Safety Verification of Output Feedback Controllers for Nonlinear Systems

Kendra Lesser[1] and Alessandro Abate[1]

***Abstract*— A high-gain observer is used for a class of feedback linearisable nonlinear systems to synthesize safety-preserving controllers over the observer output. A bound on the distance between trajectories under state and output feedback is derived, and shown to converge to zero as a function of the gain parameter of an observer. We can therefore recover safety properties under output feedback and control saturation constraints by synthesizing a controller as if the full state were available. We specifically design feedback linearising controllers that satisfy certain properties, such as stability, and then construct the associated maximal safety-invariant set, namely the largest set of all initial states that are guaranteed to produce safe trajectories over a given (possibly infinite) time horizon.**

I. Introduction

Verification is an increasingly important aspect of control system design. In particular for safety-critical systems, such as aircrafts, satellites, and autonomous vehicles, it is crucial to certify that the controllers designed for these systems will not lead to costly failures, namely that the state of the system will not reach unsafe regions of the state space. Model-based verification is gaining traction as a tool for assessing safety of control systems [1], but is only as accurate as the model being used. An often overlooked model feature is the unavailability of the full state of the system for feedback control. In such cases a controller may be designed as if the full state were available, and then a state estimate may be used in place of the actual state. In doing so, any properties the controller is designed for (linearisation, stability, safety) are no longer guaranteed.

For nonlinear systems, it is difficult to design feedback controllers to achieve desired performance, even when the full state is known. Linear systems, on the other hand, are well-studied with an array of tools available for controller design. If possible, it is therefore desirable to feedback linearise a nonlinear system to then apply techniques from linear control theory [2]. This technique, however, requires full knowledge of the state and no uncertainty in the model.

When the state is unknown, or known only partially, an observer can be used to estimate it. For systems transformable to normal form [2], a nonlinear high-gain observer (HGO) can rapidly reconstruct the actual state because the error dynamics quickly stabilize [3]. If the HGO output is used in place of the actual state in a feedback controller, the output feedback trajectory approximately recovers the trajectory that would arise under state feedback, after a brief transient period in which the estimation errors are large [4], [5]. The implication for feedback linearisable systems is that under certain conditions, linearisation (approximately) still holds under output feedback, and controller synthesis techniques for linear systems still apply [6]. These results, however, typically are used only for stabilization and reference tracking objectives (as in [3]–[6]), and have not been utilized in a verification context.

Alternately, verification and controller synthesis with safety objectives are well-studied for fully observable systems [1], [7]. A common approach for nonlinear systems is to use an optimal control formulation that relies on Hamilton-Jacobi-Bellman equations and discretization-based approximate solutions, to find the maximal set of initial conditions that lead to trajectories that do not violate safety constraints [8], [9], [10]. Feedback linearising controllers are synthesized in [11], using the above formulation to ensure that safety specifications are satisfied and that the controller does not saturate.

Controller synthesis for *partially observable* systems that are guaranteed to obey safety constraints (i.e. correct-by-design control) is much less studied. For stochastic systems, [12], [13] provide an optimal control formulation to synthesize safety maximizing controllers, which use as input a probability distribution that captures any knowledge about the state of the system. An output feedback controller linked to an observer is synthesized and proven to satisfy safety specifications for linear stochastic systems in [14] and for linear deterministic systems in [15]. Output feedback controllers are synthesized for a class of switched linear systems in [16], exploiting robustness of quantitative formulae.

This work, to the best of our knowledge, is the first to synthesize controllers with safety guarantees for a class of *nonlinear* systems. We consider full state feedback linearisable systems in normal form, and construct a HGO in order to recover the desired trajectory using output feedback. *Our contributions are a) the derivation of explicit bounds on the distance between trajectories under state feedback and trajectories under output feedback using an HGO, and b) the synthesis of feedback linearising controllers that satisfy certain specifications (such as stability, trajectory tracking, etc.) that are also provably safe and do not violate any control saturation constraints.* We achieve this by utilizing the bound we derive between trajectories to design controllers as if the state were available for feedback, and by using a more conservative safety constraint that is a function of that bound.

[1]Department of Computer Science, University of Oxford, UK, {kendra.lesser, alessandro.abate}@cs.ox.ac.uk

This work is supported by the European Commission IAPP project AMBI 324432, and by the John Fell Oxford University Press Research Fund.

II. BACKGROUND

A. Feedback Linearisable Systems

We consider a nonlinear single-input, single-output system in normal form [2] that is full state feedback linearizable:

$$\begin{aligned} \dot{x} &= Ax + B[b(x) + a(x)u] \\ y &= Cx, \end{aligned} \tag{1}$$

with

$$A = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ \vdots & \vdots & \cdots & 0 & 1 \\ 0 & 0 & \cdots & \cdots & 0 \end{bmatrix} \in \mathbb{R}^{n\times n}, \quad B = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix} \in \mathbb{R}^n,$$
$$C = \begin{bmatrix} 1 & 0 & \cdots & 0 \end{bmatrix} \in \mathbb{R}^n.$$

The state is $x \in \mathbb{R}^n$, the output is $y \in \mathbb{R}$, and the control input is $u \in \mathbb{R}$. Any system with relative degree n can be put into the form (1) [2]. For simplicity, in this work we assume that the system does not have zero dynamics, although an extension to this case is possible, following [6]. We also raise the following mild assumption on the model dynamics.

