# Run Time Assurance for Safety-Critical Systems An Introduction to Safety Filtering Approaches for Complex Control Systems

Kerianne L. Hobbs, Mark L. Mote, Matthew C. L. Abate,

Samuel D. Coogan, and Eric M. Feron POC: K. Hobbs (kerianne.hobbs@us.af.mil)

June 7, 2022

Approved for public release: distribution unlimited. Case Number AFRL-2021-2407.

More than three miles above the Arizona desert, an F-16 student pilot experienced a gravity-induced loss of consciousness (GLOC), passing out while turning at nearly 9Gs (nine times the force of gravity) flying over 400 knots (over 460 miles per hour). With its pilot unconscious, the aircraft turn devolved into a dive, dropping from over 17,000 feet to less than 8,000 feet in altitude in less than 10 seconds. An auditory warning in the cockpit called out to the pilot "altitude, altitude" just before he crossed through 11,000 feet, switching to a command to "pull up" around 8,000 feet. Meanwhile, the student's instructor was watching the event unfold from his own aircraft. As the student's aircraft passed through 12,500 feet, the instructor called over the radio "two recover," commanding the student ("two") to end the dive. As the student's aircraft passed through 11,000 feet the instructor's "two recover!" came with increased urgency. At 9,000 feet, and with terror rising in his voice the instructor yelled "TWO RECOVER!" Fortunately, at the same time as the instructor's third panicked radio call, a new Run Time Assurance (RTA) system kicked in to automatically recover the aircraft. The Automatic Ground Collision Avoidance System (Auto GCAS), an RTA system integrated on the jets less than two years earlier in the Fall of 2014, detected that the aircraft was about to collide, commanded a roll to wings level and pull up maneuver, and recovered the aircraft less than 3,000 feet above the ground. The event described here occurred in May 2016. A video from the event was declassified and publicly released in September 2016 and the footage may be found at [1]. While Auto GCAS monitored the behavior of a safety-critical cyber-physical system with a human providing the primary control functions, the same concept is gaining attention in the autonomy community looking to assure safety while integrating complex and intelligent control system designs.

RTA Systems are online verification mechanisms that filter an unverified *primary* controller output to ensure system safety. The primary control may come from a human operator, an advanced control approach, or an autonomous control approach that cannot be verified to the same level as simpler control systems designs. The critical feature of RTA systems is their ability to alter unsafe control inputs explicitly to assure safety. In many cases, RTA systems can functionally be described as containing a *monitor* that watches the state of the system and output of a primary controller, and a backup controller that replaces or modifies control input when necessary to assure safety. Note that RTA and the controllers within the architecture go by many different names, as described in the sidebar titled "RTA Aliases." RTA designs specifically to bound the behavior of neural network control systems used as the primary controller are discussed in the sidebar titled "Shielded Learning." An important quality of an RTA system is that the assurance mechanism is constructed in a way that is entirely agnostic to the underlying structure of the primary controller. By effectively decoupling the enforcement of safety constraints from performance-related objectives, RTA offers a number of useful advantages over traditional (offline) verification. Another way to think of RTA is to consider designing controllers so that there is always a plan B, as discussed in a sidebar titled "The Case for Plan B."

Verification, validation, assurance, and certification methods present the largest barriers to operational use of autonomous control in safety critical systems, such as passenger aircraft, vehicles, medical devices, and nuclear power plants. For example, the commercial aviation domain requires showing stringent safety constraint satisfaction for any failure that is more likely than "extremely improbable," defined as an event that occurs with a probability of  $10^{-9}$  [2], (more than one in a billion flight hours). While verification, validation, assurance, and certification are related, there are slight variations in their meaning. Additionally, these are also confused with concepts described in the "Safety, Reliability, and Security" sidebar. Verification is an activity that determines whether a system meets requirements [3], in effect answering the question "Did we build the system right?" Validation assesses whether a system meets the needs of the end user [4], answering the question "Did we build the right system?" Model validation refers to evaluating how well a model represents reality. Assurance is justified confidence that the system functions as intended with limited vulnerability to uncertainty, hazards, and threats based on evidence generated through development activities [5]. Certification determines whether a system conforms to a set of criteria or standards for a class of similar systems [6]-[13]. Verification, validation, certification and assurance of safety critical dynamical systems requires the development of techniques that incorporate comprehensive analysis based on rigorously developed specifications as well as techniques that continuously evaluate system behavior at run time.

A significant benefit of RTA is that it alleviates the need to design the primary controller in a way that conforms to traditional safety standards, which may not be directly applicable to the primary control design. Practical consequences of this are that RTA provides a means of testing new control algorithms on hardware platforms without compromising safety potential as well as a *near-term certification path* for autonomous controllers and safety-critical humancontrolled systems. Additionally, the verification of an assurance mechanism is generally simpler than verification of a performance-based controller as it does not require consideration of the potentially complex performance-related objectives. That is, safety verification does not become more difficult as the primary controller grows more complex, and the verification process does not need to be repeated when changes are made to the primary controller. Note that RTA is envisioned to increase safety confidence, not provide a substitute for having some level of confidence in the safety of a primary controller. Additionally, as described in the sidebar "To RTA or not to RTA? That is the Question," there are times when RTA is not appropriate.

RTA is emerging as the dominant approach for enforcing safety in real world autonomous systems and semi-autonomous systems. Though not always directly associated with the term *run time assurance*, notable autonomous systems applications are seen in road vehicles [14], [15], bipedal walking [16], [17], air traffic control [18], fixed wing aircraft collision avoidance [19]–[23] and flight envelope protection [24], VTOL aircraft [25] [26], spacecraft collision avoidance [27], manipulator arms [28], and propulsion [29] to name a few.

The rise in popularity of RTA can be attributed to many factors, including: (i) the emergence of mobile autonomous systems operating in safety critical environments, *e.g.* away from very specialized and remote applications to the vicinity of humans; (ii) the increasing complexity of the primary control system in these systems making traditional verification techniques prohibitively difficult; (iii) the growing desire to quickly update complex system software without compromising safety or repeating a costly verification process; (iv) the ability to perform fast computation online, *e.g.*, integrating trajectories or reachable sets online has only recently become a tractable task for non-trivial systems; (v) the emergence of control barrier function (CBF) and active set invariance filtering (ASIF) methods that give smooth and *minimally invasive* modifications to the desired actions. Additionally, RTA seems to provide an attractive solution to bounding the behavior of machine learning and artificial intelligence (AI) systems. While some notable efforts have been made in the direction of generating and verifying provably safe deep neural network controllers [30]–[33], the complexity and dynamic nature of these systems often precludes the ability to generate rigorous safety guarantees on the primary controller. The RTA

paradigm enables an alternate path, without a need to compromise on safety. Reinforcement learning (RL) applications on safety-critical systems often fall into the domain of safe RL, described as a process of policy learning that maximizes expected return while ensuring system performance and/or respecting safety constraints [34]. The safe RL community generally takes one of two approaches: modifying the optimality criterion to reduce risk or using an RTA mechanism during the training process, often referred to as a shield [35].

This article provides a tutorial on developing RTA systems. First a description of the basic architecture, modeling framework, and fundamental definitions are introduced. Second, categories and properties of RTA systems are identified, and systems engineering and humanmachine interaction considerations beyond the system dynamics are discussed. Third, a particular category of RTA systems called the Simplex architecture is described. Fourth, a second category of RTA systems called active set invariance filters is described. Fifth, implicit and explicit variations of both the Simplex architecture and active set invariance filter RTA approaches are described for the canonical double integrator system. Sixth, approaches for RTA system performance under uncertainty is described. Seventh, verification approaches for RTA algorithms and architectures are presented. Additional topics, applications, definitions, examples, and supplemental information are found in sidebars throughout.

# **Run Time Assurance for Safety-Critical Cyber-Physical Systems**

Cyber-physical systems feature computing devices that interact with the physical world via actuators and sensors [36], and such systems are *safety-critical* when failure would result in loss of life, damage to property, or other unacceptable damages such as environmental harm [37]. An RTA architecture acts as an online-verification and enforcement tool for cyber-physical systems that guarantees certain system-level safety requirements are met at run time. In this section, the RTA problem formulation, appropriate RTA employment in the control loop, and other important design considerations are discussed.

## The Run Time Assurance Architecture

RTA presents an approach to control design that allows a designer to sidestep the common trade-off between performance and safety. The central idea behind RTA is to decouple the task of enforcing safety constraints from all other objectives of the controller. This is achieved by splitting the control system into two components: a performance-driven *primary* controller, and a safety-driven *RTA mechanism*. The RTA mechanism preempts the desired inputs from the primary controller when necessary for ensuring the safety of the system, and otherwise lets the input pass through unaltered. This approach is associated with the feedback control architecture

shown in Figure 1b. The output of the performance-driven controller is referred to as the *desired* input and assigned the variable  $u_{des}$ , and the output of the RTA mechanism is referred to as the *actual* input (or assured input) and is assigned the variable  $u_{act}$ . At its core, an RTA mechanism is a mapping from a state  $x \in \mathbb{R}^n$  and a desired input  $u_{des} \in \mathbb{R}^m$  to an actual input  $u_{act} \in \mathbb{R}^m$ , and has the objectives of (i) enforcing safety of the system as defined by state constraints, and (ii) having  $u_{act}$  as close as possible to  $u_{des}$ , so long as it does not conflict with the first objective. A common choice for measuring how *close* the two signals are is the norm of the difference between the two variables, however, this measurement may also take other forms, such as the integral of the deviation. An RTA mechanism is said to be *assured* when the output  $u_{act}$  is always verifiably safe, as defined in later sections. Moreover, an assured RTA mechanism guarantees safety for the entire system, and for all time, and this is true regardless of the chosen primary controller. Although not considered an RTA architecture under the definitions in this paper, a related concept for an alternative safety architecture is described in the sidebar titled "Reference and Command Governors."

## **Modelling Cyber-Physical Systems**

Mathematical models that provide an abstraction of real world system behavior may be defined at varying levels of complexity based on the contrasting needs of accurately predicting system behavior and conducting mathematical analysis. This article focuses mainly on cyberphysical systems modeled as *dynamical systems*; hereafter the variable  $x \in X \subseteq \mathbb{R}^n$  denotes the state vector of the system model and the variable  $u \in U \subset \mathbb{R}^m$  denotes a bounded control input, where X is the set of all possible system states and the set U is the *admissible* set of controls, determined *a priori* by the actuation constraints of the real-world system. *Continuous-time* system models appear as systems of ordinary differential equations as in

$$\dot{x} = F_{\rm c}(x, u),\tag{1}$$

and discrete-time system models appear as state-update maps as in

$$x^+ = F_{\rm d}(x, u). \tag{2}$$

While the computational elements responsible for the control of a CPS live in a discrete world, the physical laws of motion tend to take the form of ordinary differential equations. In many cases, systems of the form (2) are obtained from discretizing equations of the form (1) over the controller update period.

A natural generalization is to consider *hybrid* systems that combine continuous-time and discrete-time elements [38], [39]. In fact, as seen below, some RTA mechanisms are naturally interpreted as inducing a hybrid system in closed-loop, even when the open-loop system is of



Figure 1: To illustrate where run time assurance sits within the control system architecture, (a) depicts a prototypical feedback control system architecture, and (b) shows the insertion of a run time assurance mechanism. The primary controller (e.g. human operator, an advanced control approach, or an autonomous control approach) is performance-driven and outputs a desired input  $u_{des}$ . The run time assurance mechanism maps the state  $x \in \mathbb{R}^n$  and a desired input  $u_{des} \in \mathbb{R}^m$  to an actual input  $u_{act} \in \mathbb{R}^m$ , with the objectives of (i) enforcing state constraints, and (ii) outputing  $u_{act}$  as close as possible to  $u_{des}$ . Run time assurance can be thought of as a safety filter on the performance controller output.

the form (1). For the purposes of this article, attention is restricted to to systems with dynamics governed by (1) or (2) and will not formally characterize hybrid systems, but we note instances where RTA tools have been specifically extended to such systems in the literature.

Throughout this article, it is discussed how one may verify or assure system models of the form (1) or (2) against safety constraints. However, (1) and (2) are only models, and real-world systems generally deviate from their models. When this deviation is significant, safety guarantees derived from the models are invalid. One method for addressing this limitation is to instead consider a *nondeterministic model* that explicitly considers uncertainties that account for the deviation between the real-world system and the *deterministic models* presented in (1) or (2). This article first focuses on the deterministic system models (1) and (2) and discusses extensions to nondeterministic models in the Section "Assurance in the Presence of Uncertainty."

#### Admissible Control of Cyber-Physical Systems

The main focus of this article is in devising *safe* controllers for cyber-physical systems modeled as in (1) or (2). Safety control schemes must conform to the boundaries of admissible control. That is, a control law can only be realized on a physical system if its outputs are bounded to a set of signals that respect the actuation constraints of that system. A feedback *control law u* denotes a mapping to  $\mathbb{R}^m$  from either the state domain or an augmentation of the state domain. For instance, in the context of RTA, many primary control laws will take the form  $u : X \to \mathbb{R}^m$ , while RTA control laws will generally take the form  $u : X \times \mathbb{R}^m \to \mathbb{R}^m$ , and either can be extended to include explicit dependence on time. Control laws are said to be *admissible* when their range is the admissible input set U, *e.g.*  $u : X \to U$  or  $u : X \times \mathbb{R}^m \to U$ .

Any control law u can be made admissible via composition with a saturation (also called clamping) function; a saturation function  $\sigma : \mathbb{R}^m \to U$  is such that  $\sigma(u)$  is approximately equal to u when u is in the interior of U and  $\sigma(u)$  is approximately the projection of u onto the boundary of U when u is outside of U. Here, there are two main approaches: (i) hard clamping, whereby the saturation function  $\sigma(u) = u$  for all  $u \in U$  and  $\sigma(u)$  is a projection onto the boundary  $\partial U$ , and (ii) smooth clamping, whereby  $\sigma$  takes the form of a differentiable function that such that  $\sigma(u) \approx u$  for  $u \in U$ . For example, for the case of scalar input and  $U = [-1, 1] \subset \mathbb{R}$ , then  $\sigma_1(u) = \max(-1, \min(1, u))$  is a hard clamping function while  $\sigma_2(u) = \tanh(u)$ ,  $\sigma_3(u) = \frac{2}{\pi} \tan(\frac{\pi}{2}u)$ ,  $\sigma_4(u) = u(1+u^2)^{-1/2}$ , and  $\sigma_5(u) = \tanh(u+0.5u^3)$  are examples of smooth clamping functions.

## Safety and Invariance of Cyber-Physical Systems

In a colloquial context, safety means freedom from harm or danger. For safety critical systems, safety is freedom from conditions that would, in a worst-case environment, lead to an unacceptable loss, such as a loss of control, physical damage to the system under control, loss of human life, human injury, property damage, environmental pollution, or failure of the mission [40], [41]. In the context of a dynamical system, safety is a characterization of the state of the system and its evolution. While alternate approaches to forming safety specifications exist (*e.g.* temporal logic), this research focuses specifically on set invariance requirements derived from static properties on the state. Specifically, safety properties are specified with state constraints, and safety relates to whether a particular initial condition will lead to the safety constraints being satisfied for all time.

In this article, safety constraints are defined with inequality constraints on the state. In particular, we consider  $\varphi_i : X \to \mathbb{R}$  for  $i \in \{1, \ldots, M\}$ , where M is the number of safety

constraints and it is required that always  $\varphi_i(x) \ge 0$  for each *i*. The set of states that satisfy all of the safety constraints is referred to as the *constraint set* (sometimes called the *allowable set* or the constraint space) and is given by

$$C_{A} := \{ x \in X \mid \varphi_{i}(x) \ge 0, i \in \{1, \dots, M\} \}.$$
(3)

*Example:* Consider a wingman aircraft with position  $(x_W, y_W)$  flying co-altitude in formation with a lead aircraft with position  $(x_L, y_L)$  and the requirement is that the two aircraft never go within 30 meters of each other. A constraint set over the states  $x = [x_L, y_L, \dot{x}_L, \dot{y}_L, x_W, y_W, \dot{x}_W, \dot{y}_W]^T$  is given by:

$$\mathcal{C}_{\rm A} = \{ x \in \mathbb{R}^8 \, | \, \sqrt{(x_{\rm L} - x_{\rm W})^2 + (y_{\rm L} - y_{\rm W})^2} - 30 \ge 0 \}.$$
(4)

 $\Diamond$ 

Importantly, there may exist states in  $C_A$  that obey the safety constraints at a given time, but that inevitably lead to violations in the future. For example, in the aircraft case it is possible for the wingman aircraft to have just more than 30 m separation with the lead aircraft initially, but be traveling too fast to be able turn, climb, or descend to avoid a collision. Though the state falls in the constraint set  $C_A$  at a given time, it is not "safe" as it will inevitably lead to a departure from the set. A meaningful definition of safety must encode additional information relating to whether the safety constraints will continue to be satisfied for all time with a particular control law and subject to a particular set of dynamics and actuation constraints.

Informally, a system is safe if the state belongs to  $C_A$  for all time. This concept is formalized with the notion of *set invariance*: A state  $x(t_0)$  is safe with respect to a closed-loop dynamical system if that state lies in a forward invariant subset of the constraint set, *i.e.* if:

$$x(t_0) \in \mathcal{C}_{\mathrm{S}} \implies x(t) \in \mathcal{C}_{\mathrm{S}} \quad \forall t \ge t_0$$

$$\tag{5}$$

where  $C_S \subseteq C_A$ . In this case,  $C_S$  is said to be a *safe set*. Furthermore, control laws are said to be safe when they render some nonempty subset of  $C_A$  forward invariant. It is generally desirable that a control law be designed such that the safe set is as large as possible. However, it is generally not possible to find an admissible control law that will render  $C_A$  itself forward invariant. The distinction between  $C_S$  and  $C_A$  is necessary as the system is subject to actuation and dynamics constraints, meaning that the control signal only has a limited amount of influence over the evolution of the state.

It is important to note that forward invariance, and by extension safety, are properties of the closed-loop system and are not defined in absence of a controller. The existence, size, and shape of  $C_S$  are all dependent on the controller. The question of whether it is possible to find



Figure 2: Safe set  $C_S$  and constraint set  $C_A$  topology for explicit (left) and implicit (right) cases. Here X represents the set of all states, the constraint set  $C_A$  is the set of states considered safe when there are no controller limitations, and the safe set  $C_S$  is a subset representing the set of states that are safe under the limits of the control system. For example, a car approaching a wall is in the constraint set  $C_A$  so long as it does not hit the wall. However, if the car is two feet from the wall and moving at 100 miles per hour, the it can be in the constraint set  $C_A$  but not in the safe set  $C_S$ . To be in  $C_S$ , the car must be able to stop before hitting the wall. In cases where it is hard to define  $C_S$ , but a predefined backup control is known (e.g. car slamming on the brakes), an implicit backup set  $C_b$  may be used instead of  $C_S$ .

a control law that will render a particular set invariant is addressed with the concept of control invariance. A set  $S \subseteq X$  is said to be *control invariant* (also called viable) if there exists an admissible control law that renders S forward invariant. The largest control invariant subset of  $C_A$  is known as the *viability kernel* [42]. Intuitively, the viability kernel is the largest achievable safe set, and it forms the boundary between the states for which it is and is not possible to find a control input that will keep the system safe for all time. A fact that will become important in later sections is that control invariant sets are forward invariant under controllers that take the form of certain optimization-based procedures.

## Identifying Safe Sets

Computing invariant sets has been the subject of a wide variety of research. A common approach to identifying an invariant region is to find a Lyapunov function. Since stability is a stronger condition than invariance, any stable region under a particular controller must also be invariant under that controller. Moreover, invariance of a set under an admissible control law  $u: X \to U$ , also implies control invariance under U. Invariant regions may also be computed by identifying barrier certificates [43]. While identifying Lyapunov functions or barrier certificates is typically a difficult task, they may be computed automatically for polynomial dynamics using sum-of-squares (SOS) optimization [15], [44], [45] at moderate to high computational costs. Other approaches include using discretized solutions the the Hamilton-Jacobi equations [46], and sampling [47]. Safe sets may also be computed for discrete time systems using approaches in computational geometry [48]. These approaches typically rely on convexity of the set  $C_A$  and linearity of the dynamics for the purposes of implementing set operations. See [49]–[51] for a more extensive overview of techniques for invariant set computation.

*Example:* Consider the double integrator system—*e.g.* a spacecraft in linear motion— described by

$$\begin{aligned} \dot{x}_1 &= x_2 \\ \dot{x}_2 &= u \end{aligned} \tag{6}$$

where the state vector  $x = [x_1, x_2]^T \in \mathbb{R}^2$  is composed of a position state  $x_1$ , and a velocity state  $x_2$ , and  $u \in [-1, 1]$  is the acceleration due to thrust. The safety constraint  $\varphi(x) = -x_1 \ge 0$  is imposed on the system, reflecting the requirement that the system avoid collision with an object located at  $x_1 = 0$ , and extending to  $x_1 \ge 0$ . The constraint set is

$$C_{\rm A} = \{ x \in \mathbb{R}^2 \, | \, -x_1 \ge 0 \}.$$
(7)

The largest control invariant subset (*i.e.* the viability kernel) of  $C_A$  is  $C_S = \{x \in \mathbb{R}^2 | h(x) \ge 0\}$ where

$$h(x) = \begin{cases} -2x_1 - x_2^2 & \text{if } x_2 > 0\\ -x_1 & \text{if } x_2 \le 0. \end{cases}$$
(8)

The unsafe states in the constraint set  $C_A \setminus C_S$  represent states for which a future collision is inevitable. Intuitively,  $x_2 = \sqrt{-2x_1}$  represents the maximum safe velocity at  $x_1$ . If this approach speed is exceeded, then there will not be sufficient distance to stop before a collision occurs. This result is made apparent by considering the flow of the system under a recovery maneuver u = -1, which applies maximum thrust force away from the obstacle. It can be seen that  $C_S$  is a forward invariant set under this control law. Figure 3 shows a depiction of the described sets, and the flow under various inputs.



Figure 3: A car moving straight towards a wall can be modeled as a double integrator system. This shows a phase plot for double integrator system with u = -1, 0, 1, where  $x_1$  is position,  $x_2$  is velocity, u = 1 represents full forward acceleration, u = 0 represents continuing at a constant velocity, and u = -1, represents full acceleration away from the wall. The viability kernel  $C_S$  is shaded green and represents the safe states where it is possible to apply full negative acceleration to avoid hitting the wall. Collision states are shaded red and represent states in which the car has hit or driven past the wall  $X \setminus C_A$ . The *inevitable* collision states  $C_A \setminus C_S$  are shaded purple and represent cases where the car has not yet hit the wall, but is traveling too fast towards the wall to be able to stop in time.

