MIT Open Access Articles

Region-of-convergence estimation for learning-based adaptive controllers

Region-of-Convergence Estimation for Learning-Based Adaptive Controllers

John F. Quindlen[1] Ufuk Topcu[2] Girish Chowdhary[3] Jonathan P. How[4]

Abstract— **Recent learning-based extensions to popular adaptive control procedures offer improved convergence, but at the cost of increased complexity. This complexity makes it difficult to analytically compute level sets that bound the system response. These level sets can be combined with the a priori known Lyapunov function for such systems to provide barrier certificates, verifying the safety of the system to maximum allowable error limits. This paper presents a complementary automated procedure for computing invariant level sets offline using simulation data. These level sets encompass combinations of safe initial conditions and parameters that will not cause the adaptive system's response to exceed constraints. First, conditions for the complete set of safe initial states and parameters, known as the region-of-convergence, are established. These conditions, coupled with the known Lyapunov functions describing the adaptation, are used to form an optimization procedure to construct verifiable level sets for the system response. These levels sets thus provide barrier certificates for safety and conservatively estimate the complete region-of-convergence. Lastly, the procedure is demonstrated on an adaptive control system.**

I. INTRODUCTION

Adaptive control procedures combine parameter estimation with control to achieve a desire level of performance while maintaining robustness. In particular, popular model reference adaptive control (MRAC) methods have been extensively studied in order to guarantee the stability of the adaptation and the ability of the uncertain system to track the desired reference model [1]–[3]. As a result, the stability and convergence is defined using well-known Lyapunov functions than can be used to prove Uniform Ultimate Boundedness of the system and determine level sets that bound the system response [3].

Recent approaches [4]–[6] have improved performance and robustness of the baseline MRAC methods, but at the cost of increased complexity. This complexity presents a challenge towards the calculation of level sets that describe the system response. For instance, concurrent learning MRAC (CL-MRAC) adds a history stack of saved data points to improve convergence of the adaptation and decrease transient errors in the response, but this history stack complicates analysis of the system. Analytical bounds on the CL-MRAC response have been shown [7], but these bounds are difficult to calculate in practice as they explicitly depend on the particular points saved in the history stack. Since CL-MRAC periodically updates the history stack, the bounds will have to be recomputed, making it difficult to provide more general certificates. Until a bounding set can be determined, these adaptive systems cannot be theoretically verified as safe.

In contrast, simulations of adaptive systems are an extremely useful tool to gain insight on the behavior of the adaptive system. Model checking procedures [8], [9] already utilize large batches of simulation traces to generate statistical guarantees on complex, stochastic, nonlinear systems. When dealing with deterministic systems, these guarantees can be tightened even further. In particular, recent work in simulation-guided region-of-attraction [10], [11] and Lyapunov function [12], [13] analysis procedures provide theoretically guaranteed bounds generated using simulation traces. These approaches intelligently sample the state space to find falsifying simulation traces and construct a linear program that describes a Lyapunov function and its maximum invariant level set that avoids such traces. The crux of these methods is that if the Lyapunov function and its corresponding maximum invariant set meets certain conditions, then they create a barrier certificate that verifies the stability of the complex nonlinear system. In short, if the system starts within the region-of-attraction identified by the procedure, then the system response will always remain bounded and eventually converge towards the equilibrium point. This provides a powerful method for automatically and provably verifying the stability of a system of interest.

In this paper, a similar simulation-based methodology is used to analyze the response of learning-based adaptive controllers. For a given deterministic adaptive control system, the system's state and adaptation trajectory is completely characterized by the initial state conditions and parameters. When constraints are placed on the allowable limits of the system response, the initial states and parameters determine whether the system response will exceed those constraints. If a certain combination of initial states and parameters do not lead to a violation of these constraints, then the combination is labeled as "safe". The complete set of safe combinations is defined as the region-of-convergence.