Assumption 1: The functions $a(x)$ and $b(x)$ are locally Lipschitz continuous in x for all $x \in \mathcal{A} \subseteq \mathbb{R}^n$, locally bounded over $\mathcal{A}$, and $a(\cdot) \geq a_0 > 0$. □

Given Assumption 1, if the state x were available for feedback and there were no restrictions on the magnitude of the control input, we could design a control law $u(t) = g(x(t))$,

$$g(x) = \frac{-b(x) + Kx + v}{a(x)} \tag{2}$$

to render the system linear, thus obtaining

$$\dot{x} = (A + BK)x + Bv, \tag{3}$$

where v is a reference signal. A control signal $u(t)$, however, is in general subject to a saturation condition such as $|u(t)| \leq u_{\max}$, i.e. $u(t) \in \mathcal{U}_{\text{safe}}$ with $\mathcal{U}_{\text{safe}} = [-u_{\max}, u_{\max}]$. This saturation condition is imposed as follows: whenever $g(x)$ lies inside the set $\mathcal{U}_{\text{safe}}$, it does not saturate, but as soon as $g(x) \notin \mathcal{U}_{\text{safe}}$, it takes the value $\pm u_{\max}$. If we let $\overline{g}(x)$ represent the controller without saturation constraints, then the actual implemented controller is defined as follows:

$$g(x) = \begin{cases} -u_{\max}, & \text{if } \overline{g}(x) \leq -u_{\max} \\ \overline{g}(x), & \text{if } -u_{\max} < \overline{g}(x) < u_{\max} \\ u_{\max}, & \text{if } \overline{g}(x) \geq u_{\max} \end{cases}. \tag{4}$$

Further, in view of (1) the state x is not available to the controller as in (2), and the output $y = Cx$ must be used instead. We would therefore like to design a state estimate $\hat{x} \in \mathbb{R}^n$ that is as close to x as possible, in order to approximately feedback-linearise the system (1).

B. High Gain Observers

We use a high-gain observer (HGO) of the form

$$\dot{\hat{x}} = A\hat{x} + B[b(\hat{x}) + a(\hat{x})u] + H(\epsilon)(y - C\hat{x}) \tag{5}$$

to construct an estimate of the state x of (1). High-gain observers have been used for nonlinear systems because of their robustness to model uncertainty and fast convergence to the true state [3]. The observer gain $H(\epsilon)$ is a function of a positive parameter ϵ, and takes the form

$$H(\epsilon) = \begin{bmatrix} \frac{\alpha_1}{\epsilon} & \frac{\alpha_2}{\epsilon^2} & \cdots & \frac{\alpha_n}{\epsilon^n} \end{bmatrix},$$

with the α_i selected such that the polynomial

$$s^n + \alpha_1 s^{n-1} + \alpha_2 s^{n-2} + \ldots \alpha_{n-1} s + \alpha_n$$

is Hurwitz (roots in the open left half plane). The selection of the α_i guarantees that the error dynamics, defined over the signal $(x - \hat{x})$, are stable. By picking ϵ arbitrarily small, the error dynamics converge to zero arbitrarily quickly. However, HGOs suffer from an initial peaking phenomenon before rapidly converging to the true state (i.e. the state estimate at early stages can be very large compared to the actual state). As such, if the controller $g(x)$ is replaced by $g(\hat{x})$, this leads to erroneous control inputs early on. One partial remedy that we employ is to saturate the controller to avoid large magnitude control inputs [3].

C. Safety Verification Under Output Feedback

We are interested in constructing feedback linearizing controllers for nonlinear systems that simultaneously satisfy both state and control input constraints (respectively for safety requirements over states, and due to input saturation), given that the state is not available for feedback. More precisely, given a set $\mathcal{X}_{\text{safe}} \subset \mathbb{R}^n$ that we would like the system trajectory $x(t)$ to remain within (and that lies within the set $\mathcal{A}$ over which Assumption 1 is valid), we would like to find the set of *all* initial states $x(0)$ (denoted as the safety-invariant set) which, under a class of feedback linearizing controllers, generate trajectories that remain within $\mathcal{X}_{\text{safe}}$ and respect the control constraints $u(t) \in \mathcal{U}_{\text{safe}}$ for all $0 \leq t \leq T$, where $T \leq \infty$.

Definition 1: The safety-invariant set Δ is the set of all initial states $x(0)$ that, under a given feedback controller $g(x)$, produce trajectories $x(t)$ that remain within the safe set $\mathcal{X}_{\text{safe}}$ without violating the control constraint $u(t) \in \mathcal{U}_{\text{safe}}$ for all $0 \leq t \leq T$, with $T \leq \infty$. □

We additionally raise the following assumptions on $\mathcal{X}_{\text{safe}}$ and the initial state $x(0)$ (the latter being non-trivial as it is required for subsequent results). The notation $\partial\mathcal{X}$ is used to denote the boundary of set $\mathcal{X}$.

