2019 American Control Conference (ACC)
Philadelphia, PA, USA, July 10-12, 2019

Realizable Set Invariance Conditions for Cyber-Physical Systems

Thomas Gurriet, Petter Nilsson, Andrew Singletary, and Aaron D. Ames

***Abstract*—There is currently a gap between control-theoretical results and the reality of robotic implementations—this makes it difficult to transfer analytical guarantees to practice. This problem is especially troubling when it comes to safety guarantees for safety-critical systems. In this paper we seek to help bridge this gap. We first make a clear theoretical distinction between a system and a model, and outline how the two need to be related for guarantees to transfer from the latter to the former. We then introduce various imperfections into the model, including uncertainty in actuation and sensing, as well as time discretization effects from digital control implementations. These assumptions lead to new criteria for controlled invariance to be realizable. We investigate these criteria and propose a digital control implementation for enforcing safety in the presence of uncertainty. Our ideas are illustrated with a numerical example where a ground robot satisfies safety constraints in the presence of perception noise.**

I. Introduction

Safety and reliability of cyber-physical systems are two of the biggest obstacles on the road towards ubiquitous robotic systems. It is widely accepted that a solution to these problems can only emerge from rigorous mathematics, methods and processes [1], [2], [3]. However, while there is a vast body of work on safety and reliability in control theory, very little of it is actually used in practice where safety margins are typically empiric and/or heuristic. Mathematically grounded approaches toward safety have several potential major benefits. In particular, systems can become adaptive and utilize large safety margins only when motivated both from a safety and from an uncertainty perspective, which improves both performance and safety compared to heuristic approaches. Furthermore, the set of safe operating conditions is made explicit, which is beneficial when a system is moved to a new environment or connected to other systems.

We believe that the gap between theory and practice is a direct echo of the discrepancy between the reality of cyber-physical systems (CPSs) and their representative models used in control theory. For the result of a theorem to apply, all the hypotheses must be satisfied. Therefore, it should not come as a surprise that system implementations based on theorems whose assumptions are impossible to satisfy in practice often turn out to be unreliable. This discrepancy needs to be addressed from both sides: with better hardware that attains performance close to idealistic assumptions, but also with mathematics capable of accounting for model imperfections. In this paper we investigate the theory side of the problem in the context of safety.

At a fundamental level, safety can be reduced to constraining system behavior. From a control perspective, safety is therefore usually associated with the concept of set invariance [4]—given a desired safety set, ensuring safe system operation is equivalent to making sure that it remains inside a safety set at all times. This is usually done in two stages. First, parts of the safety set from where future safety violations are unavoidable must be excluded, which is done by computing a viable (or controlled-invariant) subset of the safety set [5], [6]. Many approaches have been proposed to find such a subset, from discretized solutions of Hamilton-Jacobi equations [7], SOS optimization [8], sampling [9], and many others [10]. Secondly, a control law that renders the viable subset forward invariant needs to be implemented. In this second step, two different approaches are usually considered. In [11], [12], a control structure is proposed that switches between a nominal controller designed for performance and a "safe" controller designed for set-invariance. Whereas in [4], continuous optimization-based filtering of the control input is performed so as to enforce set-invariance in a minimally invasive way. At a fundamental level, both strategies reduce to enforcement of Nagumo's sub-tangentiality condition [13] with varying degrees of conservatism.

Despite solid theoretical underpinnings, the practical effectiveness of this body of work is hindered by the use of idealistic system representations and assumptions. Inclusion of more realistic representations in the context of set invariance is still fairly recent and not yet fully developed. Additive uncertainty has been considered [14], [15], and also generic parametric uncertainty [16], including for hybrid systems [17], whereas [18], [19] treat sampled data systems. But even with these "robust" approaches, systematic and unfailing set invariance has yet to be demonstrated on actual hardware, although there have been multiple cases recently of experimental success through reasonable approximations of the formal approaches to safety [16], [20].

In order to achieve the goal of verifying theory on hardware, we argue that these approaches must be revisited on the basis of a rigorous and quantifiable description of the differences between systems, and models of systems. With this approach, a formal and quantifiable link can be made between mathematical theorems and the requirements in terms of modeling and hardware implementation to achieve the results predicted by these theorems. Therefore, the main contribution of this paper is the derivation of realizable set-invariance conditions based on realistic assumptions about cyber-physical systems. In particular, we introduce *robust viability under state uncertainty*—a property that a set must satisfy in order for it to be possible to render it invariant

The authors are with the Department of Mechanical and Civil Engineering, California Institute of Technology, Pasadena, CA, 91125 USA. Emails: {tgurriet,pettni,asinglet,ames}@caltech.edu

978-1-5386-7926-5/$31.00 ©2019 AACC

under realistic assumptions. We give sufficient conditions for this property to hold, and suggest a digital control scheme that enforces invariance of such a set even in the presence of model imperfections.

The rest of the paper is laid out as follows. In Section II, a formal distinction is made between a system and a model of that system, and we introduce an approximation hierarchy of models. In Section III we first revisit Nagumo's theorem, before extending it by relaxing overly idealistic assumptions. We study robust viability under state uncertainty in Section IV, and propose an invariance-enforcing digital control scheme in Section V. A numerical example is given in Section VI before the paper is concluded in Section VII.

Notation: For a mapping $f : X \to Y$ and a subset $A \subseteq X$, we write $f(A) = \{f(x) : x \in A\}$. We write $\|\cdot\|$ to denote the standard Euclidean norm; the Euclidean ball with center x and radius r is denoted $\mathcal{B}(x, r) = \{y : \|y - x\| \leq r\}$. The diameter of a set X is defined as $\|X\| = \sup_{x,y \in X} \|x - y\|$, with this notation $\|\mathcal{B}(x, r)\| = 2r$. We write $\oplus$ and $\ominus$ for Minkowski addition and subtraction.

II. System and Model Terminology

George Box famously wrote that "all models are wrong, but some are useful" [21], which can be taken as an argument both for and against model-based control. Based on this statement alone, it seems unlikely that any sort of "guarantee" is possible to achieve in the real world, since it is impossible to foresee exactly how inputs affect the system at hand. However, guarantees can in theory be achieved even with a model that is incorrect—what is required is that the model captures (or *over-approximates*) the true behavior, which is a much less strict requirement than being exact.

While systems and system models are treated interchangeably in much of the existing literature, at the core of this paper is the clear distinction between these two fundamentally different entities.

Terminology 1. A **system** Σ is a collection of measurable physical quantities, also called "states". We refer to the number of states n as the **state dimension** of the system. The **input** to a system is a collection of physical quantities whose values affect the state of the system, and which can be (approximately) controlled via actuators to desired values. We refer to the number of inputs m as the **input dimension** of the system.

