

Adaptive Robust Control Contraction Metrics: Transient Bounds in Adaptive Control with Unmatched Uncertainties

Samuel G. Gessow and Brett T. Lopez

Abstract—This work presents a new sufficient condition for synthesizing nonlinear controllers that yield bounded closed-loop tracking error transients despite the presence of unmatched uncertainties that are concurrently being learned online. The approach utilizes contraction theory and addresses fundamental limitations of existing approaches by allowing the contraction metric to depend on the unknown model parameters. This allows the controller to incorporate new model estimates generated online without sacrificing its strong convergence and bounded transients guarantees. The approach is specifically designed for trajectory tracking so the approach is more broadly applicable to adaptive model predictive control as well. Simulation results on a nonlinear system with unmatched uncertainties demonstrates the approach.

I. INTRODUCTION

There is a growing need for nonlinear controllers that guarantee the closed-loop tracking error has bounded transients even when parametric model uncertainty is present. For example, trajectory planning or high-level decision-making can be improved by leveraging transient bounds to predict how the closed-loop tracking error will evolve in time, leading to less myopic behaviors. Furthermore, in order to limit conservatism, these controllers must be able to incorporate online model estimates without losing their strong performance guarantees. However, concurrent control and model estimation, i.e., indirect adaptive control, that preserves transient bounds is nontrivial as the controller must simultaneously account for model estimation error and the transients associated with model estimation. Traditionally, trajectory tracking with model error is addressed by establishing the input-to-state convergent (ISC) [1] or (stronger) incremental input-to-state stable (δ ISS) [2] property. If the model is also being estimated online then the controller must be robust to not only the uncertainty in the model but also the transients of the parameter estimates [3], [4]. These conditions are pragmatically difficult to establish though for general uncertain nonlinear systems. We will show that being ISC with respect to estimation transients can be bypassed and instead replaced with a simpler condition that eases the design of robust high-performance controllers without sacrificing tracking error transient bounds.

A sufficient condition for bounded tracking error transients is the existence of an input-to-state convergent control Lyapunov function (ISC-clf) where the closed-loop system is ISC with respect to the unknown model parameters. For the case where the model uncertainties are unmatched (outside the span of the control input matrix), the ISC-clf must be

model-dependent, i.e., a function of the unknown parameters. In essence one must find a *family* of ISC-clf's. An immediate consequence of this dependency is that the ISC-clf itself must change as the model is updated. This introduces significant complexity in the controller design since the transient bounds will only hold if the controller either anticipates how the model is changing or is ISC with respect to the estimation transients. Since predicting the model estimation transients is unrealistic, imposing ISC with respect to the estimation transients is the primary approach used in the literature. Aside from systems in strict-feedback form, systematically computing an ISC-clf for general uncertain nonlinear systems is difficult, and typically requires assumptions about the form of the ISC-clf, i.e., restricting it to take a quadratic form. Conversely, contraction and contraction metrics [5]–[8] is a promising new framework to systematically compute feedback policies for stabilizable nonlinear systems. A key advantage of contraction is that the metric search can be formulated as a convex optimization and is therefore applicable to a much broader class of systems. While robust [7]–[9] and direct adaptive [10], [11] control techniques based on contraction metrics have been proposed, an adaptive robust control strategy which guarantees 1) bounded tracking error transients and 2) convergence to the desired trajectory as the model estimate is improved has yet to be developed.

The main contribution of this work is a new contraction-based adaptive robust control framework that possesses desirable convergence and bounded transients guarantees for systems with unmatched uncertainties being estimated online. Specifically, we introduce the *adaptive robust control contraction metric* which, when used in feedback, can incorporate new estimates of the uncertain model parameters to improve closed-loop performance without sacrificing its strong guarantees. Central to the approach is allowing the metric to depend on the estimated model parameters. While this generally presents unique challenges in terms of ensuring the tracking error transients remain bounded as the model is changing online, we present a new sufficient condition for the metric that preserves transient bounds and can be immediately added to existing algorithms that numerically construct contraction metrics. Trajectory tracking results for a system not feedback linearizable nor in strict-feedback form with unmatched uncertainties confirm the theoretical properties and showcases the advantage of the approach when compared to existing robust techniques.

Notation: Symmetric positive definite $n \times n$ matrices are denoted S_+^n . Positive scalars are denoted \mathbb{R}_+ and strictly positive scalars are $\mathbb{R}_{>0}$. The infinity norm of a signal

Both authors are with the VECTR Laboratory, University of California, Los Angeles, Los Angeles CA, {sgessow,btlopez}@ucla.edu

$x(t)$ is $\|x(t)\|_\infty$ and the L_1 norm at time t is $|x(t)|$. The shorthand notation of a function T parameterized by a vector a with vector argument z is $T_a(z) \triangleq T(z; a)$. The directional derivative of a smooth matrix $M : \mathbb{R}^n \rightarrow \mathcal{S}_+^n$ along a vector field $v : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is $\partial_v^\circ M(x) = \sum_i^n \nabla_{x_i} M(x) \dot{x}_i$ or $\partial_v M(x) = \sum_i^n \nabla_{v_i} M(x) \dot{v}_i$.