Assumption 2: The set $\mathcal{X}_{\text{safe}}$ is a compact set contained in $\mathcal{A}$ (from Assumption 1), and $x(0) \in \mathcal{X}_{\text{safe}} \backslash \partial\mathcal{X}_{\text{safe}}$ (so $x(0)$ must lie in the interior of $\mathcal{X}_{\text{safe}}$). □

The linearising controller $g(x)$ may be designed to stabilize the system, track a given reference trajectory, etc., and the safety-invariant set tells us from which initial states

$x(0) \in \mathcal{X}_{\text{safe}}$ the trajectory must initialize to further satisfy the given state and control constraints. We could as well introduce a parametrized controller $g(x, \beta)$, as in [11], to try and maximize the size of the safety-invariant set Δ over β. If we use a parametrized controller, however, Δ will be maximal only with respect to the imposed structure on the controller that targets the other control objectives (e.g., linearisation plus stability). The interest in maximizing the size of Δ is motivated by the general absence of exact knowledge on $x(0)$ beyond the assumption that it lies within $\mathcal{X}_{\text{safe}}$, thus we seek to increase the validity of the controller we have designed over unknown initial conditions $x(0)$. Once a set Δ is computed (regardless of whether it is maximized with respect to a parameter), we have both synthesized a controller that meets performance specifications and obtained a guarantee on its safety given the set of initial states.

There are several ways to compute safety-invariant sets for dynamical systems: here we use an approach based on reachability analysis [11]. The safety-invariant set is the complement of the set of initial states that will eventually exit $\mathcal{X}_{\text{safe}}$, and therefore we can compute the *backwards reachable set* starting from all states outside of $\mathcal{X}_{\text{safe}}$. The set complement, $\mathcal{X}_{\text{safe}}^c$, is propagated backwards in time through the closed loop dynamics (3) to find the set of initial states starting inside of $\mathcal{X}_{\text{safe}}$ that end in $\mathcal{X}_{\text{safe}}^c$. This can be done for a purely deterministic system (when the reference input v in (2) is identically zero) or a non-deterministic system (reference input is allowed to vary). In the non-deterministic case, the safety-invariant set may encompass trajectories that remain within $\mathcal{X}_{\text{safe}}$ for *any possible* choice of v (i.e. even "adversarial" inputs that try to drive $x(t)$ outside of $\mathcal{X}_{\text{safe}}$), or for a *single* input v, that acts to keep $x(t)$ within $\mathcal{X}_{\text{safe}}$. For instance, the reference input v can be treated as the parametrization constant β of the linearising controller $g(x, \beta)$, and the safety-invariant set is then maximized with respect to v (while still satisfying the saturation constraints). Note that in this case we limit our focus to a *constant* parameter β, whereas in the non-deterministic case v is typically allowed to change with time.

The backwards reachable set is computed using a Hamilton-Jacobi formulation (which allows us to minimize the safety-invariant set over $\beta(t)$, if desired) [9], and while difficult to solve exactly in practice, can be approximated using well-established discretization methods, such as those implemented in the Level Set Toolbox [17]. Other computational approaches include viability kernel methods [18], which rely on propagating certain types of shapes through linear dynamics to avoid discretization.

In this work, state x is not available as input to function g, so we must compute Δ under the additional constraint that the control input is provided by $g(\hat{x})$, with $\hat{x}$ the output of an HGO. The safety-invariant set computed under output feedback, rather than state feedback, will be denoted as $\tilde{\Delta}$. While computing reachable (and hence safety-invariant) sets is well established given the feedback controller $g(x)$ and a possible reference signal v, neither linearisation nor satisfaction of the control constraints are guaranteed under the output feedback controller $g(\hat{x})$, and the computation of $\tilde{\Delta}$ becomes more difficult. First, we would need to keep track of both x and $\hat{x}$, doubling the dimensionality of the system, which severely limits the applicability of discretization methods. Second, we would no longer have linear closed-loop dynamics, and could not use the viability kernel methods that are suitable to higher dimensional systems so long as the dynamics are linear.

It is therefore desirable to first design a controller under the assumption that the state is available for feedback, compute Δ, and then *formally relate* Δ to $\tilde{\Delta}$ under output feedback. In summary:

Problem 1: Given a full state feedback linearisable system (1) whose state is not fully available for control,

1) compute a high-gain observer (5) that converges to the original dynamics (1) as quickly as possible;
2) compute a feedback linearising controller $g(x)$ (or $g(x, \beta)$), and the corresponding safety-invariant set Δ (possibly maximizing Δ as a function of β), and quantitatively relate Δ to $\tilde{\Delta}$ under the condition that $g(x)$ is replaced by $g(\hat{x})$. □

III. Main Results

In order to address Problem 1, we first provide a result on the convergence of trajectories under output and state feedback, which generates an upper bound on the distance between closed-loop trajectories. We can therefore first compute Δ, and use the bound between state and input trajectories to derive quantitative claims about $\tilde{\Delta}$.