Simulations of the adaptive system at various combinations can be employed to estimate this region with an invariant level set and create barrier certificates. The simulations help compute an absolute lower bound on the failed trajectories. Paired with the fact that the known Lyapunov function for such systems is bounded, this information forms a maximum invariant level set on safe trajectories. The com-

[1]Graduate Student, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology

[2]Assistant Professor, Department of Aerospace Engineering and Engineering Mechanics, University of Texas at Austin

[3]Assistant Professor, Department of Mechanical and Aerospace Engineering, Oklahoma State University

[4]Richard C. Maclaurin Professor of Aeronautics and Astronautics, Aerospace Controls Laboratory (Director), MIT

bination of the Lyapunov function and maximum invariant set provides a barrier certificate such that any trajectories started within some initial region will never exceed the constraints.

The paper is organized as follows. First, a background on recent learning-based model reference adaptive controllers is given. In particular, the concurrent learning MRAC algorithm is used as a case study for these approaches. This background is then used to motivate and define the region-of-convergence problem for these adaptive systems. Once the region-of-convergence has been defined, the next section describes the simulation-based procedure to estimate this set using simulation traces. The procedure uses an optimization problem to maximize an invariant set within the region-of-convergence, which also provides a barrier certificate verifying the safety. The last section demonstrates the procedure on an adaptive control system.

II. MODEL REFERENCE ADAPTIVE CONTROL

Consider the following uncertain dynamical system

$$\dot{x}(t) = Ax(t) + B(u(t) + W^{*T}\phi(x)) \quad (1)$$
$$x(0) \in X_{\text{init}}$$

where $x(t) \in \mathbb{R}^n$ is the state vector and $u(t) \in \mathbb{R}^m$ is the corresponding control input vector. Assume $x(t)$ is available for measurement and the pair of matrices $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{n \times m}$ are both known and controllable. This pair (A, B) defines the nominal system dynamics; however, uncertainties $\Delta(x) \in \mathbb{R}^m$ are also present in the actual dynamics. These uncertainties $\Delta(x)$ are assumed to have a known structure parameterized by an unknown constant weighting matrix $W^* \in \mathbb{R}^{k_n \times m}$ and a known regressor vector $\phi(x) \in \mathbb{R}^{k_n}$

$$\Delta(x) = W^{*T}\phi(x) \quad (2)$$

where $\phi(x)$ can be either a linear or nonlinear basis vector without loss of generality. Let regressor $\phi(x)$ be Lipschitz continuous to ensure the existence and uniqueness of the simulation trace [1]. While the exact values of the parameters in the system are unknown, the set of all possible parameters $W^*_{\text{feas}} \in \mathbb{R}^k$ is known, that is all possible $W^* \in W^*_{\text{feas}}$ with $k = k_n \times m$. Additionally, the initial state $x(0)$ may be uncertain, but is also assumed to fall within some known feasible region $X_{\text{init}} \in \mathbb{R}^n$.

The adaptive controller forces the system to track a known, stable reference system with desirable characteristics. In this paper, the underlying reference model is assumed to be linear and the resulting reference model dynamics are given by

$$\dot{x}_m(t) = A_m x_m(t) + B_m r(t) \quad (3)$$

where $x_m(t) \in \mathbb{R}^n$ is the reference state and $r(t) \in \mathbb{R}^{m_r}$, $m_r \leq m$ is the external reference command, assumed to be piecewise continuous and bounded. The state tracking error between the reference state trajectory and the actual state trajectory is given by $e(t) = x_m(t) - x(t)$. It is assumed that all $x_m(0)$ start at the same nominal initial condition x_{nom} and any uncertainty on $x(0)$ can be treated as disturbances to the known x_{nom}.

The control system then attempts to minimize the state tracking error $e(t)$ using control inputs $u(t)$. The control input is the combination of three components:

$$u = u_{rm} + u_{pd} - u_{ad} \quad (4)$$

a feedforward term $u_{rm} = K_r r(t)$ where $B_m = BK_r$ and $K_r \in \mathbb{R}^{m \times m_r}$, a feedback term $u_{pd} = -Kx(t)$ where $A_m = A - BK$, and an adaptive control input u_{ad}.