II. PROBLEM FORMULATION AND PRELIMINARIES

Consider the uncertain control affine nonlinear system

$$\dot{x} = f(x) + \Delta(x)^\top \theta + B(x)u, \quad (1)$$

with state $x \in \mathbb{R}^n$, control input $u \in \mathbb{R}^m$, nominal dynamics $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$, and control input matrix $B : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ with columns $b_i(x)$ for $i = 1, \dots, m$. The model parameters $\theta \in \mathbb{R}^p$ are unknown but with known dynamics $\Delta : \mathbb{R}^n \rightarrow \mathbb{R}^{p \times n}$ with rows $\Delta_i(x)$ for $i = 1, \dots, p$. No restrictions are placed on the model uncertainties for generality, hence they are allowed to be unmatched, i.e., $\Delta(x)^\top \theta \notin \text{span}\{B(x)\}$. Shorthand notation for (1) will be $\dot{x} = F_\theta(x, u)$. We make the following assumptions about the uncertain parameters θ and their estimates $\hat{\theta}$.

Assumption 1. The unknown parameters θ belong to a known compact set Θ . Furthermore, the parameter estimation error $\tilde{\theta} \triangleq \hat{\theta} - \theta$ also belongs to a known compact set $\tilde{\Theta}$.

Assumption 2. The model parameter estimate $\hat{\theta}$ is being generated online via a suitable model estimator where the estimation error $\tilde{\theta}$ is monotonically non-increasing.

Assumptions 1 and 2 are standard in the indirect adaptive control and adaptive MPC literature [12], [13]. Note that Assumption 2 is more restrictive but can be generally achieved using least squares over a sliding window [14] or set membership identification [15].

The goal of this work is to develop constructive conditions for a model-parameterized feedback policy $\kappa_\theta : \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^p \rightarrow \mathbb{R}^m$ such that the state $x(t)$ converges exponentially to a region near a desired trajectory $x_d(t)$ when the system dynamics are uncertain and being estimated online. Note that κ_θ must be model-parameterized so that the tracking error converges to zero as the parameter estimate converges to its true value. Since this work is concerned with converging to a time-varying trajectory despite the presence of model error, it is beneficial to introduce the notion of an input-to-state convergent system.

Definition 1 ([1]). A closed-loop system is *input-to-state convergent* (ISC) with respect to the input $w(t)$ if, for any feasible trajectory $x_d(t)$, the state $x(t)$ satisfies

$$|x(t) - x_d(t)| \leq \beta(|x(t_0) - x_d(t_0)|, t) + \sigma \left(\sup_{\tau \in [t_0, t]} |w(\tau)| \right),$$

where $\beta : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ is a class \mathcal{KL} function and $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is a class \mathcal{K} function.

Definition 1 ensures that the tracking error transients is bounded, and that $x(t) \rightarrow x_d(t)$ if $w(t) \rightarrow 0$. More broadly, ISC [1] is a natural extension to ISS [16] and is deeply

related to δ ISS [2]. In particular, ISC and δ ISS are equivalent if the system evolves on a compact set [17]. We will also make use of a specialized version of Definition 1, namely the notion of an exponentially convergent system, where the \mathcal{KL} function β takes an exponential form.

Definition 2. A closed-loop system is *exponentially input-to-state convergent* (eISC) with respect to the input $w(t)$ if, for any feasible trajectory $x_d(t)$, state $x(t)$ satisfies

$$|x(t) - x_d(t)| \leq \sigma_1(|x(t_0) - x_d(t_0)|)e^{-\lambda(t-t_0)} + \sigma_2 \left(\sup_{\tau \in [t_0, t]} |w(\tau)| \right),$$

where $\lambda \in \mathbb{R}_{>0}$ is the convergence rate and $\sigma_{1,2} : \mathbb{R} \rightarrow \mathbb{R}$ are class \mathcal{K} functions.

It is important to note that eISC and exponential δ ISS are not necessarily equivalent as exponential convergence to a single trajectory is a weaker condition than exponential convergence of two arbitrary trajectories. With that said, it is easy to show that exponential δ ISS implies eISC – a fact we will use later in this work.

Remark 1. Any system that satisfies Definition 1 or Definition 2 also satisfies the Universal L_∞ gain definition as presented in [9]. The key difference between the definition presented here and that in [9] is that Definitions 1 and 2 use a generic class \mathcal{K} function σ and ensure that β is a \mathcal{KL} function by definition or by making it a decaying exponential.

Definitions 1 and 2 are important characterizations of closed-loop systems but do not contain constructive conditions for a control law or control Lyapunov function (clf) that satisfy their respective mathematical definitions. While $w(t)$ is a general input, within the context of indirect adaptive control one aims to achieve ISC with respect to model estimation error while the model estimate is being updated online. A necessary and sufficient condition for a system to be eISC with respect to $\tilde{\theta}$ and $\frac{d}{dt}\tilde{\theta}$ is the existence of an exponentially ISC-control Lyapunov function (eISC-clf)¹.

Definition 3. A continuously differentiable function $V_\theta : \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^p \rightarrow \mathbb{R}_+$ where $V_\theta(x, x_d) = 0 \iff x(t) = x_d(t)$ is an *exponentially input-to-state convergent control Lyapunov function* (eISC-clf) if the following hold

$$\begin{aligned} k_1|x - x_d|^a &\leq V_\theta(x, x_d) \leq k_2|x - x_d|^a & (2) \\ |x(t) - x_d(t)| &\geq \sigma \left(\sup_{\tau \in [t_0, t]} \left| \begin{bmatrix} \tilde{\theta}(\tau) \\ \dot{\tilde{\theta}}(\tau) \end{bmatrix} \right| \right) \implies \\ \inf_{u \in \mathbb{R}^m} \left\{ \nabla_x V_\theta^\top F_{\hat{\theta}}(x, u) + \nabla_{x_d} V_\theta^\top F_{\hat{\theta}}(x_d, u_d) \right. \\ &\quad \left. - \nabla_x V_\theta^\top \Delta^\top \tilde{\theta} + \nabla_{\tilde{\theta}} V_\theta^\top \dot{\tilde{\theta}} \right\} \leq -k_3|x - x_d|^a & (3) \end{aligned}$$

where $k_1, k_2, k_3, a \in \mathbb{R}_{>0}$, $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is a class \mathcal{K} function, and $\tilde{\theta} \in \tilde{\Theta}$ and $\frac{d}{dt}\tilde{\theta} \in \tilde{\Omega}$ belong to compact sets.