## Implicit and Explicit Definitions of the Safe Set

Given a constraint set  $C_A$  and an admissible control law  $u: X \to U$ , the question arises of how one can determine if a given state  $x \in X$  is safe. In many cases, the set of safe states  $C_S$  can be identified *explicitly* with a functional representation; *e.g.* 

$$\mathcal{C}_{\mathrm{S}} = \{ x \in X \mid h(x) \ge 0 \},\tag{9}$$

with  $h: X \to \mathbb{R}$ . Checking whether  $x \in C_S$  is then equivalent to checking whether  $h(x) \ge 0$ . The safe set boundary may be determined in a number of ways, such as through Lyapunov arguments or reachability analysis. Additionally, for continuous-time systems as in (1), Nagumo's theorem states that, under appropriate regularity assumptions on f, u, and h, a necessary and sufficient condition for invariance of a candidate set  $C_S$  as in (9) is  $\frac{\partial h}{\partial x}(x)^T F_c(x, u(x)) \ge 0$  for all x such that h(x) = 0 [52], and thus Nagumo's theorem can be used to verify that a candidate safe set is invariant. Moreover, as discussed in the Section "Assurance in the Presence of Uncertainty," it is possible to consider nondeterministic effects through the identification of sets that are *robustly* forward invariant, that is, that are forward invariant under any realization of the nondeterminism.

*Example:* Consider the double integrator system (6), now paired with the constraint set  $C_A = \{x \in \mathbb{R}^2 \mid 1 - \|x\|_{\infty} \ge 0\}$ . Suppose that a control law is constructed such that for all  $x \in C_A$ ,  $u = -x_1 - x_2$ . Note that actuation constraints can be ignored in  $C_A$  if a hard clamping function is used, and that for all  $x \in C_A$ ,  $u(x) \in U$ . In this case, the closed-loop dynamics are

$$\dot{x} = \begin{bmatrix} 0 & 1\\ -1 & -1 \end{bmatrix} x,\tag{10}$$

for all states in  $C_A$ . A valid safe set can be constructed by identifying a region of attraction contained in  $C_A$ . For instance, note that

$$V(x) = 1.2x_1^2 + 0.2x_1x_2 + 1.1x_2^2$$
(11)

is a Lyapunov function for the closed-loop system and the level curve V(x) = 1 is contained in  $C_A$ . Thus,  $C_S = \{x \in \mathbb{R}^2 | h(x) \ge 0\}$ , where h(x) = 1 - V(x) defines an invariant set contained in  $C_A$ . The sets  $C_S$  and  $C_A$  are depicted along with the system flow in Figure 4(a). Note that while  $C_S$  is invariant, it is not the largest forward invariant set contained in  $C_A$ . It is said to be conservative.

Explicit identification of forward invariant subsets is typically obtained only at the expense of conservatism. In particular, the methods used to identify forward invariant sets are not generally scalable to complex and high-dimensional systems, and the safe sets obtained with these methods may greatly underapproximate the largest safe set obtainable. However, an explicit (functional) representation of the safe set boundary is not always necessary. One may *implicitly* define  $C_{\rm S}$ in terms of the closed-loop trajectories under the control law. For example, consider a *backup* control law  $u_{\rm b}: X \to U$ , and let  $\phi^{u_{\rm b}}(t;x)$  represent the state reached after starting at  $x \in X$ and applying  $u_{\rm b}$  for t units of time. Then the set

$$\mathcal{C}_{\mathrm{S}} = \{ x \in X \mid \forall t \ge 0, \ \phi^{u_{\mathrm{b}}}(t; x) \in \mathcal{C}_{\mathrm{A}} \}$$
(12)

is by definition an invariant set under  $u_b$ , and it is entirely contained in  $C_A$ . The utility of this definition arises from the observation that it is possible to check whether *individual states* are safe simply by integrating the dynamics forward from those states. For example, if  $C_A = \{x \in X \mid \varphi(x) \ge 0\}$  then the following are equivalent:

(i) 
$$\forall t \ge 0, \ \phi^{u_{\mathrm{b}}}(t;x) \in \mathcal{C}_{\mathrm{A}},$$
 (ii)  $\inf_{t \in [0,\infty)} \varphi(\phi^{u_{\mathrm{b}}}(t;x)) \ge 0$  (13)

It is interesting to note that in special cases where the infimum can be solved in closed form, the solution can be used to define an explicit safe set, as shown in (8) and in the following example.



Figure 4: Phase plots for a mass-spring-damper system (10), modeled as a double integrator with the constraint set  $C_A = \{x \in \mathbb{R}^2 | 1 - ||x||_{\infty} \ge 0\}$ , where with the complement of the constraint space  $C_A$  is shaded red, and  $C_A \setminus C_S$  is shaded purple.

Example: Consider the system,

$$\begin{bmatrix} \dot{x}_1\\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} \cos(x_2)\\ u \end{bmatrix}$$
(14)

with  $u \in [-1, 1]$ , the constraint set  $C_A = \{x \in \mathbb{R}^2 | x_1 \ge 0\}$  and the admissible control law  $u_b = 1$ . The solution for the flow from initial state  $x_0 = [x_{0,1}, x_{0,2}]^T$  is given by

$$\begin{bmatrix} \phi_1^{u_{\rm b}}(t;x_0) \\ \phi_2^{u_{\rm b}}(t;x_0) \end{bmatrix} = \begin{bmatrix} x_{0,1} - \sin(x_{0,2}) + \sin(x_{0,2} + t) \\ t + x_{0,2} \end{bmatrix}$$
(15)

and  $x_0$  is safe if  $\inf_{t \in [0,\infty)} \phi_1^{u_b}(x_0,t) = x_{0,1} - \sin(x_{0,2}) - 1 \ge 0$ . Using this information, a safe set under this maneuver can be defined explicitly as

$$C_{\rm S} = \{ x \in \mathbb{R}^2 \, | \, x_1 - \sin(x_2) - 1 \ge 0 \}.$$
(16)

The sets  $C_A$ ,  $C_S$  are depicted in Figure 5 along with the flow of the system under  $u_b = 1$ .

Importantly, the system flow  $\phi^{u_b}(t;x)$  can readily be evaluated over a *finite* time horizon  $\mathcal{T} = [0,T]$  by numerically integrating the dynamics. In this case, safety of a trajectory can be shown by ensuring that (i) the trajectory lies in  $\mathcal{C}_A$  over  $\mathcal{T}$  and (ii) the endpoint of the finite trajectory T lies in a known invariant subset  $\mathcal{C}_b \subseteq \mathcal{C}_A$ . That is,

$$\mathcal{C}_{\mathrm{S}} = \{ x \in X \mid \forall t \in \mathcal{T}, \ \phi^{u_{\mathrm{b}}}(t;x) \in \mathcal{C}_{\mathrm{A}} \land \phi^{u_{\mathrm{b}}}(T;x) \in \mathcal{C}_{\mathrm{b}} \}.$$
(17)



Figure 5: Phase plot for system (14) under  $u_{\rm b} = 1$ . The safe set  $C_{\rm S}$  is shaded green, the complement of  $C_{\rm A}$  is shaded red, and  $C_{\rm A} \setminus C_{\rm S}$  is shaded purple.

In this context, a safe terminal set  $C_b \subseteq C_A$  is referred to as a *backup set*, and the set (17) is the *safe backward image* (SBI) of  $C_b$ . Note that numerical integration over a continuous interval [0, T] requires that the trajectory be approximated with a finite amount of points sampled along this interval; *e.g.*  $\mathcal{T} = \{t_1, ..., t_K\} \subset [0, T]$ . Testing with the discrete set of points does not strictly prove invariance as it does not ensure that the trajectory will not leave between the sampled points. While the approximation is sufficient for most practical purposes, this may be addressed with a finer time discretization or by deriving bounds for an inner approximation of the constraint set as in [53].

While a single case is presented here, various forms of implicit safety may arise under modified assumptions. For instance, [54] considers the case where  $C_b$  is not forward invariant, and guarantees are in finite time. As discussed in the Section "Assurance in the Presence of Uncertainty," this concept may be extended to nondeterministic systems by approximating the forward reachable sets or invariant tubes around the recovery trajectories. The topology of the constraint and safe sets for both implicit and explicit cases is illustrated in Figure 2.

*Example:* Consider the system (10) along with the associated constraint set  $C_A = \{x \in \mathbb{R}^2 | \varphi(x) \ge 0\}$  with  $\varphi(x) = 1 - ||x||_{\infty}$ . We now consider the safe set from the previous example as the backup set. That is, consider the backup set  $C_b = \{x \in \mathbb{R}^2 | h_b(x) \ge 0\}$  with  $h_b(x) = 1 - V(x)$  and V(x) described by (11). The flow for this system can be evaluated exactly

$$\phi^{u_{\rm b}}(t;x) = e^{At}x\tag{18}$$

where A is the system matrix from (10). Consider the backup time horizon  $\mathcal{T} = [0, T]$ . A state  $x \in \mathbb{R}^2$  is in the safe backward image if the following conditions are met

(i) 
$$\forall t \in \mathcal{T}, \quad \varphi(e^{At}x) \ge 0, \qquad (ii) \ h_{\mathbf{b}}(e^{AT}x) \ge 0.$$
 (19)

The sets for this case are depicted in Figure 4(b).

Identifying Safe Backup Sets

Backup sets  $C_b$  are safe terminal sets that act as *seeds* of safety. As seen in Figure 4 the backup set may be used to implicitly identify a larger safe region via trajectories under a backup controller. This is useful because small invariant sets are generally much easier to find than large ones. Furthermore, a finite-time trajectory can be used to show safety over an infinite horizon when it is shown that it the trajectory safely reaches such a set. This concept has been frequently explored in the context of trajectory optimization and *Model-Predictive Control* (MPC) [23], [55]. In the literature related to MPC, the backup set is often referred to as a *terminal feasible invariant set*. Backup sets *are* safe sets and may be computed as any safe subset of  $C_A$  using the methods described above. However, it is common to construct these from domain knowledge of safe configurations of the particular system.

For an example of a safe backup set, consider that ground vehicles and vertical takeoff and landing (VTOL) air vehicles may bring themselves to rest in a safe location. In this case, the role of  $u_b$  is to generate a safe stopping maneuver; *e.g.* [56]. Individual rest states are invariant points in the state space, and sets of adjacent rest states form invariant surfaces or volumes. In many cases it is not possible or practical to bring a vehicle to rest, and a backup set may instead be identified from safe periodic trajectories. For example, fixed wing aircraft may compute safe loiter circles [18], [22], [23], [57] or spacecraft may compute a set of safe orbits [58]. In many cases, the invariant set is a region without volume, and only an infinitesimal perturbation is required to cause a departure from the set. In practice, it is desirable to show that such surfaces are locally attractive under some backup control law, or to explicitly identify a larger invariant region around the initial set. For example, [58] develops a control law that stabilizes the a spacecraft to a set of safe orbits.

Example: Consider the damped linear system

$$\ddot{x} = -\dot{x} + u \tag{20}$$

15

as

 $\diamond$ 

with the constraint set given by  $C_A = \{x \in \mathbb{R} \mid x \ge 0\}$ . A safe backup set is given by the subset of  $C_A$  with zero velocity:  $C_b = \{x \in \mathbb{R} \mid x \ge 0, \dot{x} = 0\}$ . Note that  $C_b \subset C_A$  and that  $C_b$  is invariant when u = 0.

*Example:* Consider two spacecraft in planar orbit around a central body: (i) a *target* spacecraft, which is in a fixed circular orbit with period  $\tau$ , and (ii) a *chaser* spacecraft of mass m. In this setting, the dynamics of the chaser spacecraft are given by the Clohessy-Wiltshire-Hill equations, as developed in [59]:

$$\dot{x}_{1} = x_{3}$$

$$\dot{x}_{2} = x_{4}$$

$$\dot{x}_{3} = 3n^{2}x_{1} + 2nx_{4} + \frac{1}{m}u_{1}$$

$$\dot{x}_{4} = -2nx_{3} + \frac{1}{m}u_{2}$$

$$\dot{x}_{5} = -|u_{1}| - |u_{2}|$$
(21)

with state  $x \in \mathbb{R}^5$  and control input  $u \in [-u_{\max}, u_{\max}]^2$ . In this setting,  $x_1, x_2$  denote the relative distances between the spacecraft,  $x_3, x_4$  denote the relative velocities, and  $x_5$  denotes a fuel state. Additionally,  $n = \frac{2\pi}{\tau}$  [59]. The constraint set is defined by the constraints that the chaser must not run out of fuel and must stay at a distance greater than  $R_{\min}$  from the target spacecraft. The constraint set is  $C_A = \{x \in \mathbb{R}^5 | x_1^2 + x_2^2 - R_{\min}^2 \ge 0, x_5 \ge 0\}$ .

Backup Set from Invariant Points: Invariant points exist as zero velocity states in the  $x_2$  $x_5$ -plane, hence  $\{x \in \mathbb{R}^5 | x_1 = x_3 = x_4 = 0\}$  is invariant for  $u \equiv 0$ . A safety set is obtained from this set as  $C_{b,1} = \{x \in \mathbb{R}^5 | x_1, x_2, x_3 = 0, x_5 \ge 0, |x_2| \ge R_{\min}\}.$ 

Backup Set from Periodic Trajectories: It can be shown that the chaser spacecraft is on an elliptical orbit around the target spacecraft whenever u = 0 and  $x_3 = \frac{n}{2}x_2$ ,  $x_4 = -2nx_1$ , with each individual orbit occupying the position states  $x_1^2 + \frac{1}{4}x_2^2 = b^2$  where  $b \ge 0$  is the semi-minor axis of the orbit. Hence, the linear subspace  $\{x \in \mathbb{R}^5 | x_3 = \frac{n}{2}x_2, x_4 = -2nx_1\}$  is an invariant manifold composed of all such closed orbits. A backup set can be constructed from this subspace by considering range of orbits that fall into the constraint set:  $C_{b,2} = \{x \in \mathbb{R}^5 | x_3 = \frac{n}{2}x_2, x_4 = -2nx_1, x_1^2 + \frac{1}{4}x_2^2 \ge R_{\min}^2\}$ 

# **Properties of Run Time Assurance Systems**

RTA has evolved over the last several decades stemming from research in control theory, computer science, electrical, mechanical, and aerospace engineering. In this section, each of several basic classifications and properties of RTA mechanisms are defined, concluding with RTA-human interaction considerations. RTA approaches can be classified as: explicit or implicit,

zero or first order, and latched or unlatched. Explicit approaches precisely define a specific safe set, while implicit RTA approaches define recovery trajectories under a predefined backup control law. Zeroth order RTA methods do not use gradients of the dynamics and are often associated with RTA systems that use methods described in the Section "The Simplex Architecture," while first order methods take advantage of gradient computations and are used in methods such as those described in the Section "Active Set Invariance Filtering." Latched RTA systems switch to a backup controller and remain under its control until a specified release condition is met, while unlatched implementations return to the primary control as soon as the system returns to a safe set. In addition to these classifications, RTA systems feature safety and performance properties described in this section discusses important concepts with respect to RTA-human interaction including: variable risk tolerance, ability to turn the RTA off, transparency, and missed detection/false alarm considerations.

## **Implicit and Explicit RTA Approaches**

Fundamentally, an RTA mechanism is a control law that renders a subset of the constraint space forward invariant, activates a recovery response near the boundary of that set, and that passes through the desired input from the primary controller everywhere else. The set made safe under the RTA is referred to as the *safe operational region* of the RTA system, and it is denoted by  $C_S$ . In cases where the recovery response relies on switching to a backup control law  $u_b$ , the safe operational region is a subset of the constraint set  $C_A$  that is forward invariant under that control law. Approaches of this type are considered in the section entitled "The Simplex Architecture." Alternatively, in cases where the recovery response is determined as the solution to an optimization program, and safety is enforced with the constraints of that program, then the safe operational region is a control invariant subset of  $C_A$ . This is because the optimization program searches over all feasible control actions — *i.e.* control invariant sets are forward invariant under certain optimization procedures. Approaches of this nature are considered in the section "Active Set Invariance Filtering."

As with the safe sets discussed previously, the safe operational region can be identified explicitly or implicitly. *Implicit* (or trajectory-based) approaches compute finite-time trajectories under a backup controller { $\phi^{u_b}(t; x) \in X | t \in T$ } online (at run time) and use the information to decide when or how to intervene. *Explicit* (or region-based) approaches may or may not switch to a backup controller, but do not rely on simulating trajectories of the backup controller online. Typically, explicit approaches identify trusted regions offline via the construction of either a large forward invariant or control invariant set. That is, implicit RTA approaches are those that utilize an online *look-ahead* and will bound the system to an implicitly defined safe set such as (17) while explicit RTA approaches determine safe states *a priori* and will bound the system to an explicitly defined safe set such as (9). This concept is demonstrated in the following example.

*Example:* Consider the mass-spring-damper system described by (10) with the constraint space  $C_A = \{x \in \mathbb{R}^2 | 1 - ||x||_{\infty} \ge 0\}$ . An RTA mechanism can be constructed by requiring that the backup controller be applied near the boundary of a set  $C_S$ , where  $C_S$  is safe under the backup controller. Let the RTA mechanism be,

$$u(x) = \begin{cases} u_{\rm b} & \text{if } \vartheta(x) \\ u_{\rm des} & \text{otherwise.} \end{cases}$$
(22)

where  $u_{\text{des}}$  is the desired control input and  $\vartheta(x)$  represents the activation condition (or switching condition). Considered below are implicit and explicit activation conditions.

## Explicit Condition

Let  $C_{\rm S} = \{x \in \mathbb{R}^2 \mid h(x) \ge 0\}$  with h(x) = 1 - V(x) and V(x) being the Lyapunov function (11). The set  $C_{\rm S}$  is said to be trusted as it is safe under  $u_{\rm b}$ . The activation condition

$$\vartheta(x):h(x) \ge \varepsilon \tag{23}$$

ensures safety of  $C_S$  with respect to  $u_{act}$  by activating  $u_b$  near the boundary where  $\varepsilon > 0$  defines the size of the RTA activation region. Since, V(x) defines a Lyapunov function for the closedloop dynamics under  $u_b$ , it is known that h(x) is increasing whenever  $u_b$  is applied. Furthermore, since  $u_b$  is activated before the boundary h(x) = 0, (22) enforces that h(x) is always positive, meaning that the system will be forward invariant in  $C_S \subset C_A$ .

#### Implicit Condition

Let  $C_{\rm b} = \{x \in \mathbb{R}^2 \mid h(x) \ge 0\}$  with h(x) = 1 - V(x) and V(x) being the Lyapunov function (11), then a condition for invariance in the safe backward image of  $C_{\rm b}$  is

$$\vartheta(x): \ [\forall t \in [0,T], \quad \varphi(e^{At}x) \ge \varepsilon_1] \land \ [h_{\mathbf{b}}(e^{AT}x) \ge \varepsilon_2]$$
(24)

where  $\varepsilon_1, \varepsilon_2 > 0$ . The safe backward image in this case defines the safe operational region for the RTA. It is depicted as the green region in Figure 4(b).

There are important tradeoffs associated with implicit and explicit approaches. While identifying invariant sets explicitly offline generally reduces the online computational burden of the algorithm, the task may become intractable for complex and high-dimensional systems. In such cases, implicit, trajectory-based approaches become favorable. The key advantage to assessing safety through trajectories online is that, in the simplest case, it only requires a finitetime simulation of the recovery maneuver, which is generally a tractable task. A secondary benefit is that it handles dynamic environments and changes in the model better; *i.e.* rather than needing to recompute an invariant subset of  $C_A$ , one only needs to update the simulation parameters. However, an important practical consideration is that trajectory-based approaches rely on the ability to accurately forecast recovery trajectories, and the predicted behavior may be increasingly sensitive to noise (*i.e.* less accurate) as the length of the backup trajectory is increased.



Figure 6: Notional depiction of the tradeoff between a longer integration horizon and a larger backup set. The thick green arcing arrow represents the desired trajectory of the primary controller, the teal region represents a backup safe set  $C_b$ , and the thin teal arcing arrows represent trajectories to the backup set. Large explicit backup sets  $C_b$ , reduce the need for online computation of trajectories to the backup set, but may be intractable for complex and high-dimensional systems. Computing online finite-time simulation of a recovery maneuver may be more tractable in some cases, but is more vulnerable to noise or model inaccuracies.

It is interesting to note that since a backup set  $C_b$  is an explicitly defined invariant subset of the constraint set  $C_A$ , an implicit approach to safety will typically involve a combination of (i) explicitly identifying an invariant set  $C_b$  offline and (ii) finding a safe trajectory to this set online. In this sense, the explicit approach can be viewed as a special case of the implicit approach where the backup trajectory consists of only a single point. Furthermore, any explicitly defined safe set can be used as part of an implicit approach. The implicit approach allows for the problem to be solved by taking full advantage of both offline and online resources. Specifically, finding a larger backup set offline reduces the online computational burden by reducing the length of the trajectory required to obtain the same operational region; likewise, integrating backup trajectories over a longer horizon reduces the need to identify large safe sets offline. An depiction of this tradeoff is illustrated in Figure 6. Generally, for simple systems, the problem can be solved offline, and as systems become more complex, one must rely more and more on online simulations.

## **Zero-Order and First-Order Methods**

Filtering approaches in RTA may be classified into zero-order and first-order methods. Zero-order algorithms are derivative free. They are typically, though not necessarily, associated with the *Simplex* architecture and choosing from a discrete set of possible control signals. For example, the RTA mechanism may choose between applying the input from the primary control law  $u_{des}$  or a backup control law  $u_b$  at any time. In this case, one may observe the effects of chatter; *i.e.* variations in the output of the RTA mechanism are not smooth with respect to variations in the initial state or desired input. This may have practical implications to real world systems. For instance, the action may cause excessive wear to the actuators, on-board components may experience interruption to their operations, or passengers may observe the intervention as being more abrupt, resulting in a less pleasant experience (*e.g.* consider an automobile that only intervenes through applying the brakes at maximum capacity). These effects may be handled with various heuristics; *e.g.* hysteresis, or interpolating between  $u_b$  and  $u_{des}$  as the state approaches the boundary of the safe operational region. Alternatively, one may turn to a first order method.

First order methods require computing derivatives, and are typically associated with methods described in the "Active Set Invariance Filtering" section. In this case, the gradient information is used in a *barrier condition*, which constrains the set of available inputs to a subset of admissible and safe inputs. An optimization framework may be used to choose the input within this set according to some cost function. A common cost specification requires that the RTA choose the control that is closest (in a 2-norm sense) to the desired input. RTA filters constructed in this way will have smooth variations in the output signal  $u_{act}$ . While this eliminates the effect of chatter, it introduces added complexity into the problem and generally requires greater computational resources. The Robotarium is an example explicit ASIF RTA implementation is discussed further in the sidebar titled "The Robotarium: An RTA Enabled Remote-Access Swarm Robotics Testbed"

## Latched and Unlatched Implementations

RTA systems based on those described in the Section "The Simplex Architecture" activate a recovery maneuver when the monitor detects that one of the boundary violation conditions has become active. At this point, the decision logic must specify when to return control to the primary controller. In an *unlatched* implementation, the decision logic releases control of the recovery controller as soon as the boundary violation condition becomes inactive. For example, an unlatched RTA implementation that switches from the primary control law  $u_{des}$  to the backup control law  $u_b$  whenever a condition  $\vartheta$  is satisfied will take the form of a hybrid system as in (22). Alternatively, a *latched* implementation will hold its activation state until a separate

	Implicit	Explicit
Simplex	spacecraft collision avoidance [27]	waypoint tracking cyber-physical system [60]
	fixed wing aircraft ground avoidance (Auto-GCAS) [61], [62]	advanced flight control systems [63]
	fixed wing aircraft obstacle avoidance [22], [23], [64]	lane keeping [65]
	safety against LTL specifications [66], [67]	aircraft conflict resolution [68]-[70]
	collision avoidance for high speed quadrotor navigation [56], [71]	decentralized collision avoidance for quadrotors [72], [73]
	collision avoidance during lane-changing maneuvers [74]	ground robot navigation around obstacles [75], [76]
ASIF	bipedal walking/ exoskeletons [77]	bipedal walking/ exoskeletons [78], [79]
	robotic manipulator arms [28]	lane keeping [14], [15]
	mobile inverted pendulum (Segway) [54]	mobile inverted pendulum (Segway) [80]
	rapid aerial exploration of unknown environments [26]	robotic grasping [81]
	multi-robot systems [82]	swarm robotics (Robotarium) [83]-[86]
	fixed wing aircraft collision avoidance [87]	motorized rehabilitative cycling system [88]

TABLE 1: Classification of RTA Applications in the Literature

release condition is met. The release condition may for example consist of following the backup trajectory for some fixed amount of time, bringing the system to a particular configuration, or having a human operator return control after assessing the problem. Auto GCAS is an example of a temporary latched system that holds the maneuver until the aircraft is on a collision free course [61]. Other implementations of latched RTAs may require a system reset before switching back, such as a computer reset supervised by a human operator. Note that this requires that the decision logic remember the current activation state.