In the absence of uncertainties, the feedforward u_{rm} and feedback u_{pd} terms would force the nominal system to track the reference trajectory. Now, the adaptive control input u_{ad} attempts to suppress the uncertainty $\Delta(x)$ by estimating the parameterization of $\Delta(x)$ and updating u_{ad} accordingly. The adaptive inputs follow the same known structure of $\Delta(x)$:

$$u_{ad} = \widehat{W}^T \phi(x) \quad (5)$$

with estimated parameters $\widehat{W}$ in place of actual W^*. The resulting estimation error is $\widetilde{W} = \widehat{W} - W^*$. Usually, the initial estimates can be set to zero, $\widehat{W}(0) = 0$, as any known quantities are assumed to be included in the nominal system (A, B).

A. Concurrent Learning Adaptive Control

Many possible techniques exist for adapting the parameter estimates [1]–[6], but MRAC-based procedures all update their estimates based upon the reference model tracking error $e(t)$. For the remainder of the paper, the concurrent learning MRAC (CL-MRAC) procedure will be used as the adaptive process, but other techniques can also be employed. The concurrent learning adaptive law is given by

$$\dot{\widehat{W}} = -\Gamma\phi(x)e^T PB - \Gamma_c \sum_{i=1}^{\overline{p}} \phi(x_i)\epsilon_i^T \quad (6)$$

where learning rates Γ and Γ_c are positive scalars. The symmetric positive definite matrix P is determined from the Lyapunov equation $A_m^T P + P A_m = -Q$ with $Q = Q^T > 0$. The innovation of the concurrent learning adaptive law is the history stack, which stores a fixed, finite number $\overline{p}$ of measurements of the state vector and estimation error $\epsilon_i = u_{ad,i} - \Delta_i = \widetilde{W}^T \phi(x_i)$ alongside the instantaneous tracking error. Central to this process are measurements of $\dot{x}$ to obtain $\Delta(x)$. The derivative $\dot{x}$ can be either measured directly or estimated using a fixed point smoother [4].

$$\Delta(x) = (B^T B)^{-1} B^T (\dot{x} - Ax - Bu) \quad (7)$$

An important aspect of this CL-MRAC procedure is the singular value maximizing (SVM) algorithm that periodically replaces the least-informative data points $(\phi(x_i), \epsilon_i)$ in the history stack with more-informative ones as they are encountered. This ensures the matrix $S(x, t) = \sum_{i=1}^{\overline{p}} \phi(x_i)\phi(x_i)^T$ is positive definite and results in a net increase in the rate of convergence. However, this periodic update of the history stack makes it difficult to compute analytical bounds, as discussed in Section III. Proofs and more detailed discussions of this algorithm can be found in [4] or in an extended version of the paper at http://tinyurl.com/jk4fzfz.

MRAC procedures use a known Lyapunov function to prove the stability of the system. This Lyapunov function V is differentiable, positive definite, and radially unbounded and describes both the stability of the reference model tracking error and parameter estimation error with $V(0) = 0$.

$$V = e^T P e + \texttt{trace}(\widetilde{W}^T \Gamma^{-1} \widetilde{W}) \qquad (8)$$

For the CL-MRAC procedure, the derivative of this Lyapunov function is a strictly negative definite function that asymptotically converges to the solution $e = 0, \widetilde{W} = 0$.

$$\dot{V} = -e^T Q e - 2\Gamma^{-1}\Gamma_c \texttt{trace}(\widetilde{W}^T \sum_{i=1}^{\overline{p}} \phi(x_i)\phi(x_i)^T \widetilde{W}) \qquad (9)$$

This Lyapunov function and its derivative was shown in [4] to prove the stability and convergence of the CL-MRAC controller. Similar proofs for other adaptive procedures establish the Lyapunov stability of the adaptation and convergence toward the stable reference model, thus also establishing stability of the adaptive system.

III. REGION-OF-CONVERGENCE OF CONSTRAINED ADAPTIVE SYSTEMS

The Lyapunov function in (8) is common to most MRAC-based adaptive procedures. If $V(e, \widetilde{W}) > 0$ and $\dot{V}(e, \widetilde{W}) \leq 0$ for all $(e, \widetilde{W}) \neq 0$, this function establishes the Lyapunov stability and convergence of the controlled system towards the desired reference trajectory. Additionally, this also indicates that in the absence of restrictions on $u(t)$, any value for W^* and $x(0)$ will not destabilize the adaptation. Because the Lyapunov function (8) is radially unbounded, the adaptive system is globally stable.