We write $(\mathbf{x}(t), \mathbf{u}(t)) \in \mathbb{R}^n \times \mathbb{R}^m$ to denote **trajectories** of the system, i.e. the measurements $\mathbf{x}(t)$ resulting from setting system inputs to $\mathbf{u}(t)$ at time t. In the following we assume that the system is such that $\mathbf{x}(t)$ is differentiable and defined for all times.

The model of a system is the mathematical representation of the evolution of the state of the system. Modeling can be done in many ways; motivated by highly dynamical robotic systems we focus on continuous-time models.

Terminology 2. A **system model** $\mathcal{M} = (f, X, U, P)$ consists of a Lipschitz continuous map $f : X \times U \times P \to \mathbb{R}^n$, where $X \subseteq \mathbb{R}^n$ is the domain of the model, $U \subseteq \mathbb{R}^m$ is a set of admissible inputs, and P is a parameter set.

Since we can not hope to obtain an exact mathematical representation of a physical system, we resort to the idea of finding models that *over-approximate* it. A model over-approximates a physical system if all possible state/input trajectories of the system are consistent with the model. That is, every possible behavior of the system must be contained in the model, but the model may contain additional behaviors that are not present in the true system. Constructing an over-approximating model for a real system is of course a very challenging task that involves a trade-off between the confidence in the model indeed being an over-approximation, and the magnitude of the "approximation gap".

Definition 1. A model $\mathcal{M} = (f, X, U, P)$ is an **over-approximation** of a physical system Σ, written $\mathcal{M} \succeq \Sigma$, if for every trajectory $(\mathbf{x}(t), \mathbf{u}(t))$ there exists a continuous parameter trajectory $\mathbf{p} : \mathbb{R} \to P$ such that

$$\frac{\mathrm{d}\mathbf{x}(t)}{\mathrm{d}t} = f(\mathbf{x}(t), \mathbf{u}(t), \mathbf{p}(t)). \tag{1}$$

The parameter signal $\mathbf{p}$ should be thought of as a way to model uncertainty and unknowns in the system and its environment, as well as any unmodeled dynamics.

It is straightforward to extend the concept of over-approximations also to pairs of system models: we say that a model $\mathcal{M}_1 = (f_1, X_1, U_1, P_1)$ over-approximates a model $\mathcal{M}_2 = (f_2, X_2, U_2, P_2)$, written $\mathcal{M}_1 \succeq \mathcal{M}_2$, if there exist (projection) maps Γ^X, Γ^U and Γ^P such that $X_1 \supseteq \Gamma^X(X_2)$, $U_1 \supseteq \Gamma^U(U_2)$, $P_1 \supseteq \Gamma^P(P_2)$, and for all $x \in X_2$ and $u \in U_2$,

$$f_1(\Gamma^X(x), \Gamma^U(u), \Gamma^P(P_2)) \supseteq T_x\Gamma^X(f_2(x, u, P_2)), \tag{2}$$

where $T_x\Gamma^X$ is the tangent map of Γ^X at x. Combined, these conditions imply that the dynamics of $\mathcal{M}_2$ via the projection maps can be embedded in the space of $\mathcal{M}_1$ in a way so that all behaviors of $\mathcal{M}_2$ are captured also by $\mathcal{M}_1$.

This partial order between a system and different models allows for exploration of the trade-off between performance and confidence. Consider two models $\mathcal{M}_1$ and $\mathcal{M}_2$ such that $\mathcal{M}_1 \succeq \mathcal{M}_2 \succeq \Sigma$. Ensuring safety of Σ with respect to a tight over-approximation $\mathcal{M}_2$ implies a comparatively small performance loss compared to a hypothetical "optimal" safe controller, while a more conservative over-approximation $\mathcal{M}_1$ implies a larger margin and more cautious behavior.

III. Invariance Conditions and Realizability

The objective of this work is to constrain a system to stay inside a closed safety set S under realistic modeling and implementation assumptions. We achieve this by enforcing forward invariance on an over-approximating model, which implies forward invariance of the system itself.

Definition 2. A set S is **forward controlled invariant** for a system Σ if the input $\mathbf{u}(t)$ can be chosen such that $\mathbf{x}(0) \in S \implies \forall t \geq 0,\ \mathbf{x}(t) \in S$.

A. Nagumo's Set-Invariance Condition

The core tool behind set invariance is the famous Nagumo's theorem, which relates invariance to derivatives taking values in the *tangent cone* $\mathcal{T}_S(x)$[1]. Here we state it in a slightly extended formulation that makes a clear distinction between a system and a system model.

Theorem 1. *Consider a system Σ, a model $\mathcal{M} = (f, X, U, P)$ such that $M \succeq \Sigma$, and a closed set $S \subseteq X$. If for every trajectory with $\mathbf{x}(0) \in S$ the input $\mathbf{u}(t)$ can be chosen Lipschitz continuous and such that*

$$f(\mathbf{x}(t), \mathbf{u}(t), P) \subseteq \mathcal{T}_S(\mathbf{x}(t)),\ \forall t \geq 0, \tag{3}$$

then the set S is forward controlled invariant for Σ.

Proof. For any system trajectory $(\mathbf{x}, \mathbf{u})$ we can find an associated continuous parameter trajectory $\mathbf{p}$ such that (1) holds. Due to continuity of f and $\mathbf{p}$, the solution of (1) for this choice of $\mathbf{p}$ is unique. Furthermore, (3) implies that $\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{x}(t) \in \mathcal{T}_S(\mathbf{x}(t))$ for all $t \geq 0$. Therefore the classical version of Nagumo's theorem [6], [13] applies, and S is forward invariant for $\mathcal{M}$ and hence also for Σ. □

With the objective of enforcing condition (3) through means of feedback control, it is fundamental to look at the realizability of this condition. In theory, the only requirement for satisfying the condition is that the set is *robustly viable*.

Definition 3. A set S is **robustly viable** for a model $\mathcal{M} = (f, X, U, P)$ if for all $x \in S$ there exists a $u \in U$ such that

$$f(x, u, P) \subseteq \mathcal{T}_S(x). \tag{4}$$

However, there are always imperfections that prevent enforcement of (3) even if (4) holds in theory. We therefore proceed by adding additional realistic assumptions that model uncertainty. We have already accounted for uncertainty in the dynamics via the unknown parameter signal $\mathbf{p}(t) \in P$. Here we introduce additional uncertainty and imperfections stemming from estimation, actuation, and digital control.

B. Actuation Uncertainty

There are many reasons as to why a desired input might not be perfectly achievable on the real system. For instance, in many cases the computed input is passed to an electronic control unit (ECU) that acts as a closed-loop controller around the property of interest (e.g. servo angle, torque, or speed). In that case, transients in the ECU control loop might cause the input to deviate from the desired value. In low-powered embedded systems there may also be quantization effects that result in small but significant discrepancies between the desired and actual inputs.