A. Convergence of State- and Output-Feedback Trajectories

The trajectory $x(t)$ under output feedback recovers the desired trajectory $\overline{x}(t)$ under state feedback, namely the solution of $\dot{\overline{x}} = (A + BK)\overline{x} + Bv$, provided Assumptions 1 and 2 are satisfied, and $\hat{x}(0) \in \mathcal{X}_{\text{safe}}$. This known result [3] is typically presented in the context of feedback stabilization, and asymptotic guarantees on convergence are provided as a function of $\epsilon \to 0$ (the parameter of the HGO). In this work, rather than only providing asymptotic results [4], we discuss precise upper bounds on the distance between $x(t)$ and $\overline{x}(t)$, as a function of ϵ. Further, whereas [4] shows that stability of the origin is preserved under output feedback, and gives an $O(\epsilon)$ bound on the distance between trajectories based on their convergence towards the origin, we show convergence of trajectories to each other as $\epsilon \to 0$ without the assumption that the control law is stabilizing. Note that without stability assumptions, the following result is only valid for finite time horizons $T < \infty$ (we will later address the infinite horizon case).

Theorem 1: Given a nonlinear system (1) satisfying Assumptions 1-2, a time horizon $T < \infty$, and a desired bound $\xi > 0$, there exists a parameter ϵ for a high-gain observer (5), such that $\|x(t) - \overline{x}(t)\|_2 \leq \xi$ for all $0 \leq t \leq T$. □

The proof follows directly from singular perturbation theory, and in particular what is known as Tikhanov's Theorem [19]. We focus on characterising the bound ξ exactly as a

function of ϵ, and show that ξ can be made arbitrarily small as $\epsilon \to 0$. For verification purposes, however, we are more interested in the value of ξ for a given ϵ, rather than the limiting behaviour as $\epsilon \to 0$, since the earlier will establish a connection between Δ and $\tilde{\Delta}$. Note that the restriction of $x(t)$ to $\mathcal{X}_{\text{safe}}$ means that in practice we only require $\mathcal{A} = \mathcal{X}_{\text{safe}}$, i.e. we do not require Lipschitz or boundedness conditions on (1) outside of $\mathcal{X}_{\text{safe}}$.

The relation to singular perturbation theory is seen by replacing the observer dynamics with a scaled version of the estimation error:

$$\eta_i = \frac{x_i - \hat{x}_i}{\epsilon^{n-i}}. \tag{6}$$

We can then write $x - \hat{x} = D(\epsilon)\eta$, with $D(\epsilon)$ a diagonal matrix with entries $[\epsilon^{n-1}, \ldots, \epsilon, 1]$. The combined state and observer dynamics can then be written as

$$\begin{aligned} \dot{x} &= Ax + B[b(x) + a(x)g(x - D(\epsilon)\eta)] \\ \epsilon\dot{\eta} &= \Lambda\eta + \epsilon B[b(x) + a(x)g(x - D(\epsilon)\eta) \\ &\qquad - b(\hat{x}) - a(\hat{x})g(x - D(\epsilon)\eta)], \end{aligned} \tag{7}$$

with

$$\Lambda = \begin{bmatrix} -\alpha_1 & 1 & 0 & \cdots & 0 \\ -\alpha_2 & 0 & 1 & \ldots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ \vdots & \vdots & \cdots & \ddots & 1 \\ -\alpha_n & 0 & \cdots & 0 & 0 \end{bmatrix}.$$

The system (7) is a standard singularly perturbed system [19]. We can think of the x dynamics as the "slow" subsystem, and the η dynamics as the "fast" subsystem. The design of the matrix $H(\epsilon)$ for the high-gain observer in (5) ensures that the matrix Λ is stable. If we scale time through the change of variable $\tau = \frac{t}{\epsilon}$, and then let $\epsilon = 0$, the corresponding boundary-layer system

$$\frac{d\eta}{d\tau} = \Lambda\eta \tag{8}$$

has an isolated, asymptotically stable root at $\eta = 0$ [19]. The reduced system that results from setting $\eta = 0$ in (7) is the system we would *like* to have (system with state feedback designed to satisfy all the properties that we have specified), whose trajectory we denote by $\overline{x}(t)$ (cf. notation at the beginning of this section).

The motivation for studying the singularly perturbed problem is as follows. First, the fast variable η varies quickly, and is approximated by the boundary-layer system in (8). Because of the difference in time scales, x varies little in this period, and remains close to its initial condition $x(0)$. Then, as η has converged to its equilibrium point, the slow variable takes over, and the behaviour of x can be approximated by the reduced system $\overline{x}$ (obtained with $\eta = 0$). We will formalize the amount that each variable can change during these time periods, in order to get bounds on the distance between $x(t)$ and $\overline{x}(t)$.

Here we do not provide the full proof of Theorem 1, which is quite lengthy and similar to that in [4], but outline the necessary steps. We first make claims on the behaviour of η, and then describe the behaviour of x as a function of η.

