derivative is given by:

$$\begin{bmatrix} \frac{\partial \lambda}{\partial a}(a, b) \\ \frac{\partial \lambda}{\partial b}(a, b) \end{bmatrix} = -\frac{1}{\frac{\partial F}{\partial p}(a, b, \lambda(a, b))} \begin{bmatrix} \frac{\partial F}{\partial a}(a, b, \lambda(a, b)) \\ \frac{\partial F}{\partial b}(a, b, \lambda(a, b)) \end{bmatrix}. \quad (8)$$

To apply this theorem and show smoothness of the function λ_{CLF} in (5), one considers the smooth function:

$$F(a, b, p) = bp^2 + 2ap - q(b), \quad (9)$$

noting that λ_{CLF} is continuous and satisfies (6) and (7) for each (a, b) in $\mathcal{S}_{\text{CLF}} := \{(a, b) \in \mathbb{R} \times \mathbb{R}_{\geq 0} : a < 0 \text{ or } b > 0\}$. In addition to being smooth, one can also verify that the controller (4) constructed from λ_{CLF} satisfies inequality (3). Indeed, this is a consequence of picking an appropriate function for F , a point which we will expand on later.

B. CBFs and Safety Filters

The main objective of this letter is to leverage Theorem 1 for constructing smooth controllers that render the resulting closed-loop system *safe*, a property that requires system trajectories to remain within a desirable subset of the state space. Formally, we say that (2) is safe on a set $\mathcal{C} \subset \mathbb{R}^n$ if $\mathcal{C}$ is forward invariant [7]. By considering sets of the form:

$$\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}) \geq 0\}, \quad (10)$$

where $h : \mathbb{R}^n \rightarrow \mathbb{R}$ is a smooth function, the existence of a safe feedback controller can be characterized using the concept of a control barrier function (CBF) [1].

Definition 1: A smooth function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ defining a set $\mathcal{C} \subset \mathbb{R}^n$ as in (10) is said to be a *control barrier function* (CBF) for (1) on $\mathcal{C}$ if zero is a regular value of h and there exists a smooth² $\alpha \in \mathcal{K}_\infty^e$ such that for all $\mathbf{x} \in \mathbb{R}^n$

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \{ \underbrace{\nabla h(\mathbf{x}) \cdot \mathbf{f}(\mathbf{x})}_{L_f h(\mathbf{x})} + \underbrace{\nabla h(\mathbf{x}) \cdot \mathbf{g}(\mathbf{x})\mathbf{u}}_{L_g h(\mathbf{x})} \} > -\alpha(h(\mathbf{x})).$$

Similar to CLFs, any feedback controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ satisfying the inequality $L_f h(\mathbf{x}) + L_g h(\mathbf{x})\mathbf{k}(\mathbf{x}) \geq -\alpha(h(\mathbf{x}))$ ensures that h remains positive along closed-loop trajectories, and therefore renders $\mathcal{C}$ forward invariant [1]. Perhaps the greatest utility of CBFs is their ability to act as a *safety filter* for a nominal feedback controller $\mathbf{k}_d : \mathbb{R}^n \rightarrow \mathbb{R}^m$. A safety filter is a controller that modifies $\mathbf{k}_d$ – preferably, in a minimally invasive fashion – so that the resulting closed-loop system is safe. The most common examples of such safety filters are instantiated via the QP:

$$\begin{aligned} \mathbf{k}(\mathbf{x}) &= \arg \min_{\mathbf{u} \in \mathbb{R}^m} \frac{1}{2} \|\mathbf{u} - \mathbf{k}_d(\mathbf{x})\|^2 \\ \text{subject to} \quad & L_f h(\mathbf{x}) + L_g h(\mathbf{x})\mathbf{u} \geq -\alpha(h(\mathbf{x})), \end{aligned} \quad (11)$$

whose solution modifies $\mathbf{k}_d$ in the $L_g h^\top$ direction:

$$\mathbf{k}(\mathbf{x}) = \mathbf{k}_d(\mathbf{x}) + \lambda(a(\mathbf{x}), b(\mathbf{x})) L_g h(\mathbf{x})^\top, \quad (12)$$

with a scalar function $\lambda : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$, where:

²A continuous function $\alpha : \mathbb{R} \rightarrow \mathbb{R}$ is an extended class $\mathcal{K}_\infty$ function ($\alpha \in \mathcal{K}_\infty^e$) if $\alpha(0) = 0$, α is increasing, and $\lim_{r \rightarrow \pm\infty} \alpha(r) = \pm\infty$.

