Global Asymptotic Stability of Nonlinear Systems while Exploiting Properties of Uncertified Feedback Controllers via Opportunistic Switching

Paul K. Wintz Ricardo G. Sanfelice João P. Hespanha

Abstract— **A hybrid control strategy is introduced that renders a compact set uniformly globally asymptotically stable for a continuous-time plant by switching between a Lyapunov-certified feedback controller and an uncertified controller. This control strategy allows for the opportunistic use of a controller that has desirable performance but lacks a Lyapunov certificate. A pair of tunable threshold functions determine conditions for switching between the controllers. To establish global uniform asymptotic stability, a nonsmooth Lyapunov function is constructed for the closed-loop hybrid system using an auxiliary memory variable and the Lyapunov certificate associated with the certified controller. Examples illustrate improvements to control effort and rate of convergence resulting from the proposed hybrid control strategy when applied to state-feedback and model-predictive control.**

I. Introduction

For some control design problems, a single continuous state-feedback controller cannot simultaneously satisfy all design requirements. In particular, the design of a globally asymptotically stabilizing controller for a nonlinear system requires multiple controllers if the system violates Brockett's conditions [1] or has certain topological obstructions, such as systems with states that live in certain topological manifolds [2], [3]. Such challenges have motivated the use of supervisory algorithms that selects between multiple controllers [4], [5], [6], [7]. For a system with one controller that renders a set-point only locally asymptotically stable, and another controller that guides the system into the vicinity of the set-point, a class of supervisors called *uniting controllers* produce global asymptotic stability of the set-point by selecting the first controller near the set-point and the second controller away from it [8], [9], [10]. Similarly, a switching strategy for a family of Lyapunov-certified controllers to achieve asymptotic stability is presented in [11]. For systems with constraints, supervisors are used to provide a backup controller that guarantees safety when the primary controller lacks such a guarantee [12]. We do not, however, know of a control strategy that allows for the opportunistic use of an uncertified controller to improve performance while preserving asymptotic stability.

In this paper, a hybrid control strategy is proposed to fill that gap by uniting a Lyapunov-certified controller with an uncertified controller via opportunistic switching. We consider an unconstrained continuous-time nonlinear plant with state space $\mathbb{R}^n$ and a given compact set $\mathcal{A} \subset \mathbb{R}^n$ that must be rendered uniformly globally asymptotically stable (UGAS). If, for a given controller, $\mathcal{A}$ is rendered UGAS and a Lyapunov function is known for the closed-loop system, then we call the controller *Lyapunov-certified*. A Lyapunov function exists for every sufficiently regular closed-loop system such that $\mathcal{A}$ is UGAS [13, Theorem 4.17], but construction of such a function is often difficult. Some controllers that do not cause $\mathcal{A}$ to be UGAS may, however, have otherwise desirable properties. A controller for which a Lyapunov function is unavailable is called *uncertified*. The novel contribution of this paper is the introduction of a hybrid control strategy, such that—given a continuous Lyapunov-certified feedback controller κ_0 and a continuous uncertified controller κ_1—the set $\mathcal{A}$ is UGAS for the resulting closed-loop system, the controller κ_1 is preferred over κ_0, and Zeno behavior does not occur.

As an example where using an uncertified controller is advantageous, suppose κ_1 is a linear quadratic regulator (LQR) for the linearization of a nonlinear system about the origin. Because an LQR feedback is an optimal control law, κ_1 is approximately optimal (by some measure) near the origin. The basin of attraction under κ_1 is an open neighborhood of the origin, but far from the origin, nonlinear dynamics dominate, so the linearization is inaccurate and κ_1 will generally not produce global stability. Our switching logic lets us use κ_1—without knowledge of the actual basin of attraction—in conjunction with a Lyapunov-certified controller to achieve global convergence to the origin and minimize costs locally. A detailed consideration of this example is given in Example 2. We envision that our switching logic could be particularly useful for reinforcement learning control, which often demonstrates good results empirically, but for which it is often difficult to produce Lyapunov certificates [14].

The remainder of the paper proceeds as follows. Section II introduces notation and preliminary concepts. Section III describes our proposed switching logic and the resulting closed-loop system. Section III-C contains theoretical results. Several examples, throughout, illustrate the behavior of the closed-loop system.

Paul K. Wintz is with the Department of Applied Mathematics, University of California, Santa Cruz, CA 95064 (pwintz@ucsc.edu), Ricardo G. Sanfelice is with the Department of Electrical and Computer Engineering, University of California, Santa Cruz, CA 95064 (ricardo@ucsc.edu), João P. Hespanha is with the Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106 (hespanha@ece.ucsb.edu).

This research was supported by the National Science Foundation under grant nos. ECS-1710621, CNS-1544396, and CNS-2039054; by the Air Force Office of Scientific Research under grant nos. FA9550-19-1-0053, FA9550-19-1-0169, and FA9550-20-1-0238; by the Army Research Office under grant no. W911NF-20-1-0253; and by the U.S. Office of Naval Research under the MURI grant no. N00014-16-1-2710.

II. Preliminaries

We denote the nonnegative real numbers by $\mathbb{R}_{\geq 0}$, and the natural numbers (0 inclusive) by $\mathbb{N}$. For $x, y \in \mathbb{R}^n$, $\langle x, y \rangle$ denotes the inner product between x and y. We write $[x^\top \; y^\top]^\top$ as (x, y). For a set S, the interior of S is denoted $\operatorname{int} S$. For a continuously differentiable function $f : \mathbb{R}^n \to \mathbb{R}$, the gradient of f at x is denoted $\nabla f(x)$. Given $x \in \mathbb{R}^n$ and a nonempty set $\mathcal{A} \subset \mathbb{R}^n$, the distance from x to $\mathcal{A}$ is $|x|_{\mathcal{A}} := \inf_{y \in \mathcal{A}} |y - x|$. A continuous function $\alpha : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is said to be in class $\mathcal{K}_\infty$ if $\alpha(0) = 0$, α is strictly increasing, and $\lim_{r \to \infty} \alpha(r) = \infty$. Given a nonempty set $\mathcal{A} \subset \mathbb{R}^n$, a function $V : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ is said to be *positive definite* with respect to $\mathcal{A}$ if $V(x) > 0$ for all $x \in \mathbb{R}^n \setminus \mathcal{A}$ and $V(x) = 0$ for all $x \in \mathcal{A}$.