1) Because the boundary-layer system is stable, there exists a Lyapunov function $W(\eta) = \eta^T P \eta$, such that $\Lambda^T P + P\Lambda = -I$. Further, there exists a constant $\rho > 0$ such that $\Omega_\rho^\eta = \{\eta : W(\eta) \leq \rho\epsilon^2\}$ is a positively invariant set, for any $\epsilon < \bar{\epsilon} = \frac{1}{4M_1\|P\|_2}$ (M_1 is the Lipschitz constant associated with the variable x for $b(x) + a(x)u$, to be further explained shortly).
2) The variable η enters Ω_ρ^η within a finite time $T(\epsilon)$, with $\lim_{\epsilon\to 0} T(\epsilon) = 0$.
3) Before η enters Ω_ρ^η, we can pick $\epsilon = \epsilon_1$, $0 < \epsilon_1 < \bar{\epsilon}$, such that $\|x(t) - \overline{x}(t)\|_2 \leq \xi$, for $0 \leq t \leq T(\epsilon)$.
4) Once η enters Ω_ρ^η, we can choose $\epsilon = \epsilon_2$, $0 < \epsilon_2 < \bar{\epsilon}$, such that $\|x(t) - \overline{x}(t)\|_2 \leq \xi$.

Without going into full details, we can elaborate upon the bounds on $\|x(t) - \overline{x}(t)\|_2$ in steps 3) and 4) above, as follows. For $0 \leq t \leq T(\epsilon)$,

$$\|x(t) - \overline{x}(t)\|_2 \leq C_1 \lambda_{\min}(P) \ln\left(\frac{k^2}{\epsilon^{2n}}\right)\epsilon, \tag{9}$$

and for $T(\epsilon) \leq t \leq T$,

$$\begin{aligned} \|x(t) - \overline{x}(t)\|_2 \leq \Bigg(&4C_1\lambda_{\min}(P)\ln\left(\frac{k^2}{\epsilon^{2n}}\right)C_2 \\ &+ \frac{M_2\gamma\|D(\epsilon)\|_2}{1+M_1}C_2 - \frac{M_2\gamma\|D(\epsilon)\|_2}{1+M_1}\Bigg)\epsilon. \end{aligned} \tag{10}$$

The constants used in (9) and (10) are as follows. The Lipschitz constant M_1 associated with the variable x for $b(x) + a(x)u$ is guaranteed to exist by Assumption 1, since both $a(\cdot)$ and $b(\cdot)$ are Lipschitz continuous in x, and therefore $b(x) + a(x)u$ is also Lipschitz continuous in x. Constant M_2 is defined as $\max_{x\in\mathcal{X}_{\text{safe}}} |a(x)|$, and γ is the Lipschitz constant associated with $g(x)$ defined in (2), which again exists by Assumption 1. The notation $\lambda_{\min}(P)$ refers to the smallest eigenvalue associated with matrix P. We define constant k as a bound on the initial distance between $x(0)$ and $\hat{x}(0)$, which we can guarantee to exist if we set $\hat{x}(0) \in \mathcal{X}_{\text{safe}}$ and impose Assumption 2. Given a neighbourhood $N_r(x(0)) \subset \mathcal{X}_{\text{safe}}$, based on Assumption 1 we claim that

$$\|Ax + B[b(x) + a(x)u]\|_2 \leq C_1,$$

for all $x \in N_r(x(0))$. We require $x(0) \in \mathcal{X}_{\text{safe}} \backslash \partial\mathcal{X}_{\text{safe}}$ so that the neighborhood $N_r(x(0))$ is contained in $\mathcal{X}_{\text{safe}}$, and therefore Assumption 1 applies. Finally, $C_2 = e^{(1+M_1)(T-T(\epsilon))}$.

For a given constant ξ, we can pick $\epsilon_1 \leq \bar{\epsilon}$ small enough such that $\|x(t) - \overline{x}(t)\|_2 \leq \xi$ for $0 \leq t \leq T(\epsilon)$ using (9). We can also pick $\epsilon_2 \leq \bar{\epsilon}$ such that $\|x(t) - \overline{x}(t)\|_2 < \xi$ for $T(\epsilon) \leq t \leq T$ using (10). By taking $\epsilon = \min\{\epsilon_1, \epsilon_2\}$, we guarantee that $\|x(t) - \overline{x}(t)\|_2 \leq \xi$ for all $0 \leq t \leq T$. The upper bound $\bar{\epsilon}$ is required to ensure the positive invariance of Ω_ρ^η. Notice that, as $\epsilon \to 0$, $\|x(t) - \overline{x}(t)\|_2 \to 0$ uniformly over $0 \leq t \leq T$. Equations (9) and (10) therefore provide both a direct way of computing ξ for a given ϵ (by taking the maximum of the right hand side of both equations) *or*

for verifying the maximum value that ϵ can take to achieve a desired bound ξ.

Theorem 1 is not specific to feedback linearisable systems, i.e. $a(x)+b(x)u$ can be replaced by a more general Lispchitz continuous function $\phi(x,u)$. A possibly tighter bound is available for linearizable systems with control input (2), because an explicit solution to (3) is available of the form $x(t) = x(0)\exp\{(A+BK)t\}$, and therefore (10) can be refined.

The next Lemma considers the scenario when the linearising controller asymptotically stabilizes the system, and $g(x)$ is given by (2) with $v=0$ and K chosen to stabilize $A+BK$. In this instance we can get a tighter bound on $\|x(t)-\overline{x}(t)\|_2$, because $A+BK$ has negative eigenvalues, thus the quantity $\|\exp\{(A+BK)t\}\|$ does not increase with t.