¹An eISC-clf is similar to an ISS-clf discussed in [3] but instead ensures exponential convergence to a time-varying trajectory $x_d(t)$. Proving the equivalence between eISC and the existence of an eISC-clf follows similarly to that in [3] and is omitted for brevity.

Synthesizing an eISC-clf that is robust to model estimation error and estimation transients is not trivial due to the presence of $\nabla_{\hat{\theta}} V(x)^\top \frac{d}{dt} \hat{\theta}$ in (3). Techniques have been developed for systems in strict-feedback form [3], [18] but a systematic constructive procedure for general nonlinear systems has yet to be proposed. A promising approach, and the one taken here, is the use of contraction analysis [5] and the control contraction metric (ccm) [6]. In [8], the robust ccm was introduced to ensure eISC despite the presence of model uncertainties or external disturbances; the robust ccm was later used for robust motion planning in [9]. The main limitation of using the robust ccm framework is that the tracking error is not guaranteed to converge to zero even if the model is perfectly known after online estimation. This is a direct result of the metric (and hence controller) being designed for a *single* model realization. In essence, picking a single model for metric / controller design eliminates the ability to incorporate new model information gathered online without violating tracking error or transient bounds. The method proposed in this work addresses this limitation while guaranteeing that the tracking error transients is bounded *and* that the tracking error will converge to zero if the model estimation error converges to zero.

III. MAIN RESULT

A. Overview

This section contains the main results of this paper. First, a brief review of control contraction metrics is presented. Then, the adaptive robust ccm is defined. Lastly, we show how an adaptive robust ccm can be constructed via convex optimization to achieve the bounds in Definition 2.

B. Adaptive Robust Control Contraction Metrics

Contraction analysis [5] is an attractive framework for synthesizing stabilizing control laws. In particular, the search for a ccm $M : \mathbb{R}^n \rightarrow \mathcal{S}_+^n$ can be formulated as a convex condition amenable to numerical methods. Furthermore, once a metric $M(x)$ is found then the corresponding control law ensures any two trajectories converge to each other exponentially. This convergence can be stated in terms of geodesics and Riemannian energy. For a smooth manifold \mathcal{M} a geodesic $\gamma : [0, 1] \times \mathbb{R} \rightarrow \mathbb{R}^n$ is the shortest path between any two points. If we impose the boundary conditions $\gamma(0, t) = x_d(t)$, and $\gamma(1, t) = x(t)$ then the Riemannian energy $E(\gamma(s, t)) = \int_0^1 \gamma_s(s, t)^\top M(\gamma(s, t)) \gamma_s(s, t) ds$ where $\gamma_s(s, t) \triangleq \nabla_s \gamma(s, t)$ can be interpreted as a form of tracking error since $E(\gamma(s, t)) = 0 \iff x(t) = x_d(t)$. Since $M(x)$ is a ccm then $\dot{E}(\gamma(s, t)) \leq -2\lambda E(\gamma(s, t))$ yielding $x(t) \rightarrow x_d(t)$ exponentially with rate λ .

To further motivate this work, and the need for a new type of ccm, consider the differential dynamics of (1) $\delta_x = A_\theta(x, u)\delta_x + \Delta(x)^\top \delta_\theta + B(x)\delta_u$, where $A_\theta(x, u) = \nabla_x f(x) + \sum_{i=1}^m \nabla_x b_i(x) u_i + \sum_{i=1}^p \nabla_x \Delta_i(x) \theta_i$ is the Jacobian of the dynamics. When the model is known, a metric $M(x)$ is constructed so that the implication $\delta_x^\top M B = 0 \implies \delta_x^\top (A^\top M + M A + \dot{M}) \delta_x \leq -2\lambda \delta_x^\top M \delta_x$ is true; the dependency on u in the Jacobian can be eliminated by

imposing $B(x)$ be a Killing vector for $M(x)$. However, the presence of the uncertainty in the differential dynamics and Jacobian requires the metric satisfy a stronger condition to achieve bounded transients. The robust ccm was proposed to meet this need but is unable to incorporate online model estimates thereby limiting its performance. Hence, a new metric is needed in order to incorporate online model estimates without sacrificing tracking error transient bounds.

Definition 4. A uniformly bounded Riemannian metric $M_\theta : \mathbb{R}^n \times \mathbb{R}^p \rightarrow \mathcal{S}_+^n$ is an *adaptive robust control contraction metric* (arccm) if for each $\theta \in \Theta$ the dual metric $W_\theta(x) \triangleq M_\theta(x)^{-1}$ satisfies

$$\begin{aligned} -\partial_x W_\theta + A_\theta W_\theta + W_\theta A_\theta^\top + B Y_\theta + Y_\theta^\top B^\top \\ \preceq -2\lambda W_\theta + \frac{1}{\alpha^2} \Delta^\top \Delta, \end{aligned} \quad (\text{C1})$$

$$\partial_{b_i}^\circ W_\theta - W_\theta \nabla_x b_i^\top - \nabla_x b_i W_\theta = 0, \quad i = 1, \dots, m \quad (\text{C2})$$

$$-\mu W_\theta \preceq \nabla_{\theta_i} W_\theta \preceq \mu W_\theta, \quad i = 1, \dots, p \quad (\text{C3})$$

where $Y_\theta : \mathbb{R}^n \times \mathbb{R}^p \rightarrow \mathbb{R}^m \times \mathbb{R}^n$ and $\lambda, \mu, \alpha \in \mathbb{R}_{>0}$.