A. Hybrid Systems

We consider hybrid systems modeled in the form [15], [9]

$$\mathcal{H} \begin{cases} \dot{x} \;\, = f(x) & x \in C \\ x^+ = g(x) & x \in D \end{cases} \tag{1}$$

with state variable $x \in \mathbb{R}^n$, flow map $f : C \to \mathbb{R}^n$, jump map $g : D \to \mathbb{R}^n$, flow set $C \subset \mathbb{R}^n$, and jump set $D \subset \mathbb{R}^n$. A solution x to $\mathcal{H}$ is defined on a hybrid time domain $\operatorname{dom} x \subset \mathbb{R}_{\geq 0} \times \mathbb{N}$, which parameterizes the solution by ordinary time $t \in \mathbb{R}_{\geq 0}$ and discrete time $j \in \mathbb{N}$. A hybrid time domain is a subset of $\mathbb{R}_{\geq 0} \times \mathbb{N}$ such that for every $(T, J) \in \operatorname{dom} x$, there exists a sequence $\{t_j\}_{j=0}^{J+1}$ such that $t_0 = 0$, $t_{j+1} \geq t_j$ for each $j \in \{0, 1, \ldots, J\}$, and $\operatorname{dom} x \cap ([0, T] \times \{0, 1, \ldots, J\}) = \cup_{j=0}^{J} ([t_j, t_{j+1}], j)$; see [15]. A solution x is said to be *complete* if $\operatorname{dom} x$ is unbounded, and is said to be *Zeno* if it is complete and the t component of $\operatorname{dom} x$ is bounded (implying $j \to \infty$ in finite ordinary time). A solution x is said to be *maximal* if there does not exist a solution y to $\mathcal{H}$ such that x is a truncation of y to a strict subset of $\operatorname{dom} y$.

B. Stability Properties

Given a differential equation $\dot{z} = f(z)$ with $f : \mathbb{R}^n \to \mathbb{R}^n$ continuous and z evolving in $\mathbb{R}^n$, and a nonempty compact set $\mathcal{A} \subset \mathbb{R}^n$, then a continuously differentiable function $V : \mathbb{R}^n \to \mathbb{R}$ is called a *Lyapunov function* if there exist $\alpha_1, \alpha_2 \in \mathcal{K}_\infty$ and a continuous positive definite function ρ such that

$$\begin{aligned} \alpha_1(|z|_{\mathcal{A}}) \leq V(z) \leq \alpha_2(|z|_{\mathcal{A}}) \qquad & \forall z \in \mathbb{R}^n, \\ \langle \nabla V(z), f(z) \rangle \leq -\rho(|z|_{\mathcal{A}}) \qquad & \forall z \in \mathbb{R}^n. \end{aligned}$$

Definition 1 ([9, Definition 3.7]): For a hybrid system $\mathcal{H}$ as in (1), a nonempty set $\mathcal{A} \subset \mathbb{R}^n$ is said to be *uniformly globally stable* for $\mathcal{H}$ if there exists a class-$\mathcal{K}_\infty$ function α such that every solution x to $\mathcal{H}$ satisfies $|x(t, j)|_{\mathcal{A}} \leq \alpha(|x(0, 0)|_{\mathcal{A}})$ for each $(t, j) \in \operatorname{dom} x$; and *uniformly globally attractive* for $\mathcal{H}$ if every maximal solution is complete and for all $\varepsilon > 0$ and $r > 0$, there exists $T > 0$ such that every solution x to $\mathcal{H}$ with $|x(0, 0)|_{\mathcal{A}} \leq r$ satisfies $|x(t, j)|_{\mathcal{A}} \leq \varepsilon$ for all $(t, j) \in \operatorname{dom} x$ such that $t + j \geq T$. If $\mathcal{A}$ is both uniformly globally stable and uniformly globally attractive for $\mathcal{H}$, then it is said to be *uniformly globally asymptotically stable* (UGAS) for $\mathcal{H}$.

Given a hybrid system $\mathcal{H}$ as in (1), a nonempty set $K \subset \mathbb{R}^n$ is said to be *forward invariant* for $\mathcal{H}$ if each maximal solution x to $\mathcal{H}$ from K is complete and satisfies $x(t, j) \in K$ for all $(t, j) \in \operatorname{dom} x$ [9, Definition 3.13].

III. Hybrid Control Strategy

We consider a nonlinear continuous-time plant

$$\dot{z} = f_P(z, u) \tag{2}$$

with state space $\mathbb{R}^n$. Let $\mathcal{A} \subset \mathbb{R}^n$ be a given nonempty compact set to asymptotically stabilize. Suppose κ_0 is a continuous Lyapunov-certified controller that renders $\mathcal{A}$ to be UGAS for $\dot{z} = f_P(z, \kappa_0(z))$ and has an associated Lyapunov function V, and suppose κ_1 is a continuous uncertified controller for (2). We write the pair of feedback control laws as $u = \kappa_q(z)$ with $q \in Q := \{0, 1\}$. The problem to solve consists of designing a switching logic for q such that $\mathcal{A}$ is UGAS for the resulting closed-loop system, Zeno behavior does not occur, and the controller κ_1 is preferred over κ_0. To solve this problem, we design a hybrid control strategy that determines when to switch between κ_0 and κ_1, as shown in Figure 1. The resulting closed-loop system is hybrid, which we model as in (1).