Lemma 1: For the feedback linearizable system in (1) satisfying Assumption 1, with a high-gain observer (5) with parameter ϵ, controller $g(x) = \frac{-b(x)+Kx}{a(x)}$ with K chosen such that $A+BK$ is stable, and time horizon T, the upper bound on $\|x(t)-\bar{x}(t)\|_2$ is given by (9) for $0 \le t \le T(\epsilon)$, and by

$$\|x(t)-\overline{x}(t)\|_2 \le \left[\left(4C_1\lambda_{\min}(P)\ln\left(\frac{k^2}{\epsilon^{2n}}\right)\right)\hat{C}_2 + (T-T(\epsilon))M_2\gamma\|D(\epsilon)\|_2\right]\epsilon$$

for $T(\epsilon) \le t \le T$. □

Notice that the matrix $A+BK$ is assumed to be diagonalisable, and if not it can be put into Jordan form. Then $A+BK = G\Theta G^{-1}$, with G the matrix of (generalised) eigenvectors, and Θ the diagonal or Jordan matrix containing the eigenvalues (with negative real part). The constant $\hat{C}_2 = \|G\|\|G^{-1}\|\left\|e^{\Theta(T-T(\epsilon))}\right\|$, and all other constants coincide with those in Theorem 1.

B. Invariance Properties of Feedback Linearisable Systems

We can now use Theorem 1 (and, by extension, Lemma 1) to show how to compute set $\tilde{\Delta}$ by calculating the corresponding set Δ under state feedback. We first provide a finite time horizon result, which does not require any stability assumptions on the controller $g(x)$. We will denote, for a given parameter $\delta > 0$, the restricted set $(\mathcal{X}_{\text{safe}} - \delta) = \{x : x \in \mathcal{X}_{\text{safe}} \cap \|x-\tilde{x}\|_2 \ge \delta, \forall\, \tilde{x} \in \partial\mathcal{X}_{\text{safe}}\}$, where $\partial\mathcal{X}_{\text{safe}}$ is the boundary of $\mathcal{X}_{\text{safe}}$.

Theorem 2: For the feedback linearisable system (1) satisfying Assumption 1, the high-gain observer (5) with parameter ϵ, safe set $\mathcal{X}_{\text{safe}}$ and control constraint $\mathcal{U}_{\text{safe}}$, and given the finite time horizon T, the safety invariant set Δ, computed using state feedback, $\mathcal{U}_{\text{safe}}$ and $(\mathcal{X}_{\text{safe}} - \xi)$ (with ξ the bound between trajectories $x(t)$ and $\bar{x}(t)$ provided in Theorem 1), is an underestimate of the safety invariant set $\tilde{\Delta}$ under output feedback for the original safe set $\mathcal{X}_{\text{safe}}$. □

Note that for some choices of ϵ, the associated ξ may exceed the diameter of $\mathcal{X}_{\text{safe}}$, in which case $\tilde{\Delta}$ will be the empty set. We assume that ϵ can be chosen small enough, and hence ξ is small enough, to render $\tilde{\Delta}$ nonempty (although an empty $\tilde{\Delta}$ is also informative). From Theorem 1, we know that for any $0 < \epsilon < \bar{\epsilon}$, there exists a corresponding ξ such that $\|x(t)-\overline{x}(t)\|_2 \le \xi$ for all $t \in [0,T]$. Therefore, for a given $\epsilon > 0$, if we can control $\overline{x}(t)$ to remain inside $(\mathcal{X}_{\text{safe}} - \xi)$ for $t \in [0,T]$, we guarantee that $x(t) \in \mathcal{X}_{\text{safe}}$.

We can extend Theorem 2 to the infinite horizon case, provided that the state feedback control law $g(x)$ asymptotically stabilizes the origin, as in Lemma 1. We again employ the control structure in (2) with $v=0$, and K chosen such that $A+BK$ has eigenvalues with negative real parts.

Theorem 3: For the feedback linearisable system (1) satisfying Assumptions 1 and 2, the high-gain observer (5) with parameter ϵ, safe set $\mathcal{X}_{\text{safe}}$, and control constraint $\mathcal{U}_{\text{safe}}$, the infinite horizon safety-invariant set Δ, computed using state feedback and constraint sets $(\mathcal{X}_{\text{safe}} - \xi)$ and $\mathcal{U}_{\text{safe}}$ (ξ is again the bound from Theorem 1), is an underestimate of the safety invariant set $\tilde{\Delta}$ under output feedback for the original constraint sets $\mathcal{X}_{\text{safe}}$ and $\mathcal{U}_{\text{safe}}$. Further, the origin remains asymptotically stable under output feedback. □

We know from Theorem 2 that, over a finite time horizon, we can underestimate the safety-invariant set $\tilde{\Delta}$ by assuming state feedback and using modified constraint sets $(\mathcal{X}_{\text{safe}} - \xi)$ and $\mathcal{U}_{\text{safe}}$. We can extend Theorem 2 to the infinite horizon by constructing a positively invariant set $\Omega_c^x \times \Omega_\rho^\eta$, with Ω_ρ^η the same as in Theorem 1, and $\Omega_c^x \subset \Delta$, the safety-invariant set using an asymptotically stabilizing state feedback controller and constraint sets $\mathcal{X}_{\text{safe}}$ and $\mathcal{U}_{\text{safe}}$. We only need Ω_c^x to be invariant while respecting the original constraint set $\mathcal{X}_{\text{safe}}$ (if $\Omega_c^x \subset \Delta$, then $x(t)$ is guaranteed to not leave $\mathcal{X}_{\text{safe}}$).