#### Properties and Systems Engineering Considerations for Run Time Assurance Design

While it is convenient to treat safety as a binary "safe" or "unsafe" property, safety is often a complex weighting of risk among primary and secondary safety constraints. This section provides a few definitions of design considerations for incorporating RTA solutions into the larger control system. In practice, RTA system design ultimately relies on domain-specific knowledge to achieve acceptable performance. However, there are common safety and performance properties of good RTA designs. In order to keep the system safe, the RTA mechanism must select control inputs that are *innocuous* and *viable*. In addition to safety, the performance of the RTA may be evaluated by its adherence to a *nuisance free* property.

## Innocuity

Ideally the entire state of the system is contained in one vector x and the entire constraint set can be explicitly defined in a set  $C_A$ . However, in practice this is rarely the case. Instead, RTA decisions are not solely based on system dynamics, but must factor in the states of interacting subsystems, human interactions, fault management, or interlock condition management. *Innocuity*  is the property that each component of the design has no unacceptable impact on the larger system design [89]. Innocuity of RTA systems implies the selected control action does not have unintended behavior, such as causing a violation of a secondary safety constraints. An interlock is a set of two mutually exclusive states such as mechanisms that prevent elevator doors from opening when the elevator is in motion, or that prevent the elevator from moving when the doors are open. This relationship between RTA-level and system-level constraints may result in a combination of RTA system modes (described as a finite state machine), if-statements, and dynamics considerations [90].

For example, in the development of Auto GCAS, the high-level requirements in order of precedence were to do no harm, do not interfere, and prevent collisions [61], [91]. The "do no harm" and "do not interfere" requirements in Auto GCAS are examples of innocuity properties. These requirements don't directly define the RTA response, but describe special considerations to prevent unintended consequences of the RTA design. This requirement ordering might seem counter-intuitive, as one might expect preventing collisions to be the top priority of a collision avoidance system. However, this order of requirements proved critical in acceptance of the system [92]. To give a little more insight, selected generalized, conceptual requirements for Auto GCAS are presented below and divided here in terms of "do no harm," "do not interfere," and "prevent collisions." The "prevent collisions" requirements are typical of what is generally considered an RTA requirement. However, the "do no harm" and "do not interfere" requirements describe how the RTA system should operate in the context of the aircraft control system that interacts with human operators, experiences failures, and has a defined set of interlock conditions.

## Do No Harm

- The automatic recovery shall not cause harm to pilot, aircraft, or components.
- The automatic recovery maneuver shall not place the aircraft in an uncontrollable state.
- The pilot shall be able to interrupt a maneuver.
- The pilot shall be able to manually engage a maneuver.
- Subsystem monitors shall determine if a failure exists and be provided to Auto GCAS.
- Auto GCAS shall not activate an automatic recovery maneuver if a failure of a subsystem supporting Auto GCAS exists.
- Auto GCAS shall leave the failed mode state if a failure does not exist.
- The automatic recovery maneuver shall not activate during aerial refueling.
- The automatic recovery maneuver shall not activate when the aircraft has an excessively high angle of attack.
- The automatic recovery maneuver shall not activate when the aircraft velocity is too low.
- Auto GCAS shall inform the pilot if the aircraft speed is too low for the automatic recovery

maneuver.

## **Do Not Interfere**

- The automatic recovery maneuver shall not activate during landings.
- The pilot shall be able to turn the system off.
- The pilot shall be able to select the protection level.
- Auto GCAS shall notify the pilot when the automatic recovery maneuver initiates and terminates.
- The system shall record data about each automatic recovery maneuver activation.
- The Digital Terrain Elevation Data used by the system shall have a resolution and accuracy that supports safe and nuisance free operation.

## **Prevent Collisions**

- The system shall conduct an automatic recovery maneuver when the projected recovery trajectory intersects the terrain profile contour with buffers.
- The system shall terminate an automatic recovery maneuver as soon as it is determined that the aircraft will clear the terrain threat.
- The recovery maneuver shall consist of a roll to wings level and a pull up.

Systems engineering, computing hardware availability, and human-machine interaction all require practical considerations that play important roles in imparting secondary safety constraints on RTA designs. From a computing hardware standpoint, assuming availability of high performance computing may not be valid. In some domains, design constraints such as size, weight, and power coupled with vibration and radiation tolerance limit the processing and memory of computing hardware onboard the dynamical system. In these cases simpler RTA solutions may be favored over optimal solutions. An example of this occurs in aviation, where RTA designs are expected to be backwards compatible with older aircraft, and simple approaches may be to climb, descend, or turn right or left at a specified rate [93]. Additionally, when the automated backup control of the RTA systems impact human operators, humans trust an easily understood maneuver consistent with human training, expectations, and preferences [94]. For Auto GCAS, familiarity of the roll to wings level and pull maneuver and its consistency with pilot training and behavior resulted in strong positive perceptions of the system [92].

#### Viability

*Viability* ensures the selected control action does not eliminate availability of a safe action downstream, *i.e.* the RTA mechanism will always be able to identify a safe input in the future.

This concept is formalized in Figure 3, where the green area under the curve describes maximum velocity for any distance from the obstacle. One way to ensure viability in an implicit RTA design approach is to define a limited set of pre-determined recovery actions. Simple, pre-defined recovery responses can be pre-verified for use, reduce online computation burden, and make the RTA system more easily understood and trusted by a human operator [94]. In Auto GCAS, rather than optimizing a trajectory for each specific collision avoidance scenario, a predefined set of verified maneuvers was created. In F-16 Auto GCAS, a single roll to wings level and 5g pull maneuver is conducted [62]. In the Automatic Air Collision Avoidance System (Auto ACAS), nine different maneuvers are considered [95]. In patented [96] and experimental Auto GCAS designs for small UAVs [97], lesser capability aircraft [98], and cargo-class aircraft [99], three or more possible maneuvers may be considered. In the case of multiple maneuvers, one popular approach that combines viability with nuisance freedom is to engage the last available maneuver right before it is too late.

#### Nuisance Freedom

In a *nuisance-free* RTA system, the constraint set  $C_A$  and corresponding operating envelope  $C_{\rm S}$  of the primary controller is as large as feasible. Other ways to describe this property is that the RTA should be minimally invasive and should not interfere with the primary function of control system  $u_{des}$  unless absolutely necessary. In Auto GCAS, a collision avoidance maneuver is considered nuisance-free if it occurs after an aware pilot would have maneuvered to avoid a collision. To formally define nuisance-free operations, a time available metric described zero time available as the point where an activation would just barely scrape the ground, and increasingly positive time available would result in maneuvers at greater altitudes above the terrain. To assign a value to an acceptable time available activation range, a 1995 study measured when pilots felt an aggressive recovery should be activated to avoid a collision [61], [91]. Pilots flew towards the ground at a variety of dive angles, bank angles, air speeds, and load factors and activated a recovery maneuver when his or her comfort threshold was met. After each run, the pilots rated the timing of the recovery initiation, their anxiety level during the maneuver, and the precision/aggressiveness of the maneuver for each run. Based on this information a 1.5 second time available design criteria was identified, and collision avoidance maneuvers that activated with less than 1.5 seconds time available were considered nuisance-free [61], [91].

More generally, the number of RTA activations, magnitude of the response, activation timing, and size of the safe set can be use to compare nuisance criteria across different RTAs.

• In many scenarios, the RTA should not come on often because activations interrupt the primary mission. An *RTA activations* metric tracks the number of seconds or discrete time

steps that a recovery maneuver is in control. This metric allows comparing different RTA approaches as it provides a measure of relative conservatism of various RTA solutions in the presence of the same, possibly faulty primary controller.

- The RTA should ideally not impart a large change in control, measurable via a *control magnitude* metric.
- An *RTA invasiveness* metric describes whether the RTA activates appropriately when near the safety constraint boundary, while refraining from activating when far from the allowable boundary. A variety of application-specific measurements of this may be applied, such as the time available metric in Auto GCAS.
- An *RTA safe set size* measures of the volume of the safe set  $C_S$ . The larger the volume, the more space the primary controller has to operate optimally, and the less invasive the RTA is. Different RTA approaches may be more or less conservative, which impacts the size of  $C_S$ .

## Run Time Assurance Integrity Monitoring

Recognizing and accommodating failures that could impact an RTA system is a critical element of RTA design. The failures could come from inside the RTA system or from externally provided information that the RTA uses to make a response decision. Possible software checks include monitoring for missing "heartbeat" signals from components and "reasonableness checks" that ensure incoming data are within a reasonable range of values [97]. For example, a failure of the inertial navigation system should not allow Auto GCAS to roll inverted and pull down into the ground instead of pulling up to avoid a collision [61]. Auto GCAS uses System Wide Integrity Management (SWIM) hardware and software tests used to mitigate single-point-of-failure risks [61], [95], [100].

## Human Interaction with Run Time Assurance Systems

When designing RTA systems, it is important to consider the interaction of the system with human operators. While the topic of human-machine interaction can swell quickly, this article focuses on a few key design considerations in RTA.

## Variable Risk Tolerance and an "Off Switch"

One of the challenges of developing an RTA system is the need to balance nuisancefreedom with the risk tolerance of a particular task or operator. Adding the ability to tune risk corresponding to safety buffers or other factors facilitates flexibility and acceptability of the RTA in operational use. For example, Auto GCAS features two pilot selectable modes that enable variable risk tolerance: a "norm" mode with adequate safety buffers for most cases, and a "min" mode which minimizes buffer size for nuisance-free low-level flying at the trade-off of reduced protection [61], [91]. In addition, the ability to turn the system off was an important feature that engendered pilot trust [92].

## Transparency and Trust

The primary basis of trust in automation comes from dependability and predictability [101]. However additional factors with a measurable impact on human trust include the past performance (*e.g.* system pedigree or heritage), simplified and understandable performance, system intent, and reliability [94]. In Auto GCAS transparency comes through pilot training before operations and through visual indicators on the aircraft's head-up-display (HUD) [92], [102]. A discussion of the difference between safety, security, and reliability is provided in the sidebar titled "Safety, Reliability, and Security." Reliability is important for human trust in RTA systems as humans initially assume that machines are perfect and any errors rapidly deteriorate trust [103]. Knowledge of Auto GCAS's 98% reliability increased test pilot trust of the system [92].

#### Missed Detections and False Alarms

It is important to consider an acceptable rate of missed detections and false alarms in the design of an RTA system. *Missed detections* are a failure of the monitor to detect an imminent safety violation. For example, if Auto GCAS fails to detect an imminent collision with terrain, it could lead to loss of life and the aircraft. *False alarms* occur when the RTA monitor believes the system is entering an unsafe state, when in reality no safety violation is imminent. False alarms erode confidence in the RTA system. For example, during psychological study on the development and deployment of Auto GCAS, one pilot indicated that a single false fly-up (automatically maneuvering to avoid the ground when there was no imminent collision) would likely cause pilots to turn the system off and lose the protection it provided [61], [91]. Developing pilot trust in Auto GCAS hinged on its nuisance avoidance criteria, *i.e.* keeping the number of false alarms very low [102].

In statistical hypothesis testing, missed detections and false alarms are Type I and Type II errors. Hypothesis testing is a formal technique to verify that a system meets a specification or requirement [104]. A hypothesis can have two possible outcomes: a null hypothesis  $H_0$ , such as a collision is not imminent, and an alternative hypothesis  $H_1$ , such as a collision is imminent. As summarized in Table 2, a Type I error is a false alarm (*i.e.* activate  $u_b$  when not near the constraint boundary  $\partial C_A$ ), while a Type II error is a missed detection (*i.e.* failing to activate

 $u_{\rm b}$  near  $\partial C_{\rm A}$ , resulting in a departure from the safe set. When considering a collision avoidance system for example, a Type I error occurs when a collision avoidance system maneuvers to avoid a collision that would not have happened, and a Type II error occurs when a collision avoidance system does not maneuver to prevent a collision.

	$H_0$	$H_1$
	Collision	Collision
	Not Imminent	Imminent
Do Not Reject $H_0$		Type II Error
	Correct	False Negative
System Decides		Missed Detection
Not Imminent		$\beta$ probability
Reject H <sub>0</sub>	Type I Error	
	False Positive	Correct
System Decides	False Alarm	
Collision Imminent	$\alpha$ Probability	

TABLE 2: False Alarms and Missed Detections Type I and Type II Errors, with Collision Avoidance as an Example.

In some systems, like Auto GCAS, the acceptable rates of false alarms  $\alpha$  are very low because of the consequence of an unnecessary interruption of the system mission is high. In other systems like elevators, the acceptable rate of a missed detection  $\beta$  (not detecting a human is in the way of the closing doors) would be very low, and false alarms would be more acceptable (there is almost always an option to take the stairs). During the development of RTA systems, studies are required to determine acceptable  $\alpha$  and  $\beta$  rates.

# The Simplex Architecture

Simplex Architecture RTA designs [60], [105]–[107] feature a *monitor* or watchdog that watches for imminent violations of safety constraints, and a backup *safety controller* that provides a guaranteed safe control output to remedy the unsafe situation. A critical feature here is that the backup control is applied via a switch, rather than gradually as in ASIF described next. One way to describe Simplex RTAs is as a hybrid dynamical system [38] described by Eq. (22). The monitor and safety controller are inserted before the plant as demonstrated in Fig. 7. In this model, the following functional components describe specific functions and interactions that may be implemented on one or more physical components:



Figure 7: Simplex Architecture with a physical plant, a verified safety controller, a verified decision module and switch, and an unverified complex controller.

- *Plant*: This functional component is the system under control and could be a spacecraft, aircraft, boat, car, or other robotic system.
- *Primary controller*: This functional component is the primary, high performance controller of the system that may not be fully verified with traditional offline verification techniques. Depending on the application this could be a human pilot, neural network-based control, or complex control software.
- *Backup Controller*: This functional component continuously computes a safe control response to the state of the plant. For example, in Auto GCAS the safety controller is a predefined roll and pull maneuver. One school of thought places heavy verification burden on the safety controller, boundary monitor, and decision logic as a way to certify advanced controllers that cannot be verified to the current certification standards [108]. Another school of thought requires no assumptions on the safety or verification of the backup controller. That is, safety can be guaranteed through online integration of the backup controller, and the system may revert at any time to the most recently computed safe trajectory. Formal guarantees are obtainable in both cases.
- *Boundary (a.k.a. safety) monitor*: This functional component monitors the state of the plant and the nominal controller for predetermined safety boundary violations. Examples of boundary violations might be a trajectory prediction that intersects the terrain (indicating an imminent ground collision) or a control input that would cause excessive acceleration (risking damage to the structure or components).
- *Decision Logic*: The functional component takes inputs from all the other components pictured and determines which control output, complex or safety, to output to the plant.

## Variations on Simplex

Variations on Simplex include *nested RTA*, *multi-recovery RTA*, and *multi-monitor RTA* designs. *Nested RTA* designs implement safety controls at each layer of the control architecture as shown in Figure 8. *Multi-recovery RTA* designs have multiple reversionary controllers at the



Figure 8: Nested feedback RTA architecture for an Unmanned Aerial System (UAS) plant [109]. In a nested RTA architecture, the inner-loop control law RTA ensures aircraft stability, the guidance control law RTA ensures safety of commands generated to follow waypoints, the flight management system RTA checks that the waypoints along the path are safe, and the mission planning system RTA ensures the safety of task allocations to meet mission goals of the plant. In each RTA layer, a monitor watches for unsafe conditions and decision logic determines switching between the primary controller and backup controller.

same level of the control architecture and a switch that considers all possible recovery function options. The concept is discussed in a recent ASTM standard for the use of RTA in unmanned aircraft [108] and depicted in Figure 9. The best approach to designing multiple recoveries is



Figure 9: ASTM F3269 RTA architecture with multiple recovery control functions at the same level in the control hierarchy [108]. An example of this might be an aircraft that avoids collisions with other aircraft, avoids collisions with the ground, and avoids flying into no-fly zones. Each of these boundaries are disjoint; however, it is possible that violations of multiple boundaries could be predicted to occur at the same time.

an area of active research. One approach is to combine multiple RTA systems as an integrated solution. For example, the Automatic Integrated Collision Avoidance System (Auto ICAS) [110], [111] integrates ground and midair collision avoidance into a comprehensive collision avoidance system. The advantage of this integrated approach is that each system can consider possible conflicts between boundary violations and select a better solution than either might select on their own. For instance, in a situation where an aircraft flying close to the ground is on a collision course with another aircraft, a ground-aware midair collision avoidance system will filter from the possible maneuvers to select a maneuver that avoids the midair collision while at the same time ensuring none of those maneuvers will fly the aircraft into the ground. An integrated solution has the benefit of generating a better solution than either separate system might otherwise choose. However, this increase in performance also comes at a large increase in cost and schedule. An alternative multiple recovery controller integration approach is to use an RTA network architecture where each boundary monitor and recovery controller function is separate [112]. An example of an architecture with multiple safety monitors and recovery functions is the Expandable Variable Autonomy Architecture (EVAA) RTA Network [112]. On the one hand, this modular approach eases verification of each individual component and facilitates quick integration of new components and upgrades. On the other hand, additional stress is placed on the verification of the decision module/switching component to determine which action to take in the case where multiple safety boundaries may be violated at the same time. The EVVA [112] answer to this challenge is to approach the decision like a human pilot, who responds to the most critical boundary violation first before moving onto the next violation. A disadvantage of the modular approach is that it is possible to violate a safety boundary. For instance, if a ground collision and geofence collision were both imminent, the system might decide to engage ground collision avoidance to first clear the threat of ground collision before engaging a geofence controller and violating the geofence boundary. Depending on the scenario, this may or may not be acceptable. An integrated solution could adjust the boundary of a geofence controller based on awareness of ground and midair collision avoidance, or could select a collision avoidance maneuver that also kept the aircraft within the geofence. Multi-monitor RTA designs feature checks and balances to ensure integrity of the RTA system. These designs may include monitors beyond a boundary or safety monitor. These monitors may include a *failure monitor*, an *interlock* monitor, and a human supervisor, among others [90]. The failure monitor watches for unreliable information coming into the RTA such as a failed or stale sensor signal and prevents the RTA from acting on incorrect state information. The interlock monitor watches for conditions where it may be safe to compute a backup control solution, but unsafe to engage a maneuver. For example, an interlock condition might occur when a robotic system is physically connected to another system and rapid maneuver may result in damage to one or both robots. Finally, there are many circumstances where a cyber-physical system may be acting with a human teammate

or under supervision of a human. In these cases, a human supervisor may have roles such as turning the RTA off, changing the RTA risk or sensitivity level, or manually activating the RTA recovery control function.

## Simplex RTA Safety Monitor Approaches

As discussed previously, approaches may be considered implicit or explicit based on whether they rely on simulating backup trajectories online. Approaches to the design of a Simplex-based RTA safety monitor tend to fall in categories of *boundary violation* (i.e. explicit), *trajectory prediction* (i.e. implicit), or *reach-tube prediction* (i.e. nondeterministic). Boundary violation monitors activate a recovery controller after a safety boundary has been violated. These designs often feature a buffer to allow a recovery response time to complete before violating the actual safety boundary. Trajectory prediction monitors predict the trajectory that the recovery



Figure 10: The RTA "Ameoba" diagram depicts the entire state space X, containing the set of allowable states  $C_A$  in purple, a safety buffer region  $C_{S2}$ , and a safe zone with without switching  $C_{S1}$ . The trajectory of a primary controller is depicted in red. When the trajectory crosses the switching condition boundary between  $C_{S1}$  and  $C_{S2}$ , the RTA activates the backup controller and the green trajectory shows backup control to a safe region  $C_b$ . Image inspired by [109].

would take if it were to activate and if that trajectory prediction intersects the safe boundary with a small buffer, the recovery maneuver will activate [113], [114]. Reach-tube-based monitors study a set of possible future states, arising from uncertainty in the dynamics, and in instances where this prediction intersects the unsafe set the RTA recovery function will activate to ensure safety [115], [116].

#### Simplex RTA as a Near-Term Certification Path

Within the last five years progress has been made towards developing certification criteria for RTA [117] based on the Simplex architecture resulting in the publication of the first ASTM standard for the use of RTA in unmanned aircraft [108]. A second iteration [118] of this standard opens the door for alternative filtering RTA designs.

#### **Provably Safe RTA with Black Box Backup Controllers**

Implicit RTA systems rely on online simulations of a backup controller in order to assess safety. An important consequence of this architectural choice is that the backup controller need not be guaranteed to always produce a safe solution in order to guarantee safety of the RTA system. That is, in the event that the backup controller fails to compute a new safe trajectory at any time, it can follow the most recently found safe trajectory until it reaches the backup set. In this sense, the job of the backup controller is to search for and *propose* candidate safe trajectories. A monitor can be used to determine whether this candidate trajectory is both feasible and safe. A practical consequence of this is that high performing black-box controllers can safely be used to generate backup trajectories. This concept has been used in practice [18] and is thoroughly explored for the context of Simplex-based systems in [119]. Furthermore, the idea extends to other methods of RTA. For example, in [54] an ASIF is constructed which utilizes a backup controller taking the form of a neural network approximation of an optimal control policy.

## **Canonical Algorithms**

Included in this section are a set of canonical algorithms for Simplex-based RTA. In this case, we consider the deterministic discrete time system,

$$x(i+1) = F(x(i), u(i))$$
(25)

where  $x \in X \subseteq \mathbb{R}^n$  denotes the state,  $u \in U \subset \mathbb{R}^m$  denotes the control input and  $i \in \mathbb{Z}$  denotes the time index. A constraint set is defined as  $\mathcal{C}_A \subseteq X$ . It is assumed that a backup control law  $u_b: X \to U$  is defined and we let  $\phi^{u_b}(i; x)$  be the state of (25) reached at  $i \ge 0$  when beginning at state x at time i = 0 and evolving under  $u_b$ .