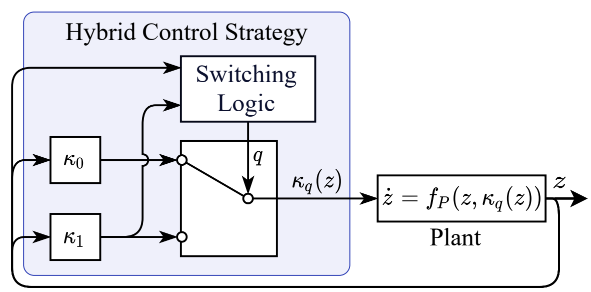

Fig. 1: The switching logic passes q as an output to a switch, which determines whether κ_0 or κ_1 is applied to the plant.

A. Outline of Hybrid Control Strategy

Our hybrid control strategy uses the plant state variable z, the logic variable q described above, and a memory variable $v \in \mathbb{R}_{\geq 0}$. The purpose of each variable is summarized here:

- $z \in \mathbb{R}^n$ is the state of the plant. Our goal is to steer z asymptotically to $\mathcal{A}$.
- $q \in Q$ determines the current feedback controller. When $q = 0$, controller κ_0 is used and when $q = 1$, κ_1 is used.
- $v \in \mathbb{R}_{\geq 0}$ records the value of $V(z)$ at each switch, and then decreases along flows, converging to zero (the dynamics of v are designed in Section III-B). When using the κ_1 controller, $V(z)$ can increase because κ_1 is uncertified, so v is used as an upper bound for $V(z)$, restricting how much $V(z)$ can grow before triggering a switch to $q = 0$. Because v converges to zero, $V(z)$ will be squeezed to zero as well.

Hence, the state of the closed-loop system is

$$x := (z, v, q) \in \mathcal{X} := \mathbb{R}^n \times \mathbb{R}_{\geq 0} \times Q,$$

and we aim to uniformly globally asymptotically stabilize the compact set

$$\mathcal{A}_{\mathcal{X}} := \{x \in \mathcal{X} \mid z \in \mathcal{A}, v = 0\} = \mathcal{A} \times \{0\} \times Q. \quad (3)$$

The rate of change of $V(z)$ is central to our discussion, so for each $z \in \mathbb{R}^n$ and each $q \in Q$, we define

$$\dot{V}_q(z) := \langle \nabla V(z), f_P(z, \kappa_q(z)) \rangle.$$

Because V is a Lyapunov function for $\dot{z} = f_P(z, \kappa_0(z))$, there exists a continuous positive definite function ρ such that $\dot{V}_0(z) \leq -\rho(|z|_A)$ for all $z \in \mathbb{R}^n$.

The basic idea of our hybrid control strategy is as follows. Our strategy implements a switching logic that uses two continuous functions $\sigma_0, \sigma_1 : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ chosen such that σ_1 is positive definite and $\sigma_0(s) > \sigma_1(s)$ for all $s \geq 0$.[1] These functions define thresholds on $\dot{V}_1(z)$ for switching between the feedback controllers κ_0 and κ_1.

(S0) While the feedback controller κ_0 is applied to the plant, due to q being equal to 0, we monitor $\dot{V}_1(z)$. We say that $\dot{V}_1$ is "small enough to switch to $q = 1$" at $z_0 \in \mathbb{R}^n$ if

$$z_0 \in \mathcal{Z}_{0 \mapsto 1} := \{z \in \mathbb{R}^n \mid \dot{V}_1(z) \leq -\sigma_0(|z|_{\mathcal{A}})\}. \quad (4)$$

If $\dot{V}_1(z)$ is small enough to switch to $q = 1$, then κ_1 will produce convergence toward $\mathcal{A}$, so the switching logic updates q from 0 to 1 and records the value of $V(z)$ in v. Conversely, we say that $\dot{V}_1$ is "large enough to hold $q = 0$" at $z_0 \in \mathbb{R}^n$ if

$$z_0 \in \mathcal{Z}_0 := \{z \in \mathbb{R}^n \mid \dot{V}_1(z) \geq -\sigma_0(|z|_{\mathcal{A}})\}. \quad (5)$$

The system is allowed to flow if $q = 0$ and $z \in \mathcal{Z}_0$.

(S1) While the feedback controller κ_1 is applied, due to q being equal to 1, the values of v, $V(z)$, and $\dot{V}_1(z)$ are monitored. We say that $\dot{V}_1$ is "large enough to switch to $q = 0$" at $z_1 \in \mathbb{R}^n$ if

$$z_1 \in \mathcal{Z}_{1 \mapsto 0} := \{z \in \mathbb{R}^n \mid \dot{V}_1(z) \geq -\sigma_1(|z|_{\mathcal{A}})\}, \quad (6)$$

and "small enough to hold $q = 1$" if

$$z_1 \in \mathcal{Z}_1 := \{z \in \mathbb{R}^n \mid \dot{V}_1(z) \leq -\sigma_1(|z|_{\mathcal{A}})\}. \quad (7)$$

If $\dot{V}_1(z)$ is large enough to switch to $q = 0$, then κ_1 is performing poorly. Rather than switching immediately, however, we wait to switch until $V(z) \geq v$. This provides leeway in case κ_1 briefly causes a small increase to $V(z)$. (The dynamics of v are designed, below, such that if z remains in $\mathcal{Z}_{1 \mapsto 0}$ long enough, then $V(z)$ will eventually equal v.) While $q = 1$ and either $z \in \mathcal{Z}_1$ or $V(z) < v$, the system flows and we continue to use κ_1.