While the ultimate stability of the adaptation is invariant to the particular values of $x(0)$ and W^*, the transient response of $e(t)$ and $\widetilde{W}(t)$ during the adaptation must also be considered. Many systems place additional constraints on the system response, such as maximum allowable thresholds on the tracking error $e(t)$. Stability alone does not guarantee these performance constraints will be met. For instance, consider Figure 1, which depicts a sample trajectory taken from the CL-MRAC controlled system in Section V-A. The adaptation is known to be stable, but the transient tracking error $e(t)$ response exceeds the allowable threshold. The error dynamics are dependent upon the particular $x(0)$, W^* values used; therefore, it is desirable to identify whether a set of initial conditions $x(0)$ and parameters W^* will cause the adaptive system to exceed any performance thresholds.

The set of all possible combinations of initial conditions $x(0)$ and parameters W^* is the feasible search region $R_{\text{feas}} \in \mathbb{R}^{n+k}$, where $R_{\text{feas}} := X_{\text{init}} \cup W^*_{\text{feas}}$. Since different combinations of $x(0)$ and W^* will either cause the system to exceed the constraints or not, the feasible region R_{feas} can be broken up into two independent subsets. In order to define these subsets, first define the space outside these constraint thresholds, called the failure region R_{fail}.


Fig. 1. State space diagram of the response of a CL-MRAC system. Even though the system is stable [4], the actual state exceeds the allowable error.

Definition 1 The failure region $R_{\text{fail}} \in \mathbb{R}^n$ is the set of infeasible tracking errors $e_{\text{fail}} \in \mathbb{R}^n$ outside the allowable bounds.

$$R_{\text{fail}} := \{||e|| > e_{\text{limit}}\}$$

From this, the region-of-convergence (ROC) $R_0 \subset R_{\text{feas}}$ is defined below.

Definition 2 The region-of-convergence $R_0 \in \mathbb{R}^{n+k}$ is the set of $W^* \in W^*_0$ and $x(0) \in X_0$ such that the solution to (1), trajectory $\Phi(t : x_0, w_0)$, never enters the region R_{fail}. The trajectory $\Phi(t : x_0, w_0)$ is initialized at values $x_0 \in X_0$ and $w_0 \in W^*_0$, where $W^*_0 \subseteq W^*_{\text{feas}}$ and $X_0 \subseteq X_{\text{init}}$.

This region-of-convergence is generally unknown ahead of time, so the task is to estimate this region. Due to the nonlinear nature of the adaptive system, the shape of this region in $\mathbb{R}^{n+k}$ space can be complex and even non-convex, as will be seen in Section V-A.

The second subset of R_{feas} is the region-of-failure, R_F. Just like R_0, R_F considers initial conditions $x(0)$ and W^* rather than the space of all $e(t)$ like R_{fail}.

Definition 3 The region-of-failure $R_F \in \mathbb{R}^{n+k}$ is the set of $W^* \in W^*_0$ and $x(0) \in X_0$ such that the trajectory $\Phi(t : x_0, w_0)$ enters the region R_{fail}, again where $W^*_0 \subseteq W^*_{\text{feas}}$ and $X_0 \subseteq X_{\text{init}}$. Note that the intersection $R_F \cap R_0 = \emptyset$.

Previous work has provided some analytical tools for estimating the ROC. More specifically, extensions to the respective Lyapunov convergence proofs can be used to calculate upper bounds on the state tracking error $||e(t)||$ for each MRAC method under consideration [3], [7], as well as other approaches such as L1 adaptive control [14]. These upper bounds will then determine whether the system will exceed any of the indicated constraints. However, with more complex procedures, these analytical bounds on $e(t)$ become increasingly difficult and unwieldy to calculate offline, thus complicating a priori safety verification.