$$\begin{aligned} a(\mathbf{x}) &:= L_f h(\mathbf{x}) + L_g h(\mathbf{x}) \mathbf{k}_d(\mathbf{x}) + \alpha(h(\mathbf{x})) \\ b(\mathbf{x}) &:= \|L_g h(\mathbf{x})\|^2. \end{aligned} \quad (13)$$

For the QP in (11), λ is the Lagrange multiplier associated with the constraint and is given by:

$$\lambda(a, b) = \lambda_{\text{QP}}(a, b) := \begin{cases} 0 & b = 0 \\ \text{ReLU}(-a/b) & b > 0, \end{cases} \quad (14)$$

where $\text{ReLU}(y) := \max\{0, y\}$. This QP-based controller has the advantage of being pointwise optimal, but is not smooth even if the problem data itself is smooth. Before proceeding, we make precise the notion of a smooth safety filter.

Definition 2: Given a CBF $h : \mathbb{R}^n \rightarrow \mathbb{R}$ and nominal controller $\mathbf{k}_d : \mathbb{R}^n \rightarrow \mathbb{R}^m$ for (1), a controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ of the form (12) is said to be a *smooth safety filter* for (1) with respect to $\mathcal{C}$ if $\mathbf{k}$ is smooth and for all $\mathbf{x} \in \mathbb{R}^n$

$$\underbrace{a(\mathbf{x}) + b(\mathbf{x})\lambda(a(\mathbf{x}), b(\mathbf{x}))}_{L_f h(\mathbf{x}) + L_g h(\mathbf{x}) \mathbf{k}(\mathbf{x}) + \alpha(h(\mathbf{x}))} \geq 0. \quad (15)$$

In the following section, we present our characterization of smooth safety filters.

III. TOWARDS SMOOTH SAFETY FILTERS

A Sontag-like safety filter can also be derived using Theorem 1. Here, we use the same function as in (9), taking the other root to the quadratic function:

$$\lambda_S(a, b) := \begin{cases} 0 & b = 0 \\ \frac{-a + \sqrt{a^2 + q(b)b}}{b} & b > 0. \end{cases} \quad (16)$$

One can verify λ_S is continuous and satisfies (6) and (7) on

$$\mathcal{S} = \{(a, b) \in \mathbb{R} \times \mathbb{R}_{\geq 0} : a > 0 \text{ or } b > 0\}, \quad (17)$$

implying λ_S is smooth on (17) by Theorem 1 and thus produces a smooth safety filter via (12) since $a + b\lambda_S(a, b) \geq 0$ for all $(a, b) \in \mathcal{S}$. Nevertheless, the resulting safety filter is overly conservative as illustrated by the following example.

Example 1: Consider a single integrator $\dot{\mathbf{x}} = \mathbf{u}$ equipped with the nominal controller $\mathbf{k}_d(\mathbf{x}) = -\mathbf{x}$, which stabilizes the system to the origin. Our task is to additionally avoid a circular obstacle of radius $r_o \in \mathbb{R}_{>0}$ located at $\mathbf{x}_o \in \mathbb{R}^2$, which can be accomplished by designing a safety filter for $\mathbf{k}_d$. This is achieved using $h(\mathbf{x}) = \|\mathbf{x} - \mathbf{x}_o\|^2 - r_o^2$ as a CBF to construct the QP safety filter in (11) and the Sontag safety filter defined by (12) where $\lambda = \lambda_S$ as in (16) with $q(b) = \sigma b$, $\sigma \in \mathbb{R}_{>0}$. The trajectory of the single integrator under each safety filter is provided in Fig. 2(a) for different σ . Decreasing σ decreases the conservatism of the controller, but it is still overly invasive even for arbitrarily small σ .

Motivated by the previous example, we set out to develop a general class of smooth safety filters that are less invasive by analyzing how the choice of F impacts the behavior of the resulting safety filter. Towards this development, we aim to answer the following: What properties should a function $(a, b, p) \mapsto F(a, b, p)$ satisfy so that one of its roots produces a smooth safety filter, and how may this function be selected such that the resulting safety filter is minimally invasive?




Fig. 1. Illustration of methodology for generating smooth safety filters.