Several remarks are in order. Firstly, Eqs. (C1) and (C2) are similar to those in [8] but the metric is now allowed to be parameter-dependent. In the absence of parameter adaptation Eqs. (C1) and (C2) are sufficient conditions for the closed-loop system to be ISC with respect to $\hat{\theta}$. Note (C2) is only needed if the control input matrix depends on x thereby ensuring $B(x)$ is a Killing vector field for $M(x)$. Secondly, since the metric is parameter-dependent, a valid metric is immediately available for any $\theta \in \Theta$. Hence, $M_\theta(x)$ represents a *family of metrics* that can immediately incorporate online estimates of θ , thereby addressing the fundamental limitations of the robust ccm. Since the metric will change as θ is continuously being replaced with $\hat{\theta}$ from the estimator, (C1) would usually include another term related to the transients of $\hat{\theta}$ thereby complicating the search for $M_\theta(x)$. Introducing (C3), which is part of the novelty of this work, we bypass the difficulty of incorporating model estimation transients in (C1) while still achieving tracking error transient bounds; this will be shown in Theorem 1. Condition (C3) can be equivalently stated (see Proposition 1) as a bound on the log gradient of the Riemannian energy when $x(t) \neq x_d(t)$. This interpretation of condition (C3) provides a useful insight into how (C3) influences the metric and subsequently the Riemannian energy, namely it limits how fast the log of the Riemannian energy can change.

Proposition 1. When $x(t) \neq x_d(t)$ then condition (C3) holds for any curve $c(s, t)$ if and only if

$$|\nabla_{\theta_i} \log(E_\theta(c(s, t)))| \leq \mu, \quad i = 1, \dots, p \quad (4)$$

where $E_\theta(c(s, t))$ is the Riemannian energy of curve $c(s, t)$.

Proof. Since $x(t) \neq x_d(t)$ then $E_\theta(c) > 0$. Suppose $|\nabla_{\theta_i} \log(E_\theta(c))| \leq \mu$ (omitting the arguments of curve c). Applying the chain rule and Liebnitz integral rule to the

definition of the Riemannian energy, we obtain

$$\begin{aligned} |\nabla_{\theta_i} \log(E_\theta(c))| &= |\nabla_{\theta_i} E_\theta(c)| / E_\theta(c) \leq \mu, \\ \implies -\mu E_\theta(c) &\leq \nabla_{\theta_i} E_\theta(c) \leq \mu E_\theta(c), \\ \implies -\mu M_\theta(c) &\preceq \nabla_{\theta_i} M(c) \preceq \mu M_\theta(c), \end{aligned}$$

where the last implication follows from the definition of E_θ . Letting $W_\theta(x) = M_\theta(x)^{-1}$ then one arrives to (C3). Conversely, if (C3) holds then $|c_s^\top \nabla_{\theta_i} M(c) c_s| \leq \mu c_s^\top M_\theta(c) c_s$ for any curve c . Integrating with respect to s and applying Cauchy-Schwarz inequality yields

$$|\nabla_{\theta_i} E_\theta(c)| \leq \mu E_\theta(c) \implies |\nabla_{\theta_i} \log(E_\theta(c))| \leq \mu. \quad \square$$

For numerically constructing an arccm, we see that Eqs. (C1) to (C3) are jointly convex in W_θ , Y_θ and $\frac{1}{\alpha^2}$ so they can be efficiently computed using numerical methods. However, they are non-convex in μ or λ so a line search method is needed to optimize over these parameters. Picking an appropriate cost function that balances the competing objectives of convergence rate, overshoot, and control effort is also critical to obtain the desirable closed-loop response.

We now present the main result of the paper, which shows that the tractable conditions laid out in Definition 4 are sufficient to achieve definition Definition 2. The methodology of the proof is similar to the ones in [8], [9] with the key difference being how to handle $M_{\hat{\theta}}$ being a function of $\hat{\theta}$. If $x(t)$ happens to be in the cut locus then the proof can be modified with the technique used in [7]. We make the following assumption to ensure the uniqueness of geodesics and smoothness of the Riemannian energy.

Assumption 3. For the control system (1) the set of times for which $x(t)$ is in the cut locus of $x_d(t)$ has zero measure. As shown in [7] this assumption is not strictly necessary and the upper Dini derivative can be used to remove the assumption.

Theorem 1. Let $\hat{\theta}$ be the current estimate of the unknown model parameters θ . If an adaptive robust ccm $M_\theta(x)$ exists, then the closed-loop system is eISC with respect to the estimation error $\tilde{\theta}$ with the control law

$$u(t) = u_d(t) + \int_0^1 K_{\hat{\theta}}(\gamma(s, t)) \gamma_s(s, t) ds, \quad (5)$$

where $K_{\hat{\theta}}(x) \triangleq Y_{\hat{\theta}}(x) M_{\hat{\theta}}(x)$ computed via the conditions in Definition 4 and $\gamma(s, t)$ is a geodesic connecting $x(t)$ and $x_d(t)$. Hence, if $\tilde{\theta} \rightarrow 0$ then $x(t) \rightarrow x_d(t)$ exponentially.