For instance, consider the following bound on tracking error for a CL-MRAC system. This bound is taken from the proof for Corollary 2 in [7]. It states that the tracking error $e(t)$ is bounded by an exponential function

$$||e||_{L_\infty} \leq (k_1||e(0)||_2^2 + k_2||W^*||_2^2)^{\frac{1}{2}} e^{-k_3 t} \qquad (10)$$

where k_1 and k_2 are functions of known constants P, Γ, and Γ_c. The difficulty in computing the analytical bounds arises from the fact that function k_3 depends explicitly on the saved values of $\phi(x_i)$ and ϵ_i stored in the history stack from the adaptive law in (6) and is periodically updated with new information. As mentioned in Section II-A, these values are gradually replaced with more-informative ones as the system evolves, changing the value of k_3. Since k_3 is not constant and evolves as the system executes, it is extremely difficult to compute the analytical bounds for $e(t)$ offline and ensure the system is safe before execution.

In comparison, it is rather straightforward to generate large numbers of simulations of the adaptive systems of interest and these simulations give insight into the sensitivity of the performance to X_{init} and W^*_{feas}. While simulation traces of the system at various samples of X_{init} and W^*_{feas} can provide statistical bounds on the system response, they generally lack the guarantees of the difficult, but verifiable, analytically determined bounds. Simulation-guided verification methods provide a bridge between the two. Simulation traces can be used to construct verifiable invariant level sets to estimate the region-of-convergence and form barrier certificates.

IV. ESTIMATION OF THE REGION-OF-CONVERGENCE

Large numbers of simulations are used to estimate the region-of-convergence of learning-based MRAC systems. The following procedure selectively samples R_{feas} and performs simulations at these values to find the smallest boundary between R_0 and R_F. The simulation traces are then used to construct barrier certificates to ensure a system will never enter R_{fail} if the system is initialized within set R_{init}.

This approach is an extension of recent work in simulation-guided region-of-attraction estimation, which addresses a closely related problem. These procedures [10], [11] attempt to find which system trajectories converge to the origin in asymptotically stable, but not globally attractive, nonlinear systems. Not surprisingly, these approaches can be readily adapted to address the region-of-convergence estimation problem. In particular, the robust estimation procedure [10], [11] provides a good groundwork for a region-of-convergence estimation procedure. The ROC procedure will utilize many of the same functions. A much more detailed analysis and discussion of these functions can be found in [15].

While the region-of-attraction (ROA) procedures and region-of-convergence (ROC) procedure are very similar, they do address different problems. Unlike ROA procedures, the ROC procedure already starts with a known Lyapunov function. The procedure does not need to search for a suitable Lyapunov function to characterize the system as it is provided by the respective adaptive control method under consideration. For MRAC systems, this will look like (8). Additionally, the ROC method must enforce stricter conditions (boundedness) in place of attraction to the origin. These differences require a slightly different approach for region-of-convergence estimation. The following subsections describe the ROC estimation procedure.

A. Invariant Subsets of the Region-Of-Convergence

The estimation of the region-of-convergence relies upon invariant sublevel sets of R_0 to form the basis of an optimization procedure. These invariant sublevel sets describe a set of $(e, \widetilde{W})$ for which the system trajectory will fall below a set threshold and ensure the system remains out of R_{fail}. For MRAC systems, sublevel sets can be written in terms of the Lyapunov function as follows.

Definition 4 For $\eta > 0$ and Lyapunov function $V : \mathbb{R}^{n+k} \to \mathbb{R}$ from (8), the η-sublevel set of V is:

$$\Omega_{V,\eta} := \{e \in \mathbb{R}^n, \widetilde{W} \in \mathbb{R}^k | V(e, \widetilde{W}) \leq \eta\}. \quad (11)$$

Sublevel sets of R_{feas} can then be used to define a set of initial values $x(0)$ and W^*, R_{init}, such that the Lyapunov function V of the corresponding response falls below some bound β. If all trajectories that enter R_{fail} have $V(e, \widetilde{W}) > \beta$, then the initial set R_{init} is a subset of R_0 and provides a barrier certificate $B(e, \widetilde{W})$ to verify the safety of the system.