A. Minimally Invasive Smooth Safety Filters

To answer the question posed in the previous subsection, let $(a, b, p) \mapsto F(a, b, p)$ be a smooth function and suppose there exists a continuous function $\lambda : \mathcal{S} \rightarrow \mathbb{R}$ satisfying the conditions of Theorem 1, which implies that λ is smooth on $\mathcal{S}$. We are also interested in ensuring that:

$$a + b\lambda(a, b) \geq 0, \quad (18)$$

for all $(a, b) \in \mathcal{S}$ so that λ may be used to construct a smooth safety filter satisfying (15). Note that F in (9) is constructed such that $\frac{\partial F}{\partial p} = 2(a + b\lambda_S(a, b)) > 0$ directly implies the satisfaction of (18). However, we demonstrate that this is not the only path towards certifying that one of F 's roots produces a smooth safety filter. In what follows, we introduce more general sufficient conditions on F so that one of its roots produces a smooth safety filter.

Our starting point is one of the direct consequences of Theorem 1 – the function λ must satisfy:

$$\frac{\partial \lambda}{\partial a}(a, b) = -\frac{\frac{\partial F}{\partial a}(a, b, \lambda(a, b))}{\frac{\partial F}{\partial p}(a, b, \lambda(a, b))}, \quad (19)$$

for all $(a, b) \in \mathcal{S}$, implying it is a solution to the ODE:

$$\frac{dp}{da} = -\frac{\frac{\partial F}{\partial a}(a, b, p)}{\frac{\partial F}{\partial p}(a, b, p)}, \quad (20)$$

from an appropriate initial condition.³ For a fixed $b = b^* \geq 0$, the trajectory $a \mapsto \lambda(a, b^*)$ defines a curve $\{(a, b, p) \in \mathcal{S} \times \mathbb{R} : p = \lambda(a, b), b = b^*\}$. As b is varied, this curve defines an entire surface $\{(a, b, p) \in \mathcal{S} \times \mathbb{R} : p = \lambda(a, b)\}$, thereby recovering $(a, b) \mapsto \lambda(a, b)$ as a function of both a and b (cf. Fig. 1). Although working with an ODE is generally more challenging than working with an algebraic equation, this new perspective allows us to reformulate the goal of satisfying

³In (20), one may think of p as the state, a as “time”, and b as a parameter.

inequality (18) as a set invariance problem. To do this, we define, for each $b \geq 0$, the set-valued map⁴

$$\mathcal{H}_b(a) := \{p \in \mathbb{R} : h_b(p, a) := a + bp \geq 0\}. \quad (21)$$

Ultimately, we would like to ensure that for each fixed $b = b^* \geq 0$, the solution to (19) satisfies $\lambda(a, b^*) \in \mathcal{H}_{b^*}(a)$ for all $(a, b^*) \in \mathcal{S}$ so that we can conclude that λ satisfies (18) for all $(a, b) \in \mathcal{S}$. The following proposition constitutes the first result of this letter, establishing conditions on F such that the flow of (20) satisfies inequality (18).

Proposition 1: Let $F : \mathbb{R}^3 \rightarrow \mathbb{R}$ be a smooth function that defines the dynamics in (20). Suppose that for all $b > 0$:

$$\left. \frac{\partial F}{\partial p} \right|_{p=-\frac{a}{b}} = b \left. \frac{\partial F}{\partial a} \right|_{p=-\frac{a}{b}}. \quad (22)$$

Then, for all $b > 0$, the set $\mathcal{H}_b$ is invariant for (20).

Proof: Note that the dynamics (20) may not be well-defined for a given initial condition $p_0(a_0) \in \mathcal{H}_b(a_0)$ since we may have $\frac{\partial F}{\partial p}(a_0, b, p_0) = 0$. Nevertheless, there exists no solution⁵ from such initial conditions, and the empty solution is contained in $\mathcal{H}_b$. For all other initial conditions in $\mathcal{H}_b(a_0)$, the dynamics are well-defined and thus generate a unique smooth solution $\lambda(\cdot, b) : I(p_0(a_0)) \rightarrow \mathbb{R}$ defined on a maximal interval of existence $I(p_0(a_0)) \subset \mathbb{R}$. We now show that when $b > 0$, $\mathcal{H}_b$ is invariant with respect to each of these solutions in the sense that $\lambda(a, b) \in \mathcal{H}_b(a)$ for all $a \in I(p_0(a_0))$. To this end, our proof leverages Nagumo's Theorem [7, Ch. 4]. We first verify that all points along the boundary of the set $\mathcal{H}_b(a)$ are regular points of h_b in (21), which follows directly from $\frac{\partial h_b}{\partial p}(p, a) = b \neq 0$. Then, $\mathcal{H}_b$ is invariant with respect to each $\lambda(\cdot, b)$ provided that:

$$\frac{dh_b}{da} = \frac{\partial h_b}{\partial a} + \frac{\partial h_b}{\partial p} \frac{dp}{da} = 1 - b \left[\frac{\partial F}{\partial p} \right]^{-1} \frac{\partial F}{\partial a} = 0,$$

whenever $h_b(a, p) = 0$, or, equivalently, when $p = -\frac{a}{b}$. It follows from (22) that the above holds so long as $\frac{\partial F}{\partial p} \neq 0$, which must be true for the solution to exist, and, by Nagumo's Theorem, implies the invariance of $\mathcal{H}_b$. ■

Proposition 1 provides a simple condition (22), for which one can use to help assess if an implicit function λ to a given smooth function F will satisfy inequality (18). The key idea is that with Proposition 1, if λ satisfies (18) for some $(a^*, b^*) \in \mathcal{S}$, it will satisfy (18) for all a in its interval of existence for the same $b = b^*$. The following lemma establishes this idea formally and additional conditions on λ so that it satisfies (18) for all $(a, b) \in \mathcal{S}$.

Lemma 1: Let $F : \mathbb{R}^3 \rightarrow \mathbb{R}$ be a smooth function satisfying (22) for all $b > 0$, and suppose there exists a continuous function $\lambda : \mathcal{S} \rightarrow \mathbb{R}$ satisfying (6) and (7) for all $(a, b) \in \mathcal{S}$ as in (17). Then, if $\lambda(0, b) > 0$ for all $b > 0$, λ is smooth on $\mathcal{S}$ and satisfies inequality (18) for all $(a, b) \in \mathcal{S}$.

Proof: The proof is divided into two cases: i) $b > 0$ and ii) $b = 0$. For each $b > 0$, consider an initial condition of (20), with $a_0 = 0$, satisfying $p_0^b = \lambda(0, b) > 0$. Since F and λ satisfy (6) and (7), the conditions of Theorem 1 hold, which implies λ is smooth on $\mathcal{S}$. Moreover, $\lambda(a, b)$ is the unique

solution to the ODE (20) by definition since, by Theorem 1, it must satisfy (8) for all $(a, b) \in \mathcal{S}$. Thus, the maximal interval of existence of this solution must be the domain of λ , which is equal to $\mathbb{R}$ when $b > 0$. Moreover, since $p_0^b \in \mathcal{H}_b(0)$ by definition and F satisfies (22), the conditions of Proposition 1 hold, which implies $\lambda(a, b) \in \mathcal{H}_b(a)$ for all $a \in \mathbb{R}$ for each $b > 0$. When $b = 0$ any value of $\lambda(a, b)$ satisfies $\lambda(a, b) \in \mathcal{H}_b(a)$ for all $a > 0$. Thus, $\lambda(a, b) \in \mathcal{H}_b(a)$ for all $(a, b) \in \mathcal{S}$, which implies λ satisfies (18), as desired. ■

Lemma 1 suggests that with condition (22) from Proposition 1, F only needs to be constructed so that λ is positive when $a = 0$ for it to generate a smooth safety filter. Note that the above result does not guarantee the existence of a continuous $\lambda : \mathcal{S} \rightarrow \mathbb{R}$ satisfying (6) and (7), but states that, if such a function exists, then it is smooth on $\mathcal{S}$ and satisfies inequality (19). We will provide examples of F satisfying these conditions shortly. First, we combine Proposition 1 and Lemma 1 to establish the main result of this letter, which formalizes the construction of smooth safety filters.

Theorem 2: Consider system (1) with a smooth nominal controller $\mathbf{k}_d : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and let $h : \mathbb{R}^n \rightarrow \mathbb{R}$ be a CBF for (1) on a set $\mathcal{C} \subset \mathbb{R}^n$ as in (10). Let $\lambda : \mathcal{S} \rightarrow \mathbb{R}$ satisfy the conditions of Lemma 1 for some smooth function $F : \mathbb{R}^3 \rightarrow \mathbb{R}$. Then, the controller:

$$\mathbf{k}_s(\mathbf{x}) = \mathbf{k}_d(\mathbf{x}) + \lambda(a(\mathbf{x}), b(\mathbf{x}))L_g h(\mathbf{x})^\top, \quad (23)$$

where $a : \mathbb{R}^n \rightarrow \mathbb{R}$ and $b : \mathbb{R}^n \rightarrow \mathbb{R}$ are as in (13), is a smooth safety filter for (1).