Proof. Let $\gamma(s, t)$ be a geodesic where $\gamma(0, t) = x_d(t)$ and $\gamma(1, t) = x(t)$. Moreover, let $c(s, t)$ be a curve that evolves according to (1) where $c(s, t') = \gamma(s, t')$ for some time t' . If the control law $u(s, t) = u_d(t) + \int_0^s K_{\hat{\theta}}(c) c_s ds$ where $\delta_u = u_s(s, t) = K_{\hat{\theta}}(c) c_s$ is applied for some interval $t \in [t', t' + \epsilon]$, then the differential dynamic of c becomes

$$\dot{c}_s = [A_{\hat{\theta}}(c, u) + B(c) K_{\hat{\theta}}(c)] c_s + \Delta(c)^\top \tilde{\theta},$$

where $\theta(s, t) = s\hat{\theta}(t) + (1-s)\theta(t)$ so $\delta_\theta = \theta_s = \tilde{\theta}(t)$.

Letting $V_{\hat{\theta}}(c, c_s) = c_s^\top M_{\hat{\theta}}(c) c_s$ then

$$\begin{aligned} \frac{d}{dt}(c_s^\top M_{\hat{\theta}} c_s) &= c_s^\top \left(A^\top M_{\hat{\theta}} + M_{\hat{\theta}} A + \dot{M}_{\hat{\theta}} + \tilde{\theta}^\top \Delta M_{\hat{\theta}} \right. \\ &\quad \left. + M_{\hat{\theta}} \Delta^\top \tilde{\theta} + K_{\hat{\theta}}^\top g^\top M_{\hat{\theta}} + M_{\hat{\theta}} K_{\hat{\theta}} B \right) c_s, \end{aligned}$$

where $\dot{M}_{\hat{\theta}}(c) = \partial_c M_{\hat{\theta}}(c) + \partial_{\hat{\theta}} M_{\hat{\theta}}(c)$. Making use of (C1) in Definition 4, taking the Schur compliment and multiplying both sides by $[\eta_x, \delta_\theta]$ and $[\eta_x, \delta_\theta]^\top$ gives

$$\begin{bmatrix} \eta_x \\ \delta_\theta \end{bmatrix}^\top \begin{bmatrix} \mathcal{W} & -\Delta \\ -\Delta & \alpha^2 I \end{bmatrix} \begin{bmatrix} \eta_x \\ \delta_\theta \end{bmatrix} \geq 0,$$

where $\mathcal{W} = \partial_c W_\theta - \lambda W_\theta - A_\theta W_\theta - W_\theta A_\theta^\top - g Y_\theta - Y_\theta^\top g^\top$. Expanding the above expression, making the substitutions $\eta_x = M_{\hat{\theta}} c_s$ and $Y_{\hat{\theta}} = K_{\hat{\theta}} M_{\hat{\theta}}^{-1}$, and applying condition (C3) one can show that $\dot{V}_{\hat{\theta}}$ becomes

$$\frac{d}{dt}(c_s^\top M_{\hat{\theta}} c_s) \leq -\rho(t) c_s^\top M_{\hat{\theta}} c_s + \alpha^2 |\tilde{\theta}|^2,$$

where $\rho(t) \triangleq \lambda - p\mu \left| \frac{d}{dt} \hat{\theta} \right|$. Integrating along $c(s, t)$ leads to

$$\int_0^1 \frac{d}{dt}(c_s^\top M_{\hat{\theta}} c_s) ds \leq -\rho(t) \int_0^1 c_s^\top M_{\hat{\theta}} c_s ds + \alpha^2 \int_0^1 |\tilde{\theta}|^2 ds.$$

Switching the order of integration and differentiation yeilds

$$\frac{d}{dt} E(c(s, t)) \leq -\rho(t) E(c(s, t)) + \alpha^2 |\tilde{\theta}|^2.$$

For $t \in [t', t' + \epsilon]$ taking the limit as ϵ goes to 0 gives

$$\begin{aligned} \lim_{\epsilon \rightarrow 0} \frac{d}{dt} E(c(s, t)) &= \frac{d}{dt} \lim_{\epsilon \rightarrow 0} E(c(s, t)) \\ &\leq \lim_{\epsilon \rightarrow 0} \left(-\rho(t) E(c(s, t)) + \alpha^2 |\tilde{\theta}|^2 \right). \end{aligned}$$

Since $c(s, t') = \gamma(s, t')$ and t' was arbitrary we obtain

$$\frac{d}{dt} E(\gamma(s, t)) \leq -\rho(t) E(\gamma(s, t)) + \alpha^2 |\tilde{\theta}|^2. \quad (6)$$

Recalling $\tilde{\theta}$ and $\frac{d}{dt} \hat{\theta}$ are bounded by assumption, letting $\bar{\rho} \triangleq \lambda - p\mu \left\| \frac{d}{dt} \hat{\theta} \right\|_\infty \leq \rho$ then (6) becomes

$$E(\gamma(s, t)) \leq E(\gamma(s, 0)) e^{-\bar{\rho} t} + \frac{\alpha^2}{\bar{\rho}} \|\tilde{\theta}\|_\infty^2 (1 - e^{-\bar{\rho} t}), \quad (7)$$

which can be put into the same form as that in Definition 2 by noting $M_{\hat{\theta}}(x)$ is uniformly bounded so there exists $\underline{a}, \bar{a} \in \mathbb{R}_{>0}$ such that $\underline{a}I \preceq M_{\hat{\theta}}(x) \preceq \bar{a}I$ which leads to

$$|x(t) - x_d(t)| \leq \sqrt{\frac{\bar{a}}{\underline{a}}} |x(0) - x_d(0)| e^{-\bar{\rho} t} + \frac{\alpha^2}{\bar{\rho}} \|\tilde{\theta}\|_\infty^2 (1 - e^{-\bar{\rho} t}).$$