Definition 5 Given an initial set $R_{\text{init}} \subset R_{\text{feas}}$ and a failure region $R_{\text{fail}} \subset R_{\text{feas}}$, a function $B: R_{\text{feas}} \to \mathbb{R}$ is a barrier certificate if:

$$B(e, \widetilde{W}) \leq 0 \qquad \text{for all } (e, \widetilde{W}) \in R_{\text{init}} \quad (12)$$

$$B(e, \widetilde{W}) > 0 \qquad \text{for all } (e, \widetilde{W}) \in R_{\text{fail}} \quad (13)$$

$$[\dot{e}; \dot{\widetilde{W}}]^T \, \nabla B(e, \widetilde{W}) < 0 \quad \text{for all } (e, \widetilde{W}) \in R_{\text{feas}} \quad (14)$$

$$\text{such that } B(e, \widetilde{W}) = 0.$$

With this definition, the following Lemma provides a description of barrier certificates for the system response in terms of sublevel sets of the Lyapunov function.

Lemma 1 Given 1) a continuously differentiable Lyapunov function $V : \mathbb{R}^{n+k} \to \mathbb{R}$ such that

$$V(0) = 0 \text{ and } V(e, \widetilde{W}) > 0 \text{ for all } (e, \widetilde{W}) \neq 0 \quad (15)$$

$$\dot{V}(e, \widetilde{W}) < 0 \text{ for all } (e, \widetilde{W}) \neq 0, \quad (16)$$

2) a η-sublevel set $\Omega_{V,\eta}$ of V, and 3) an initial set $R_{\text{init}} \subset \Omega_{V,\eta}$ and $R_{\text{init}} \subset R_{\text{feas}}$, if set R_{init} meets these conditions

$$V(e, \widetilde{W}) \leq \eta \text{ for all } (e, \widetilde{W}) \in R_{\text{init}} \quad (17)$$

$$V(e, \widetilde{W}) > \eta \text{ for all } (e, \widetilde{W}) \in R_{\text{fail}}, \quad (18)$$

then $B(e, \widetilde{W}) = V(e, \widetilde{W}) - \eta$ defines a barrier certificate and $R_{\text{init}} \subset R_0$.

Proof: A proof for Lemma 1 is omitted here for space, but can be found online in an extended version of this paper at http://tinyurl.com/jk4fzfz. ∎

Remark 1 If two additional weak assumptions are met, then the η-sublevel set can be increased slightly from the result in Lemma 1. This is further discussed and proven in the extended paper.

B. Computation of Barrier Certificates

The barrier certificate verifies the safety of the system as long as the system is initialized within the initial set R_{init}. It identifies a region-of-convergence of the system; however, as an estimate of the full ROC, R_{init} may be conservative. It is therefore advantageous to not only identify a barrier certificate of the system, but also attempt to maximize the region to include as much of R_0 as possible.

The barrier certificate computation process is broken into two steps: generate samples and compute bounds. These two steps naturally extend from the earlier region-of-attraction work. As such, the following overview focuses on the particular modifications required for the ROC. The details of the SimLFG and CWOpt algorithms discussed next can be found with full descriptions in the earlier works [10], [15].

This first step involves selectively sampling $x(0)$ and W^* in the R_{feas} space, looking for points on either side of the R_0 boundary. This algorithm uses a search parameter β_{sim} to define a ball $\in R_{\text{feas}}$ and randomly selects points in that region to run simulations. If the simulations return failures or converged trajectories, β_{sim} will shrink or expand accordingly. Once the procedure has simulated N_{conv} converged trajectories, it terminates.

After the procedure terminates, the C and D datasets of trajectories that lead to success or failure can be used to determine a candidate level set and barrier certificate. At a minimum, the failed trajectories in D provide a strict upper bound on the allowable size of the η-sublevel set, since $V(e, \widetilde{W}) \geq \beta_{UB}$ for all $(x(0), W^*) \in D$, and therefore $\eta < \beta_{UB}$.

Next, the invariant sublevel set for the Lyapunov function can be computed using the following optimization problem.

Problem 1 Given the Lyapunov function V and a positive definite function $l_2 \in \mathbb{R}$, define

$$\gamma_l^* = \underset{\gamma, s_2, s_3}{\text{maximize}} \; \gamma \;\; \text{subject to} \tag{19}$$

$$\gamma > 0 \tag{20}$$

$$s_2, s_3 \in \Sigma[e, \widetilde{W}] \tag{21}$$

$$- [(\gamma - V)s_2 + \nabla V[\dot{e}; \dot{\widetilde{W}}]s_3 + l_2] \in \Sigma[e, \widetilde{W}] \tag{22}$$

where $\Sigma[e, \widetilde{W}]$ refers to a positive semidefinite function $\mathbb{R}^{n+k} \to \mathbb{R}$. Notice that since $\dot{V} < 0$ for CL-MRAC systems, the resulting level set Ω_{V, γ_l^*} can also be directly used to form a barrier certificate as in Lemma 1.