The idea behind the algorithms in this section is to assess the safety of a *probe step*. That is, at any state  $x_{curr} \in X$ , whether to accept or reject the desired control input  $u_{des} \in U$  is based on whether the candidate next state reached under this input  $x_{cand} := F(x_{curr}, u_{des})$  is safe. In the first case, an explicit safe set is assumed to be defined. The second case considers an implicit approach, which relies on online simulations under the backup dynamics. A deterministic discrete time system is assumed for clarity of the algorithms and their associated guarantees. However, the

Algorithm T Region-Dased Shiplex Ther (RDSI)		
input : Current State $x_{curr} \in X$		
: Desired Input $u_{des} \in U$		
<b>output</b> : Safe Control Input $u_{act} \in U$		
<b>predefined</b> : Constraint set $C_A \subset X$		
: Invariant Set $\mathcal{C}_{\mathrm{S}} \subset \mathcal{C}_{\mathrm{A}}$		
: Backup Control Law $u_{\rm b}: X \to U$ .		
1: function $u_{\text{act}} = \text{RBSF}(x_{\text{curr}}, u_{\text{des}})$		
2: $x_{\text{cand}} \leftarrow F(x_{\text{curr}}, u_{\text{des}})$		
3: <b>if</b> $x_{\text{cand}} \in \mathcal{C}_{\text{S}}$ <b>then</b>		
4: return $u_{des}$		
5: else		
6: return $u_{\rm b}$		

Algorithm 1 Region-Based Simpley Filter (RBSE)

algorithms may be adapted to continuous time systems, to multiple backup controllers, latched implementations, etc.

#### Explicit Simplex Filter

Algorithm 1 describes an explicit Simplex-based RTA algorithm, the Region-Based Simplex Filter (RBSF). In addition to a backup control law  $u_{\rm b}$ , it is assumed that a safe set  $C_{\rm S} \subseteq C_{\rm A}$ is defined such that  $C_{\rm S}$  is invariant under  $u_{\rm b}$ . The algorithm renders the RTA system forward invariant in  $C_{\rm S}$ . This is apparent through the observations that (i) by nature of the invariance property, applying  $u_{\rm b}$  anywhere in  $C_{\rm S}$  will lead to a next state that is still in the set (ii) the sample step allows for inputs that would cause a departure from  $C_{\rm S}$  to be detected and replaced with  $u_{\rm b}$ . In practice, one may choose to add conservatism to the approach by working with any underapproximation  $\tilde{C}_{\rm S} \subseteq C_{\rm S}$ . The region  $C_{\rm S} \setminus \tilde{C}_{\rm S}$  is often referred to as a *buffer*. A practical consideration is that the backup controller should locally attract solutions outside of  $\tilde{C}_{\rm S}$  to  $\tilde{C}_{\rm S}$ . That is, it is favorable to ensure that a recovery is possible in the case where the state is perturbed outside of the safe region. One way to do this is to ensure that all initial conditions in the buffer will lead to  $\tilde{C}_{\rm S}$  if  $u_{\rm b}$  is applied; *e.g.* this is the case when  $C_{\rm S}$  and  $\tilde{C}_{\rm S}$  are both sublevel sets of a Lyapunov function.

*Example:* Consider the double integrator system described by (6) discretized over a period of 0.1 s, and consider the constraint set (7). The backup control law  $u_{\text{des}} = -1$  renders the set  $C_{\text{S}} = \{x \in \mathbb{R}^2 \mid -2x_1 - x_2^2 \ge 0\}$  forward invariant and it is clear that  $C_{\text{S}} \subseteq C_{\text{A}}$ . Figure

15(a) shows a simulation of the RBSF algorithm with these parameters under the desired control  $u_{\text{des}} = 1$ .

#### Implicit Simplex Filter

Algorithm 2 describes an implicit Simplex-based algorithm, the Simulation-based Simplex Filter (SBSF). It is assumed that a backup set  $C_b \subseteq C_A$  is defined such that  $C_b$  is invariant under a backup control law  $u_b : X \to U$ . No assumptions are made on the structure of  $u_b$ , it may come from a closed form control expression, a black box function, an MPC, a motion primitive, etc. In contrast to the RBSF algorithm, which constrains the system to evolve in the explicitly defined invariant region, SBSF uses the smaller invariant set  $C_b$  as a means to access a larger implicitly defined invariant set. Specifically, SBSF constrains the system to the safe backward image of  $C_b$ :

$$\mathcal{C}_{\rm S} = \{ x \in X \mid \phi^{u_{\rm b}}(i;x) \in \mathcal{C}_{\rm A} \; \forall i \in \{0,...,N-1\}, \; \phi^{u_{\rm b}}(N;x) \in \mathcal{C}_{\rm b} \}$$
(26)

where,  $\phi^{u_{\rm b}}(i;x)$  is the *i*<sup>th</sup> point in an N + 1 point trajectory obtained from simulating under the dynamics (25) and backup controller  $u_{\rm b}$ . Since  $C_{\rm S}$  is forward invariant under  $u_{\rm b}$ , the algorithm renders the RTA system forward invariant in  $C_{\rm S}$ . As in the case of the RBSF algorithm, this becomes apparent from the observations that (i) applying  $u_{\rm b}$  anywhere in  $C_{\rm S}$  will lead to a state that is also in  $C_{\rm S}$  and (ii) the sample step under  $u_{\rm des}$  can reliably be used to detect inputs that would cause a departure from  $C_{\rm S}$ . Conservatism can be added to this approach by using underapproximations of  $C_{\rm A}$  and  $C_{\rm b}$ . Note that the above algorithm assumes a deterministic system and relies on simulating single trajectories under the backup dynamics. This can be extended to the nondeterministic case by simulating an overapproximation of the forward reachable set under the backup dynamics and requiring that  $C_{\rm b}$  is robustly forward invariant.

*Example:* Consider the double integrator system described by (6) discretized over a period of 0.1 s, and consider the constraint set (7). The backup control law  $u_{des} = -1$  renders the set  $C_{b} = \{x \in \mathbb{R}^{2} \mid -x_{1} \geq 0, -x_{2} \geq 0\}$  forward invariant and it is clear that  $C_{b} \subseteq C_{A}$ . Figure 16(a) shows a simulation of the SBSF algorithm with these parameters under the desired control  $u_{des} = 1$ .

*Example:* (From [58]) Consider the system (21) discretized in time to have a 1 s resolution, and with  $C_A = \{x \in \mathbb{R}^5 | ||r||_{\infty} - r_{\min} \ge 0, \kappa_1 + \kappa_2 ||r||_{\infty} - ||v||_{\infty} \ge 0\}$  where  $r := [x_1, x_2]^T$ , and  $v := [x_3, x_4]^T$  are the relative position and speed vectors respectively. The backup controller  $u_b$  is taken as the first input in the solution to a mixed integer quadratic program (MIQP); solved subject to the constraints that (i) all intermediate points lie in  $C_A$  and (ii) the endpoint lies in a set  $C_b$ , that is defined from the set of natural motion trajectories lying outside

Algorithm 2 Simulation-Based Simplex Filter (SBSF)				
input :	Current State $x_{\text{curr}} \in \mathbb{R}^n$			
:	Desired Input $u_{\text{des}} \in U$			
output :	Safe Control Input $u_{act} \in U$			
<b>predefined</b> : Constraint Set $C_A \subset X$				
:	Invariant Backup Set $\mathcal{C}_{\mathrm{b}} \subseteq \mathcal{C}_{\mathrm{A}}$			
:	Backup Control Law $u_{\rm b}: X \to U$ .			
1: function $u_{act} = SBSF(x_{curr}, u_{des})$ 2: $x_{curr} \leftarrow F(x_{curr}, u_{des})$				
3: <b>compute:</b> $\phi^{u_b}(i; x_{cand})  \forall i \in \{0, \dots, N\}$				
4: <b>if</b> $\phi^{u_{\mathrm{b}}}(i;x)$	$(x_{\text{cand}}) \in \mathcal{C}_{A}  \forall i \in \{0, \dots, N-1\} \text{ AND } \phi^{u_{b}}(N; x_{\text{cand}}) \in \mathcal{C}_{b} \text{ then }$			
5: return	n $u_{ m des}$			
6: else				
7: retur	<b>n</b> $u_{\rm b}$			

of  $C_A$ . In this case  $x \in C_S$  is true whenever the MIQP has a feasible solution with initial state x. Figure 11 shows a simulation of the RTA system with a primary controller that is designed to drive the system to an invariant point along the y-axis. The simulation uses parameters m = 50 kg, n = 0.001027 rad/s, and  $u_{\text{max}} = 0.5$  N,  $r_{\text{min}} = 0.5$  km,  $\kappa_1 = 0.5$  m/s, and  $\kappa_2 = 2 n \text{ s}^{-1}$ .

# Active Set Invariance Filtering

Active set invariance filtering (ASIF) methods are RTA approaches associated with a class of first-order, optimization-based algorithms. As with many Simplex-based techniques they are system agnostic, provably correct, and exhibit a number of robustness properties which make them attractive in practice. In addition, ASIFs are minimally invasive with respect to the safety constraints. This results in a smoother and more gradual intervention than with approaches that rely on switching to backup controllers directly.

This section considers systems of the form

$$\dot{x} = f(x) + g(x)u. \tag{27}$$



Figure 11: Projected trajectories of a chaser spacecraft relative to a target spacecraft at the origin under Clohessy-Wiltshire-Hill dynamics, described by (21) over a 2400 s period. For safety, the chaser spacecraft is required to stay at least 0.5 km away from the target in both the x and y directions (shown by dashed red lines on the left), and is required to stay below a dynamic velocity limit (shown by dashed red lines in the figure on the right) where the spacecraft's relative velocity must slow down as it gets closer. The chaser spacecraft trajectories without RTA are grey. The trajectories with a simulation-based Simplex filter RTA are blue when  $u_{act} = u_{des}$ , and red when  $u_{act} = u_{b}$ . The backup trajectories (cyan) are shown at 8 s intervals.

When (1) is of the form (27), the dynamics are said to be *control affine*, and it is assumed throughout the following that  $f(x) : \mathbb{R}^n \to \mathbb{R}^n$  and  $g(x) : \mathbb{R}^n \to \mathbb{R}^{n \times m}$  are locally Lipshitz continuous in their inputs so that, in particular, solutions to (27) are unique when they exist. ASIFs are constructed as quadratic programs, where the objective function is typically the  $l^2$ norm of the difference between the desired input  $u_{des}$  and actual input  $u_{act}$ , and where the constraints on the program, which are known as *barrier constraints*, take the form  $BC_i(x, u) =$  $a_i(x)u + b_i(x) \ge 0$ ,  $i \in \{1, ..., N\}$ . The canonical ASIF controller is given below as ASIF-QP:

# **ASIF-QP**

$$u_{\text{act}}(x, u_{\text{des}}) = \operatorname{argmin}_{u \in U} \|u - u_{\text{des}}\|_2^2$$
(28)

s.t. 
$$BC_i(x, u) = a_i(x)u + b_i(x) \ge 0, \qquad i = 1, ..., N$$
 (29)
The ASIF-QP above guarantees system safety with respect to an operational region  $C_S$  when the properties  $\vartheta_1^{BC}$  and  $\vartheta_2^{BC}$  are satisfied by the barrier constraints:

•  $\vartheta_1^{BC}$ : if (29) is satisfied then the system is forward invariant in some set  $C_S \subseteq C_A$ . *i.e.* 

$$\vartheta_1^{\text{BC}} : BC_i(x, u) \ge 0, \ \forall i \in \{1, ..., N\} \implies \mathcal{C}_{\text{S}} \text{ is forward invariant}$$
(30)

where  $C_{S} \subseteq C_{A}$ .

•  $\vartheta_2^{BC}$ : it is possible to satisfy (29) from any state in the set  $C_S$ . *i.e.*  $\forall x \in C_S, \forall i \in \{1, \dots, N\}$ ,

$$\vartheta_2^{\mathrm{BC}} : \sup_{u \in U} [BC_i(x, u)] \ge 0 \tag{31}$$

The first property ensures that the RTA mechanism will bound trajectories to the set  $C_S$ , as long as there exists a feasible input satisfying the barrier constraint. The second property ensures that for all states in  $C_S$ , it is possible to find a feasible input that satisfies the barrier constraint. Since quadratic programs can be solved in such a way that feasible solutions are found whenever they exist, guaranteeing the existence of a solution is sufficient for ensuring feasibility of the ASIF-QP. When both properties are satisfied,  $C_S$  defines a safe set with respect to the ASIF-QP. Moreover when the barrier constraints satisfying the properties above, the ASIF-QP is *minimally invasive* in the sense that the  $l^2$  distance between  $u_{act}$  and  $u_{des}$  is minimized, subject to the barrier constraints.

As with the case of Simplex, implicit and explicit approaches emerge for ASIF based on how the safe operational region  $C_S$  is identified. First, *explicit* ASIFs are studied. In this case,  $C_S$  obtained by identifying a control invariant set explicitly as the super zero level set of a smooth function. Next, in the case of *implicit* ASIFs, a control invariant set is identified implicitly through the trajectories of a backup control law. A more thorough review of explicit ASIF approaches is provided in [120] and implicit ASIF approaches in [53]. See also [49] for a review of both implicit and explicit ASIF approaches.

## Barrier Constraints from an Explicitly Defined Safe Set

Consider a set  $C_S$  that is defined as the super zero-level set of a continuously differentiable function  $h : \mathbb{R}^n \to \mathbb{R}$ , *i.e.*,

$$\mathcal{C}_{\mathrm{S}} = \{ x \in \mathbb{R}^n \,|\, h(x) \ge 0 \}$$
(32)

$$\partial \mathcal{C}_{\mathrm{S}} = \{ x \in \mathcal{C}_{\mathrm{S}} \, | \, h(x) = 0 \}$$
(33)

and suppose that 0 is a regular value of h; *i.e.* the gradient of h does not vanish on the boundary. The goal is to design a barrier constraint BC(x, u) that satisfies  $\vartheta_1^{BC}$  and  $\vartheta_2^{BC}$ , and consequently makes the set  $C_S$  forward invariant under the ASIF-QP. One method for ensuring the forward invariance of  $C_S$  involves checking subtangentality condition over the boundary of  $C_S$ . Sometimes called Nagumo's condition, this result is as follows: if

$$L_f h(x) + L_g h(x)u(x) \ge 0 \tag{34}$$

holds for all  $x \in \partial C_S$ , then  $C_S$  is forward invariant for the closed-loop dynamics of (27) under u(x), where  $L_f$ ,  $L_g$  in (34) denote Lie derivatives of h along f and g, respectively. Informally, (34) ensures that  $\dot{h}(x) \ge 0$  for all  $x \in \partial C_S$  and, thus, any control law u(x) satisfying (34) ensures the forward invariance of  $C_S$  when applied to (27).



Figure 12: Barrier geometry for constraint function  $h : \mathbb{R}^2 \to \mathbb{R}$ .

The constraint (34) is not very practical for use directly in an optimization framework as it is activated exactly on the boundary of  $C_S$ , which is a region without volume. The solution presented in [121] is to use similar condition that is active everywhere in  $C_S$ , and that includes a strengthening term which relaxes the constraint away from the boundary: for all  $x \in C_S$ 

$$L_f h(x) + L_g h(x)u(x) + \alpha(h(x)) \ge 0$$
(35)

where  $\alpha : \mathbb{R} \to \mathbb{R}$  is an extended class  $\kappa_{\infty}$  function. A continuous function  $\alpha : \mathbb{R} \to \mathbb{R}$  is class- $\kappa_{\infty}$  if  $\alpha$  is strictly increasing and  $\alpha(0) = 0$ . Since  $\alpha(h(x)) = 0$  on the boundary, satisfaction of (35) implies satisfaction of (34). Hence, (35) implies  $C_{\rm S}$  is forward invariant for the closed-loop dynamics of (27) under u(x). For this reason, an attractive method for constructing the ASIFs (28)–(29) is to form barrier constraints using

$$BC(x,u) = L_f h(x) + L_g h(x)u + \alpha(h(x)).$$
(36)

Applying the previous theoretical results, the resulting ASIF ensures the forward invariance of  $C_S$  and therefore satisfies  $\vartheta_1^{BC}$ , as discussed above.

In order to satisfy the requirement  $\vartheta_2^{BC}$ , the functions h,  $\alpha$  must be chosen in such a way that ensures that there always exists  $u \in U$  satisfying the barrier constraint; equivalently,

$$\sup_{u \in U} [L_f h(x) + L_g h(x)u] \ge -\alpha(h(x)).$$
(37)

must hold for all  $x \in C_S$ . Though it is on occasion possible to obtain h,  $\alpha$  directly through intuitive reasoning, satisfying this requirement is accomplished in general through the computation of a *control invariant* subset of  $C_A$ . Specifically, if  $C_S$  is a control invariant set, then it is always possible to find a strengthening function  $\alpha(x)$  such that the viability property is satisfied [80]. In practice,  $\alpha$  often takes the form of a linear or odd polynomial function, and can to an extent be chosen in a way that shapes the response of the RTA mechanism to have some desired behavior. For example, a common trade-off that occurs when shaping this function is in choosing between a more gradual response that activates further inside the safe set, and a more abrupt intervention activating closer to the boundary of  $C_S$ . Alternatively, one may bypass choosing  $\alpha$  altogether with a relaxed QP formulation (see [80], [49] for more details). Identifying a large control invariant subset  $C_S \subseteq C_A$  is the key challenge in implementing this method.

*Example:* Consider the double integrator system described by (6) and the constraint set (7). A control invariant set is given by the super zero level set of  $h(x) = -2x_1 - x_2^2$  and given a class  $\kappa_{\infty}$  function  $\alpha$ , the barrier constraint is  $-2x_2(1+u) + \alpha(-2x_1 - x_2^2) \ge 0$ . Note that  $\forall x \in C_S$ ,  $\alpha(h(x)) \ge 0$  and  $\sup_{u \in U} [L_f h(x) + L_g h(x)u] = 0$ , hence the barrier constraint is viable.

Example: Consider the dynamics of a spacecraft in unconstrained rotation,

$$\dot{x} = J^{-1}(-x \times Jx) + J^{-1}u \tag{38}$$

where  $x \in \mathbb{R}^3$  represents the angular velocity,  $J = \text{diag}(J_1, J_2, J_3) \in \mathbb{R}^{3 \times 3}$  is a diagonal inertia matrix and  $u \in [-1, 1]^3$ . A control invariant set for this system is given by  $C_S = \{x \in \mathbb{R}^3 \mid h(x) \ge 0\}$  with,

$$h(x) = K - x^T J x \tag{39}$$

for any  $K \in \mathbb{R}_+$ . The fact that this set is control invariant is made apparent with the observation that the rotational kinetic energy  $x^T J x$  is constant when u = 0. Choosing  $\alpha(x) = x$  the barrier constraint is  $\nabla h^T(x)\dot{x} + h(x) \ge 0$ , which simplifies to,

$$K - x^T J x - 2x^T u \ge 0 \tag{40}$$

It is easy to show that this constraint is viable everywhere within  $C_S$ . Namely,  $K - x^T J x \ge 0$ everywhere in  $C_S$ , and one can always find  $u \in U$  such that  $-2x^T u \ge 0$ .

 $\Diamond$ 

A	lgorithm	3	Explic	cit Ac	tive	Set	Invari	iance	Filte	r
---	----------	---	--------	--------	------	-----	--------	-------	-------	---

input :	Current State $x_{curr} \in X$
:	Desired Input $u_{des} \in U$
output :	Safe Control Input $u_{\text{act}} \in U$
predefined:	Allowable Set $\mathcal{C}_A \subset X$
:	Invariant Set $\mathcal{C}_{\mathrm{S}} \subset \mathcal{C}_{\mathrm{A}}$
:	Extended Class $\kappa_{\infty}$ Function $\alpha : \mathbb{R} \to \mathbb{R}$

- 1: function  $u_{\text{act}} = \text{EASIF}(x_{\text{curr}}, u_{\text{des}})$
- 2: solve: (28) subject to the constraint (36).
- 3: return  $u_{\text{act}}(x_{\text{curr}}, u_{\text{des}})$ .

If there exists an extended class  $\kappa_{\infty}$  function satisfying (37) for all  $x \in X$  then h(x) is said to be a *control barrier function* [120], [122], [123]. Barrier functions were first developed in the context of nonlinear and hybrid system verification [43], [124], [125]. The idea was first extended to the RTA problem with quadratic programs in [126], and it was developed more in [122]. In recent years, there has been a surge in interest in this topic. The theory has been extended to relax assumptions on smoothness [127], to discrete time [128] stochastic [129] and nondeterministic systems, and to enforce temporal logic specifications [130] [131] [132]. Likewise, the practicality of the approach is demonstrated in a variety of real world applications, as shown in Table 1. It is important to note that in order for (36) to depend on the control input u, the constraint function h(x) must have a relative degree of one; *i.e.*  $L_g h(x) \neq 0$ . Safety constraint functions with higher relative degrees are addressed in the literature on exponential control barrier functions [133], and high order control barrier functions [134], [135].

*Example:* Consider the system  $\ddot{x} = u$  with  $u \in U = (-\infty, \infty)$  and the constraint function  $h(x) = -x_1$ . Note that while actuation is unlimited in this case, it is not possible to enforce the barrier constraint  $L_f h(x) + L_g h(x) u \ge \alpha(h(x))$  as  $L_g h(x) = 0$ .

## Barrier Constraints from an Implicitly Defined Safe Set

As discussed in the previous section, the barrier constraint is viable when it restricts the system to a control invariant set. If an explicit representation of a large control invariant subset of  $C_A$  can be obtained, then with the appropriate choice of  $\alpha$ , a valid barrier constraint is obtained through (36). The key idea behind the implicit approach to ASIF [53], [54], [80] is to filter with respect to an implicitly defined control invariant set. Specifically, given a backup control law  $u_b$  and a backup set  $C_b$  that is invariant under  $u_b$ , one can develop barrier constraints that activate

near the boundary of the safe backward image of  $C_b$ ; see (17). The idea was first proposed in [80] and further developed in [53], [54]. A tutorial on the topic is provided in [87], and some notable applications are listed in Table 1. A key advantage to this approach is that it does not require the computation of a large control invariant set.