Therefore, the closed-loop system is eISC with respect to $\tilde{\theta}$ with the adaptive robust ccm $M_\theta(x)$ and control law (5). Moreover, if $\tilde{\theta} \rightarrow 0$ then $x(t) \rightarrow x_d(t)$ exponentially. \square

It is important to note that although the system no longer has to be ISC / eISC with respect to the estimation transients, the convergence rate of the closed-loop system will change during estimation transients. It is critical to ensure that $\frac{d}{dt} \hat{\theta}$ never violates the condition $\rho(t) = \lambda - p\mu \left| \frac{d}{dt} \hat{\theta} \right| > 0$. Pragmatically this can be achieved by selecting or mod-

ulating the parameter estimation rate. Furthermore, there is an inherent trade-off between short-term tracking error convergence and long-term control performance from having a more accurate model estimate. In other words, the effects of fast model estimation must be carefully considered because large estimation transients will result in slower short-term convergence but better long-term control since the model is more accurately known.

C. Offline / Online Computation

As mentioned previously, the conditions in Definition 4 are jointly convex in $W_\theta(x)$, $Y_\theta(x)$, and $\frac{1}{\alpha^2}$. The values of these are determined offline using existing state-dependent LMI solvers. The controller (5) needs to be computed online, and requires a geodesic at each time. Geodesic computation does involve a nonlinear optimization but can be performed in real-time [19] and is less expensive than nonlinear MPC. The parameter estimates $\hat{\theta}$ should be computed online by any suitable model estimator to maximize performance.

IV. ILLUSTRATIVE EXAMPLE

Consider the system with state $x = [x_1, x_2, x_3]^\top$ unknown parameters $\theta = [\theta_1, \theta_2, \theta_3, \theta_4]^\top$ and dynamics

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \end{bmatrix} = \begin{bmatrix} x_3 - \theta_1 x_1 \\ -x_2 - \theta_2 x_1^2 \\ \tanh(x_2) - \theta_3 x_3 - \theta_4 x_1^2 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} u, \quad (8)$$

where $\theta_1 \in [-1, 1]$, $\theta_2 \in [0.5, 1.5]$, $\theta_3 \in [-0.6, 0.75]$ and $\theta_4 \in [-1.75, 0.5]$. The true model parameters are $\theta^* = [-0.3, 0.8, -0.25, -0.75]^\top$. This system has been used before in [11] (modified from that in [6]) because it demonstrates the versatility of contraction since the systems is not feedback linearizable, is not in strict feedback form, and has unmatched uncertainties. A desired trajectory x_d was generated by specifying $x_{1d}(t) = \sin(t)$; the remaining states $x_{2d}(t)$ and $x_{3d}(t)$ were generated based on an updating model of the system. An arccm was computed via sum-of-squares programming with YALMIP [20] where W_θ and Y_θ were parameterized by fourth-order matrix polynomials in state x and parameter θ . In order to ensure the conditions in Definition 4 were satisfied for all $\theta \in \Theta$, Eqs. (C1) and (C3) were imposed at each point in a uniform 21 x 21 grid for the unmatched parameters θ_1 and θ_2 . The control input was computed as outline in Section III-C.

Two controllers were tested to verify the properties of the proposed approach: one with model estimation and the other without. The adaptive controller was initialized with an estimate $\hat{\theta}$ set to the midpoint of Θ . At $t = 3$ s the parameters began to linearly change to the true parameters with convergence occurring at $t = 7$ s, i.e., $\tilde{\theta} = 0$; this is indicated by green vertical lines in Figs. 1 and 2. The non-adaptive controller used the same metric but $\hat{\theta}$ was fixed to the midpoint of Θ . The Riemannian energy — an indicator of the closed-loop tracking error — was calculated for the two controllers and is plotted in Fig. 1. As seen in the figure both controllers obey (7), as the energy is always upper bounded by the predicted conservative bound (black line).

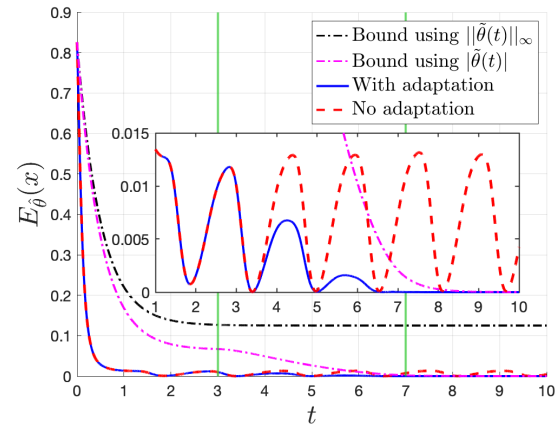


Fig. 1: Riemannian energy of the new method is lower than that of a controller with a static reference model. The model with adaptation begins adapting at $t = 3$ s and finishes at $t = 7$ s.