Figure 2 shows a representative plot of $-\sigma_0$, $-\sigma_1$, $\mathcal{Z}_{0 \mapsto 1}$, and $\mathcal{Z}_{1 \mapsto 0}$. Note that $\mathcal{Z}_{0 \mapsto 1} \subset \operatorname{int} \mathcal{Z}_1$, which ensures that if $\dot{V}_1(z_0)$ is small enough to switch to $q = 1$, then there is a neighborhood of z_0 where $\dot{V}_1$ is small enough to hold $q = 1$. Similarly, $\mathcal{Z}_{1 \mapsto 0} \subset \operatorname{int} \mathcal{Z}_0$, so if $\dot{V}_1(z_1)$ is large enough to switch to $q = 0$, then there is a neighborhood of z_1 where $\dot{V}_1$ is large enough to hold $q = 0$. Furthermore, $\mathcal{Z}_0 \cup \mathcal{Z}_1 = \mathbb{R}^n$, so either holding 0 or holding 1 is possible everywhere. The sets $\mathcal{Z}_{1 \mapsto 0}$ and $\mathcal{Z}_{0 \mapsto 1}$ are closed and disjoint, precluding Zeno solutions (see Proposition 1).

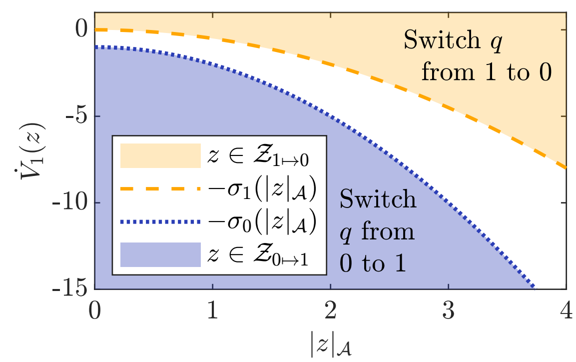

Fig. 2: A plot of $-\sigma_0(|z|_{\mathcal{A}}), -\sigma_1(|z|_{\mathcal{A}}), \mathcal{Z}_{0 \mapsto 1}$, and $\mathcal{Z}_{1 \mapsto 0}$.

Before we formulate the hybrid closed-loop system, we demonstrate the switching logic with an example.

Example 1 (Switching Logic): Consider the plant $\dot{z} = u$ with $z, u \in \mathbb{R}$, controllers $\kappa_0(z) := -z$, $\kappa_1(z) := -z^3$, and pick $\sigma_1(s) := s^2$ and $\sigma_0(s) := 1.5s^2 + 10^{-3}$ for all $s \geq 0$. Figure 3 shows plots of solutions to $\dot{z} = \kappa_0(z)$, $\dot{z} = \kappa_1(z)$, and $\dot{z} = \kappa_q(z)$ with q switching according to our hybrid control strategy.[2] Initially, the solution with the feedback κ_1 converges quickly but slows as z approaches zero. On the other hand, the solution with the feedback κ_0 converges slowly far from the origin, but accelerates relative to the solution with the feedback κ_1, becoming smaller than it at $t = 1.7\,\mathrm{s}$. The switched solution uses κ_1 far from the origin and κ_0 near the origin, producing overall faster convergence.

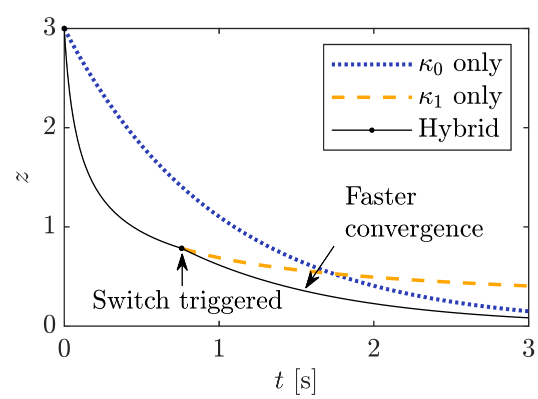

Fig. 3: Solutions from Example 1 using κ_0 only, κ_1 only, and opportunistic switching between κ_0 and κ_1.

Plots of V, v, $-\sigma_0$, $-\sigma_1$, $\dot{V}_1$, and q for a solution from initial values $(z_0, v_0, q_0) = (2, 0, 0)$ are shown in Figure 4. At $t = 0\,\mathrm{s}$, $\dot{V}_1(z_0) < -\sigma_0(|z_0|_{\mathcal{A}})$, so, per (S0), the system

[1]The function σ_0 is strictly positive—not positive definite—because $\sigma_0(0) > \sigma_1(0) = 0$.

[2]Simulations are computed in MATLAB with the *HyEQ Toolbox* [16].

immediately switches to $q = 1$. As time progresses, $\dot{V}_1(z)$ increases until it surpasses $-\sigma_1(|z|_{\mathcal{A}})$ at $t = 0.4\,\mathrm{s}$. This indicates $V(z)$ is not decreasing fast enough to hold $q = 1$, but because $V(z)$ is less than v, the switch to $q = 0$ is delayed until v equals $V(z)$ at $t \approx 0.9\,\mathrm{s}$, as required in (S1).