Proof: The proof follows directly from Lemma 1 since λ is smooth and satisfies (18), implying (23) satisfies (15). ■

Example 2: We use our results to construct the functions:

$$F_1(a, b, p) = bp^2 + ap - \frac{1}{4}q(b), \quad (24a)$$

$$F_2(a, b, p) = e^{\frac{p}{\sigma}} - e^{-\frac{a}{\sigma b}} - 1, \quad (24b)$$

where q is as in (9) and $\sigma \in \mathbb{R}_{>0}$. These functions satisfy (22) and the conditions of Lemma 1 where:

$$F_1(0, b, p) = bp^2 - \frac{1}{4}q(b) = 0 \implies \lambda(0, b) = \pm \frac{1}{2} \sqrt{\frac{q(b)}{b}},$$

$$F_2(0, b, p) = e^{\frac{p}{\sigma}} - 2 = 0 \implies \lambda(0, b) = \sigma \log(2).$$

Hence, for each F there exists a continuous λ satisfying $\lambda(0, b) > 0$ for all $b > 0$ (recall that $q(b) > 0$ for all $b > 0$). After verifying these conditions, we proceed to compute the implicit functions satisfying (6):

$$\lambda_1(a, b) = \frac{1}{2}\lambda_S(a, b) = \begin{cases} 0 & b = 0 \\ \frac{-a + \sqrt{a^2 + bq(b)}}{2b} & b > 0 \end{cases} \quad (25a)$$

$$\lambda_2(a, b) = \begin{cases} 0 & b = 0 \\ \sigma \log\left(1 + e^{-\frac{a}{\sigma b}}\right) & b > 0, \end{cases} \quad (25b)$$

which are continuous on $\mathcal{S}$, and therefore smooth on $\mathcal{S}$ as one can verify that $\frac{\partial F}{\partial p}(a, b, \lambda(a, b)) \neq 0$ for all $(a, b) \in \mathcal{S}$. These functions are plotted in Fig. 2(b) for a fixed b . As guaranteed by Lemma 1, each λ satisfies inequality (18). Moreover, when $q(b) = \sigma b$ both (25a) and (25b) approach λ_{QP} in (14) in the limit as $\sigma \rightarrow 0$. Indeed, when $b > 0$ both (25a) and (25b) are smooth approximations of the ReLU function, corresponding to the Squareplus approximation (yielding a ‘‘Half-Sontag’’ formula) and the Softplus approximation, respectively [20].

⁴One may think of (21) as a collection of ‘‘time-varying’’ safe sets (recall that a is our ‘‘time’’ variable) for the dynamics in (19).

⁵We consider solutions in the classical sense.




Fig. 2. (a) Comparison between the trajectory of a single integrator generated by the QP controller (blue) in (11) and Sontag safety filter (red) from (16) for Example 1. Each safety filter is implemented from the initial condition $x_0 = (-4, 3.9)$ with $\alpha = 2$. The values of σ are varied from 0.2 to 0.001. (b) Illustration of the sets $\mathcal{H}_b$ (dotted black line) and $\mathcal{H}_b^\epsilon$ (dashed black line) from (21) and (26), respectively, for a fixed b . (c-d) Trajectories of the single integrator under the influence of the smooth safety filters from (25a) (c) and (25b) (d). The trajectories are generated with σ varying from 1 to 0.01 for (c) and from 0.7 to 0.01 for (d). In each plot, more transparent curves correspond to smaller values of σ .

The functions in (25) are used to construct smooth safety filters via Theorem 2, and are applied to the scenario from Example 1, cf. Fig. 2(c-d).