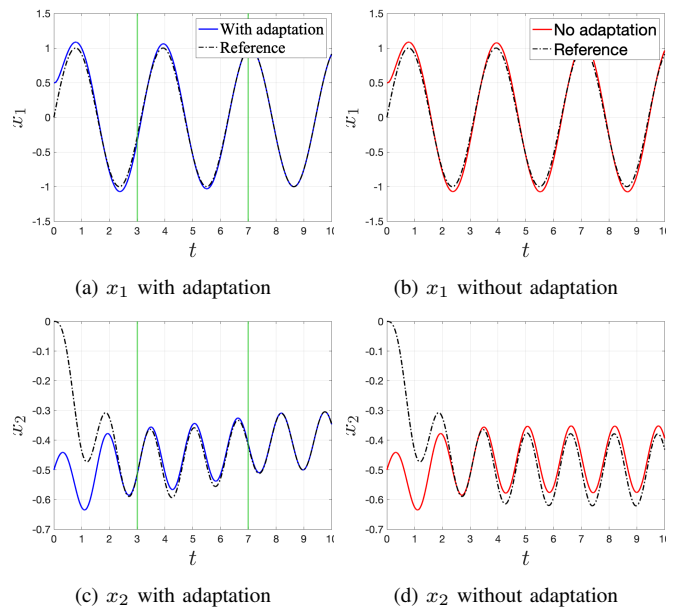


Fig. 2: States with and without model updates. The reference for x_1 is the same for both controllers but differs for other states due to model updates. The model updates start at $t = 3$ s and converges at $t = 7$ s (vertical green lines). State x_3 omitted for brevity as it exhibits identical behavior.

The black line is the bound based on the most aggressive adaption and worst case $\tilde{\theta}$, which can be used to construct a worst-case error bound from the start and therefore be employed for predictive control. Fig. 1 also shows the bound calculated by integrating (6) (magenta line). In practice, the bound from (6) would not be known as it relies on knowing $\tilde{\theta}(t)$ (and hence knowing θ at all time) but is useful here to confirm the theoretical result that $x \rightarrow x_d$ as $\tilde{\theta} \rightarrow 0$. As shown in the inset of Fig. 1, the tracking error with the robust (non-adaptive) controller satisfies the conservative bound but does not converge to zero. Although the magenta bound is not an achievable worst case bound, the worst case bound could be improved if $\tilde{\Theta}$ is reduced via, e.g., set membership identification, to achieve zero tracking error. This is in contrast to only improving $\tilde{\theta}$ which yields better performance but the conservative bound is unchanged. Fig. 2 shows the time history of x_1 and x_2 for reference.

V. CONCLUSION

We presented a contraction-based framework for robust trajectory tracking despite the presence of uncertain dynamics that are being estimated online. After introducing the adaptive robust control contraction metric (arccm), which is allowed to depend on the unknown parameters, a feedback control law was proposed and shown to yield an exponentially input-to-state convergent closed-loop system. By allowing the metric to depend on the unknown parameters we are able to guarantee the trajectory tracking error will converge to zero if the parameter estimates converge to their true value. The key mechanism for allowing the metric to be parameter dependent is a novel gradient condition that only changes the convergence rate, as opposed to sacrificing stability or transient bounds. Future work includes using arccm's in adaptive MPC, as well as developing new methods for computing contraction metrics in high-dimensions.

VI. APPENDIX

A Lyapunov generalization of the result presented in Section III are derived here. Note that generalizing comes at the expense of losing convexity of the constructive conditions.

Definition 5. A continuously differentiable function $V_\theta : \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^p \rightarrow \mathbb{R}_+$ where $V_\theta(x, x_d) = 0 \iff x(t) = x_d(t)$ is an *adaptive* eISC-clf if the following hold

$$k_1|x - x_d|^a \leq V_\theta(x, x_d) \leq k_2|x - x_d|^a \quad (2)$$

$$|x(t) - x_d(t)| \geq \sigma \left(\sup_{\tau \in [t_0, t]} |\tilde{\theta}(\tau)| \right) \implies$$

$$\inf_{u \in \mathbb{R}^m} \left\{ \nabla_x V_\theta^\top F_{\tilde{\theta}}(x, u) + \nabla_{x_d} V_\theta^\top F_{\tilde{\theta}}(x_d, u_d) - \nabla_x V_\theta^\top \Delta^\top \tilde{\theta} \right\} \leq -k_3|x - x_d|^a \quad (9)$$

$$-\mu V_\theta \leq \nabla_{\theta_i} V_\theta \leq \mu V_\theta, \quad i = 1, \dots, p \quad (10)$$

where $k_1, k_2, k_3, a, \mu \in \mathbb{R}_{>0}$, $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is a class \mathcal{K} function, and $\tilde{\theta} \in \Theta$ and $\frac{d}{dt}\tilde{\theta} \in \Omega$ belong to compact sets.

Theorem 2. Let $\hat{\theta}$ be the current estimate of the model parameters θ . If an adaptive eISC-clf $V_{\hat{\theta}}(x, x_d)$ exists, then the closed-loop system is eISC with respect to the estimation error $\tilde{\theta}$. Hence, if $\tilde{\theta} \rightarrow 0$ then $x(t) \rightarrow x_d(t)$ exponentially.

Proof. Let $V_{\hat{\theta}}$ be an aISC-clf. Consider when $|x - x_d| \geq \sigma(\sup_{\tau \in [t_0, t]} |\tilde{\theta}(\tau)|)$. Differentiating along (1) yields

$$\dot{V}_{\hat{\theta}} \leq -k_3|x - x_d|^a + \nabla_{\tilde{\theta}} V_{\hat{\theta}}^\top \dot{\tilde{\theta}} \leq -\frac{k_3}{k_2}V_{\hat{\theta}} + p\mu V_{\hat{\theta}}|\dot{\tilde{\theta}}|$$

where we make use of Eqs. (2), (9) and (10) to obtain the second inequality. If $\frac{d}{dt}\tilde{\theta}$ is designed so $\rho(t) = \frac{k_3}{k_2} - p\mu|\frac{d}{dt}\tilde{\theta}| > 0$ and since $\tilde{\theta}$ and $\frac{d}{dt}\tilde{\theta}$ are bounded then letting $\bar{\rho} \triangleq \frac{k_3}{k_2} - p\mu|\frac{d}{dt}\tilde{\theta}|_\infty \leq \rho(t)$ yields $\dot{V}_{\hat{\theta}} \leq -\bar{\rho}V_{\hat{\theta}} \implies |x(t) - x_d(t)| \leq \sigma_1(|x(0) - x_d(0)|)e^{-\lambda t}$ via (2). Since this holds only when the implication (9) is true, we can conclude $|x(t) - x_d(t)| \leq \sigma_1(|x(0) - x_d(0)|)e^{-\lambda t} + \sigma_1(\sup_{\tau \in [t_0, t]} |\tilde{\theta}(\tau)|)$. Therefore, the closed-loop system is eISC with respect to $\tilde{\theta}$. Moreover, if $\tilde{\theta} \rightarrow 0$ then $x(t) \rightarrow x_d(t)$ exponentially. \square