Consider a smooth backup control law  $u_b : \mathbb{R}^n \to U$  and suppose that the constraint set is given by,

$$\mathcal{C}_{\mathcal{A}} := \{ x \in \mathbb{R}^n \, | \, \varphi(x) \ge 0 \}$$
(41)

and a backup set  $\mathcal{C}_{\mathrm{b}} \subseteq \mathcal{C}_{\mathrm{A}}$  is given by,

$$\mathcal{C}_{\mathbf{b}} = \{ x \in \mathbb{R}^n \,|\, h_{\mathbf{b}}(x) \ge 0 \}.$$

$$\tag{42}$$

The safe backward image of  $\mathcal{C}_{\mathrm{b}}$  is

$$\mathcal{C}_{\mathrm{S}} = \{ x \in \mathbb{R}^n \, | \, \forall t \in [0, T_{\mathrm{b}}], \ \varphi(\phi^{u_{\mathrm{b}}}(t; x)) \ge 0 \land h_{\mathrm{b}}(\phi^{u_{\mathrm{b}}}(T_{\mathrm{b}}; x)) \in \mathcal{C}_{\mathrm{b}} \}.$$
(43)

where  $\phi^{u_b}(t;x)$  is the flow of (27) under  $u_b$ , evaluated t units of time from x. Given an extended class  $\kappa_{\infty}$  function  $\alpha$ , it can be shown that the following constraints are sufficient for invariance in  $C_s$ :

$$\frac{\mathrm{d}\varphi(t_{\mathrm{b}};x))}{\mathrm{d}t} + \alpha(\varphi(\phi(t_{\mathrm{b}};x))) \ge 0 \tag{44}$$

$$\frac{\mathrm{d}h_{\mathrm{b}}(\phi(T_{\mathrm{b}};x))}{\mathrm{d}t} + \alpha(h_{\mathrm{b}}(\phi(T_{\mathrm{b}};x)) \ge 0$$
(45)

for all  $t_b \in [0, T_b]$ . The constraint (44) enforces that the points along the backup trajectory stay in the constraint space  $C_A$ , and (45) enforces that the endpoint of the backup trajectory stay in  $C_b$ . Expanding the derivative terms yields the implicit ASIF barrier constraints:

$$\nabla\varphi(\phi^{u_b}(t_{\mathbf{b}};x))D\phi^{u_b}(t_{\mathbf{b}};x)[f(x)+g(x)u] + \alpha(\varphi(\phi^{u_b}(t_{\mathbf{b}};x))) \ge 0$$
(46)

$$\nabla h_{\rm b}(\phi^{u_b}(T_{\rm b};x))D\phi^{u_b}(T_{\rm b};x)[f(x) + g(x)u] + \alpha(h_{\rm b}(\phi^{u_b}(T_{\rm b};x))) \ge 0$$
(47)

for all  $t_{\rm b} \in [0, T_{\rm b}]$ . Note that since  $t_{\rm b}$  lives in the interval  $[0, T_{\rm b}]$ , (46) represents an uncountable number of constraints. In practice one may approximate the constrained set of states by numerically integrating (27) forward under  $u_{\rm b}$  and evaluating  $\phi^{u_{\rm b}}(t;x)$  at discrete times  $t_{\rm b} \in \{0, t_1, ..., t_{\rm N-1}, T_{\rm b}\}$ . It is shown in [53] how the finite set of intermediate constraints may be tightened in such a way that is sufficient for satisfying the constraints with an infinite number of trajectory points. One may compute  $D\phi^{u_{\rm b}}(t_{\rm b};x)$  by integrating along with (27) a *sensitivity* matrix  $Q(t_{\rm b}, x) = D\phi^{u_{\rm b}}(t_{\rm b};x)$ . As explained in [136],  $Q(t_{\rm b}, x)$  is obtained as the solution to the following differential equation,

$$\frac{\mathrm{d}Q(t_{\mathrm{b}},x)}{\mathrm{d}t_{\mathrm{b}}} = Df_{\mathrm{cl}}(\phi^{u_{\mathrm{b}}}(t_{\mathrm{b}};x))Q(t_{\mathrm{b}},x)$$
(48)

Algorithm 4 Implicit Active Set Invariance Filter				
input	: Current State $x_{curr} \in \mathbb{R}^n$			
	: Desired Input $u_{des} \in U$			
outpu	<b>it</b> : Safe Control Input $u_{act} \in U$			
<b>predefined</b> : Allowable Set $C_A \subset X$				
	: Invariant Backup Set $\mathcal{C}_{\mathrm{b}} \subseteq \mathcal{C}_{\mathrm{A}}$			
	: Smooth Backup Control Law $u_{\rm b}: X \to U$ .			
	: Extended Class $\kappa_{\infty}$ Function $\alpha : \mathbb{R} \to \mathbb{R}$			
1: funct	ion $u_{\text{act}} = \text{IASIF}(x_{\text{curr}}, u_{\text{des}})$			
2: co	<b>pupute:</b> $\phi^{u_{\rm b}}(t_{\rm b}; x_{\rm curr})  \forall t_{\rm b} \in \{0, t_1,, t_{{\rm N}-1}, T_{\rm b}\}$ by integrating (27) under $u_{\rm b}$			
3: ev	<b>valuate:</b> $Df_{cl}(\phi^{u_{b}}(t_{b}; x_{curr}))  \forall t_{b} \in \{0, t_{1},, t_{N-1}, T_{b}\}$			
4: <b>c</b>	<b>pupute:</b> $D\phi^{u_b}(t_b, x_{curr})  \forall t_b \in \{0, t_1,, t_{N-1}, T_b\}$ by integrating (48)			
5: so	lve: (28) subject to the constraints (46)-(47)			

6: **return**  $u_{\text{act}}(x_{\text{curr}}, u_{\text{des}})$ .

with  $Q(0, x) = I_{n \times n}$  and where  $Df_{cl}(\phi^{u_b}(t_b; x))$  is the Jacobian of the closed-loop dynamics of (27) under the control law  $u_b(x)$ , evaluated at  $\phi^{u_b}(t_b; x)$ . In order for this term to be defined, it is necessary that the control law  $u_b$  be smooth. As such, it is often necessary to consider smooth saturation functions for bounding the control law to the admissible domain U.

*Example:* Considered here is finite-time safety (collision avoidance) for two vehicles having combined dynamics,

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \\ \dot{x}_4 \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & -b_1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & -b_2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \end{bmatrix}$$
(49)

where  $x_1, x_3$  denote positions and  $x_2, x_4$  denote velocities for the vehicles,  $b_1 = 0.1$  and  $b_2 = 0.25$ . The safety constraint is collision avoidance  $\varphi(x) = x_3 - x_1$  and the constraint set is  $C_A = \{x \in \mathbb{R}^4 | \varphi(x) \ge 0\}$ . To simplify the analysis, finite-time safety is considered, with the endpoint constraint (47) being omitted. The backup controller  $u_b = [-1, 1]^T$  is integrated over a 10 s horizon and  $\alpha(x) = 2x$ . Figure 13 shows the result of a simulation with  $u_{des}(t) = [1 - \exp(-0.1t), \exp(-0.25t) - 1]^T$ .



Figure 13: Simulation of Implicit ASIF algorithm for cooperative collision avoidance of two vehicles.

*Example:* Consider the system given by (38) and the constraint set  $C_A = \{x \in \mathbb{R}^3 | \varphi(x) \ge 0\}$  with,

$$\varphi(x) = \omega_{\max}^2 - x_1^2 - x_2^2 - x_3^3 \tag{50}$$

where  $\omega_{\max} \in \mathbb{R}_+$  represents the maximum allowable angular speed. It can be shown that this set is not itself forward invariant in general. The backup control law

$$u_{\rm b} = \tanh((x \times Jx) - k_{\rm d}Jx) \tag{51}$$

with  $k_d \ge 0$  is bounded to [-1, 1] and stabilizes system (38) to the origin. Furthermore, it can be shown that  $C_b = \{x \in \mathbb{R}^3 | K - x^T J x \ge 0\}$  is invariant under  $u_b(x)$  (see from previous example that this is invariant under the control law  $u(x) \equiv 0$ ). The largest ellipsoid of this form that is contained in  $C_A$  is obtained by choosing  $K = \omega_{\max}^2 / \min(J_1, J_2, J_3)$ .

Figure 14 shows an illustration of the filtered trajectory under the (unsafe) primary controller  $u_d(x) = [\sin(\frac{t}{2}), \sin(\frac{t}{2} - \frac{\pi}{4}), \sin(\frac{t}{4} + \frac{\pi}{4})]^T$  and with  $\omega_{\max} = 1 m/s$ ,  $k_d = 1$ ,  $J_1 = 12 \text{ kg m}^2$ ,  $J_2 = 12 \text{ kg m}^2$ ,  $J_3 = 5 \text{ kg m}^2$  and the backup horizon  $t_b \in \{0, 0.05, ..., 3\}$ . The trajectory is colored blue where  $u_{\text{des}} = u_{\text{act}}$  and red where the RTA mechanism is active.  $\Diamond$ 

#### **Additional Design Considerations**

Recall that when the properties  $\vartheta_1^{BC}$  and  $\vartheta_2^{BC}$  are satisfied, a set  $C_S$  is forward invariant under the RTA control law  $u_{act}$  given by the ASIF-QP. While guarantees exist under the assumptions of the model, and indeed the formulations presented exhibit some robustness properties, it is desirable in practice to consider the behavior of the system in the case where an unmodelled effect perturbs the state outside of  $C_S$ . A simple solution is to make  $C_S$  a locally attractive set by switching to a stabilizing backup controller whenever  $x \notin C_S$ . Alternatively, for explicit ASIF, *if h is a concave function*, then the barrier constraint will cause  $u_{act}$  to attract solutions to the



Figure 14: Simulation of implicit ASIF algorithm for spacecraft angular velocity dynamics. The constraint set  $C_A$  is represented by the purple sphere, the backup set  $C_b$  is represented by the blue ellipsoid. The surface of this ellipsoid has constant kinetic energy. The bottom right plot projects this information into normed space, and the top right plot shows the desired and actual control values.

interior of  $C_S$  whenever the ASIF-QP is feasible. A backup controller should be present for the case where a feasible solution is not found. While outside of the scope of this article, it is interesting to note that attracting solutions outside of  $C_S$  to  $C_S$  doesn't necessarily correspond to the best response of the physical system. For instance [137] (Section V.B) presents a method for minimizing a function of *damage* in the presence of an inevitable collision.

# **Discussion and Comparison of Approaches on Double Integrator System**

Figure 16 shows a comparison of the four algorithms presented in this article on the double integrator system given by (6) with  $u \in [-1, 1]$ ,  $C_A = \{x \in X \mid -x_1 \ge 0\}$ , and a controller update period of 10 Hz. For the case of the Simplex algorithms, the continuous dynamics are discretized to fit this period. The simulations begin at state  $x(0) = [-1.75, 0]^T$  and have desired input  $u_{des} = 1$ . Here it can be seen that both Simplex-based algorithms exhibit near identical behavior, and likewise both ASIF algorithms exhibit similar behavior. In spite of the implicit methods lacking explicit knowledge of the boundary, all cases exhibit similar operational regions. Note that the backup trajectory under the implicit Simplex filter violates the safety constraint. This is by design, as the trajectories are integrated forward after taking a theoretical probe step. A buffer could be added to the switching region to assure safety. The ASIF approaches are slightly more conservative due to the presence of the  $\alpha$  term in the barrier constraints. This



((b)) The explicit active set invariance filter offers a more gradual intervention that solves a quadratic program to minimize the difference of the actual control signal from the desired control signal of the primary controller, while assuring safety according to Algorithm 3.

Figure 15: Comparison of explicit run time assurance algorithms on double integrator system (6) with constraint set (7) and  $U \in [-1, 1]$ . The viability kernel is shaded green, collision states are shaded red, and the *inevitable* collision states are shaded purple. In the figures on the left,  $x_1$  represents distance to the collision state boundary,  $x_2$  represents velocity relative to the collision state under the desired control from the primary controller is blue, the state under the filtered control is red, and safety constraint boundaries are green. Figures on the right show the actual control output of the filtered signal.

function may be modified to shape the response.

While the Simplex and ASIF methods are most distinctive in terms of behavior, the implicit and explicit methods are most distinctive in terms of requirements. For instance, the key challenge in using the implicit methods in the design of the backup controller and the identification of the backup set. Once these have been identified the it is relatively straightforward to implement either approach. Likewise, the key challenge in implementing the explicit algorithms is identifying an invariant set. Importantly, the explicit ASIF approach is the only algorithm that does not require any backup controller to be defined.



((a)) The implicit simplex filter switches from primary control (full acceleration) to a backup control (full deceleration), as predicted by the path of the system under the backup control according to Algorithm 2.



((b)) The implicit active set invariance filter computes a backup trajectory by integrating under a backup control law and solving the quadratic program under implicit barrier constraints according to Algorithm 4.

Figure 16: Comparison of implicit run time assurance algorithms on double integrator system (6) with constraint set (7) and  $U \in [-1, 1]$ . The viability kernel is shaded green, collision states are shaded red, and the *inevitable* collision states are shaded purple. In the figures on the left,  $x_1$  represents distance to the collision state boundary,  $x_2$  represents velocity relative to the collision state boundary, the state under the desired control from the primary controller is blue, the state under the filtered control is red, implicit backup trajectories are cyan, and safety constraint boundaries are green. Figures on the right show the actual control output of the filtered signal versus simulation time.

# Assurance in the Presence of Uncertainty

While CPS may be *modeled* using dynamical systems as in (1)–(2), real world systems rarely —if ever— obey the dynamics of their model precisely. Furthermore, deterministic models do not capture the effects of external disturbances, which are universal in real-world systems.

One way to address model error is to construct a higher fidelity system model, *i.e.* a higher order parameterized model that can be tuned in system identification. However, producing

higher fidelity modes can become an exhaustive procedure and can be unfruitful when a newly suggested complex model performs less accurately in comparison to its simpler predecessor. These hindrances have attracted the attention of numerous fields of study, *e.g.*, robust control [138], Sim2Real [139].

A second way to address model error and external disturbances is to construct a *nondeterministic* system model, *i.e.* a dynamical model that accounts for environmental, internal disturbances, errors, by incorporating one or more noise terms. In continuous-time, one such a model is represented by

$$\dot{x} = f(x) + g_1(x)u + g_2(x)w, \tag{52}$$

where  $x \in X$  and  $u \in U$  maintain their definitions as the system state and control input, and where now the term  $w(t) \in W \subset \mathbb{R}^p$  is included to represent a bounded nondeterministic input signal to the system. Employing a nondeterministic system model as in (52) reduces the design burden involved in forming an accurate representation of real world phenomena. When designed correctly, a nondeterministic model *will* in general describe the real world behavior accurately.

This section presents run time assurance mechanisms in the context of uncertain systems, as in (52). First, an *explicit* RTA formulation is presented in which *robust control barrier functions* are employed to assure system trajectories evolve in a given safe set. Next, an implicit solution to this problem is presented, where run time safety is assessed through the online computation of a reachable sets.

## **Explicit Run Time Assurance for Uncertain Systems**

In this setting of (52), an ASIF provides a control policy  $u_{act}$  that renders a given safe set  $C_S = \{x \in X | h(x) \ge 0\}$  forward invariant, regardless to disturbance input w(t) chosen. For this reason, ASIFs are naturally constructed from *robust control barrier functions*, as discussed next.

Robust control barrier functions extend the general barrier function theory presented above to the nondeterministic setting of (52). A continuously differentiable function  $h : \mathbb{R}^n \to \mathbb{R}$  is a robust control barrier function for (52) if there exists a class- $\kappa$  function  $\alpha : \mathbb{R} \to \mathbb{R}$  such that for all  $x \in C_S := \{x \in X \mid h(x) \ge 0\}$  there exists a  $u \in \mathbb{R}^m$  satisfying

$$BC(x, u, w) = L_f h(x) + L_{g_1} h(x)u + L_{g_2} h(x)w + \alpha(h(x)) \ge 0$$
(53)

for all  $w \in W$ , where  $L_f h(x)$ ,  $L_{g_1} h(x)$  and  $L_{g_2} h(x)$  are naturally taken to be the lie derivatives of h along f,  $g_1$  and  $g_2$ . The barrier constraint (53) is a linear inequality on the variable u for any  $x \in \mathbb{R}^n$  and any  $w \in W$ , and when a candidate controller  $u_{\text{des}}(x)$  satisfies (53) for all  $x \in C_S$  and all  $w \in W$ , the set  $C_S$  will be robustly forward invariant for the closed-loop dynamics of (52) under  $u_{des}$ .

When instead  $u_{des}(x)$  does not satisfy (53), an ASIF can be employed to assure the forward invariance of  $C_S$ . One such filter, given below as RASIF-QP, is posed as quadratic program where the constraints in this program are constructed from (53).

# **RASIF-QP**

$$u_{\rm act}(x) = \arg\min_{u \in U} ||u - u_{\rm des}(x)||_2^2$$
(54)

s.t. 
$$L_f h(x) + L_{g_1} h(x)u + L_{g_2} h(x)w + \alpha(h(x)) \ge 0$$
  
for all  $w \in W$  (55)

When h is a robust control barrier function for (52), the RCBF-QP is always feasible and the set  $C_S$  is robustly forward invariant for closed-loop dynamics of (52) under  $u_{act}$ . However, this program contains an infinite number of linear constraints and, thus, is not always practically implementable. In the special instance where the disturbance set W is a polytope, the program RCBF-QP can be reduced to include only a finite number of linear constraints. Assume, for instance, that

$$W = \mathbf{Conv}\{w^1, \cdots, w^q\}$$
(56)

for some  $w^1, \dots, w^q \in \mathbb{R}^p$ , where **Conv** denotes the convex hull operator. Then an explicit ASIF is constructed using

## **RASIF-QP** (polytope W)

$$u_{\rm act}(x) = \arg\min_{u \in \mathbb{R}^m} ||u - u_{\rm des}(x)||_2^2$$
(57)

s.t. 
$$L_f h(x) + L_{g_1} h(x)u + L_{g_2} h(x)w + \alpha(h(x)) \ge 0$$
  
for all  $w \in \{w^1, \cdots, w^q\}$  (58)

and this program contains only q linear constraints.

#### **Implicit Active Set Invariance for Uncertain Systems**

In this section, an implicit RTA formulation is introduced for the nondeterministic setting of (52). As in the case of deterministic systems, the implicit ASIF allows the system to leave the safe set  $C_S$  when safety is verified *a priori* via the assessment of a known safe backup control policy.

In the deterministic setting of (27), the safety of a given state  $x \in X$  with respect to the backup controller  $u_{\rm b}$  is assessed via a system simulation; however, in the nondeterministic setting of (52), it is necessary to assess the effects of the disturbance, perhaps through, *e.g.*, computing reachable sets [67], [140], [141]. For initial set  $A \subseteq X$  and time  $t \ge 0$ , the time-treachable set of (52) under  $u_{\rm b}$  is

$$R(t;A) := \left\{ \Phi^{u_{\mathsf{b}}}(t;x,\mathbf{w}) \in X \mid x \in A, \, \mathbf{w} : [0,t) \to \mathcal{W} \right\}$$
(59)

where  $\Phi^{u_{\rm b}}(t; x, \mathbf{w})$  denotes the state of (52) reached at time t when beginning at state  $x \in X$ at time 0 and evolving subject to the backup control input  $u_{\rm b}$  and the disturbance signal  $\mathbf{w}$ :  $[0,t] \to \mathcal{W}$ . In particular, an implicit ASIF ensures  $R(T_{\rm b}; x) \subseteq C_{\rm b}$  for all states x along the system trajectory, where  $T_{\rm b} \ge 0$  maintains its definition and function as the horizon time of the backup controller.

In general, it can be difficult to compute reachable sets in closed form. However numerous efficient computational methods exists for over-approximating reachable sets, with computational speeds suitable for, *e.g.*, computing reachable sets in the control-loop and enforcing safe behavior online from these predictions [140]–[143]. Of particular interest to this work, the mixed monotonicity property of dynamical systems can be applied to compute a hyperrectangular over-approximation of R(t; x) using a single simulation of a related 2n-dimensional deterministic embedding system; a technique that have been applied in domains including transportation system [144], and biological systems [145]. Further details on the mixed monotonicity property are provided in the sidebar titled "Mixed Monotonicity for Efficient Reachability" in this work.

In the remainder of this section, we present one possible implicit ASIF construction, where over-approximations of reachable sets are computed online using the mixed monotonicity property. The basic idea is as follows: when the system is at state x, the ASIF computes a hyperrectangular overapproximation of  $R(T_b; x)$  using the mixed monotonicity procedure. A barrier condition is then enforced for each vertex of the approximation, as to ensure  $R(T_b; x) \subseteq C_b$  for all states x along the system trajectory. In this way safety is ensured. Specific details are as follows.

For a finite set  $S \subset \mathbb{R}$ , and some fixed parameter p > 0, the Log-Sum-Exponential (LSE)

of S is given by

$$LSE(\mathcal{S}) = -\frac{1}{p} \log \sum_{s \in \mathcal{S}} \exp(-p \cdot s).$$
(60)

The LSE has several useful properties: namely, LSE(S, p) is differentiable with respect to the elements of S, and LSE(S, p) approximates min S, *i.e.*,

$$\min \mathcal{S} - \frac{n}{p} \log 2 \le \mathsf{LSE}(\mathcal{S}, p) < \min \mathcal{S}$$
(61)

for all p > 0, and this approximation can be made arbitrarily tight by choosing p large enough. We use the LSE, in this section, to approximate the minimum evaluation of h over a hyperrectangular subset of the statespace. Define the following:

$$LSE_{h}(a) := LSE(\{h(z) \mid z \in \langle\!\langle a \rangle\!\rangle\}, p)$$
(62)

$$\gamma(t; x) := \mathsf{LSE}_h(\Phi^E(t; x)) \tag{63}$$

$$\Psi(x) := \sup_{0 \le \tau \le T_{\rm b}} \gamma(\tau; x) \tag{64}$$

where  $\langle\!\langle x, y \rangle\!\rangle$  in (62) denotes the set of  $2^n$  corners of a rectangle [x, y] and is given by

$$\langle\!\langle x, y \rangle\!\rangle := \{ z \in X \mid z_i \in \{x_i, y_i\} \}.$$
 (65)

From the above definitions we have that  $\Psi(x) \ge 0$  when there exists a  $t \in [0, T_b]$  so that  $[\Phi^E(t; x)] \subseteq C_b$  and in this case  $R(T_b; x) \subseteq C_b$  as  $C_b$  is assumed robustly forward invariant for (52) under the backup controller  $u_b$ . Thus, an implicit ASIF for (52) is constructed using Algorithm 5. See [140], [142], [143] for further details.

# Verification of Run Time Assurance Algorithms and Architectures

A key question in trusting RTA systems is "who checks the checker?" In many applications, RTA systems serve a safety critical role and therefore must be rigorously verified. One can imagine a second RTA to check the RTA, and another third to check that RTA until it is turtles all the way down. Proponents of the Simplex architecture advocate simple switching logic and preplanned backup control actions that reduce the verification burden [108]. Other approaches may rely on mathematical proofs for verification [146]. Nearly all approaches see the RTA as a modular component within the closed-loop control system that can be verified sub-component. It is worth noting that this verification can be done on the algorithms or the actual software implementation. This section focuses on verification of RTA algorithms and architectures in the early design phases using formal and compositional reasoning. Algorithm 5 Implicit ASIF for Nondeterministic Control Systems

input	: Current State $x_{curr} \in X$
	: Desired Input $u_{des} \in U$
output	: Safe Control Input $u_{\text{act}} \in U$
predefine	<b>d</b> : Constraint set $\mathcal{C}_A \subset X$
	: Invariant Set $\mathcal{C}_{\mathrm{S}} \subset \mathcal{C}_{\mathrm{A}}$
	: Backup Control Law $u_{\rm b}: X \to U$ .