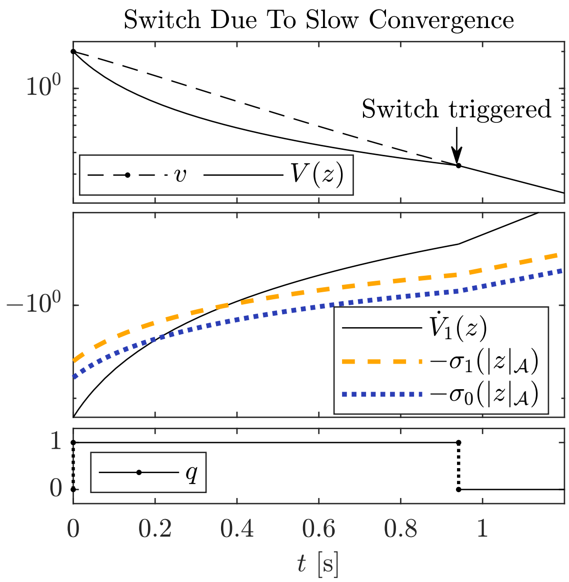

Fig. 4: In Example 1, the controller κ_1 initially has good performance, causing $V(z)$ to decrease quickly, but after $\dot{V}_1(z)$ moves above $-\sigma_1(|z|_{\mathcal{A}})$, v starts to catch up with $V(z)$ and a switch is triggered when $V(z) = v$. While $q = 0$, v does not affect switching, so it is hidden from plots.

B. Construction of the Closed-Loop System

We are now equipped to define the hybrid closed-loop system. Following (S0), if $q = 0$, then the system flows while z is in $\mathcal{Z}_0$ and jumps when z enters $\mathcal{Z}_{0\mapsto 1}$. Thus, if $q = 0$, jumps occur when $x = (z, v, q)$ belongs to

$$D_0 := \mathcal{Z}_{0\mapsto 1} \times \mathbb{R}_{\geq 0} \times \{0\} \tag{8}$$

and flows occur when x belongs to

$$C_0 := \overline{\mathcal{X}_0 \setminus D_0} = \mathcal{Z}_0 \times \mathbb{R}_{\geq 0} \times \{0\} \tag{9}$$

where $\mathcal{X}_0 := \mathbb{R}^n \times \mathbb{R}_{\geq 0} \times \{0\}$. Similarly, following (S1), when $q = 1$, jumps occur only when $z \in \mathcal{Z}_{1\mapsto 0}$, and $V(z) \geq v$, and flows occur if either $z \in \mathcal{Z}_1$ or $V(z) \leq v$. Hence, if $q = 1$, then the system jumps when x is in

$$D_1 := \{x \in \mathcal{X}_1 \mid V(z) \geq v\} \cap (\mathcal{Z}_{1\mapsto 0} \times \mathbb{R}_{\geq 0} \times \{1\}) \tag{10}$$

and flows when x is in

$$C_1 := \overline{\mathcal{X}_1 \setminus D_1} = \{x \in \mathcal{X}_1 \mid V(z) \leq v\} \cup (\mathcal{Z}_1 \times \mathbb{R}_{\geq 0} \times \{1\}) \tag{11}$$

where $\mathcal{X}_1 := \mathbb{R}^n \times \mathbb{R}_{\geq 0} \times \{1\}$. Then, the jump set is $D := D_0 \cup D_1$ and the flow set is $C := C_0 \cup C_1$. Note that the flow set is the closed complement of the jump set: $C = \overline{\mathcal{X} \setminus D}$.

Next, we define the discrete and continuous dynamics of the hybrid closed-loop system. At each jump, z is constant, since the plant state is continuous in time; v is set equal to $V(z)$ to record its value; and q is toggled to the opposite value in $\{0, 1\}$. During flows, z evolves according to $f_P(z, \kappa_q(z))$ and the logic variable q is held constant. The continuous dynamics of v are designed next.

We write the continuous dynamics for v as $\dot{v} = f_v(z, v, q)$. When $q = 0$, the value of v does not affect the switching scheme; we simply pick $f_v(z, v, 0) = -v$ so that v exponentially converges to zero. When $q = 1$, however, the behavior of v is crucial to ensuring that solutions to the closed-loop system converge to $\mathcal{A}_{\mathcal{X}}$. We design $f_v(z, v, 1)$ to satisfy the following rules:

(R1) If $V(z) = v = 0$, then $f_v(z, v, 1) = 0$ because $x \in \mathcal{A}_{\mathcal{X}}$ has already been achieved.

(R2) If $V(z) \leq v \neq 0$, then $f_v(z, v, 1) < 0$ is such that v converges to zero. The motivation for this choice is that $V(z)$ is allowed to increase while $V(z) < v$, so by making v converge to zero, $V(z)$ is squeezed from above, forcing either convergence or a switch to $q = 0$.

If $V(z) > v$, then v is allowed to increase because, eventually, one of the following must occur:

- z enters $\mathcal{Z}_{1\mapsto 0}$ prompting a switch to $q = 0$;
- $V(z) = v$, in which case (R1) or (R2) applies; or
- $z \in \mathcal{Z}_1$ and $V(z) > v$ hold for the rest of time, so $V(z)$ converges to 0 and v is squeezed to 0 as well.

The behavior of v while $V(z) \leq v \neq 0$ is crucial to the performance of the closed-loop system. We prescribe the following cases:

(R3) If $z \in \mathcal{Z}_1$ and $V(z) < v$, then v remains greater than $V(z)$ as long as z remains in $\mathcal{Z}_1$. This guarantees that the leeway above $V(z)$ is maintained while $V(z)$ is decreasing fast enough to hold $q = 1$.

(R4) If $z \in \mathcal{Z}_{1\mapsto 0}$ and $V(z) < v$, then $\dot{V}_1(z) > f_v(z, v, 1)$ holds and, furthermore, if z remains in $\mathcal{Z}_{1\mapsto 0}$, then v decreases until it reaches $V(z)$ in finite time, causing a switch to $q = 0$. This acts as a fail-safe in case $V(z)$ otherwise fails to converge to zero.