We see that imposing a gradient bound on $V_{\hat{\theta}}(x, x_d)$ in (10) similar to (C3) again bypasses requiring ISC with respect to the parameter estimation transients. When $x(t) \neq x_d(t)$ then (10) can also be interpreted as a bound on the log gradient of $V_{\hat{\theta}}(x, x_d)$. Despite the similarity, it is well-known that computing clf's for general nonlinear system presents several difficulties that contraction theory circumvents [6].

REFERENCES

- [1] A. Pavlov, N. Van De Wouw, and H. Nijmeijer, *Uniform output regulation of nonlinear systems: a convergent dynamics approach*. Springer, 2006, vol. 205.
- [2] D. Angeli, "A lyapunov approach to incremental stability properties," *IEEE Transactions on Automatic Control*, vol. 47, no. 3, pp. 410–421, 2002.
- [3] M. Krstic, P. V. Kokotovic, and I. Kanellakopoulos, *Nonlinear and adaptive control design*. John Wiley & Sons, Inc., 1995.
- [4] M. Zamani, N. van de Wouw, and R. Majumdar, "Backstepping controller synthesis and characterizations of incremental stability," *Systems & Control Letters*, vol. 62, no. 10, pp. 949–962, 2013.
- [5] W. Lohmiller and J.-J. E. Slotine, "On contraction analysis for nonlinear systems," *Automatica*, vol. 34, no. 6, pp. 683–696, 1998.
- [6] I. R. Manchester and J.-J. E. Slotine, "Control contraction metrics: Convex and intrinsic criteria for nonlinear feedback design," *IEEE Transactions on Automatic Control*, vol. 62, no. 6, pp. 3046–3053, 2017.
- [7] S. Singh, A. Majumdar, J.-J. Slotine, and M. Pavone, "Robust online motion planning via contraction theory and convex optimization," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 5883–5890.
- [8] I. R. Manchester and J.-J. E. Slotine, "Robust control contraction metrics: A convex approach to nonlinear state-feedback H^∞ control," *IEEE Control Systems Letters*, vol. 2, no. 3, pp. 333–338, 2018.
- [9] P. Zhao, A. Lakshmanan, K. Ackerman, A. Gahlawat, M. Pavone, and N. Hovakimyan, "Tube-certified trajectory tracking for nonlinear systems with robust control contraction metrics," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5528–5535, 2022.
- [10] B. T. Lopez and J.-J. E. Slotine, "Adaptive nonlinear control with contraction metrics," *IEEE Control Systems Letters*, vol. 5, no. 1, pp. 205–210, 2020.
- [11] —, "Universal adaptive control of nonlinear systems," *IEEE Control Systems Letters*, vol. 6, pp. 1826–1830, 2021.
- [12] H. K. Khalil, "Adaptive output feedback control of nonlinear systems represented by input-output models," *IEEE transactions on Automatic Control*, vol. 41, no. 2, pp. 177–188, 1996.
- [13] J. Köhler, P. Kötting, R. Soloperto, F. Allgöwer, and M. A. Müller, "A robust adaptive model predictive control framework for nonlinear uncertain systems," *International Journal of Robust and Nonlinear Control*, vol. 31, no. 18, pp. 8725–8749, 2021.
- [14] G. Chowdhary, T. Yucelen, M. Mühlegg, and E. N. Johnson, "Concurrent learning adaptive control of linear systems with exponentially convergent bounds," *International Journal of Adaptive Control and Signal Processing*, vol. 27, no. 4, pp. 280–301, 2013.
- [15] R. L. Kosut, M. K. Lau, and S. P. Boyd, "Set-membership identification of systems with parametric and nonparametric uncertainty," *IEEE Transactions on Automatic Control*, vol. 37, no. 7, pp. 929–941, 1992.
- [16] E. D. Sontag and Y. Wang, "On characterizations of the input-to-state stability property," *Systems & Control Letters*, vol. 24, no. 5, pp. 351–359, 1995.
- [17] B. S. Rüffer, N. Van De Wouw, and M. Mueller, "Convergent systems vs. incremental stability," *Systems & Control Letters*, vol. 62, no. 3, pp. 277–285, 2013.
- [18] M. Zamani and P. Tabuada, "Backstepping design for incremental stability," *IEEE Transactions on Automatic Control*, vol. 56, no. 9, pp. 2184–2189, 2011.
- [19] K. Leung and I. R. Manchester, "Nonlinear stabilization via control contraction metrics: A pseudospectral approach for computing geodesics," in *2017 American Control Conference (ACC)*. IEEE, 2017, pp. 1284–1289.
- [20] J. Lofberg, "Yalmip: A toolbox for modeling and optimization in matlab," in *2004 IEEE international conference on robotics and automation (IEEE Cat. No. 04CH37508)*. IEEE, 2004, pp. 284–289.