1: function  $u_{act}(x) = ASIF(x_{curr}, u_{des})$ 

2: compute:

- 3:  $u^* = \arg\min_{u \in \mathbb{R}^m} ||u u_{des}||_2^2$
- 4: s.t.  $\frac{\partial \Psi}{\partial x}(x_{\text{curr}})(f(x_{\text{curr}}) + g_1(x_{\text{curr}})u + g_2(x_{\text{curr}})w) \ge -\alpha(\Psi(x_{\text{curr}}))$
- 5:  $\forall w \in \langle\!\langle \underline{w}, \, \overline{w} \rangle\!\rangle$
- 6: **if** Program feasible **then return**  $u^*$
- 7: else return  $u_{\rm b}(x_{\rm curr})$

#### Formal Specification and Analysis of RTA System Requirements, Architecture and Design

Formal methods tools from computer science for automated verification of hardware and software [147] can be brought to the problem of verifying RTA. In addition to the ability to verify software implementations, formal methods can be used to aid in modeling and automated analysis of algorithm and decision logic with mathematical rigor. At one level, Lyapunov proofs can be used to guarantee safety of control systems. Formal methods can supplement these proofs, or perhaps be used where Lyapunov proofs do not exist. Three goals of formal methods are specification, analysis, and synthesis.

*Formal specification* facilitates a common, unambiguous understanding of the system between users, designers, programmers, and testers. "Specification is difficult, unglamorous, and arguably the biggest bottleneck facing verification and validation of aerospace, and other, autonomous systems" [148]. RTA systems and formal specification and analysis have been applied successfully in a variety of applications [60], [63], [66], [97], [106], [107], [149]–[161]. Formal RTA requirements specifications have also been developed and analyzed [162]–[164].

*Formal analysis* helps to find errors and increase reliability of the system. Formal analysis can be divided [165] into axiomatic (*e.g.* . Hoare logic [166], theorem proving [167]), semantic approaches (*e.g.* . model checking [147]), and static analysis methods (*e.g.* . abstract interpretation [168]). Many formal analysis approaches focus on the satisfiability modulo theories (SMT)

where sets of variables are substituted with binary-valued functions of non-binary values called predicates [169]. An example predicate is an inequality that is evaluated as true or false, such as  $x+y \le c$  with variables x and y, and constant c. These predicates can be combined in a variety of logics such as first order and temporal logics such as linear temporal logic (LTL), computational tree logic (CTL), timed computational tree logic (TCTL), metric temporal logic (MTL), and signal temporal logic (STL) [147], [170]–[172]. Formal verification of linear temporal logic specifications has been applied to PID attitude control of spacecraft [173], a Simplex RTA system for spacecraft attitude [160], [174], [175], switching logic for automatic maneuvering [161], an LTL specification monitor automaton [67], and many other applications.

*Formal synthesis* integrates formal methods into the development process to convert a design to implementation with some level of automation. While many controller designs are synthesized in some fashion, formal synthesis typically involves automatic generation of a controller, or a run time assurance filter from formal specifications [176].

#### **Compositional Verification**

Compositional verification, sometimes referred to as a "divide and conquer" approach to verification [177], helps to verify or test portions of algorithms or systems architectures separately and then make conclusions about the system as a whole. This verification approach is helpful in modern engineering practices where engineering tasks are divided over large teams [178]. Compositional verification includes statements such as: if property  $p_1$  for subsystem A and property  $p_2$  for subsystem B both hold, it implies property p holds for the entire system. A compositional verification approach can be used to apply formal methods to verify components of an RTA system, especially in Simplex-based RTA systems where multiple simple components (monitors, backup controllers, and switching functions) work together to ensure safety [66], [90], [160], [174], [175], [179].

Compositional verification often relies on *assume-guarantee reasoning* of the architecture [180] or the hybrid switching [181]. Assume-guarantee contracts make guarantees on the output of a component based on assumptions about the input. Assumptions are expectations that a component has about it's environment, while guarantees are behavior properties provided by the component. Verification techniques are applied to individual components with explicit, formal, assumptions and guarantees, as depicted in Figure 17.



Figure 17: Depiction of assume-guarantee reasoning for an automatic collision avoidance system, were assumptions properties are guaranteed by a component based on assumptions about their environment. Guarantees on the output of one component translate to assumptions on the input of another component until larger properties of the overall system are proven.

# Conclusions

RTA Systems are growing in popularity as a way to assure safety by altering unsafe control inputs from a primary controller. The assurance mechanism of RTA systems is constructed agnostic to the underlying structure of the primary controller, whether the primary control comes from a human operator, an advanced control approach, or an autonomous control approach. By effectively decoupling the enforcement of safety constraints from performance-related objectives, RTA offers a number of useful advantages over traditional (offline) verification.

This article provided a tutorial on developing RTA systems with particular emphasis on implicit and explicit safety definitions as well as Simplex architecture and active set invariance filtering RTA approaches. Implicit (or trajectory-based) safety definitions computed finite-time trajectories under a backup controller online, while explicit (or region-based) approaches identified trusted regions offline via the construction of either a large forward invariant or control invariant set. Simplex Architecture RTA designs *monitor* for imminent violations of safety constraints, and switch from the primary controller to a backup controller when an imminent unsafe condition is detected. By contrast, backup control intervenes gradually in ASIF approaches, rather than with a hard switch.

Given the increasing complexity of modern control systems, RTA in its various forms is

one tool to assure safety and a potential path to certification in safety critical domains.

# Sidebar: Run Time Assurance

Run Time Assurance (RTA) systems guarantee safety of increasingly complex and intelligent control system designs by monitoring the state of the system and intervening when necessary. RTA acts as a safety filter that sits between the primary controller and the plant to assure the state of the system adheres to safety properties. The design of the RTA safety mechanism is decoupled from the design of the primary controller, allowing for RTA to focus on safety while the primary controller optimizes performance. Historically, verification techniques for novel control approaches trail years to decades behind the development of the approaches themselves [66]. RTA provides a path to introduce the use of these advanced controllers faster, even in safety-critical applications. This article provides a tutorial on many ways to design RTA, including explicit and implicit monitoring, Simplex and Active Set Invariance Filter (ASIF) intervention approaches, and uncertainty with a variety of examples. Explicit monitors define forward invariant or control invariant safe sets offline that are enforced online at run time, while implicit monitors project the a finite-time trajectory of a backup controller online. Simplex interventions switch from the primary controller to a backup controller, while ASIF interventions gradually optimally modify the control signal subject to safety constraints. A conceptual and a mathematical description are provided for each. Where possible, the Automatic Ground Collision Avoidance System (Auto GCAS), an implicit Simplex RTA design, is used as an example to illustrate several RTA concepts. The article assumes an introductory understanding of control theory and state space concepts.

# Sidebar: RTA Aliases

RTA provides a conceptually simple and highly effective approach to enforcing safety constraints that has frequently been explored, utilized, and independently reinvented in a wide variety of areas. As such, RTA and its applications go by many names in the literature, including: active set invariance, active collision avoidance, advanced driver assistance, active safety, safety filtering, emergency braking, sandboxing, conflict resolution, and constraint-based planning. The primary controller goes by several names in the literature including advanced controller [182], advanced system [41], experimental control module [152], baseline controller [152], unverified controller [174], [183], nominal controller [184], complex subsystem [154], [183], and complex controller [60]. Finally, the backup controller goes by many names in the literature, including the recovery mechanism, reversionary controller [182], reversionary system [41], safety control module [152], verified controller [174], backup controller [184], safety remediation controller

[154], recovery controller [61], [108], [159], [183], and safety controller [60].

Part of the different naming conventions comes from the different intentions and the resulting implementation of the RTA intervention, which are grouped here as: recovery and reversionary controllers. *Recovery Controllers* intervene in emergency situations and return control authority to the primary controller as soon as possible. Recovery controllers are not intended to perform the wide range of tasks a dynamical system is designed to perform, but instead focuses solely on recovery maneuvers from an unsafe condition (e.g. collision avoidance, adherence to keep out zones, or maintaining state limits). By contrast, *reversionary controllers* detect inappropriate behavior of the primary controller and revert to a simpler, verified controller to complete the rest of the tasks. The reversionary controller may sacrifice performance for assurance. This article focuses on the design of recovery controller-based RTA systems.

## Sidebar: Shielded Learning

No discussion of RTA for complex and autonomous control systems would be complete without discussing the emerging challenges of using machine learning, and in particular, reinforcement learning (RL), where a dynamical system learns from experience collected over many episodes [185]. The use of RL has attracted attention for its ability to tackle complex tasks with high-dimensional state spaces by world experts in Go [186], [187] and StarCraft II [188]. In addition, it is becoming popular in robotics research to operate in unknown environments and learn state representations for many tasks [189]. Both control theory and RL share a common concept of a system state (usually  $x \in \mathcal{X}$  in control theory and  $s \in \mathcal{S}$  in RL). A control input ( $u \in \mathcal{U}$  in control theory) is formulated as an *action* a or  $A \in \mathcal{A}$  in RL, and the actions are determined by a *policy* (i.e. controller) that maps states to actions. State-action sequences  $(S_0, A_0, S_1, A_1, \dots, S_n, A_n)$  are called *trajectories* in RL. Both fields have a notion of partial observability. A measurement y may come from sensors in control theory, or an observation O in RL may be used to estimate the system state (e.g.  $\hat{x} \in \mathcal{X}$ ). In control theory, a controller is optimized to minimize a cost function, while in RL, a policy is optimized to maximize a reward function. In both control theory and RL, dynamics describe how the state evaluates in the environment. RL approaches may be model-based, where the agent learns a model of the environment, or *model-free* where dynamics are not explicitly modeled and the action is learned directly through interactions with the environment.

A challenge in RL, especially when it takes place on a physical agent, is ensuring safety of that agent. Safe RL is defined as the process of learning a policy that maximizes expected return while ensuring adherence to safety constraints [34]. Safe RL usually employs either *reward shaping* or *shielding* (*i.e.* RTA) to incorporate safety in the training process. Reward shaping

factors safety into the reward function to reduce risk. However, it is difficult to properly weight safety rewards and reward shaping can sometimes have unintended impacts on performance. In addition, reward shaping does not guarantee safety during operations. By contrast, RTA can be used to filter the actions of the RL algorithms to ensure safety. A variety of RTA approaches have been explored in the literature including human-like intervention [190], Lyapunov-based approaches [191], [192], barrier functions [193], and formal verification of safety constraints [194], [195]. Determining the appropriate way to include RTA in the training process is still an area of active research.

## Sidebar: The Case for Plan B

A fundamental concept underlying RTA is: always maintaining a safe "Plan B" during mission execution, with the purpose of executing "Plan B" as soon as the nominal activity appears to be dangerous. Considered from that perspective, maintaining a "Plan B" makes sense and hundreds of examples can be found in everyday life. For example, financial advisers always recommend their clients keep a few weeks' worth of salary in a savings account so as to be ready and face unexpected financial expenses, such as the medical costs following an accident, the replacement of a commuting vehicle following a collision, a few weeks at a hotel if the home is damaged by a hurricane, and so on. Defensive driving is another example, whereby a driver is trained to always be aware of her surroundings and be ready to take appropriate action in case of unexpected events. Currently U.S. regulations specify three levels of "circuit breakers," that are used to halt trading in stock exchanges when the S&P 500 index drops below certain critical levels. One of the oldest examples of an engineered "Plan B" can be found in Greek mythological the story of Ariadne's thread [196]–[198]. In the story, Ariadne, daughter of Minos the king of Crete, conceived of using a thread as a "Plan B" mechanism that would allow her lover to exit Daedalus' Labyrinth after killing the Minotaur. Similar themes appear in folktales such as "Le Petit Poucet" (Hop-o'-My-Thumb) [199], [200] and "Hansel and Gretel."

One may also observe this paradigm in nature. It is now known that acacia trees, whose leaves are one common form of food for giraffes and other herbivorous fauna in Africa, are capable of exercising a kind of "Plan B" making these leaves lethal by raising tannin-C in their leaves to deadly levels in a matter of minutes if they are over-grazed [201]. Likewise, rodents very often create burrows with emergency exits or "bolt holes" to escape predators, flooding or other causes that may block their main entrances [202]. The sympathetic nervous system in mammals acts as a physiological "Plan B" by preparing the body for a "fight or flight" response. Upon detection of a threat, glucose is transmitted from storage sites, blood is shifted to the organs that are most essential for physical exertion, heart rate is increased, adrenaline is released, and

cognition is sharpened [203].

# Sidebar: Safety, Reliability, and Security

The primary function of an RTA system is to ensure safety (and in some cases security), not reliability. Safety and reliability are two different concepts. Reliability is freedom from errors and failures, while safety is freedom from harm. *Reliability* of a system is its ability to perform consistently over time. For example, physical components may fatigue and degrade over time, and reliability focuses on how long the performance is consistent and failure is unlikely [40], [204]. Reliability of software, which doesn't degrade over time (unless it changes over time, in the case of intelligent control that continues to learn online), is based on it's freedom from design errors. In a related sense, *security* is freedom from harm from others. The difference in safety and security is in intent, severity, and likelihood of harm.

## Sidebar: To RTA or not to RTA? That is the Question

RTA is not a panacea for controller safety. Employing RTA incurs a cost-benefit tradeoff. Implementing RTA can increase system cost, complexity, the amount of testing or evidence required for certification, the number of hardware and software components that can fail, and the attack surface when cyber security is a concern. Designers must ask themselves critical questions before deciding to use an RTA system. First, *could a simple controller meet design requirements?* If a simple control design is sufficient, traditional simulation and testing can provide adequate confidence, or the primary controller can be exhaustively verified, an RTA system should not be used. Second, *is the system safety or mission critical?* If the system is not safety or mission critical, an RTA could be an unnecessary design complication.

If the system under consideration is safety or mission critical and the control functions can only be achieved by a complex or neural network control scheme, and certification of that high performance controller is not possible via traditional simulation, test, or exhaustive analysis, then an RTA system may be appropriate.

# Sidebar: Reference and Command Governors

While this article focuses entirely on RTA designs that monitor the primary controller and intervene before the control signal reaches the plant, it is worth noting that other approaches, such as *reference and command governors* can be used for safety assurance. Instead of modifying the control signal of the primary controller before it reaches the plant, reference and governors wrap around the closed loop system and change the reference signal [205]. The reference governor

acts as a pre-filter on the desired reference command r(t) and state measurement or estimate x(t) under noise w(t), resulting in a modified reference command v(t). A command or reference governor architecture is depicted in Figure S1.



Figure S1: Depiction of a reference governor which modifies the input to the primary controller, rather than the primary controller output to assure safety.

# Sidebar: The Robotarium: An RTA Enabled Remote-Access Swarm Robotics Testbed

The Robotarium [83]–[86] is a remotely accessible swarm robotics testbed with the stated mission to *democratize robotics by providing remote access to a state-of-the-art multi-robot research facility*. The lab consists of a 12ft×14ft custom arena and a number of differential drive robots, known as GritsBots. The testbed exhibits features such as Vicon-based real time motion-capture, wireless inductive charging, video capture of experiments, and an above mounted projector that allows users to project time-varying environmental projections onto the testbed during experiments. Users may submit code to control the robots through MATLAB or Python scripts, that are submitted through the Robotarium website.

The open-access nature of the lab presents unique challenges in terms of safety. Experiments must remain as faithful as possible to the user specified behavior while being provably safe, so as to prevent damage to the lab equipment. Offline verification of the user code would be a prohibitively difficult and time consuming task for the operators of the lab, and would place an unnecessary burden of correctness on its users. Consequently, an RTA solution to safety is utilized, and user submitted algorithms are treated as black box functions that can be probed at run time. Specifically, control barrier function based quadratic programs (explicit ASIF) are used to prevent collisions between the various robots involved in an experiment. An explicit ASIFbased approach is attractive in this case as it is minimally invasive to the desired inputs and provides a smooth overriding behavior. Since the planning is conducted under a single-integrator dynamics model, explicit identification of viable sets is fully-tractable problem offline.

# Sidebar: Mixed Monotonicity for Efficient Reachability

A dynamical system, possibly subject to a nondeterministic disturbance input, is mixed monotone when there exists a related decomposition function that separates the system dynamics into cooperative and competitive state interactions. Mixed monotonicity applies to continuoustime systems [206], [207], discrete-time systems [145], as well as systems with disturbances [208], [209], and it generalizes the *monotonicity* property of dynamical systems for which trajectories maintain a partial order over states [210], [211].

For an *n*-dimensional mixed monotone system with a disturbance input, it is possible to construct a 2n-dimensional monotone embedding system from the decomposition function. This enables one to apply the powerful theory of monotone dynamical systems to the embedding system and can be used to compute useful properties of the initial system. In particular, such approaches are useful for efficiently computing reachable sets and robustly forward invariant sets for systems with disturbances, techniques which have been successfully demonstrated in online safety applications [140]–[143].

A dynamical system

$$\dot{x} = F(x, w),\tag{S1}$$

with state  $x \in X \subseteq \mathbb{R}^n$  and disturbance  $w \in \mathcal{W} \subseteq \mathbb{R}^m$ , is a *mixed monotone system* when there exists a related decomposition function  $d: X \times \mathcal{W} \times X \times \mathcal{W} \to \mathbb{R}^n$  so that for all  $x, \hat{x} \in X$  and all  $w, \hat{w} \in \mathcal{W}$  the following hold:

- $\begin{array}{l} \bullet \ d(x,w,x,w) = F(x,w). \\ \bullet \ \frac{\partial d_i}{\partial x_j}(x,w,\widehat{x},\widehat{w}) \geq 0 \ \text{for all } i,j \ \text{with } i \neq j. \\ \bullet \ \frac{\partial d_i}{\partial \widehat{x}_j}(x,w,\widehat{x},\widehat{w}) \leq 0 \ \text{for all } i,j. \\ \bullet \ \frac{\partial d_i}{\partial w_k}(x,w,\widehat{x},\widehat{w}) \geq 0 \ \text{and } \ \frac{\partial d_i}{\partial \widehat{w}_k}(x,w,\widehat{x},\widehat{w}) \leq 0 \ \text{for all } i,k. \end{array}$

Large classes of systems have been shown to be mixed monotone, with decomposition functions constructed from, e.g., bounds on the systems Jacobian matrix [208] of domain specific knowledge [212], [213]. In certain instances decomposition functions can also be computed by solving an optimization problem [214].

An important feature of mixed monotone systems is that hyperrectangular overapproximations of reachable sets can be efficiently computed using a single simulation of a related 2n dimensional embedding system constructed from the decomposition function. Assume X is

an extended hyperrectangle and  $W := [\underline{w}, \overline{w}]$  is a hyperrectangle. When (S1) is mixed monotone with decomposition function d,

$$\begin{bmatrix} \dot{x} \\ \dot{\hat{x}} \end{bmatrix} = E(x,\,\hat{x}) := \begin{bmatrix} d(x,\,\underline{w},\,\hat{x},\,\overline{w}) \\ d(\hat{x},\,\overline{w},\,x,\,\underline{w}) \end{bmatrix}$$
(S2)

is the *embedding system* relative to d, and  $\Phi^E(t; a)$  denotes the state of (S2) reached at time  $t \ge 0$  when beginning at state  $a \in X \times X$  at time 0. For any hyperrectangular set of states  $[\underline{x}, \overline{x}] \subset X$  the following is true: if  $\Phi^E(\tau; (\underline{x}, \overline{x})) \in X \times X$  for all  $0 \le \tau \le t$ , then

$$R(t; [\underline{x}, \overline{x}]) \subseteq [\Phi_{1:n}^{E}(t; (\underline{x}, \overline{x})), \Phi_{n+1:2n}^{E}(t; (\underline{x}, \overline{x}))].$$
(S3)

The results of (S3) provide an efficient procedure for over-approximating reachable sets for systems with disturbances; in particular, a hyperrectangular over-approximation of  $R(t; [\underline{x}, \overline{x}])$  is computable from a single simulation of the embedding system (S2), where the bottom and top corners of the approximation are the first n and last n coordinates of  $\Phi^E(t; (\underline{x}, \overline{x}))$ . This result is demonstrated in the following numerical example.

Example: Consider the system

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = F(x) = \begin{bmatrix} x_2^2 + 2 \\ x_1 \end{bmatrix}$$
(S4)

with state space  $X = \mathbb{R}^2$ . The system (S4) is mixed monotone with respect to the *decomposition* function  $d: X \times X \to \mathbb{R}^2$  given by

$$d_{1}(x,\hat{x}) = \begin{cases} x_{2}^{2} + 2 & \text{if } x_{2} \ge \max\{0, -\hat{x}_{2}\}, \\ \hat{x}_{2}^{2} + 2 & \text{if } \hat{x}_{2} \le \min\{0, -x_{2}\}, \\ 2 & \text{if } x_{2} \le 0 \le \hat{x}_{2}, \end{cases}$$

$$d_{2}(x,\hat{x}) = x_{1}.$$
(S5)

Reachable sets for (S4) are now approximated from a single simulation of the embedding system (S2), as described in (S3). An example is shown in Figure S2 where  $R(1; X_0)$  is approximated for  $X_0 = [-1/2, 1/2]^2$ .

# References

- [1] Wikimedia Commons, "Auto-gcas saves f-16," Website, https://upload.wikimedia.org/ wikipedia/commons/7/75/Auto-GCAS\_saves\_F-16.webm.
- [2] A. Circular, "Ac 25.1309-1a, system design and analysis," *Washington, DC: Federal Aviation Administration*, 1988.



Figure S2: Approximating reachable sets of (S4) from the set of initial conditions  $X_0 = [-1/2, 1/2] \times [-1/2, 1/2]$ . (a)  $X_0$  is shown in red.  $R^F(1; X_0)$  is shown in green. The hyperrectangular over-approximation of  $R^F(1; X_0)$ , which is computed from the embedding system (S2) as described in (S3), is shown in light green. (b) Visualization of the bounding procedure from (S3). The trajectory of (S2) that yields Figure S2(a) is shown in blue, where  $\Phi^E$  is projected to the  $x_1, \hat{x}_1$  plane. The southeast cones corresponding to  $X_0$  and the hyperrectangular over-approximation of  $R^F(1; X_0)$  are shown in red and green, respectively.

- [3] Defense Acquisition University, "Verification," https://www.dau.mil/glossary/pages/2860.aspx, 2018, accessed 14 November 2018.
- [4] —, "Validation," https://www.dau.mil/glossary/pages/2852.aspx, 2018, accessed 14 November 2018.
- [5] J. William R. Nichols and T. Scanlon, "DoD Developer's Guidebook for Software Assurance," Carnegie Mellon University Software Engineering Institute, DoD Developer's Guidebook for Software Assurance, 12 2018.
- [6] Federal Aviation Administration (FAA), "Standard airworthiness certification regulations: Title 14, code of federal regulations," Website, 10 2017, https://www.faa.gov/aircraft/air\_ cert/airworthiness\_certification/std\_awcert/std\_awcert\_regs/regs/.
- [7] "MIL-HDBK-516C: Airworthiness Certification Criteria," Department of Defense, Guidance, 12 2014.
- [8] "MIL-STD-882E: System Safety," Department of Defense, Standard Practice, May 2012.
- [9] "Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment," SAE International, Guidance, 1996.
- [10] "Guidelines For Development Of Civil Aircraft and Systems," SAE International, Guidance, 2010.
- [11] "Software Considerations in Airborne Systems and Equipment Certification," RTCA Inc., Guidance, 2011.