(R5) If $z \in \mathcal{Z}_{0\mapsto 1}$ and $V(z) = v$ (as is the case immediately after every switch to $q = 1$), then $\dot{V}_1(z) < f_v(z, v, 1)$ must hold. This condition, in combination with (R3), ensures that $z \in \mathcal{Z}_{0\mapsto 1}$ and $V(z) = v$ only occur simultaneously immediately after a switch to $q = 1$, and the switch is immediately followed by an open interval I of ordinary time such that $V(z) < v$ for all $t \in I$. During I, a switch to $q = 0$ is impossible, due to the design of D_1.

To satisfy (R1)–(R5), we define f_v at each $(z, v, 1)$ as

$$f_v(z, v, 1) = -\sigma_1(|z|_{\mathcal{A}}) + \mu(V(z) - v) \tag{12}$$

with $\mu > 0$. Clearly, (12) satisfies (R1) and (R2). Inspecting $\mathcal{Z}_1$, $\mathcal{Z}_{1\mapsto 0}$, and $\mathcal{Z}_{0\mapsto 1}$, we see that (12) also satisfies (R3)–(R5). The term $\mu(V(z) - v)$ pushes v toward $V(z)$ at a rate proportional to the difference $V(z) - v$, which helps v to "catch up" if $V(z)$ has dropped quickly.

Combining the cases for $q = 0, 1$, the dynamics of v are

$$\dot{v} = f_v(z, v, q) := \begin{cases} -v, & \text{if } q = 0, \\ -\sigma_1(|z|_{\mathcal{A}}) + \mu(V(z) - v), & \text{if } q = 1. \end{cases} \tag{13}$$

The system parameters σ_1 and μ affect the rate at which v converges toward zero while $q = 1$. Larger choices of σ_1 and μ cause v to decay faster, which reduces the amount $V(z)$ can increase before switching back to $q = 0$, whereas smaller choices of σ_1 and μ correspond with a stronger preference for κ_1 (see Example 3).

The construction above leads to the hybrid closed-loop system $\mathcal{H} = (C, f, D, g)$ with state $x = (z, v, q) \in \mathcal{X}$ and data given by

$$\begin{cases} f(x) := (f_P(z, \kappa_q(z)), f_v(x), 0) & \forall x \in C := C_0 \cup C_1 \\ g(x) := (z, V(z), 1 - q) & \forall x \in D := D_0 \cup D_1 \end{cases} \tag{14}$$

with f_v given in (13) and D_0, D_1, C_0, C_1 in (8)–(11). The parameters of our hybrid control strategy are $\mu > 0$, and continuous functions $\sigma_0, \sigma_1 : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ such that σ_1 is positive definite and $\sigma_1(s) < \sigma_0(s)$ for all $s \geq 0$.

Example 2 (LQR): Consider the nonlinear plant

$$\dot{z} = A_1 z + h(\|z\|) A_2 z + u \tag{15}$$

with $A_1 := \left[\begin{smallmatrix} 0 & 2 \\ -2 & 0 \end{smallmatrix}\right]$, $A_2 := 4I$, and $h(s) = \min\{s, 1\}$ for $s \geq 0$. This system behaves like $\dot{z} = A_1 z$ near the origin and like $\dot{z} = (A_1 + A_2)z$ far from it. The origin of (15) is UGAS for $\kappa_0(z) := \left[\begin{smallmatrix} -5 & 0 \\ 0 & -6 \end{smallmatrix}\right] z$. For κ_1, we linearize (15) about the origin and use the linear quadratic regulator (LQR) feedback that solves the following infinite-horizon optimal control problem:

$$\begin{aligned} \underset{u}{\text{minimize}} \quad & \int_0^\infty \|z(t)\|^2 + \|u(t)\|^2 \, dt \\ \text{subject to} \quad & \dot{z} = A_1 z + u. \end{aligned} \tag{16}$$

The LQR feedback is $u = \kappa_1(z) := -z$. Figure 5 shows a solution to the hybrid closed-loop system with $\mu = 1$, $\sigma_0(s) := 0.5s^2$, and $\sigma_1(s) := 0.8s^2 + 10^{-3}$. The switching logic uses κ_1 near the origin, significantly reducing $\|u\|$. The leeway between v and $V(z)$ allows $\dot{V}_1(z)$ to be briefly larger than $-\sigma_1(|z|_{\mathcal{A}})$ without triggering a switch to $q = 0$. In contrast, if $\mu = 4$ (not shown), then v decreases faster, causing v to reach $V(z)$ and triggering a switch to $q = 0$. The switch is followed by a spike in control effort, a period of faster convergence, and a subsequent switch back to $q = 1$.

C. Uniform Global Asymptotic Stability of $\mathcal{A}_{\mathcal{X}}$

Our results require the following assumption.

Assumption 1: The functions f_P, κ_0, κ_1, and V satisfy the following properties.

(B1) f_P, κ_0, and κ_1 are continuous;

(B2) V is continuously differentiable.

For solutions to $\mathcal{H}$, every jump is followed by an interval of flow. The following proposition states that for every solution to $\mathcal{H}$, the lengths of all intervals of flow have a strictly positive lower bound. As a consequence, $\mathcal{H}$ does not have Zeno solutions.

Proposition 1: Suppose Assumption 1 holds. Then, for each solution x to $\mathcal{H}$ in (14), there exists $\gamma > 0$ such that $t_{j+1} - t_j > \gamma$ for all $(t_j, j), (t_{j+1}, j+1) \in \operatorname{dom} x$.