Assumption 1 (**Actuation**)**.** For a desired input $\bar{\mathbf{u}}(t)$ to the system, the physical input $\mathbf{u}(t)$ is such that:

$$\mathbf{u}(t) \in \{\bar{\mathbf{u}}(t)\} \oplus \Delta_u, \quad \Delta_u \subseteq \mathbb{R}^m. \tag{5}$$

This assumption reflects noise and unmodeled dynamics between the controller output $\bar{\mathbf{u}}(t)$ and the actual value of the system input $\mathbf{u}(t)$. The following result shows that input uncertainty can be incorporated into the model as an additional parameter variation.

Proposition 1. *A model $\mathcal{M} = (f, X, U, P)$ with input uncertainty Δ_u is equivalent to $\bar{\mathcal{M}} = (\bar{f}, X, U, \bar{P})$ with*

$$\bar{P} = P \times \Delta_u, \quad \bar{f}(x, u, (p, \delta_u)) = f(x, u + \delta_u, p). \tag{6}$$

Thus, actuation uncertainty can be accounted for by enforcing the robust sub-tangentiality condition (3) on an over-approximating model $\bar{\mathcal{M}} \succeq \mathcal{M}$, which is possible if the safety set is robustly viable with respect to $\bar{\mathcal{M}}$.

C. Sensing Uncertainty

Condition (3) assumes knowledge of the current state x, which is typically estimated via filters and hence not exactly known. To obtain overall guarantees on system safety with inaccurate state estimates, a measure of the degree of inaccuracy must be known so that the safety controller can account for the worst-case scenario.

Assumption 2 (**State Estimation**)**.** If the state of the system is $\mathbf{x}(t) \in S$, we have access to an estimate $\bar{\mathbf{x}}(t)$ of that state such that:

$$\mathbf{x}(t) \in \{\bar{\mathbf{x}}(t)\} \oplus \Delta_x, \quad \Delta_x \subseteq \mathbb{R}^n. \tag{7}$$

This assumption relates directly to the accuracy of the sensors, as well as to the accuracy of the state estimation method. We have opted for a robust formulation of state uncertainty assumptions for simplicity and consistency with other assumptions. For linear systems it is known how this type of robust state-valued observers can be synthesized [23], but optimal observers can be arbitrarily complex whereby more practical alternatives have been proposed [24]. However, most state estimation filters operate under probabilistic assumptions. In this setting, the resulting probability distribution can be converted into a chance constraint set Δ_x with $\mathbb{P}\left[\mathbf{x}(t) - \bar{\mathbf{x}}(t) \in \Delta_x\right] \geq 1 - \delta$, and Δ_x can be taken as the state estimate. This results in safety guarantees that hold with high probability.

As opposed to input uncertainty, state uncertainty can not be directly incorporated into the model since also the right-hand side in (3) depends on the state. Consequently, we need to strengthen the sub-tangentiality condition as follows.

Proposition 2. *Let $\mathcal{M} = (f, X, U, P)$ be a model that over-approximates a system Σ and assume that $\bar{\mathbf{x}}(t)$ satisfies Assumption 2. If for every trajectory with $\mathbf{x}(0) \in S$ the input $\mathbf{u}(t)$ can be chosen such that*

$$\forall x \in S \cap \{\bar{\mathbf{x}}(t)\} \oplus \Delta_x,\ f(x, \mathbf{u}(t), P) \subseteq \mathcal{T}_S(x), \tag{8}$$

then S is forward controlled invariant for Σ.

Proof. Assume for contradiction that (8) holds and that there exists a trajectory $(\mathbf{x}, \mathbf{u})$ with $\mathbf{x}(0) \in S$ for which S is not forward invariant. By Theorem 1 there exists a time τ such

[1] The tangent cone $\mathcal{T}_S(x)$ of S at x can be defined as the set of all directions v such that for every sequence $\{x_i\} \to x$ in S and every monotone sequence $\{t_i\} \to 0$, there exists a sequence $\{v_i\} \to v$ such that $x_i + t_i v_i \in S$ for all i [22, Prop. 5.2].

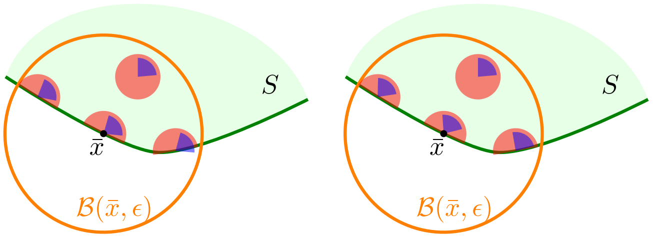
S
$\bar{x}$
$\mathcal{B}(\bar{x},\epsilon)$
S
$\bar{x}$
$\mathcal{B}(\bar{x},\epsilon)$

Fig. 1. Graphical illustration of the ϵ-robust viability condition. Tangent cones for the set S are shown in red, and cones $\bar{f}(x) = f(x, \bar{u}, P)$ in blue. To the left, $\bar{f}(\bar{x}) \subseteq \mathcal{T}_S(\bar{x})$, but this inclusion property does not hold in the whole set $S \cap \mathcal{B}(\bar{x}, \epsilon)$ which presents a problem when state uncertainty is present. To the right however, the same input $\bar{u}$ satisfies the sub-tangentiality condition in all of $S \cap \mathcal{B}(\bar{x}, \epsilon)$ so it can safely be applied even in the presence of state uncertainty.

that $\mathbf{x}(\tau) \notin \mathcal{T}_S(\mathbf{x}(\tau))$. This is however a contradiction of (8) and Assumption 2. □

Essentially, the result states that in order to enforce set invariance under imperfect sensing, it is sufficient to enforce condition (3) for all possible values of the state around the nominal sensed value. However, this stricter condition also implies more requirements on the safe set S; robust viability is no longer enough. We therefore extend the robust viability concept to allow for uncertainty in the state.

Definition 4. A set S is **robustly viable under state ϵ-uncertainty** (abbreviated **ϵ-robustly viable**) for a model $\mathcal{M} = (f, X, U, P)$ if for all $\bar{x} \in S$ there exists a $u \in U$ such that

$$\forall x \in S \cap \mathcal{B}(\bar{x}, \epsilon), \; f(x, u, P) \subseteq \mathcal{T}_S(x). \tag{9}$$

This condition is stronger than that in Definition 3 since it requires existence of inputs that guarantee the sub-tangentiality condition in a *set*, rather than at a single point as in (4), as illustrated in Fig. 1.

D. Digital Control

Proposition 2 allows a set invariance condition to be enforced even if the state is not exactly known. However, one would still need access to a continuous estimation of the states, and the capability to continuously modify the inputs. This is potentially possible through means of analog electronics; however, modern systems overwhelmingly rely on digital electronics. In this case, estimates of the states are available only at discrete time instances, and the input is modified at discrete points in time. Furthermore, the time between receiving a state estimate and setting a computed control command is non-negligible. Therefore, it is important to derive a set invariance condition that addresses this reality. For that, we consider a typical timing structure of a digital control system as depicted in Fig. 2.