Under state feedback (where the control law renders the origin asymptotically stable), there exists a Lyapunov function $V(\bar{x}) = \bar{x}^T Q \bar{x} \le -\|\bar{x}\|_2^2$ for $x \in \Delta$. By definition, the set $\Omega_c^x = \{\bar{x} : V(\bar{x}) \le c\} \subset \Delta$ is positively invariant. For $\eta \in \Omega_\rho^\eta$, the set Ω_c^x remains positively invariant under output feedback.

We can show that for $t \ge T_1$, where

$$T_1 = \frac{\lambda_{\max}(Q)}{2}\ln\left(\frac{\lambda_{\max}(Q)\|x(0)\|_2^2}{c}\right), \qquad (11)$$

$x(t)$ (using output feedback) will have reached Ω_c^x. Note that we would like to make c as large as possible, to reduce the time it takes for $x(t)$ to reach Ω_c^x.

Therefore, after time T_1, we can guarantee that $(x,\eta) \in \Omega_c^x \times \Omega_\rho^\eta$, with $\Omega_c^x \subset \Delta$, and that (x,η) will remain inside $\Omega_c^x \times \Omega_\rho^\eta$ for all $t \ge T_1$. Hence for all $t \ge T_1$, we are guaranteed that $x(t)$ does not violate the safety constraints, and $g(\hat{x})$ does not violate the control constraints. Before time T_1, we do not have the guarantee on the state constraints, and apply Theorem 2 over the finite time horizon $[0,T_1]$. We compute ξ as in Theorem 2, using (9) and (10) from Theorem 1, and shrink $\mathcal{X}_{\text{safe}}$ appropriately, then compute $\tilde{\Delta} \subset \Delta$, which is the infinite horizon safety-invariant set.

In order to ensure the invariance of Ω_c^x and Ω_ρ^η, we require an upper bound on the value that ϵ can take. From Theorem 1, we know that $\epsilon < \bar{\epsilon} = \frac{1}{4M_1\|P\|}$ to ensure that Ω_ρ^η is invariant.

To ensure that Ω_c^x is invariant and that all trajectories outside of Ω_c^x enter Ω_c^x in finite time requires $\epsilon < \min\{\epsilon_3, \epsilon_4\}$, with

$$\epsilon_3 = \frac{\beta}{2\lambda_{\max}(Q)x_{\max}L\gamma},$$

and ϵ_4 chosen such that

$$c \geq 16\lambda_{\max}(Q)^3 L^2\gamma^2\|D(\epsilon)\|_2^2\epsilon_2^2$$

is satisfied, while at the same time $\Omega_c^x \subset \Delta$. The constant L is equal to $\max_{x\in\Omega_c^x}|a(x)|$, and $\beta = \min_{x\in\partial\Omega_c^x}\|x\|_2^2$. Therefore we require $\epsilon < \hat{\epsilon} = \min\{\bar{\epsilon}, \epsilon_1, \epsilon_2\}$.

The recovery of asymptotic stability of the origin can be shown in the same manner as in [6].

To summarize, we can compute a finite- or infinite-horizon safety-invariant set under output feedback with the following procedure:

1) Finite horizon case:

- Compute ξ as a function of ϵ and time horizon T according to (9) and (10) from Theorem 1.
- Compute $\tilde{\Delta}$ using standard reachability techniques for fully observable systems, with state constraints $(\mathcal{X}_{\text{safe}} - \xi)$, control constraints $\mathcal{U}_{\text{safe}}$, and time horizon T.

2) Infinite horizon case:

- Compute Δ using $\mathcal{X}_{\text{safe}}$ and $\mathcal{U}_{\text{safe}}$.
- Find the largest c such that $\Omega_c^x \subset \Delta$.
- Compute T_1 according to (11), and ξ as a function of ϵ and T_1 according to Theorem 1.
- Compute $\tilde{\Delta}$ using standard reachability techniques for fully observable systems, with state constraints $(\mathcal{X}_{\text{safe}} - \xi)$, control constraints $\mathcal{U}_{\text{safe}}$, and time horizon T_1.

The obtained set $\tilde{\Delta}$ is the safety-invariant set for either the finite or infinite horizon using a controller designed for state feedback, but whose input is an estimate of the state produced by a high-gain observer. A simple case study that demonstrates the outlined techniques for computing $\tilde{\Delta}$ is provided in an extended version of this paper [20].