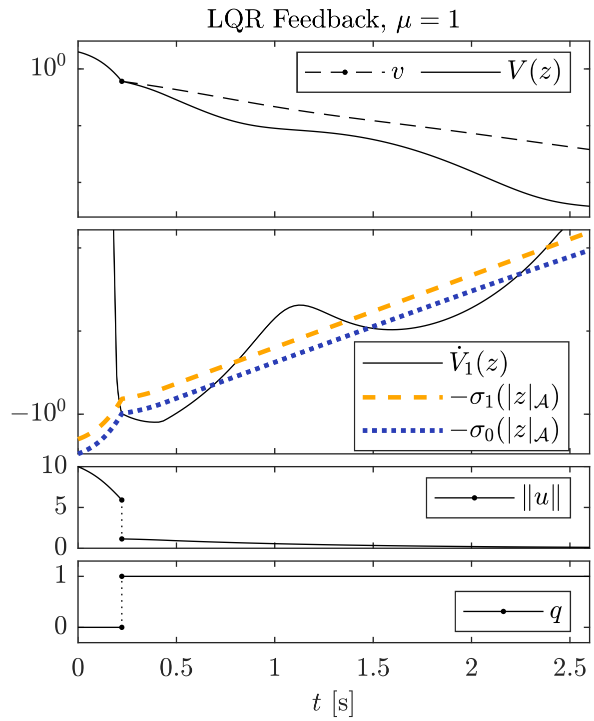

Fig. 5: In Example 2, the switching logic uses the LQR controller κ_1 near the origin, significantly reducing $\|u\|$. The leeway between v and $V(z)$ allows $\dot{V}_1(z)$ to be briefly greater than $-\sigma_1(|z|_{\mathcal{A}})$ without triggering a switch to $q = 0$.

The next result asserts that every maximal solution to $\mathcal{H}$ is complete and converges to $\mathcal{A}_{\mathcal{X}}$; thus, the hybrid controller successfully steers the plant state z to $\mathcal{A}$.

Theorem 1: Suppose Assumption 1 holds. Then, the set $\mathcal{A}_{\mathcal{X}}$ in (3) is UGAS for $\mathcal{H}$ given in (14).

Proof sketch. Consider the function

$$\widetilde{V}(x) := \max\{V(z), v\}. \tag{17}$$

Outside $\mathcal{A}_{\mathcal{X}}$, $\widetilde{V}(x)$ decreases along flows because if $q = 0$, then both $V(z)$ and v are decreasing; if $q = 1$ and $V(z) \geq v$, then $V(z)$ is decreasing; and if $q = 1$ and $V(z) \leq v$, then v is decreasing. To show that $\widetilde{V}(x)$ is decreasing, let $\widetilde{V}^\circ(x, f(x))$ be the Clarke generalized directional derivative [17] of $\widetilde{V}$ at $x \in \mathcal{X}$ in the direction of $f(x)$ and let

$$\mathcal{B} := \{x \in \mathcal{X} \mid q = 1, V(z) = v, \dot{V}_1(z) > -\sigma_1(|z|_{\mathcal{A}})\}.$$

For all $x \in C \setminus (\mathcal{B} \cup \mathcal{A}_{\mathcal{X}})$, $\widetilde{V}^\circ(x, f(x)) < 0$. For each $x \in \mathcal{B}$, the flow map $f(x)$ points out of C, so flows are impossible. Thus, $\widetilde{V}(x)$ decreases along flows in $C \setminus \mathcal{A}_{\mathcal{X}}$. Within $\mathcal{A}_{\mathcal{X}}$, $\widetilde{V}$ is zero. Furthermore, $\widetilde{V}(x)$ does not increase at jumps and Zeno behavior does not occur (see Proposition 1). Therefore, $\widetilde{V}$ is a (nonsmooth) Lyapunov function for $\mathcal{H}$, and $\mathcal{A}_{\mathcal{X}}$ is UGAS [18, Theorem 7.8], [9, Theorem 3.22]. ■

It can also be shown that $\mathcal{H}$ satisfies the hybrid basic conditions [9, Definition 2.18], so the asymptotic stability of $\mathcal{A}_{\mathcal{X}}$ is robust to small disturbances [9, Theorem 3.26].

The next example shows how μ affects performance.

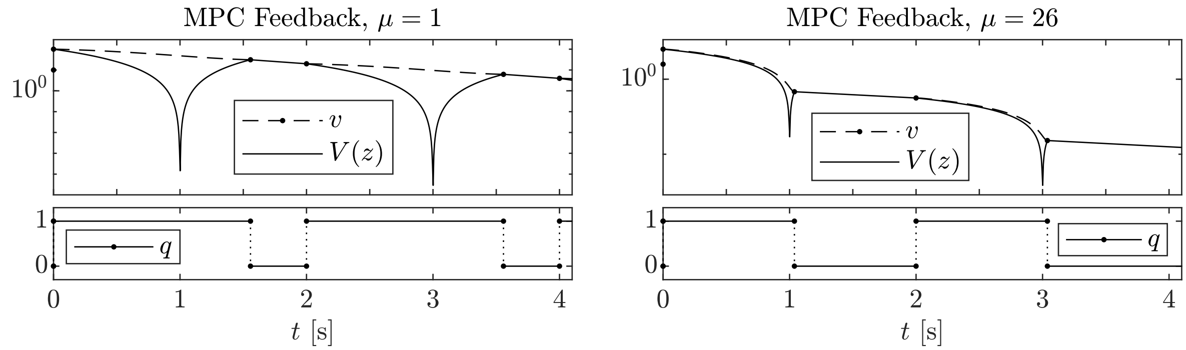

Fig. 6: In Example 3 with $\mu = 1$, a slow decrease in v allows $V(z)$ to return almost to its value at the start of each $q = 1$ interval. When $\mu = 26$, however, v decreases faster, limiting the increase in $V(z)$ after each drop.