Assumption 3 (Hard Real-Time Implementation). The digital control system operates at a fixed loop frequency $1/\Delta_t$. Each control loop starts at time t_k, $k \in \mathbb{N}$ with an estimate $\bar{x}_k = \bar{\mathbf{x}}(t_k)$ of the state as in Assumption 2. This estimate is then used to calculate the subsequent control action $\bar{u}_k$ that is activated at time $t_k + \alpha\Delta_t$ for $\alpha \in [0, 1]$.

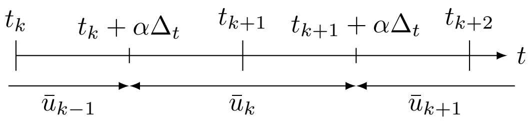

Fig. 2. Time visualization of a real-time control loop. A sensor reading is performed at times t_k which is the basis for computation of an input signal that changes at times $t_k + \alpha\Delta_t$.

Here the continuous time assumption is replaced by a more realistic correctness assumption on the timing of the control loop. We refer the reader to the real-time computing literature for more details on the complex issue of guaranteeing a bounded computing time [25]. Incorporating the digital control module into the analysis makes the overall system cyber-physical (i.e. hybrid). Rather than incorporating these discrete components into an overall model, we set out to design the digital controller in a way that guarantees safety with respect to the continuous-time model.

Having outlined Assumptions 1–3 that we argue accurately capture realistic implementations, we define the two problems that must be solved to enforce invariance under these assumptions: finding sets that are robustly viable under state uncertainty, and enforcing uncertain cone inclusion constraints through the means of digital real-time control.

Problem 1. Given a model $\mathcal{M} = (f, X, U, P)$, a safety set S, and a maximal magnitude of state uncertainty ϵ, find a set $\bar{S} \subseteq S$ that is ϵ-robustly viable.

Problem 2. Given a model $\mathcal{M} \succeq \Sigma$ and a set $\bar{S}$ that is ϵ-robustly viable, construct a digital controller such that if Assumptions 1–3 are satisfied, then $\bar{S}$ is invariant for Σ.

We contribute towards a solution of the first problem in Section IV by giving different sufficient conditions for robust viability under state uncertainty. We then propose a solution to Problem 2 in Section V. Both these problems are however challenging and we do not possess complete solutions at this point in time. We elaborate further in the conclusion (Section VII) on future work that we plan to undertake.

IV. On Robust Viability Under State Uncertainty

The condition in Definition 4 is of type "for a neighborhood of every point", and is therefore cumbersome to verify. Here we give sufficient conditions that imply ϵ-state-robust viability: one tightened formulation of the standard robust viability property that opens the door for known viability algorithms to be adapted to the case with state uncertainty, and one condition that hinges on the existence of a function certificate.

A. Cone Algebra

The results in this section are based on studying expansions and contractions of cones. Let $K \subset \mathbb{R}^n$ be a cone, i.e. a set that contains 0 and that is closed under scaling. For a positive

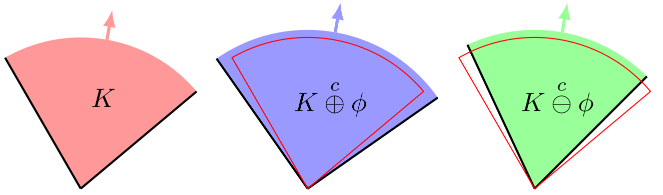

Fig. 3. Illustration of expansion (blue) and contraction (green) operations on a cone K (red).

angle $\phi > 0$ we define the inflated and deflated cones $K \overset{c}{\oplus} \phi$ and $K \overset{c}{\ominus} \phi$ as follows:

$$K \overset{c}{\oplus} \phi = \left\{ tv : t \in \mathbb{R}_+, \exists k \in K \text{ s.t. } \frac{k \cdot v}{\|k\|\|v\|} \geq \cos(\phi) \right\},$$
$$K \overset{c}{\ominus} \phi = \left\{ tv : t \in \mathbb{R}_+, \frac{k \cdot v}{\|k\|\|v\|} \geq \cos(\phi) \implies k \in K \right\}.$$

The operations are illustrated in Fig. 3. It can be verified that $(K \overset{c}{\oplus} \phi_1) \overset{c}{\oplus} \phi_2 = K \overset{c}{\oplus} (\phi_1 + \phi_2)$ and that $(K \overset{c}{\ominus} \phi) \overset{c}{\oplus} \phi \subseteq K$; two facts that are used below. For a set $X \subset \mathbb{R}^n$ we define:

$$\text{cone}(X) = \{tx : t \in \mathbb{R}_+, x \in X\}.$$

B. Conditions for State-Robust Viability

Our results rely on assumptions on how the tangent cone varies along the boundary of the safe set. To express these assumptions in a convenient manner we restrict attention the class of *practical sets* [13], which is broad enough to address most applications and enjoys a simple expression of the tangent cone. A practical set is described by a collection of r continuously differentiable functions $h_j : \mathbb{R}^n \to \mathbb{R}$:

$$S = \left\{ x \in \mathbb{R}^n \mid \min_{j \in [\![1,r]\!]} h_j(x) \geq 0 \right\},$$
$$\partial S = S \cap \left\{ x \in \mathbb{R}^n : \min_{j \in [\![1,r]\!]} h_j(x) = 0 \right\}, \tag{10}$$

where $[\![1,r]\!] = \{1, \ldots, r\} \subset \mathbb{N}$. For such a set, the tangent cone to S at x is expressed by:

$$\mathcal{T}_S(x) = \bigcap_{j \in Act(x)} \{ z \in \mathbb{R}^n \mid \langle \nabla h_j(x), z \rangle \geq 0 \}, \tag{11}$$

where the *active constraint set* at x is $Act(x) = \{ j \in [\![1,r]\!] \mid h_j(x) = 0 \}$.

Sub-tangentiality conditions become more cumbersome at "corners" of the safe set, i.e. at points x where $Act(x) > 1$. To illustrate our ideas clearer, we therefore start with a simplified condition that only applies for $r = 1$, i.e. to sets that are defined by a single inequality and therefore do not have corners.

Theorem 2. *Consider a practical set $S = \{x : h(x) \geq 0\}$ and assume that $x \mapsto cone(\nabla h(x))$ is μ-Lipschitz in a cone sense along ∂S, i.e. that for $x, \bar{x} \in \partial S$,*

$$\nabla h(\bar{x}) \subseteq cone(\nabla h(x)) \overset{c}{\oplus} \mu \|x - \bar{x}\|. \tag{12}$$

Under these conditions, if for all $x \in S$ there exists a $u \in U$ such that 1) the mapping $x \mapsto cone(f(x,u,P))$ is λ-Lipschitz in x locally in $\mathcal{B}(x,\epsilon)$:

$$cone(f(\bar{x},u,P)) \subseteq cone(f(x,u,P)) \overset{c}{\oplus} \lambda \|x - \bar{x}\|, \tag{13}$$

and 2):

$$f(x,u,P) \subseteq T_S(x) \overset{c}{\ominus} 2(\mu + \lambda)\epsilon, \tag{14}$$

then S is robustly viable under ϵ-uncertainty.