IV. Conclusions and Extensions

The use of a high-gain observer design allows us, for a specific class of feedback linearisable models, a) to design controllers to satisfy certain specifications (in this work, safety) as if the full state were available for feedback; and b) to derive a bound on the distance between the trajectories under state and output feedback, as a function of the observer parameter ϵ. Given this quantitative bound, if we can select ϵ as small as needed, we can completely recover whatever properties are satisfied under full state feedback.

Possible extensions include feedback linearisable systems with zero dynamics and model uncertainty, which are discussed in [6] and are aligned with the approach discussed in this work. On the other hand, convergence of $\hat{x}$ to x using a high-gain observer in the presence of output disturbances unfortunately does not hold. However, we may still be able to obtain a bound on the distance between trajectories, even if we cannot claim that the bound converges to zero with time. Another straightforward extension is to consider more complex temporal properties, such as reach-avoid or other LTL specifications: the difficulty lies in the construction of the state feedback controller. This is typically done by resorting to formal (finite-state) abstractions [1], [7], [21], possibly coupled downstream with the design of an incrementally stabilizing controller, which may violate the Lipschitz continuity assumptions required for the feedback controller in this work.

References

[1] P. Tabuada, *Verification and Control of Hybrid Systems: A Symbolic Approach*. Springer, 2009.

[2] S. Sastry, *Nonlinear Systems: Analysis, Stability, and Control*. Springer, 1999.

[3] H. K. Khalil and L. Praly, "High-gain observers in nonlinear feedback control," *International Journal of Robust and Nonlinear Control, Special Issue on High-Gain observers in nonlinear feedback control*, vol. 24, pp. 993 – 1015, 2014.

[4] A. N. Atassi and H. K. Khalil, "A separation principle for the stabilization of a class of nonlinear systems," *IEEE Transactions on Automatic Control*, vol. 44, no. 9, pp. 1672 – 1687, 1999.

[5] A. Teel and L. Praly, "Tools for semiglobal stabilization by partial state and output feedback," *SIAM Journal on Control and Optimization*, vol. 33, no. 5, p. 14431488, 1995.

[6] L. B. Freidovich and H. K. Khalil, "Performance recovery of feedback-linearization-based designs," *IEEE Transactions on Automatic Control*, vol. 53, no. 10, pp. 2324 – 2334, 2008.

[7] A. Abate, M. Prandini, J. Lygeros, and S. Sastry, "Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems," *Automatica*, vol. 44, no. 11, pp. 2724–2734, 2008.

[8] C. Tomlin, G. J. Pappas, and S. Sastry, "Conflict resolution for air traffic management: A study in multiagent hybrid systems," *IEEE Transaction on Automatic Control*, vol. 43, no. 4, pp. 509 – 521, 1998.

[9] I. Mitchell, A. Bayen, and C. 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.

[10] K. Margellos and J. Lygeros, "Hamilton-Jacobi formulation for reach-avoid problems with an application to air traffic management," in *American Control Conference*, 2010, pp. 3045–3050.

[11] M. Oishi, I. Mitchell, C. Tomlin, and P. Saint-Pierre, "Computing viable sets and reachable sets to design feedback linearizing control laws under saturation," in *IEEE Conference on Decision and Control*, 2006, pp. 3801 – 3807.

[12] K. Lesser and M. Oishi, "Reachability for partially observable discrete time stochastic hybrid systems," *Automatica*, vol. 50, no. 8, pp. 1989–1998, 2014.

[13] J. Ding, A. Abate, and C. Tomlin, "Optimal control of partially observable discrete time stochastic hybrid systems for safety specifications," in *American Control Conference*, 2013, pp. 6231–6236.

[14] K. Lesser and A. Abate, "Controller synthesis for probabilistic safety specifications using observers," in *IFAC Conference on Analysis and Design of Hybrid Systems*, 2015, pp. 329–334.

[15] S. Haesaert, A. Abate, and P. Van den Hof, "Correct-by-design output feedback of LTI systems," in *54th IEEE Conference on Decision and Control*, 2015, pp. 6159–6164.

[16] O. Mickelin, N. Ozay, and R. M. Murray, "Synthesis of correct-by-construction control protocols for hybrid systems using partial state information," in *American Control Conference*, 2014, pp. 2305–2311.

[17] I. Mitchell, "The flexible, extensible and efficient toolbox of level set methods," *Journal of Scientific Computing*, vol. 35, pp. 300–329, 2008.

[18] J. Maidens, S. Kaynama, I. Mitchell, M. Oishi, and G. Dumont, "Lagrangian methods for approximating the viability kernel in high-dimensinoal systems," *Automatica*, vol. 49, no. 7, pp. 2017–29, 2013.

[19] W. R. Wasow, *Asymptotic Expansions for ordinary differential equations*. Interscience Publishers, 1965.

[20] K. Lesser and A. Abate, "Safety verification of output feedback controllers for nonlinear systems," 2016, preprint, http://arxiv.org/abs/1603.06627.

[21] M. Zamani, M. Esfahani, R. Majumdar, A. Abate, and J. Lygeros, "Symbolic control of stochastic systems via approximately bisimilar finite abstractions," *IEEE Transactions on Automatic Control*, vol. 59, no. 12, pp. 3135–3150, 2014.