- [12] "Design Assurance Guidance for Airborne Electronic Hardware," RTCA Inc., Guidance, 2000.
- [13] "Formal Methods Supplement to DO-178C and DO-278A," RTCA Inc., Guidance, 2011.
- [14] C. Hu, Z. Wang, Y. Qin, Y. Huang, J. Wang, and R. Wang, "Lane keeping control of autonomous vehicles with prescribed performance considering the rollover prevention and input saturation," *IEEE Transactions on Intelligent Transportation Systems*, 2019.
- [15] X. Xu, J. W. Grizzle, P. Tabuada, and A. D. Ames, "Correctness guarantees for the composition of lane keeping and adaptive cruise control," *IEEE Transactions on Automation Science and Engineering*, vol. 15, no. 3, pp. 1216–1229, 2017.
- [16] T. Gurriet, S. Finet, G. Boeris, A. Duburcq, A. Hereid, O. Harib, M. Masselin, J. Grizzle, and A. D. Ames, "Towards restoring locomotion for paraplegics: Realizing dynamically stable walking on exoskeletons," in 2018 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2018, pp. 2804–2811.
- [17] O. Harib, A. Hereid, A. Agrawal, T. Gurriet, S. Finet, G. Boeris, A. Duburcq, M. E. Mungai, M. Masselin, A. D. Ames *et al.*, "Feedback control of an exoskeleton for paraplegics: Toward robustly stable, hands-free dynamic walking," *IEEE Control Systems Magazine*, vol. 38, no. 6, pp. 61–87, 2018.
- [18] O. Sanni, M. Mote, D. Delahaye, M. Gariel, T. Khamvilai, E. Feron, and S. Saber, "Ariadne: A common-sense thread for enabling provable safety in air mobility systems with unreliable components," 2021. [Online]. Available: http://hdl.handle.net/10754/666807
- [19] E. Squires, P. Pierpaoli, and M. Egerstedt, "Constructive barrier certificates with applications to fixed-wing aircraft collision avoidance," in 2018 IEEE Conference on Control Technology and Applications (CCTA). IEEE, 2018, pp. 1656–1661.
- [20] E. Squires, P. Pierpaoli, R. Konda, S. Coogan, and M. Egerstedt, "Composition of safety constraints with applications to decentralized fixed-wing collision avoidance," *arXiv* preprint arXiv:1906.03771, 2019.
- [21] E. Squires, R. Konda, P. Pierpaoli, S. Coogan, and M. Egerstedt, "Safety with limited range sensing constraints for fixed wing aircraft," *arXiv preprint arXiv:2010.10883*, 2020.
- [22] T. Schouwenaars, M. Valenti, E. Feron, and J. How, "Implementation and flight test results of milp-based uav guidance," in 2005 IEEE Aerospace Conference. IEEE, 2005, pp. 1–13.
- [23] T. Schouwenaars, "Safe trajectory planning of autonomous vehicles," Ph.D. dissertation, Massachusetts Institute of Technology, 2006.
- [24] C. Hanson, "Capability description for nasa's f/a-18 tn 853 as a testbed for the integrated resilient aircraft control project," 2009.
- [25] 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). IEEE, 2017, pp. 3293–3298.

- [26] A. Singletary, T. Gurriet, P. Nilsson, and A. D. Ames, "Safety-critical rapid aerial exploration of unknown environments," in 2020 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2020, pp. 10270–10276.
- [27] M. Mote, C. W. Hays, A. R. Collins, E. Feron, and K. L. Hobbs, "Natural motion-based trajectories for automatic spacecraft collision avoidance during proximity operations," in *IEEE Aerospace*, 2021.
- [28] A. Singletary, P. Nilsson, T. Gurriet, and A. D. Ames, "Online active safety for robotic manipulators," in 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 2019, pp. 173–178.
- [29] E. Wong, J. D. Schierman, T. Schlapkohl, and A. Chicatelli, "Towards run-time assurance of advanced propulsion algorithms," in 50th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, 2014, p. 3636.
- [30] G. Katz, C. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer, "Reluplex: An efficient smt solver for verifying deep neural networks," in *International Conference on Computer Aided Verification*. Springer, 2017, pp. 97–117.
- [31] H.-D. Tran, X. Yang, D. M. Lopez, P. Musau, L. V. Nguyen, W. Xiang, S. Bak, and T. T. Johnson, "Nnv: The neural network verification tool for deep neural networks and learning-enabled cyber-physical systems," in *International Conference on Computer Aided Verification*. Springer, 2020, pp. 3–17.
- [32] S. Bak, H.-D. Tran, K. Hobbs, and T. T. Johnson, "Improved geometric path enumeration for verifying relu neural networks," in *International Conference on Computer Aided Verification.* Springer, 2020, pp. 66–96.
- [33] S. Gokulanathan, A. Feldsher, A. Malca, C. Barrett, and G. Katz, "Simplifying neural networks using formal verification," in NASA Formal Methods Symposium. Springer, 2020, pp. 85–93.
- [34] J. Garcia and F. Fernández, "A comprehensive survey on safe reinforcement learning," *Journal of Machine Learning Research*, vol. 16, no. 1, pp. 1437–1480, 2015.
- [35] M. Alshiekh, R. Bloem, R. Ehlers, B. Könighofer, S. Niekum, and U. Topcu, "Safe reinforcement learning via shielding," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.
- [36] R. Alur, Principles of Cyber-Physical Systems. Cambridge, MA: The MIT Press, 2015.
- [37] J. C. Knight, "Safety critical systems: Challenges and directions," in *Proceedings of the* 24th International Conference on Software Engineering, 2002, pp. 547–550.
- [38] R. Goebel, R. G. Sanfelice, and A. R. Teel, *Hybrid dynamical systems*. Princeton University Press, 2012.
- [39] P. Tabuada, *Verification and control of hybrid systems: a symbolic approach*. Springer Science & Business Media, 2009.

- [40] N. Leveson, Engineering a Safer World: Systems Thinking Applied to Safety. MIT press, 2011.
- [41] J. D. Schierman, M. D. DeVore, N. D. Richards, N. Gandhi, J. K. Cooper, K. R. Horneman, S. Stoller, and S. Smolka, "Runtime assurance framework development for highly adaptive flight control systems," Barron Associates, Inc. Charlottesville, Tech. Rep., 2015.
- [42] J.-P. Aubin, A. M. Bayen, and P. Saint-Pierre, *Viability theory: new directions*. Springer Science & Business Media, 2011.
- [43] S. Prajna and A. Jadbabaie, "Safety verification of hybrid systems using barrier certificates," in *International Workshop on Hybrid Systems: Computation and Control*. Springer, 2004, pp. 477–492.
- [44] J. Anderson and A. Papachristodoulou, "Advances in computational lyapunov analysis using sum-of-squares programming," *Discrete & Continuous Dynamical Systems-B*, vol. 20, no. 8, p. 2361, 2015.
- [45] S. Prajna, A. Jadbabaie, and G. J. Pappas, "A framework for worst-case and stochastic safety verification using barrier certificates," *IEEE Transactions on Automatic Control*, vol. 52, no. 8, pp. 1415–1428, 2007.
- [46] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin, "A time-dependent hamilton-jacobi formulation of reachable sets for continuous dynamic games," *IEEE Transactions on automatic control*, vol. 50, no. 7, pp. 947–957, 2005.
- [47] J. H. Gillula, S. Kaynama, and C. J. Tomlin, "Sampling-based approximation of the viability kernel for high-dimensional linear sampled-data systems," in *Proceedings of the* 17th international conference on Hybrid systems: computation and control, 2014, pp. 173–182.
- [48] D. P. Bertsekas and I. B. Rhodes, "On the minimax reachability of target sets and target tubes," *Automatica*, vol. 7, no. 2, pp. 233–247, 1971.
- [49] T. Gurriet, "Applied safety critical control," Ph.D. dissertation, California Institute of Technology, 2020.
- [50] I. Mitchell, "A summary of recent progress on efficient parametric approximations of viability and discriminating kernels." in *SNR*@ *CAV*, 2015, pp. 23–31.
- [51] F. Borrelli, A. Bemporad, and M. Morari, *Predictive control for linear and hybrid systems*. Cambridge University Press, 2017.
- [52] M. Nagumo, "Über die lage der integralkurven gewöhnlicher differentialgleichungen," Proceedings of the Physico-Mathematical Society of Japan. 3rd Series, vol. 24, pp. 551– 559, 1942.
- [53] T. Gurriet, M. Mote, A. Singletary, P. Nilsson, E. Feron, and A. D. Ames, "A scalable safety critical control framework for nonlinear systems," *IEEE Access*, vol. 8, pp. 187 249– 187 275, 2020.

- [54] T. Gurriet, M. Mote, A. Singletary, E. Feron, and A. D. Ames, "A scalable controlled set invariance framework with practical safety guarantees," in 2019 IEEE 58th Conference on Decision and Control (CDC). IEEE, 2019, pp. 2046–2053.
- [55] D. Mayne, "An apologia for stabilising terminal conditions in model predictive control," *International Journal of Control*, vol. 86, no. 11, pp. 2090–2095, 2013.
- [56] B. T. Lopez and J. P. How, "Aggressive 3-d collision avoidance for high-speed navigation," in 2017 IEEE International Conference on Robotics and Automation (ICRA), 2017, pp. 5759–5765.
- [57] T. Schouwenaars, J. How, and E. Feron, "Receding horizon path planning with implicit safety guarantees," in *Proceedings of the 2004 American control conference*, vol. 6. IEEE, 2004, pp. 5576–5581.
- [58] U. Borrmann, L. Wang, A. D. Ames, and M. Egerstedt, "Natural motion-based trajectories for automatic spacecraft proximity operation collision avoidance," *IFAC-PapersOnLine*, vol. 48, no. 27, pp. 68–73, 2015.
- [59] W. Clohessy and R. Wiltshire, "Terminal guidance system for satellite rendezvous," *Journal of the Aerospace Sciences*, vol. 27, no. 9, pp. 653–658, 1960.
- [60] S. Bak, K. Manamcheri, S. Mitra, and M. Caccamo, "Sandboxing controllers for cyberphysical systems," in *International Conference on Cyber-physical systems*, ser. ICCPS, 2011.
- [61] D. E. Swihart, A. F. Barfield, E. M. Griffin, R. C. Lehmann, S. C. Whitcomb, B. Flynn, M. A. Skoog, and K. E. Prosser, "Automatic ground collision avoidance system design, integration, & flight test," *IEEE Aerospace and Electronic Systems Magazine*, vol. 26, no. 5, pp. 4–11, 2011.
- [62] E. M. Griffin, R. M. Turner, S. C. Whitcomb, D. E. Swihart, J. M. Bier, K. L. Hobbs, and A. C. Burns, "Automatic ground collision avoidance system design for pre-block 40 f-16 configurations," in *Asia-Pacific International Symposium on Aerospace Technology*, 2012.
- [63] M. Aiello, J. Berryman, J. Grohs, and J. Schierman, "Run-time assurance for advanced flight-critical control systems," in AIAA Guidance, Navigation, and Control Conference, 2010, p. 8041.
- [64] A. Burns, K. Hobbs, J. Bier, and S. Whitcomb, "Advanced capabilities for the analog flight control f-16," in NATO Sensors and Electronics Technology Panel 168 Symposium, 2012.
- [65] N. Minoiu Enache, S. Mammar, M. Netto, and B. Lusetti, "Driver steering assistance for lane-departure avoidance based on hybrid automata and composite lyapunov function," *IEEE Transactions on Intelligent Transportation Systems*, vol. 11, no. 1, pp. 28–39, 2010.
- [66] K. L. Hobbs, "Elicitation and formal specification of run time assurance requirements

for aerospace collision avoidance systems," Ph.D. dissertation, Georgia Institute of Technology, 2020.

- [67] M. Abate, E. Feron, and S. Coogan, "Monitor-based runtime assurance for temporal logic specifications," in 2019 IEEE 58th Conference on Decision and Control (CDC), 2019, pp. 1997–2002.
- [68] C. Tomlin, G. J. Pappas, and S. Sastry, "Conflict resolution for air traffic management: A study in multiagent hybrid systems," *IEEE Transactions on Automatic Control*, vol. 43, no. 4, pp. 509–521, 1998.
- [69] C. J. Tomlin, J. Lygeros, and S. S. Sastry, "A game theoretic approach to controller design for hybrid systems," *Proceedings of the IEEE*, vol. 88, no. 7, pp. 949–970, 2000.
- [70] C. Tomlin, I. Mitchell, and R. Ghosh, "Safety verification of conflict resolution manoeuvres," *IEEE Transactions on Intelligent Transportation Systems*, vol. 2, no. 2, pp. 110–120, 2001.
- [71] S. Liu, M. Watterson, S. Tang, and V. Kumar, "High speed navigation for quadrotors with limited onboard sensing," in 2016 IEEE International Conference on Robotics and Automation (ICRA), 2016, pp. 1484–1491.
- [72] J. H. Gillula, G. M. Hoffmann, H. Huang, M. P. Vitus, and C. J. Tomlin, "Applications of hybrid reachability analysis to robotic aerial vehicles," *The International Journal of Robotics Research*, vol. 30, no. 3, pp. 335–354, 2011.
- [73] G. M. Hoffmann and C. J. Tomlin, "Decentralized cooperative collision avoidance for acceleration constrained vehicles," in 2008 47th IEEE Conference on Decision and Control. IEEE, 2008, pp. 4357–4363.
- [74] C.-F. Lin, J.-C. Juang, and K.-R. Li, "Active collision avoidance system for steering control of autonomous vehicles," *IET Intelligent Transport Systems*, vol. 8, no. 6, pp. 550–557, 2014.
- [75] V. Muthukumaran, R. G. Sanfelice, and G. H. Elkaim, "A hybrid control strategy for autonomous navigation while avoiding multiple obstacles at unknown locations," in 2019 IEEE 15th International Conference on Automation Science and Engineering (CASE). IEEE, 2019, pp. 1042–1047.
- [76] D. Phan, J. Yang, R. Grosu, S. A. Smolka, and S. D. Stoller, "Collision avoidance for mobile robots with limited sensing and limited information about moving obstacles," *Formal Methods in System Design*, vol. 51, no. 1, pp. 62–86, 2017.
- [77] T. Gurriet, M. Tucker, A. Duburcq, G. Boeris, and A. D. Ames, "Towards variable assistance for lower body exoskeletons," *IEEE Robotics and Automation Letters*, vol. 5, no. 1, pp. 266–273, 2019.
- [78] Q. Nguyen and K. Sreenath, "Safety-critical control for dynamical bipedal walking with precise footstep placement," *IFAC-PapersOnLine*, vol. 48, no. 27, pp. 147–154, 2015.

- [79] Q. Nguyen, A. Hereid, J. W. Grizzle, A. D. Ames, and K. Sreenath, "3d dynamic walking on stepping stones with control barrier functions," in 2016 IEEE 55th Conference on Decision and Control (CDC). IEEE, 2016, pp. 827–834.
- [80] 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 2018 ACM/IEEE 9th International Conference on Cyber-Physical Systems (ICCPS). IEEE, 2018, pp. 98–106.
- [81] W. S. Cortez, D. Oetomo, C. Manzie, and P. Choong, "Control barrier functions for mechanical systems: Theory and application to robotic grasping," *IEEE Transactions on Control Systems Technology*, 2019.
- [82] Y. Chen, A. Singletary, and A. D. Ames, "Guaranteed obstacle avoidance for multi-robot operations with limited actuation: a control barrier function approach," *IEEE Control Systems Letters*, vol. 5, no. 1, pp. 127–132, 2020.
- [83] D. Pickem, L. Wang, P. Glotfelter, Y. Diaz-Mercado, M. Mote, A. Ames, E. Feron, and M. Egerstedt, "Safe, remote-access swarm robotics research on the robotarium," *arXiv* preprint arXiv:1604.00640, 2016.
- [84] D. Pickem, P. Glotfelter, L. Wang, M. Mote, A. Ames, E. Feron, and M. Egerstedt, "The robotarium: A remotely accessible swarm robotics research testbed," in 2017 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2017, pp. 1699– 1706.
- [85] S. Wilson, P. Glotfelter, L. Wang, S. Mayya, G. Notomista, M. Mote, and M. Egerstedt, "The robotarium: Globally impactful opportunities, challenges, and lessons learned in remote-access, distributed control of multirobot systems," *IEEE Control Systems Magazine*, vol. 40, no. 1, pp. 26–44, 2020.
- [86] S. Wilson, P. Glotfelter, S. Mayya, G. Notomista, Y. Emam, X. Cai, and M. Egerstedt, "The robotarium: Automation of a remotely accessible, multi-robot testbed," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2922–2929, 2021.
- [87] Y. Chen, M. Jankovic, M. Santillo, and A. D. Ames, "Backup control barrier functions: Formulation and comparative study," *arXiv preprint arXiv:2104.11332*, 2021.
- [88] A. Isaly, B. C. Allen, R. G. Sanfelice, and W. E. Dixon, "Zeroing control barrier functions for safe volitional pedaling in a motorized cycle," *IFAC-PapersOnLine*, vol. 53, no. 5, pp. 218–223, 2020.
- [89] C. M. Holloway, *Understanding the Overarching Properties*. National Aeronautics and Space Administration, Langley Research Center, 2019.
- [90] K. L. Hobbs, J. Davis, L. Wagner, and E. M. Feron, "Formal specification and analysis of spacecraft collision avoidance run time assurance requirements," in 2021 IEEE Aerospace Conference. IEEE, 2021, pp. 1–19.

- [91] D. Swihart, A. Barfield, E. Griffin, R. Lehmann, S. Whitcomb, M. Skoog, B. Flynn, and K. Prosser, "Design, integration and flight test of an autonomous ground collision avoidance system," *Gyroscopy and Navigation*, vol. 2, no. 2, pp. 84–91, 2011.
- [92] J. B. Lyons, N. T. Ho, W. E. Fergueson, G. G. Sadler, S. D. Cals, C. E. Richardson, and M. A. Wilkins, "Trust of an automatic ground collision avoidance technology: A fighter pilot perspective," *Military Psychology*, vol. 28, no. 4, pp. 271–277, 2016.
- [93] W. A. Olson, "Airborne collision avoidance system x," Massachusetts Institute of Technology, Lincoln Laboratory, Tech. Rep., 2015.
- [94] J. D. Lee and K. A. See, "Trust in automation: Designing for appropriate reliance," *Human Factors*, vol. 46, no. 1, pp. 50–80, 2004.
- [95] R. Turner, R. Lehmann, J. Wadley, D. Kidd, D. Swihart, J. Bier, and K. Hobbs, "Automatic aircraft collision avoidance algorithm design for fighter aircraft," in Asia-Pacific International Symposium on Aerospace Technology, 2012.
- [96] M. A. Skoog, K. Prosser, and L. Hook, "Ground collision avoidance system (iGCAS)," 4 2017, uS Patent 9,633,567.
- [97] P. Sorokowski, M. Skoog, S. Burrows, and S. Thomas, "Small UAV automatic ground collision avoidance system design considerations and flight test results," 2015.
- [98] J. D. Carpenter, "Simulation and piloted simulator study of an automatic ground collision avoidance system for performance limited aircraft," Air Force Institute of Technology, Wright-Patterson AFB, OH, Tech. Rep., 2019.
- [99] A. W. Suplisson, "Optimal Recovery Trajectories for Automatic Ground Collision Avoidance Systems (Auto GCAS)," Air Force Institute of Technology, Tech. Rep., 2015.
- [100] A. Burns, D. Harper, A. F. Barfield, S. Whitcomb, and B. Jurusik, "Auto gcas for analog flight control system," in 2011 IEEE/AIAA 30th Digital Avionics Systems Conference. IEEE, 2011, pp. 8C5–1.
- [101] P. Madhavan and D. A. Wiegmann, "Similarities and differences between human-human and human-automation trust: an integrative review," *Theoretical Issues in Ergonomics Science*, vol. 8, no. 4, pp. 277–301, 2007.
- [102] N. T. Ho, G. G. Sadler, L. C. Hoffmann, J. B. Lyons, and W. W. Johnson, "Trust of a military automated system in an operational context," *Military Psychology*, vol. 29, no. 6, pp. 524–541, 2017.
- [103] K. A. Hoff and M. Bashir, "Trust in automation: Integrating empirical evidence on factors that influence trust," *Human Factors*, vol. 57, no. 3, pp. 407–434, 2015.
- [104] J. Milton and J. C. Arnold, *Introduction to Probability and Statistics*. The McGraw-Hill Companies, Inc, New York, NY, 2003.
- [105] L. Sha, R. Rajkumar, and M. Gagliardl, "A software architecture for dependable and

evolvable industrial computing systems." Carnegie-Mellon University Software Engineering Institute, Pittsburgh, PA, Tech. Rep., 1995.

- [106] J. G. Rivera, A. A. Danylyszyn, C. B. Weinstock, L. R. Sha, and M. J. Gagliardi, "An architectural description of the simplex architecture." Carnegie-Mellon University Software Engineering Institute, Pittsburgh, PA, Tech. Rep., 1996.
- [107] C. Reis, A. Barth, and C. Pizano, "Browser security: lessons from google chrome," *Commun. ACM.*, vol. 52, pp. 45–49, 2009.
- [108] "Standard Practice for Methods to Safely Bound Flight Behavior of Unmanned Aircraft Systems Containing Complex Functions," ASTM International, Standard, 2017.
- [109] J. D. Schierman, M. D. DeVore, N. D. Richards, and M. A. Clark, "Runtime assurance for autonomous aerospace systems," *Journal of Guidance, Control, and Dynamics*, vol. 43, no. 12, pp. 2205–2217, 2020.
- [110] S. E. Jones, A. K. Petry, C. A. Eger, R. M. Turner, E. M. Griffin *et al.*, "Automatic integrated collision avoidance system," in *17th Australian International Aerospace Congress: AIAC 2017.* Engineers Australia, Royal Aeronautical Society, 2017, p. 13.
- [111] R. M. Turner, A. J. Albert, E. M. Griffin, A. C. Burns, F. Barfield, K. L. Price et al., "Automatic collision avoidance technology," in AIAC18: 18th Australian International Aerospace Congress (2019): HUMS-11th Defence Science and Technology (DST) International Conference on Health and Usage Monitoring (HUMS 2019): ISSFD-27th International Symposium on Space Flight Dynamics (ISSFD). Engineers Australia, Royal Aeronautical Society., 2019, p. 495.
- [112] M. A. Skoog, L. R. Hook, and W. Ryan, "Leveraging astm industry standard f3269-17 for providing safe operations of a highly autonomous aircraft," in 2020 IEEE Aerospace Conference. IEEE, 2020, pp. 1–7.
- [113] B. Eller, P. Stanfill, R. Turner, S. Whitcomb, D. Swihart, A. Burns, and K. Hobbs, "Test and evaluation of a modified f-16 analog flight control computer," in *AIAA Infotech@ Aerospace Conference*, 2013, p. 4726.
- [114] J. Wadley, S. Jones, D. Stoner, E. Griffin, D. Swihart, K. Hobbs, A. Burns, and J. Bier, "Development of an automatic aircraft collision avoidance system for fighter aircraft," in *AIAA Infotech@ Aerospace Conference*, 2013, p. 4727.
- [115] S. Bak and K. Hobbs, "Efficient n-to-n collision detection for space debris using 4d aabb trees," *arXiv preprint arXiv:1901.10475*, 2019.
- [116] K. Hobbs, P. Heidlauf, A. Collins, and S. Bak, "Space debris collision detection using reachability," in 5th International Workshop on Applied Verification of Continuous and Hybrid Systems, ser. EPiC Series in Computing. EasyChair, 2018.
- [117] L. R. Hook, M. Clark, D. Sizoo, M. A. Skoog, and J. Brady, "Certification strategies using run-time safety assurance for part 23 autopilot systems," in 2016 IEEE Aerospace

Conference. IEEE, 2016, pp. 1–10.