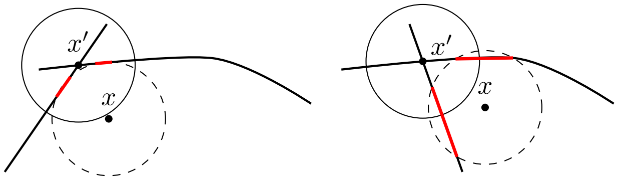

Fig. 4. Examples where (15) are satisfied (left) and violated (right). The condition requires that for any ball $\mathcal{B}(x,\epsilon)$ that intersects ∂S, there must exists a point $x' \in \partial S$ where all constraints that are active anywhere in the ball are active, and has distance at most ϵ from the intersection of the ball with ∂S. To the left, all intersection segments (in red) lie within ϵ-vicinity of x' where both constraints are active. To the right however, this does not hold.

Proof. By (11) cone Lipschitz continuity of ∇h is equivalent to cone Lipschitz continuity of $\mathcal{T}_S(x)$ for the same Lipschitz constant. Take any $x \in S$ and consider two cases: $B(x,\epsilon) \subseteq \text{int}(S)$, and $B(x,\epsilon) \not\subseteq \text{int}(S)$. In the first case the result evidently holds since $T_S(x) = \mathbb{R}^n$ in the interior of S. Consider therefore the second case; we can find a point $x' \in \partial S \cap B(x,\epsilon)$ and a $\bar{u}$ such that by (14), $\text{cone}(f(x',\bar{u},P)) \subseteq T_S(x') \overset{c}{\ominus} 2(\mu+\lambda)\epsilon$ since the right-hand side is a cone. Take any $\bar{x} \in \partial S \cap \mathcal{B}(x,\epsilon)$, then $\|x' - \bar{x}\| \leq 2\epsilon$ and consequently,

$$\begin{aligned} f(\bar{x},\bar{u},P) &\subseteq \text{cone}(f(\bar{x},\bar{u},P)) \\ &\subseteq \text{cone}(f(x',\bar{u},P)) \overset{c}{\oplus} 2\lambda\epsilon \\ &\subseteq (T_S(x') \overset{c}{\ominus} 2(\mu+\lambda)\epsilon) \overset{c}{\oplus} 2\lambda\epsilon \\ &\subseteq (T_S(\bar{x}) \overset{c}{\ominus} 2(\mu+\lambda)\epsilon) \overset{c}{\oplus} 2\lambda\epsilon \overset{c}{\oplus} 2\mu\epsilon \subseteq T_S(\bar{x}). \end{aligned}$$

Since x was arbitrary condition (9) therefore holds. □

For a condition that applies to more general sets, it is also necessary to ensure that corners are not too close to each other, and that the angles at the corners are not too sharp.

Proposition 3. *Consider a set $S = \{x : \min_{j \in [\![1,r]\!]} h_j(x) \geq 0\}$ and assume that for all $x \in S$ with $\mathcal{B}_\epsilon(x) \cap \partial S \neq 0$,*

$$\exists x' \in \partial S \text{ s.t. } \begin{cases} Act(x') = \bigcup_{\bar{x} \in \mathcal{B}(x,\epsilon)} Act(\bar{x}), \\ \max_{\bar{x} \in \partial S \cap \mathcal{B}(x,\epsilon)} \|x' - \bar{x}\| \leq \epsilon. \end{cases} \tag{15}$$

Then, if (12) holds for all h_j, $j \in [\![1,r]\!]$, and for each $x \in X$ a $u \in U$ can be found such that (13) and (14) hold, then S is robustly viable under ϵ-uncertainty.

Proof. Consider the proof of Theorem 2 and let $S_j = \{x : h_j(x) \geq 0\}$ so that $S = \cap_{j \in [\![1,r]\!]} S_j$. Due to (15) it is possible to pick $x' \in \partial S \cap B(x,\epsilon)$ so that all inequalities that are active anywhere in $\mathcal{B}(x,\epsilon)$ are active at x'. Thus, satisfaction of (14) at x' implies that $f(x',u,P) \subset T_{S_j}(x') \overset{c}{\ominus} 2(\mu + \lambda)\epsilon$, i.e. satisfaction of the subtangentality constraint for each $j \in Act(x')$. The remainder of the proof of Theorem 2 then applies for each j separately. □

Examples where the additional condition (15) is satisfied or violated are shown in Fig. 4. In essence, the requirement prohibits corners that are too sharp: if an ϵ-ball can be positioned such that it touches two separate boundaries $\{x : h_i(x) = 0\}$ and $\{x : h_j(x) = 0\}$ without being close to the set $\{x : h_i(x) = h_j(x) = 0\}$, then robust viability can not be verified via pointwise conditions.

We conclude this section by giving a certificate-based condition that serves to verify tightened sub-tangentiality conditions such as (14). In addition, it can be implemented as a control barrier function (CBF) condition [4] in an online control scheme to ensure that a robust sub-tangentiality condition is satisfied. An advantage of the CBF condition over the sub-tangentiality condition is that the $\mathcal{K}_\infty$-function α can be tuned to achieve a smooth response, as opposed to the often abrupt intervention caused by the sub-tangentiality condition.

Theorem 3. *Let $S = \{x : h(x) \geq 0\}$ for a continuously differentiable control barrier function h. Suppose that for all $p \in P$*

$$\begin{aligned}\langle \nabla h_j(x), f(x,u,p)\rangle - \|\nabla h_j(x)\|\|f(x,u,p)\| \sin(\epsilon) \\ \geq -\alpha(h(x))\end{aligned} \tag{16}$$

for a $\mathcal{K}_\infty$ function α. Then the following tightened sub-tangentiality condition holds:

$$f(x,u,P) \subseteq T_S(x) \overset{c}{\ominus} \epsilon. \tag{17}$$

Proof. Consider $x \in \partial S$. Since $\alpha(h(x)) = 0$,

$$\langle \nabla h_j(x), f(x,u,p)\rangle \geq \|\nabla h_j(x)\|\|f(x,u,p)\| \sin(\epsilon). \tag{18}$$

Let $d(u,v) = \arccos(\langle u, v\rangle / (\|u\|\|v\|))$ be the angular distance pseudometric, then (18) is equivalent to $d(\nabla h_j(x), f(x,u,p)) \leq \pi/2 - \epsilon$. Take a v such that $d(f(x,u,p), v) \leq \epsilon$, the triangle inequality in this metric then gives $d(v, \nabla h_j(x)) \leq \pi/2 - \epsilon + \epsilon = \pi/2$. This is equivalent to $\langle v, \nabla h_j(x)\rangle \geq \cos(\pi/2) = 0$, i.e. $v \in \mathcal{T}_S(x)$ by (11), which completes the proof by definition of $\overset{c}{\ominus}$. □

V. Digital Safety Control

Assuming knowledge of a set that is robustly viable under state uncertainty, we turn to the problem of enforcing sub-tangentiality conditions via digital control. Let $\mathcal{M} = (f, X, U, P)$ be a model where any potential actuation uncertainty is included in the parameter set. Since digital control only allows for control actions to be taken at discrete time instances, it is necessary to account for system evolution in between these instances to ensure safety, i.e., reachability computations are required.