B. Robust Smooth Safety Filters

The results in the previous subsection provide sufficient conditions under which the solution $(a, b) \mapsto \lambda(a, b)$ of (20) satisfies inequality (18) for all $(a, b) \in \mathcal{S}$. Such conditions require $\mathcal{H}_b$ to be invariant for (20), which precludes the consideration of safety filters that remain in a strict subset of $\mathcal{H}_b$. For example, even though Sontag's formula in (16) satisfies inequality (18) for all $(a, b) \in \mathcal{S}$, one can verify that the F producing this formula (9) does not satisfy (22). This is because the formula is contained within the set:

$$\mathcal{H}_b^\epsilon(a) := \{p \in \mathbb{R} : a + \frac{1}{\epsilon}bp \geq 0\}, \quad (26)$$

for $\epsilon = 2$, cf. Fig. 2(b). The above set satisfies $\mathcal{H}_b^\epsilon(a) \subseteq \mathcal{H}_b(a)$ for all $a \leq 0$ and $\epsilon \geq 1$. When $a > 0$, $\mathcal{H}_b^\epsilon(a) \not\subseteq \mathcal{H}_b(a)$ for any $\epsilon \geq 1$. In this situation, the fact that Sontag's formula satisfies inequality (18) relies on the fact that it is always positive, $\lambda(a, b) \geq 0$. Motivated by this observation, in this subsection we study the invariance of $\mathcal{H}_b^\epsilon \cap \mathbb{R}_{\geq 0}$, noting that:

$$\epsilon \geq 1 \implies \mathcal{H}_b^\epsilon(a) \cap \mathbb{R}_{\geq 0} \subset \mathcal{H}_b(a),$$

for all $(a, b) \in \mathcal{S}$. Taking this intersection imposes the additional requirement that $\lambda(a, b) \geq 0$ for all $(a, b) \in \mathcal{S}$, which is not restrictive since negative values of λ imply the resulting controller is pushing in the wrong direction (i.e., toward the constraint boundary) and is attempting to violate (18). As illustrated in Fig. 1 and Fig. 2(b), increasing ϵ lifts the boundary of $\mathcal{H}_b^\epsilon$ off that of $\mathcal{H}_b$ leading to a more restricted set of values that $\lambda(a, b)$ may achieve. Although increasing ϵ imposes a more conservative condition on λ , it adds an additional robustness margin to the resulting safety filter, which, as demonstrated in Section IV, may be useful in practice. The following result establishes conditions on F so that the flow of (19) satisfies the tightened condition in (26).

Proposition 2: Let $F : \mathbb{R}^3 \rightarrow \mathbb{R}$ be a smooth function that defines the dynamics in (20). Suppose that for all $b > 0$:

$$\epsilon \frac{\partial F}{\partial p} \Big|_{p=-\frac{\epsilon a}{b}} = b \frac{\partial F}{\partial a} \Big|_{p=-\frac{\epsilon a}{b}}, \quad \frac{\partial F}{\partial a} \Big|_{p=0} = 0. \quad (27)$$

Then, for each $b > 0$, the set $\mathcal{H}_b^\epsilon \cap \mathbb{R}_{\geq 0}$ is invariant for (20).

Proof: Showing the invariance of $\mathcal{H}_b^\epsilon$ follows the same steps as that of Proposition 1 by replacing $\mathcal{H}_b$ from (21) with

$\mathcal{H}_b^\epsilon(p, a) := a + \frac{1}{\epsilon}bp$, which defines $\mathcal{H}_b^\epsilon(a)$ as its zero superlevel set. To show that $\mathbb{R}_{\geq 0} = \{p \in \mathbb{R} : p \geq 0\}$ is invariant, we define $h_p(p) := p$, which defines $\mathbb{R}_{\geq 0}$ as its zero superlevel set and satisfies $\frac{\partial h_p}{\partial p} \neq 0$, implying zero is a regular value of h_p . Hence, $\mathbb{R}_{\geq 0}$ is invariant provided that:

$$\frac{dh_p}{da} = \frac{dp}{da} = -\frac{\frac{\partial F}{\partial a}(a, b, p)}{\frac{\partial F}{\partial p}(a, b, p)},$$

evaluates to zero when $p = 0$, which follows from (27) provided $\frac{\partial F}{\partial p} \neq 0$ when $p = 0$. Using a similar argument to that in the proof of Proposition 1, we may exclude points satisfying $\frac{\partial F}{\partial p} = 0$ in the above, since such points cannot lie along any trajectory produced by the dynamics in (20). Thus, since both $\mathcal{H}_b^\epsilon$ and $\mathbb{R}_{\geq 0}$ are invariant for (20) and the intersection of invariant sets is invariant [7, Prop. 4.13], it follows that $\mathcal{H}_b^\epsilon \cap \mathbb{R}_{\geq 0}$ is invariant for (20), as desired. ■

Similar to Proposition 1, the above result provides a simple condition, (27), that one may use to help determine if an implicit function λ satisfying (6) and (7) will satisfy inequality (18). Note that results similar to Lemma 1 and Theorem 2 can be stated for Proposition 2, the formal statements of which we omit here in the interest of space. As noted earlier, Sontag's F in (9) satisfies the conditions of Proposition 2, with $\epsilon = 2$. The following example introduces a safety filter that satisfies such conditions for any $\epsilon \geq 1$.