Example 3 (MPC): Let $\dot{z} = u$ with $z, u \in \mathbb{R}^2$, let $\kappa_0(z) := \frac{1}{2}\left[\begin{smallmatrix} -1 & 2 \\ -2 & -1 \end{smallmatrix}\right] z$, and let κ_1 be a model predictive controller (MPC) with periodic updates. Between updates, zero-order hold (ZOH) is used to generate the control signal. The switching logic for $\mathcal{H}$ preserves the stability properties of $\mathcal{A}_{\mathcal{X}}$ regardless of the choice of κ_1, so solutions to the hybrid closed-loop system converge even if the MPC algorithm fails to compute updated control values by the next ZOH sample time. Suppose $T = 1\,\mathrm{s}$ is the ZOH sample-time used in the computation of the MPC feedback (that is, the duration that each input value is designed to be applied for) and suppose $T_c = 2\,\mathrm{s}$ is the actual time required to compute the MPC feedback. Because $T_c > T$, a new MPC feedback value is not available at every sample time, in which case the feedback values from the previous interval are reused.

Figure 6 shows solutions to $\mathcal{H}$ with κ_1 computed using the MATLAB MPC Toolbox and with $\sigma_1(s) := 0.3s^2$, and $\sigma_0(s) := 0.36s^2 + 0.5$. For $\mu = 1$, we see that $V(z)$ decreases quickly as t approaches $1\,\mathrm{s}$ (the end of the interval when the MPC feedback value is designed to be applied), but since an updated value is unavailable, κ_1 holds the same value until $t = 2\,\mathrm{s}$. After $t = 1\,\mathrm{s}$, $V(z)$ rises quickly until it hits v, causing a switch to $q = 0$. Consequently, despite κ_1's poor performance, the closed-loop system achieves convergence. Note, however, that $V(z)$ increases significantly after $t = 1\,\mathrm{s}$ and again after $t = 3\,\mathrm{s}$. The increase is limited by v, but a smaller increase may be desirable. To reduce the amount that $V(z)$ can increase, a larger value of μ should be chosen, causing v to follow $V(z)$ more closely as $V(z)$ decreases.

IV. Conclusion

Future work includes analyzing our hybrid control strategy when applied to systems with disturbances and noisy measurements and inputs. We are also interested in integrating our strategy with existing supervisory control strategies that ensure constraint satisfaction by switching between a primary controller that is not provably safe and a backup controller with safety guarantees to create a hybrid closed-loop system that is provably safe and convergent.

References

[1] R. W. Brockett, "Asymptotic stability and feedback stabilization," in *Differential Geometric Control Theory*, pp. 181–191, Birkhäuser, 1983.

[2] C. G. Mayhew, R. G. Sanfelice, and A. R. Teel, "Quaternion-based hybrid control for robust global attitude tracking," *IEEE Transactions on Automatic Control*, vol. 56, pp. 2555–2566, Nov. 2011.

[3] R. G. Sanfelice, M. J. Messina, S. E. Tuna, and A. R. Teel, "Robust hybrid controllers for continuous-time systems with applications to obstacle avoidance and regulation to disconnected set of points," in *2006 American Control Conference*, pp. 3352–3357, IEEE, 2006.

[4] D. Liberzon, *Switching in Systems and Control*. Birkhäuser, 2012.

[5] G. Battistelli, J. P. Hespanha, and P. Tesi, "Supervisory control of switched nonlinear systems," *International Journal of Adaptive Control and Signal Processing*, vol. 26, pp. 723–738, Aug. 2012.

[6] J. P. Hespanha, *Logic-Based Switching Algorithms in Control*. PhD Thesis, Yale University, New Haven, CT, 1998.

[7] J. P. Hespanha and A. S. Morse, "Supervision of families of nonlinear controllers," in *Proceedings of 35th IEEE Conference on Decision and Control*, vol. 4, pp. 3772–3773 vol.4, Dec. 1996. ISSN: 0191-2216.

[8] C. Prieur, "Uniting local and global controllers with robustness to vanishing noise," *Mathematics of Control, Signals and Systems*, vol. 14, pp. 143–172, May 2001.

[9] R. G. Sanfelice, *Hybrid Feedback Control*. Princeton Series in Applied Mathematics, Princeton University Press, 2021.

[10] A. R. Teel and N. Kapoor, "Uniting local and global controllers," in *1997 European Control Conference (ECC)*, pp. 3868–3873, July 1997.

[11] N. H. El-Farra, P. Mhaskar, and P. D. Christofides, "Output feedback control of switched nonlinear systems using multiple Lyapunov functions," *Systems & Control Letters*, vol. 54, pp. 1163–1182, Dec. 2005.

[12] 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)*, pp. 3504–3508 vol.6, IEEE, 1998.

[13] H. K. Khalil, *Nonlinear Systems*. Pearson, third ed., 2014.

[14] H. Dai, B. Landry, L. Yang, M. Pavone, and R. Tedrake, "Lyapunov-stable neural-network control," in *Robotics: Science and Systems XVII*, Robotics: Science and Systems Foundation, July 2021.

[15] R. Goebel, R. G. Sanfelice, and A. R. Teel, *Hybrid Dynamical Systems: Modeling, Stability, and Robustness*. Princeton University Press, 2012.

[16] R. G. Sanfelice, D. A. Copp, and P. Nanez, "A toolbox for simulation of hybrid systems in MATLAB/Simulink: hybrid equations (HyEQ) toolbox," in *Proceedings of the 16th international conference on Hybrid systems: computation and control - HSCC '13*, pp. 101–106, ACM Press, 2013.

[17] F. H. Clarke, *Optimization and Nonsmooth Analysis*. Classics in Applied Mathematics, Society for Industrial and Applied Mathematics, Jan. 1990.

[18] R. G. Sanfelice, R. Goebel, and A. R. Teel, "Invariance principles for hybrid systems with connections to detectability and asymptotic stability," *IEEE Transactions on Automatic Control*, vol. 52, pp. 2282–2297, Dec. 2007.