Definition 5. The **reachable set** for a model $\mathcal{M}$ under input signal $\mathbf{u}$, denoted $\mathcal{R}^{\mathbf{u}}_{=t}(\Delta)$, is the set of states reachable from Δ in time t. Similarly, the **interval reachable set** for a model $\mathcal{M}$, denoted $\mathcal{R}^{\mathbf{u}}_{\leq t}(\Delta)$, is the set of states reachable in time up to t, i.e.

$$\mathcal{R}^{\mathbf{u}}_{\leq t}(\Delta) = \bigcup_{\tau \in [0,t]} \mathcal{R}^{\mathbf{u}}_{=t}(\Delta). \tag{19}$$

We use non-bold letters to denote constant input signals, i.e., for $\mathbf{u}(t) \equiv \bar{u}$ we write $\mathcal{R}^{\bar{u}}_{\leq t} = \mathcal{R}^{\mathbf{u}}_{\leq t}$.

We propose the digital control scheme in Algorithm 1 that accounts for state uncertainty and time quantization effects as a solution to Problem 2. The essence of the algorithm is to select an input that enforces the robust sub-tangentiality condition everywhere in the set where the system state may be during the time interval $[\alpha\Delta_t, (\alpha+1)\Delta_t]$.

Algorithm 1: Digital safety controller

```
At $t = t_k$, obtain state estimate
    $X_{k-1} = \{\bar{\mathbf{x}}(t_k)\} \oplus \Delta_x$;
Propagate estimate via forward reachability to time
    $t_k + \alpha\Delta_t$: $X'_{k-1} = R^{\bar{u}_{k-1}}_{=\alpha\Delta_t}(X_{k-1})$;
Select $\bar{u}_k$ such that $f(\bar{x}, \bar{u}_k, P) \subseteq T_S(\bar{x})$ for all
    $\bar{x} \in S \cap R^{\bar{u}_k}_{\leq\Delta_t}(X'_{k-1})$;
```

Theorem 4. *Let λ_x and λ_p be Lipschitz constants of f in x and p, respectively. and let*

$$V = \sup_{x \in X, u \in U, p \in P} \|f(x,u,p)\| \tag{20}$$

be an upper bound on vector field velocity. Suppose that S is robustly viable under ϵ-uncertainty and that state estimates Δ_x are such that

$$\left(\|\Delta_x\| + \lambda_p \Delta_t \left(\alpha + \frac{1}{2}\right) \|P\|\right) e^{\lambda_x \Delta_t \left(\alpha + \frac{1}{2}\right)} + V\Delta_t \leq \epsilon.$$

Suppose furthermore that one iteration of Algorithm 1 can be computed in less than time $\alpha\Delta_t$. Then Algorithm 1 solves Problem 2.

This result exhibits the trade off between system uncertainty and possible performance. With significant uncertainties and/or long sampling times the set of safe inputs shrinks. Remark that when the sampling time $\Delta_t \to 0$, the left-hand side goes to $\|\Delta_x\|$, i.e. all the uncertainty stemming from the digital implementation disappears. The inequality is the result of over-approximations of the size of reachable sets. The over-approximations are general but fairly crude; if tighter a priori over-approximations can be obtained the result can be strengthened accordingly. Before proving Theorem 4 we state two lemmas that bound the size of reachable sets.

Lemma 1. *Assume that f is λ_x-Lipschitz in x and λ_p-Lipschitz in p. If $\|\Delta\| \leq \epsilon$, then*

$$\|\mathcal{R}^{\mathbf{u}}_{=t}(\Delta)\| \leq (\epsilon + \lambda_p t \|P\|) e^{\lambda_x t}. \tag{21}$$

Proof. Consider the function $\phi(t) = \|\mathbf{x}(t; \mathbf{u}, \mathbf{p}, x_0) - \mathbf{x}(t, \mathbf{u}, \mathbf{p}', x'_0)\|$, where $\mathbf{x}(t; \mathbf{u}, \mathbf{p}, x_0)$ is the (unique) solution of (1) with initial condition $\mathbf{x}(0) = x_0$, input $\mathbf{u}$ and parameter trajectory $\mathbf{p}$. Then elementary calculations show that $\left|\frac{\mathrm{d}}{\mathrm{d}t}\phi(t)\right| \leq \lambda_x \phi(t) + \lambda_p t \|P\|$. Consequently we have $\phi(t) \leq \phi(0) + \int_0^t \left|\frac{\mathrm{d}\phi(\tau)}{\mathrm{d}\tau}\right| \mathrm{d}\tau \leq \phi(0) + \lambda_p t \|P\| + \int_0^t \lambda_x \phi(\tau) \mathrm{d}\tau$ and Gronwall's lemma [26, p. 310] gives that $\phi(t) \leq (\phi(0) + \lambda_p t \|P\|) e^{\lambda_x t}$. □

Lemma 2. *Assume that $\|f(x,u,p)\| \le V$ everywhere. Then $\|\mathcal{R}^{\mathbf{u}}_{\le t}(\Delta)\| \le \|\mathcal{R}^{\mathbf{u}}_{=t/2}(\Delta)\| + Vt$.*

Proof. Follows from remarking that the state for a time $\tau \in [0,t]$ must be within distance $Vt/2$ of the set $\mathcal{R}^{u}_{=t/2}(\Delta)$. □

Proof of Theorem 4. From the two lemmas it follows that for any $u \in U$ with the input $\mathbf{u}(t) = \bar{u}_{k-1}$ for $t \in [0, \alpha\Delta_t]$ and $\mathbf{u}(t) = u$ for $t > \alpha\Delta_t$ we have

$$\begin{aligned}\|R^{\mathbf{u}}_{\le\Delta_t}(X'_{k-1})\| &\le \|R^{\mathbf{u}}_{=\Delta_t/2}(X'_{k-1})\| + V\Delta_t\\ &\le \left(\delta + \lambda_p \Delta_t \left(\alpha + \frac{1}{2}\right)\|P\|\right) e^{\lambda_x \Delta_t (\alpha + \frac{1}{2})} + V\Delta_t,\end{aligned}$$

which by assumption is smaller than ϵ. As a consequence, robust viability under ϵ-uncertainty of S implies that a $\bar{u}_k \in U$ that satisfies the condition on Line 3 in Algorithm 1 always exists provided that $R^{\bar{u}_k}_{\le\Delta_t}(X'_{k-1})$ intersects S.