Example 3: The smooth safety filters from Example 2 can approximate the QP-based safety filter (12) arbitrarily closely. Generally, this is desirable; however, in certain situations, such as when handling uncertainty, one may wish to modulate how conservative the resulting safety filter is by tuning the value of ϵ (cf. Fig. 1). For this, we introduce:

$$F(a, b, p) = bp^2 + \epsilon ap - \frac{\epsilon^2}{4}q(b), \quad (28)$$

which generalizes both (9) and (24a) so that F satisfies Proposition 2 for any $\epsilon \geq 1$ and produces a robust version of Sontag's formula $\lambda(a, b) = \frac{\epsilon}{2}\lambda_S(a, b)$ for any $\epsilon \geq 1$.

Remark 1: Our approach allows for characterizing safety filters via ϵ . Although various choices of F may contain tuning parameters, the behavior of such a safety filter is limited by the value of ϵ . When $\epsilon = 1$, $\lambda(a, b)$ may approach the boundary $a + b\lambda(a, b) = 0$ of constraint (18), whereas when $\epsilon > 1$, $\lambda(a, b)$ may only approach the boundary of the tightened constraint $a + \frac{1}{\epsilon}b\lambda(a, b) \geq 0$, resulting in a more conservative but also more robust safety filter (cf. Fig. 2(b)).




Fig. 3. In each plot, the green/purple curves correspond to the trajectory obtained by tracking the smooth safe velocity generated by the robust Sontag safety filter for different ε with $q(b) = 0.1b$. Visualizations of trajectories from (c) are provided in (d) and (e), where the reduced-order trajectory corresponds to that of a single integrator $\dot{x} = u$ with $u = k_0(x)$ that the full-order model seeks to track.

IV. NUMERICAL EXAMPLES

This section illustrates the practical benefits of smooth safety filters by applying them to the model-free safety-critical control paradigm introduced in [10]. Here, we design a safety filter for a reduced-order model, the safe trajectory of which is tracked by the full-order dynamics in a model-free fashion. We consider the same setting as in [10, Ex. 2], which involves designing a controller for a planar Segway with configuration $(x, \varphi) \in \mathbb{R}^2$, where x is the position and φ the pitch angle, with the objective of driving forward at a desired velocity $\dot{x}_d$ and stopping before colliding with a wall located at $x_{\max}$. This leads to the safety constraint $h(x) = x_{\max} - x$, which is used as a CBF for a one-dimensional single integrator to construct a safety filter $k_0 : \mathbb{R} \rightarrow \mathbb{R}$ that produces a safe velocity for the Segway. On the full-order dynamics, this velocity is tracked by the PD controller:

$$\mathbf{k}(\mathbf{x}) = K_p(x - k_0(x)) + K_\varphi\varphi + K_{\dot{\varphi}}\dot{\varphi},$$

that also attempts to keep the Segway upright, where $\mathbf{x} = (x, \varphi, \dot{x}, \dot{\varphi}) \in \mathbb{R}^4$ is the state and $K_p, K_\varphi, K_{\dot{\varphi}} \in \mathbb{R}_{>0}$ are gains. Implementation of this controller does not require knowledge of the full-order dynamics, which may be uncertain or difficult to compute, and allows for enforcing *input-to-state* safety [13] of the closed-loop system [10, Prop. 1].

We now compare the response of the Segway when the safe velocity is generated nonsmoothly by a QP-based safety filter, used in [10], and smoothly by the robust Sontag safety filter from (28) for different values of ε . We begin by using the same parameters for the controller as in [10, Ex. 2], the results of which are shown in Fig. 3(a). Here, both controllers safely track the reference velocity, and the response of the smooth controller approaches that of the QP controller as $\varepsilon \rightarrow 1$. Although this approach does not directly rely on model knowledge, it relies on tuning the gains of the tracking controller to achieve safety. In general, safety can be achieved by increasing the proportional gain K_p to track the reference velocity more aggressively. When increasing K_p too much, however, we observe that the controller attempting to track a non-differentiable reference signal causes instabilities and safety violation.⁶ In contrast, the same controller attempting

to track a smooth reference velocity is more oscillatory, but maintains safety. On the other extreme, taking K_p too low (see Fig. 3(c)) results in safety violation for the QP controller and smooth controller with $\varepsilon = 1$, whereas controllers with $\varepsilon > 1$ maintain safety due to their increased robustness.