- [118] "Standard Practice for Methods to Safely Bound Behavior of Aircraft Systems Containing Complex Functions Using Run-Time Assurance," ASTM International, Standard, 2021.
- [119] U. Mehmood, S. Bak, S. A. Smolka, and S. D. Stoller, "Safe cps from unsafe controllers," *arXiv preprint arXiv:2102.12981*, 2021.
- [120] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in 2019 18th European Control Conference (ECC). IEEE, 2019, pp. 3420–3431.
- [121] J. Aubin, "Viability theory. modern birkhäuser classics," 2009.
- [122] 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*, vol. 62, no. 8, pp. 3861–3876, Aug 2017.
- [123] X. Xu, P. Tabuada, J. W. Grizzle, and A. D. Ames, "Robustness of control barrier functions for safety critical control," *IFAC-PapersOnLine*, vol. 48, no. 27, pp. 54–61, 2015.
- [124] S. Prajna, "Barrier certificates for nonlinear model validation," *Automatica*, vol. 42, no. 1, pp. 117–126, 2006.
- [125] M. Maghenem and R. G. Sanfelice, "Multiple barrier function certificates for weak forward invariance in hybrid inclusions," in 2019 IEEE 58th Conference on Decision and Control (CDC). IEEE, 2019, pp. 6319–6324.
- [126] A. D. Ames, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs with application to adaptive cruise control," in 53rd IEEE Conference on Decision and Control. IEEE, 2014, pp. 6271–6278.
- [127] P. Glotfelter, J. Cortés, and M. Egerstedt, "Nonsmooth barrier functions with applications to multi-robot systems," *IEEE control systems letters*, vol. 1, no. 2, pp. 310–315, 2017.
- [128] A. Agrawal and K. Sreenath, "Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation." in *Robotics: Science and Systems*, 2017.
- [129] C. Santoyo, M. Dutreix, and S. Coogan, "A barrier function approach to finite-time stochastic system verification and control," *Automatica*, vol. 125, p. 109439, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/ S0005109820306415
- [130] M. Srinivasan, S. Coogan, and M. Egerstedt, "Control of multi-agent systems with finite time control barrier certificates and temporal logic," in 2018 IEEE Conference on Decision and Control (CDC), 2018, pp. 1991–1996.
- [131] L. Lindemann and D. V. Dimarogonas, "Control barrier functions for signal temporal logic tasks," *IEEE control systems letters*, vol. 3, no. 1, pp. 96–101, 2018.
- [132] G. Yang, C. Belta, and R. Tron, "Continuous-time signal temporal logic planning with

control barrier functions," in 2020 American Control Conference (ACC). IEEE, 2020, pp. 4612–4618.

- [133] Q. Nguyen and K. Sreenath, "Exponential control barrier functions for enforcing high relative-degree safety-critical constraints," in 2016 American Control Conference (ACC). IEEE, 2016, pp. 322–328.
- [134] X. Tan, W. Shaw Cortez, and D. V. Dimarogonas, "High-order barrier functions: Robustness, safety and performance-critical control," *IEEE Transactions on Automatic Control*, pp. 1–1, 2021.
- [135] W. Xiao and C. Belta, "Control barrier functions for systems with high relative degree," in 2019 IEEE 58th Conference on Decision and Control (CDC), 2019, pp. 474–479.
- [136] H. Seywald and R. Kumar, "Desensitized optimal trajectories," *Analytical Mechanics Associates Rept*, pp. 03–16, 2003.
- [137] M. Mote, M. Egerstedt, E. Feron, A. Bylard, and M. Pavone, "Collision-inclusive trajectory optimization for free-flying spacecraft," *Journal of Guidance, Control, and Dynamics*, pp. 1–12, 2020.
- [138] K. Zhou and J. C. Doyle, *Essentials of robust control*. Prentice hall Upper Saddle River, NJ, 1998, vol. 104.
- [139] S. Höfer, K. Bekris, A. Handa, J. C. Gamboa, M. Mozifian, F. Golemo, C. Atkeson, D. Fox, K. Goldberg, J. Leonard *et al.*, "Sim2real in robotics and automation: Applications and challenges," *IEEE Transactions on Automation Science and Engineering*, vol. 18, no. 2, pp. 398–400, 2021.
- [140] M. Abate and S. Coogan, "Enforcing safety at runtime for systems with disturbances," in 2020 59th IEEE Conference on Decision and Control (CDC), 2020, pp. 2038–2043.
- [141] M. Abate, M. Mote, E. Feron, and S. Coogan, "Verification and runtime assurance for dynamical systems with uncertainty," in *Proceedings of the 24th International Conference on Hybrid Systems: Computation and Control*, ser. HSCC '21. New York, NY, USA: Association for Computing Machinery, 2021. [Online]. Available: https://doi.org/10.1145/3447928.3456656
- [142] C. Llanes, M. Abate, and S. Coogan, "Safety from in-the-loop reachability for cyberphysical systems," in 2021 Workshop on Computation-Aware Algorithmic Design for Cyber-Physical Systems (CAADCPS). CPS-IoT week, 2021.
- [143] —, "Safety from fast, in-the-loop reachability with application to UAVs," 2021, in submission.
- [144] S. Coogan, M. Arcak, and A. A. Kurzhanskiy, "Mixed monotonicity of partial first-infirst-out traffic flow models," in 2016 IEEE 55th Conference on Decision and Control (CDC), 2016, pp. 7611–7616.

- [145] H. L. Smith, "The discrete dynamics of monotonically decomposable maps," *Journal of Mathematical Biology*, vol. 53, no. 4, p. 747, 2006.
- [146] C. Muñoz, A. Narkawicz, G. Hagen, J. Upchurch, A. Dutle, M. Consiglio, and J. Chamberlain, "Daidalus: detect and avoid alerting logic for unmanned systems," in 2015 IEEE/AIAA 34th Digital Avionics Systems Conference (DASC). IEEE, 2015, pp. 5A1–1.
- [147] C. Baier, J. Katoen, and K. Larsen, Principles of Model Checking. MIT Press, 2008.
- [148] K. Y. Rozier, "Specification: The biggest bottleneck in formal methods and autonomy," in Working Conference on Verified Software: Theories, Tools, and Experiments. Springer, 2016, pp. 8–26.
- [149] R. Avram, X. Zhang, J. A. Muse, and M. Clark, "Nonlinear adaptive control of quadrotor UAVs with run-time safety assurance," in AIAA Guidance, Navigation, and Control Conference, 2017, p. 1896.
- [150] M. Bodson, J. Lehoczky, R. Rajkumar, L. Sha, and J. Stephan, "Analytic redundancy for software fault-tolerance in hard real-time systems," in *Foundations of Dependable Computing*. Springer, 1994, pp. 183–212.
- [151] L. Sha, R. Rajkumar, and M. Gagliardi, "Evolving dependable real-time systems," in 1996 IEEE Aerospace Applications Conference. Proceedings, vol. 1. IEEE, 1998, pp. 335–346.
- [152] D. Seto, B. Krogh, L. Sha, and A. Chutinan, "The simplex architecture for safe online control system upgrades," in *Proceedings of the 1998 American Control Conference. ACC* (*IEEE Cat. No. 98CH36207*), vol. 6. IEEE, 1998, pp. 3504–3508.
- [153] L. Sha, "Using simplicity to control complexity," *IEEE Software*, vol. 18, no. 4, pp. 20–28, 2001.
- [154] S. Bak, D. K. Chivukula, O. Adekunle, M. Sun, M. Caccamo, and L. Sha, "The systemlevel simplex architecture for improved real-time embedded system safety," in 2009 15th IEEE Real-Time and Embedded Technology and Applications Symposium. IEEE, 2009, pp. 99–107.
- [155] S. Bak, A. Greer, and S. Mitra, "Hybrid cyberphysical system verification with simplex using discrete abstractions," in 2010 16th IEEE Real-Time and Embedded Technology and Applications Symposium. IEEE, 2010, pp. 143–152.
- [156] R. Alur, C. Courcoubetis, N. Halbwachs, T. A. Henzinger, P.-H. Ho, X. Nicollin, A. Olivero, J. Sifakis, and S. Yovine, "The algorithmic analysis of hybrid systems," *Theoretical computer science*, vol. 138, no. 1, pp. 3–34, 1995.
- [157] D. K. Kaynar, N. Lynch, R. Segala, and F. Vaandrager, "The theory of timed i/o automata," *Synthesis Lectures on Distributed Computing Theory*, vol. 1, no. 1, pp. 1–137, 2010.
- [158] S. Mitra, "A verification framework for hybrid systems," Ph.D. dissertation, Massachusetts Institute of Technology, 2007.
- [159] L. R. Hook, M. Skoog, M. Garland, W. Ryan, D. Sizoo, and J. VanHoudt, "Initial
considerations of a multi-layered run time assurance approach to enable unpiloted aircraft," in *2018 IEEE Aerospace Conference*. IEEE, 2018, pp. 1–11.

- [160] K. H. Gross, "Evaluation of verification approaches applied to nonlinear system control," Master's thesis, Air Force Institute of Technology, 3 2016.
- [161] K. L. Hobbs, I. Perez, A. Fifarek, and E. M. Feron, "Formal verification of system states for spacecraft automatic maneuvering," in *AIAA Scitech 2019 Forum*, 2019, p. 1187.
- [162] I. Lee, H. Ben-Abdallah, S. Kannan, M. Kim, O. Sokolsky, and M. Viswanathan, "A monitoring and checking framework for run-time correctness assurance," 1998.
- [163] I. Lee, S. Kannan, M. Kim, O. Sokolsky, and M. Viswanathan, "Runtime assurance based on formal specifications," *Departmental Papers (CIS)*, p. 294, 1999.
- [164] C. Torens and F.-M. Adolf, "Formal requirements and model-checking for V&V automation of a RPAS mission management system," in AIAA Infotech @ Aerospace, Kissimmee, FL, 1 2015, pp. 1–13.
- [165] M. Wooldridge, "Agents and software engineering," AI\* IA Notizie, vol. 11, no. 3, pp. 31–37, 1998.
- [166] C. A. R. Hoare, "An axiomatic basis for computer programming," *Communications of the ACM*, vol. 12, no. 10, pp. 576–580, 1969.
- [167] M. Davis, G. Logemann, and D. Loveland, "A machine program for theorem-proving," *Communications of the ACM*, vol. 5, no. 7, pp. 394–397, 1962.
- [168] P. Cousot and R. Cousot, "Abstract interpretation frameworks," *Journal of logic and computation*, vol. 2, no. 4, pp. 511–547, 1992.
- [169] L. De Moura and N. Bjørner, "Satisfiability modulo theories: Introduction and applications," *Communications of the ACM*, vol. 54, no. 9, pp. 69–77, 2011.
- [170] E. M. Clarke and E. A. Emerson, "Design and synthesis of synchronization skeletons using branching time temporal logic," in *Workshop on Logic of Programs*. Springer, 1981, pp. 52–71.
- [171] R. Koymans, "Specifying real-time properties with metric temporal logic," *Real-time systems*, vol. 2, no. 4, pp. 255–299, 1990.
- [172] O. Maler and D. Nickovic, "Monitoring temporal properties of continuous signals," in Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems. Springer, 2004, pp. 152–166.
- [173] K. Gross, J. Hoffman, M. Clark, E. Swenson, R. Cobb, M. Whalen, and L. Wagner, "Evaluation of formal methods tools applied to a 6u cubesat attitude control system," in *AIAA SPACE 2015 Conference and Exposition*, 2015, p. 4529.
- [174] K. H. Gross, M. A. Clark, J. A. Hoffman, E. D. Swenson, and A. W. Fifarek, "Run-time assurance and formal methods analysis applied to nonlinear system control," *Journal of Aerospace Information Systems*, vol. 14, no. 4, pp. 232–246, 2017.

- [175] K. H. Gross, M. Clark, J. A. Hoffman, A. Fifarek, K. Rattan, E. Swenson, M. Whalen, and L. Wagner, "Formally verified run time assurance architecture of a 6U cubesat attitude control system," in *AIAA Infotech@ Aerospace*, 2016, p. 0222.
- [176] B. Könighofer, M. Alshiekh, R. Bloem, L. Humphrey, R. Könighofer, U. Topcu, and C. Wang, "Shield synthesis," *Formal Methods in System Design*, vol. 51, no. 2, pp. 332– 361, 2017.
- [177] C. S. Păsăreanu, D. Giannakopoulou, M. G. Bobaru, J. M. Cobleigh, and H. Barringer, "Learning to divide and conquer: applying the l\* algorithm to automate assume-guarantee reasoning," *Formal Methods in System Design*, vol. 32, no. 3, pp. 175–205, 2008.
- [178] D. A. Peled, Software Reliability Methods. Springer Science & Business Media, 2013.
- [179] K. H. Gross, A. W. Fifarek, and J. A. Hoffman, "Incremental formal methods based design approach demonstrated on a coupled tanks control system," in *Proceedings of the 2016 IEEE 17th International Symposium on High Assurance Systems Engineering (HASE)*. IEEE Computer Society, 2016, pp. 181–188.
- [180] M. W. Whalen, A. Gacek, D. Cofer, A. Murugesan, M. P. Heimdahl, and S. Rayadurgam, "Your" what" is my" how": Iteration and hierarchy in system design," *IEEE software*, vol. 30, no. 2, pp. 54–60, 2012.
- [181] T. A. Henzinger, M. Minea, and V. Prabhu, "Assume-guarantee reasoning for hierarchical hybrid systems," in *International Workshop on Hybrid Systems: Computation and Control*. Springer, 2001, pp. 275–290.
- [182] J. D. Schierman, D. Neal, E. Wong, and A. K. Chicatelli, "Runtime assurance protection for advanced turbofan engine control," in 2018 AIAA Guidance, Navigation, and Control Conference, 2018, p. 1112.
- [183] M. Clark, X. Koutsoukos, J. Porter, R. Kumar, G. Pappas, O. Sokolsky, I. Lee, and L. Pike, "A study on run time assurance for complex cyber physical systems," Air Force Research Laboratory Aerospace Systems Directorate, Tech. Rep., 2013.
- [184] K. Hobbs, J. Davis, L. Wagner, and E. Feron, "Formal specification and analysis of spacecraft collision avoidance run time assurance requirements," in 2021 IEEE Aerospace Conference, vol. 1. IEEE, 2021, pp. 1–16.
- [185] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. MIT press, 2018.
- [186] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot *et al.*, "Mastering the game of go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [187] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton *et al.*, "Mastering the game of go without human knowledge,"

Nature, vol. 550, no. 7676, pp. 354-359, 2017.

- [188] O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev *et al.*, "Grandmaster level in starcraft ii using multi-agent reinforcement learning," *Nature*, vol. 575, no. 7782, pp. 350–354, 2019.
- [189] J. Ibarz, J. Tan, C. Finn, M. Kalakrishnan, P. Pastor, and S. Levine, "How to train your robot with deep reinforcement learning: lessons we have learned," *The International Journal of Robotics Research*, p. 0278364920987859.
- [190] W. Saunders, G. Sastry, A. Stuhlmüller, and O. Evans, "Trial without error: Towards safe reinforcement learning via human intervention," in *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2018, pp. 2067–2069.
- [191] T. J. Perkins and A. G. Barto, "Lyapunov design for safe reinforcement learning," *Journal of Machine Learning Research*, vol. 3, no. Dec, pp. 803–832, 2002.
- [192] Y. Chow, O. Nachum, E. Duenez-Guzman, and M. Ghavamzadeh, "A lyapunov-based approach to safe reinforcement learning," in *Advances in Neural Information Processing Systems*, 2018, pp. 8092–8101.
- [193] 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*, vol. 33, 2019, pp. 3387–3395.
- [194] A. Murugesan, M. Moghadamfalahi, and A. Chattopadhyay, "Formal methods assisted training of safe reinforcement learning agents," in NASA Formal Methods Symposium. Springer, 2019, pp. 333–340.
- [195] N. Fulton and A. Platzer, "Safe reinforcement learning via formal methods: Toward safe control through proof and learning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, 2018, pp. 6485–6492.
- [196] S. Cristello, Ariadne's Thread: The Needle, The Haystack, The Thread // Arts Club of Chicago, 2018 (accessed December 7, 2020). [Online]. Available: https: //theseenjournal.org/ariadnes-thread/
- [197] N. M. Jensen, *Philosophy in labyrinths*, (accessed December 7, 2020). [Online]. Available: http://www.lavigne.dk/labyrinth/e1a\_phil.htm
- [198] A. Stieger, *Myth and Creativity: Ariadne's Thread and a Path Through the Labyrinth*, 2014 (accessed December 7, 2020). [Online]. Available: https://www.creativitypost.com/ article/myth\_and\_creativity\_ariadnes\_thread\_and\_a\_path\_through\_the\_labyrinth
- [199] *Contes de Perrault / Le Petit Poucet*, (accessed November 9, 2020). [Online]. Available: https://fr.wikisource.org/wiki/Contes\_de\_Perrault\_(%C3%A9d.\_1902)/Le\_Petit\_Poucet
- [200] I. Opie and Peter, "The classic fairy tales," Oxford University Press, p. 21, 1974.
- [201] J. Flowers and P. Shalgar, The incredible killer tree that

SOS 2015 broadcasts neighbours. (accessed an to its Decem-7. 2020). [Online]. Available: https://uk.blastingnews.com/world/2015/10/ ber the-incredible-killer-tree-that-broadcasts-an-sos-to-its-neighbours-00610971.html

- [202] R. G. Pisano and T. I. Storer, "Burrows and feeding of the norway rat," *Journal of Mammalogy*, vol. 29, no. 4, pp. 374–383, 1948. [Online]. Available: http://www.jstor.org/stable/1375126
- [203] R. M. Sapolsky, "Stress in the wild," *Scientific American*, vol. 262, no. 1, pp. 116–123, 1990.
- [204] K. Wiegers and J. Beatty, Software requirements. Pearson Education, 2013.
- [205] E. Garone, S. Di Cairano, and I. Kolmanovsky, "Reference and command governors for systems with constraints: A survey on theory and applications," *Automatica*, vol. 75, pp. 306–328, 2017.
- [206] S. Coogan, "Mixed monotonicity for reachability and safety in dynamical systems," in 2020 59th IEEE Conference on Decision and Control (CDC), 2020, pp. 5074–5085.
- [207] G. Enciso, H. Smith, and E. Sontag, "Nonmonotone systems decomposable into monotone systems with negative feedback," *Journal of Differential Equations*, vol. 224, no. 1, pp. 205 – 227, 2006. [Online]. Available: http://www.sciencedirect.com/science/ article/pii/S002203960500210X
- [208] P.-J. Meyer, A. Devonport, and M. Arcak, "Tira: Toolbox for interval reachability analysis," in *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control*, 2019, pp. 224–229.
- [209] S. Coogan and M. Arcak, "Efficient finite abstraction of mixed monotone systems," in Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control, 2015, pp. 58–67.
- [210] H. Smith, Monotone Dynamical Systems: An Introduction to the Theory of Competitive and Cooperative Systems, ser. Mathematical surveys and monographs. American Mathematical Society, 2008. [Online]. Available: https://books.google.com/books?id= vOfNAwAAQBAJ
- [211] D. Angeli and E. D. Sontag, "Monotone control systems," *IEEE Transactions on Automatic Control*, vol. 48, no. 10, pp. 1684–1698, Oct 2003.
- [212] M. Abate and S. Coogan, "Computing robustly forward invariant sets for mixed-monotone systems," in 2020 59th IEEE Conference on Decision and Control (CDC), 2020, pp. 4553– 4559.
- [213] S. Coogan and M. Arcak, "Stability of traffic flow networks with a polytree topology," *Automatica*, vol. 66, pp. 246–253, 2016.
- [214] M. Abate, M. Dutreix, and S. Coogan, "Tight decomposition functions for continuous-

time mixed-monotone systems with disturbances," *IEEE Control Systems Letters*, vol. 5, no. 1, pp. 139–144, 2021.

## **Author Biographies**



**Kerianne Hobbs** is the Safe Autonomy Lead at the Autonomy Capability Team (ACT3) at the Air Force Research Laboratory. There she investigates rigorous specification, analysis, and bounding techniques to enable certification of autonomous and learning controllers for aircraft and spacecraft applications. Her previous experience includes work in automatic collision avoidance at AFRL from 2011-2014, and Autonomy Verification and Validation research

from 2012-2020. Kerianne has a BS in Aerospace Engineering from Embry-Riddle Aeronautical University, an MS in Astronautical Engineering from the Air Force Institute of Technology, and a Ph.D. in Aerospace Engineering from the Georgia Institute of Technology.



**Mark L. Mote** is a co-founder of Pytheia. He received his Ph.D. in Robotics from the Georgia Institute of Technology in 2021, and his B.S. and M.S. degrees in Aerospace Engineering from the Georgia Institute of Technology in 2015 and 2018. His current research interests are in safe autonomy, perception, optimization, trajectory planning, and run time assurance for safety-critical systems.



**Matthew C. L. Abate** (Student Member, IEEE) received a B.A. degree in engineering sciences and a B.E degree in mechanical engineering from Dartmouth College, Hanover, NH, USA, in 2017. He received a M.S. degree in Mechanical Engineering from Georgia Tech, Atlanta, GA, USA, in 2018, and a M.S. degree in Electrical and Computer Engineering from the same institution in 2020. He is currently pursuing a Ph.D. degree in Robotics at Georgia Tech,

supervised by Dr.'s Samuel Coogan and Eric Feron. His research interest is in the area of safety and verification for uncertain dynamical systems.



**Sam Coogan** is assistant professor at Georgia Tech in the School of Electrical and Computer Engineering and the School of Civil and Environmental Engineering. Prior to joining Georgia Tech in 2017, he was an assistant professor in the Electrical Engineering Department at UCLA. He received the B.S. degree in Electrical Engineering from Georgia Tech and the M.S. and Ph.D. degrees in Electrical Engineering from the University of California, Berkeley. His research

is in the area of dynamical systems and autonomy and focuses on developing scalable tools for verification and control of networked, cyber-physical systems with an emphasis on transportation systems. He received the Donald P Eckman Award from the American Automatic Control Council in 2020, a Young Investigator Award from the Air Force Office of Scientific Research in 2019, a CAREER award from NSF in 2018, and the Outstanding Paper Award for the IEEE Transactions on Control of Network Systems in 2017.



**Eric Feron** is professor of Electrical, Computer, and Mechanical Engineering at King Abdullah University of Science and Technology (KAUST), Saudi Arabia. Eric Feron's interests are with aerospace systems, ranging from aircraft to space systems, with an emphasis on robotic autonomy and artificial intelligence. Prior to his KAUST appointment, Eric Feron was a professor of Aeronautics and Astronautics at the Massachusetts Institute of Technology (MIT) and a professor

of Aerospace Engineering at the Georgia Institute of Technology (Georgia Tech), United States. Eric Feron also works on occasion with the Ecole Nationale de l'Aviation Civile (ENAC) and the Institut Supérieur de l'Aéronautique et de l'Espace (ISAE-Supaéro), France, where he was born.