The result now follows by noting that the condition on Line 3 in Algorithm 1 necessarily implies that (8) holds for all $\tau \in [\alpha\Delta_t, (\alpha+1)\Delta_t]$. An induction argument over k then implies that S is forward invariant and that $R^{\bar{u}_k}_{\le\Delta_t}(X'_{k-1})$ intersects S at every step. □

We remark that Algorithm 1 is still sound if over-approximations of reachable sets are used rather than exact reachable sets themselves. In that case an additional margin must however be included in the robust viability condition to ensure that the digital control implementation remains feasible.

VI. Ground Robot with State Uncertainty

We illustrate some of the theory from Section IV and how robust set invariance conditions can achieve adaptive safety under varying perception conditions. Consider a ground robot model with state $\mathbf{x}(t) \in \mathbb{R}^2$ and dynamics

$$\frac{\mathrm{d}^2\mathbf{x}}{\mathrm{d}t^2} = \mathbf{u}. \tag{22}$$

The robot is tasked with tracking a desired trajectory, which is done with a standard PD controller. However, it is also subject to a safety constraint of remaining inside the circular region $\{x \in \mathbb{R}^2 : \|x\| \le \bar{r}\}$. We enforce this constraint via a radial barrier function

$$h(x,\dot{x}) = \bar{r} - r(x) - \frac{(\dot{r}(x,\dot{x}))^2}{2\bar{u}}, \tag{23}$$

where $\bar{u}$ is a maximal radial acceleration and $r(x) = \|x\|$, $\dot{r}(x,\dot{x}) = \frac{\langle x,\dot{x}\rangle}{\|x\|}$.

The zero super-level set S of h is shown in Fig. 5 in $(r,\dot{r})$ coordinates. The safety-enforcing control signal is equal to $u = -\bar{u}$ along the whole boundary of the set. Since the control does not change, it follows that S is robustly invariant under state uncertainty without needing to apply the results in Section IV. However, for illustration purposes we discuss how a modified set that satisfies a tightened sub-tangentialty condition can be constructed. Consider the red curve in Fig. 5; it demarcates the boundary of the set $\tilde{S} = \{(r,\dot{r}) : \tilde{h}(r,\dot{r}) \ge 0\}$ for $\tilde{h}(r,\dot{r}) = \bar{r} - r - \frac{(\dot{r}+\delta)^2}{2\bar{u}}$—a barrier obtained from the artificial dynamics $\dot{v} = u$, $\dot{r} = v + \delta$.

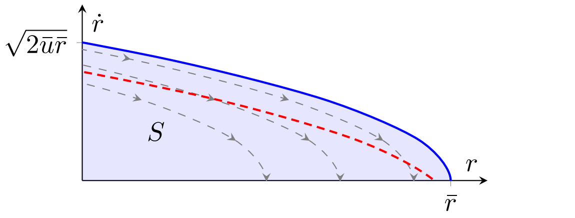

Fig. 5. Viable set for system $\ddot{r} = u$ with constraints $r \le \bar{r}$ and $u \ge -\bar{u}$.

For $f(\dot{r},u) = \begin{bmatrix}\dot{r} & u\end{bmatrix}^T$ denoting the dynamics on state space form, the modified barrier $\tilde{h}$ has the property that

$$\frac{\left\langle \nabla\tilde{h}(\dot{r}), f(\dot{r},-\bar{u})\right\rangle}{\|\nabla\tilde{h}(\dot{r})\|\|f(\dot{r},-\bar{u})\|} = \frac{\delta}{\sqrt{\bar{u}^2 + (\dot{r}+\delta)^2}\sqrt{\bar{u}^2+\dot{r}^2}} \ge \frac{\delta}{\bar{u}^2}.$$

By calculations analogous to the proof of Theorem 3 it follows that

$$f(\dot{r},-\bar{u}) \in \mathcal{T}_{\tilde{S}} \overset{c}{\ominus} \arcsin\left(\frac{\delta}{\bar{u}^2}\right), \tag{24}$$

i.e., $f(\dot{r},-\bar{u})$ satisfies a tightened sub-tangentiality condition with respect to $\tilde{h}$. This can be seen in Fig. 5 by noting that the gray curves (integral curves of the flow $f(\dot{r},-\bar{u})$) cross the red curve (zero level set of $\tilde{h}$) at an angle.

We now demonstrate the benefit of robust barriers in a simulation with sensor uncertainty: we assume that the state can not be perfectly measured, and that the quality of state measurements depends on the distance from the point $[1,-1]$ where a sensor is located (the measurement noise is Gaussian with state-dependent variance). The estimated state and variance are given by a Kalman filter, and this is the basis for both PD control and barrier enforcement. The Kalman filter gives a Gaussian distribution which does not have finite support; we therefore convert the Gaussian probability distribution into a chance constraint by selecting an appropriate confidence region X such that $\mathbb{P}[x \in X] \ge 1-\delta$, and implement a robust barrier for the state uncertainty set X via interval analysis. This results in a probabilistic safety guarantee with low probability of constraint violation. Here we chose $\delta = 0.05$.

Fig. 6 shows simulations of the system with and without noise, and with normal and robust barrier enforcement. As can be seen, the safety constraint is violated for the normal barrier in the presence of noise, but this does not happen with the barrier that is robust to state uncertainty. Furthermore, the barrier that is robust to state uncertainty adjusts its degree of cautiousness depending on the state estimate quality—the robot moves closer to the boundary when state estimates are accurate, but stays further away when uncertainty is increased, which showcases the advantage of controllers that are aware of the interplay between safety and uncertainty.

VII. Conclusions and Future Work

We have provided partial answers to the question of how safety can be achieved in the presence of uncertainty. In particular, we generalized viability to viability *under state uncertainty* which applies in situations when the state estimate is not perfectly known, and presented various results regarding properties of state-robustly viable sets. Secondly, we proposed a digital control implementation for enforcing