V. CONCLUSION

We presented a general characterization of smooth safety filters based on the Implicit Function Theorem, leading to smooth universal formulas for safety-critical control that enable quantifying the conservatism of the resulting safety filter. The practical benefits of such smooth safety filters were showcased through their application to safety-critical control with reduced-order models [10]. Future efforts will focus on extending our approach to multiple safety constraints and showcasing the benefits of smooth safety filters on hardware.

REFERENCES

- [1] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. Autom. Control*, vol. 62, no. 8, pp. 3861–3876, Aug. 2017.
- [2] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proc. Eur. Control Conf.*, 2019, pp. 3420–3431.
- [3] W. Xiao, C. G. Cassandras, and C. Belta, *Safe Autonomy with Control Barrier Functions: Theory and Applications*. Cham, Switzerland: Springer Nat., 2023.
- [4] J. Breeden and D. Panagou, "Robust control barrier functions under high relative degree and input constraints for satellite trajectories," *Automatica*, vol. 155, Sep. 2023, Art. no. 111109.
- [5] M. Jankovic, "Robust control barrier functions for constrained stabilization of nonlinear systems," *Automatica*, vol. 96, pp. 359–367, Oct. 2018.
- [6] M. Alyaseen, N. Atanasov, and J. Cortes, "Continuity and boundedness of minimum-norm CBF-safe controllers," 2023, *arXiv:2306.07398*.
- [7] F. Blanchini and S. Miani, *Set-Theoretic Methods in Control*. Cham, Switzerland: Springer, 2008.
- [8] A. J. Taylor, P. Ong, T. G. Molnar, and A. D. Ames, "Safe backstepping with control barrier functions," in *Proc. Conf. Decis. Control*, 2022, pp. 5775–5782.
- [9] M. Krstić, I. Kanellakopoulos, and P. Kokotović, *Nonlinear and Adaptive Control Design*. Hoboken, NJ, USA: Wiley, 1995.
- [10] T. G. Molnar, R. K. Cosner, A. W. Singletary, W. Ubellacker, and A. D. Ames, "Model-free safety-critical control for robotic systems," *IEEE Robot. Autom. Letters*, vol. 7, no. 2, pp. 944–951, Apr. 2022.
- [11] T. G. Molnar and A. D. Ames, "Safety-critical control with bounded inputs via reduced order models," in *Proc. Am. Control Conf.*, 2023, pp. 1414–1421.
- [12] E. D. Sontag, "A 'universal' construction of Artstein's theorem on nonlinear stabilization," *Syst. Control Lett.*, vol. 13, no. 2, pp. 117–123, Aug. 1989.
- [13] S. Kolathaya and A. D. Ames, "Input-to-state safety with control barrier functions," *IEEE Control Syst. Lett.*, vol. 3, no. 1, pp. 108–113, Jan. 2019.
- [14] M. Krstić, "Inverse optimal safety filters," *IEEE Trans. Autom. Control*, early access, May 22, 2023, doi: [10.1109/TAC.2023.3278788](https://doi.org/10.1109/TAC.2023.3278788).
- [15] M. Li and Z. Sun, "A graphical interpretation and universal formula for safe stabilization," in *Proc. Am. Control Conf.*, 2023, pp. 3012–3017.
- [16] P. Mestres and J. Cortes, "Feasibility and regularity analysis of safe stabilizing controllers under uncertainty," 2023, *arXiv:2301.04603*.
- [17] P. Ong and J. Cortes, "Universal formula for smooth safe stabilization," in *Proc. Conf. Decis. Control*, 2019, pp. 2373–2378.
- [18] M. Romdlony and B. Jayawardhana, "Stabilization with guaranteed safety using control Lyapunov–Barrier function," *Automatica*, vol. 66, pp. 39–47, Apr. 2016.
- [19] L. H. Loomis and S. Sternberg, *Advanced Calculus*. Boston, MA, USA: Addison-Wesley, 1968.
- [20] J. T. Barron, "Squareplus: A softplus-like algebraic rectifier," 2021, *arXiv:2112.11687*.

⁶Note that the results in [10] rely on differentiability of k_0 .