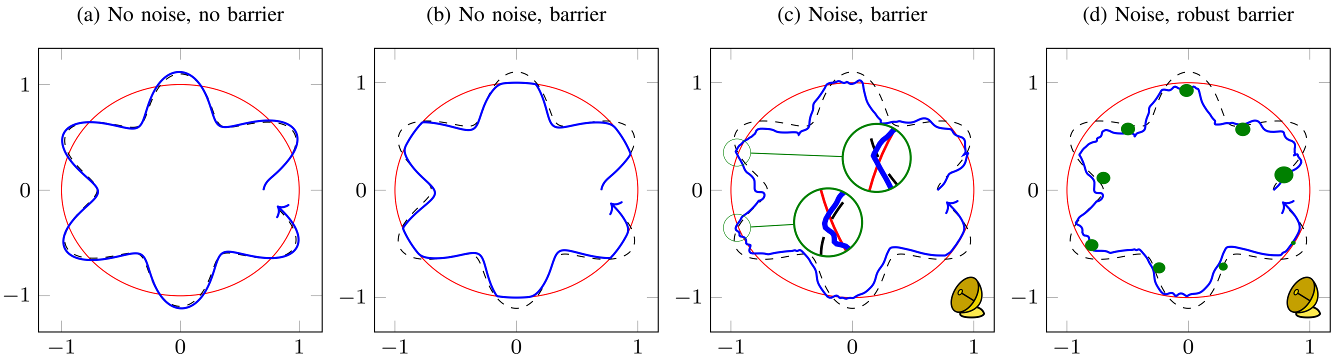

Fig. 6. Barrier performance in different situations for an example where a robot tracks the dashed black line but is constrained to stay inside of the red circle. When localization is noise-free (b), a barrier enforces the constraint $h \geq 0$, but this fails in (c) where localization is subject to error, as shown in the magnifications. However, enforcing a robust barrier constraint (8) implies safety even in the presence of noise (d). Furthermore, the degree of conservatism is adaptive with respect to the noise magnitude: the robot moves closer to the boundary when it is close to the sensor location $[1, -1]$ where the perception quality is higher. Green circles show 95% confidence regions for the estimated state.

invariance of such sets, and gave conditions for such an implementation to be equipped with correctness guarantees.

However, the introduction of state-robust viability opens up additional new problems that remain to be addressed. Firstly, there is no known general way to compute state-robustly viable sets. The results in Section IV suggest that algorithms for finding "standard" robustly viable sets under certain conditions can be extended with robustness margins to obtain sets that are viable under state uncertainty, but it may also be possible to find state-robustly viable sets directly. Secondly, the implementation in Section V relies on reachability computations, which in general can not be done exactly, and even if reachable sets can be computed it is potentially challenging to find an input on Line 3 in Algorithm 1. In future work we aim to explore over-approximations of reachable sets, and how robust inputs can be synthesized via linear over-approximations of the constraint sets analogously to the method in [16]. However, both these steps involve additional over-approximations that have to be accounted for in order for a result like Theorem 4 to be valid. This is all subject to future work, in conjunction with hardware demonstrations of these ideas.

Acknowledgment: This work was supported by NSF CPS award #1724457 and a grant from the NASA JPL President's and Director's Fund.

References

[1] Ragunathan Raj Rajkumar, Insup Lee, Lui Sha, and John Stankovic. Cyber-physical systems: the next computing revolution. In *Proc. ACM DAC*, pages 731–736, 2010.

[2] Guy André Boy. From automation to tangible interactive objects. *Annual Reviews in Control*, 38(1):1–11, 2014.

[3] Edward A Lee. Cyber physical systems: Design challenges. In *Proc. IEEE ISORC*, pages 363–369, 2008.

[4] Aaron D. Ames, Xiangru Xu, Jessy W. Grizzle, and Paulo Tabuada. Control Barrier Function Based Quadratic Programs for Safety Critical Systems. *IEEE Trans. Autom. Control*, 62(8):3861–3876, 2017.

[5] Franco Blanchini. Set invariance in control. *Automatica*, 35(11):1747–1767, 1999.

[6] Jean-Pierre Aubin. *Viability theory*. Springer Science, 2009.

[7] Ian M Mitchell, Alexandre M Bayen, and Claire J Tomlin. A time-dependent hamilton-jacobi formulation of reachable sets for continuous dynamic games. *IEEE Trans. Autom. Control*, 2005.

[8] Xiangru Xu, Jessy W Grizzle, Paulo Tabuada, and Aaron D Ames. Correctness guarantees for the composition of lane keeping and adaptive cruise control. *IEEE Trans. Autom. Sci. Eng.*, 2017.

[9] Jeremy H Gillula, Shahab Kaynama, and Claire J Tomlin. Sampling-based approximation of the viability kernel for high-dimensional linear sampled-data systems. In *Proc. ACM HSCC*, pages 173–182, 2014.

[10] Ian Mitchell. A summary of recent progress on efficient parametric approximations of viability and discriminating kernels. In *Proc. SNR@CAV*, pages 23–31, 2015.

[11] Tom Schouwenaars. *Safe trajectory planning of autonomous vehicles*. PhD thesis, Massachusetts Institute of Technology, 2005.

[12] Stanley Bak, Deepti K Chivukula, Olugbemiga Adekunle, Mu Sun, Marco Caccamo, and Lui Sha. The system-level simplex architecture for improved real-time embedded system safety. In *Proc. IEEE RTAS*, pages 99–107, 2009.

[13] Franco Blanchini and Stefano Miani. *Set-theoretic methods in control*. Springer, 2008.

[14] Quan Nguyen and Koushil Sreenath. Optimal robust safety-critical control for dynamic robotics. *International Journal of Robotics Research*, in review, 2016.

[15] Muhammad Zakiyullah Romdlony and Bayu Jayawardhana. On the new notion of input-to-state safety. In *Proc. IEEE CDC*, pages 6403–6409, 2016.

[16] Thomas Gurriet, Andrew Singletary, Jake Reher, Laurent Ciarletta, Eric Feron, and Aaron Ames. Towards a framework for realizable safety critical control through active set invariance. In *Proc. ACM/IEEE ICCPS*, pages 98–106, 2018.

[17] Jun Chai and Ricardo G. Sanfelice. On robust forward invariance of sets for hybrid dynamical systems. *Proc. ACC*, pages 1199–1204, 2017.

[18] Ian M Mitchell, Shahab Kaynama, Mo Chen, and Meeko Oishi. Safety preserving control synthesis for sampled data systems. *Nonlinear Analysis: Hybrid Systems*, 10:63–82, 2013.

[19] Charles Dabadie, Shahab Kaynama, and Claire J Tomlin. A practical reachability-based collision avoidance algorithm for sampled-data systems: Application to ground robots. In *Proc. IEEE/RSJ IROS*, pages 4161–4168, 2014.

[20] Li Wang, Aaron D. Ames, and Magnus Egerstedt. Safety Barrier Certificates for Collisions-Free Multirobot Systems. *IEEE Trans. Robotics*, 33(3):661–674, 2017.

[21] G.E.P. Box. Robustness in the Strategy of Scientific Model Building. In *Robustness in Statistics*, pages 201–236. 1979.

[22] F H Clarke, Y S Ledyaev, R J Stern, and P R Wolenski. *Nonsmooth analysis and control theory*. Springer-Verlag, 1998.

[23] Jeff S. Shamma and Kuang Yang Tu. Set-valued observers and optimal disturbance rejection. *IEEE Trans. Autom. Control*, 44(2):253–264, 1999.

[24] Franco Blanchini and Mario Sznaier. A Convex Optimization Approach to Synthesizing Bounded Complexity ℓ^∞ Filters. *IEEE Trans. Autom. Control*, 57(1):216–221, 2012.

[25] Xiaocong Fan. *Real-Time Embedded Systems: Design Principles and Engineering Practices*. Newnes, 2015.

[26] Wolfgang Walter. *Ordinary Differential Equations*. Springer, 1998.