The s_2, s_3 functions can be determined according to the procedure described in [15]. In practice, these functions are usually ϵ-scaled versions of the Lyapunov function, ϵV, for some small real ϵ or another quadratic function of the system dynamics.

The resulting γ_l^* yields a suitable lower bound for the η-sublevel set used to form a barrier certificate, so that $\gamma_l^* \leq \eta < \beta_{UB}$. Therefore, by setting $\eta = \gamma_l^*$, the process returns a theoretically verified, albeit conservative, estimate R_{init} of the full region-of-convergence R_0. A second algorithm from

[15], CWOpt, can also be used to attempt to maximize γ_l^* even further is available, but was not used here. Note, it is also possible to use the β_{UB} upper bound to generate barrier certificates using counter-example based methods from [12], [13], but requires additional toolboxes and software.

C. Discussion

In practice, this procedure will result in a conservative estimate of the region-of-convergence. Even though the estimate is conservative, it is still valuable as a certified bound. Additionally, the process generates a number of samples of $x(0)$ and W^* in R_0 that did not fall within the level set. A parallel research effort has focused on using all the converged trajectories to compute a tighter estimate of the region-of-convergence. However, this approach replaces the strong theoretical guarantees with statistical bounds. The region-of-convergence estimation procedure within this paper is intended as a complementary approach.

V. Examples

The following example demonstrates the region-of-convergence estimation process for concurrent learning MRAC system [4] described in Section II. For these types of systems, the form of the Lyapunov function $V(e, \widetilde{W})$ is already known from (8). Additional examples are included online in an extended form of the paper at `http://tinyurl.com/jk4fzfz`. Those examples also cover a more complex nonlinear wingrock model [4], but were omitted here due to space.

A. Second order linear system

Consider the second oder linear system in (23) with two unknown parameters, W_1^* and W_2^*.

$$\dot{x}_1 = x_2 \tag{23}$$
$$\dot{x}_2 = (-0.2 + W_1^*)x_1 + (-0.2 + W_2^*)x_2 + u$$

This uncertain system tracks a stable, second order linear reference model with poles $\lambda_{1,2} = -0.65 \pm 1.038j$. The reference model is excited by commanded x_{m_1} positions $z_{cmd} = \pm 1.5$. The resulting feedforward control input is given by $u_{rm} = -\omega_{rm}^2 x_{m_1} - 2\zeta \omega_{rm} x_{m_2} + \omega_{rm}^2 z_{cmd}$ where $\omega_{rm} = 1$ rad/s and $\zeta = 0.5$. As mentioned earlier, the actual uncertain system is controlled by a CL-MRAC adaptive control law. The adaptive parameters are set to $\Gamma = 2$, $\Gamma_c = 0.2$, and $\bar{p} = 20$ saved data points in the history stack.

In this example, there is no initial tracking error so that $e(0) = x_m(0) - x(0) = 0$, leaving only initial weight estimation error $\widetilde{W}(0) = -W^*$ to consider. This restricts the search space to a two-dimensional problem in W_{feas}^*, but even with this restriction, the region-of-convergence estimation problem is complex. The performance thresholds defining R_{fail} were set to $e_{1,\text{limit}} = 1.0$ and $e_{2,\text{limit}} = 20$. Essentially, the limit on $e_1(t)$ is the driving factor. The uncertain adaptive system can then be simulated using the estimation procedure.

The results of the simulation samples are shown in Figure 2. This plot demonstrates the process from Section IV. The algorithm uses the failed trajectories to calculate a

Fig. 2. Sampled region-of-convergence for the adaptive system in (23). The initial tracking error is removed, $e(0) = 0$, leaving only $\widetilde{W}(0) = -W^*$. The pairs (W_1^*, W_2^*) which lead to failure are labeled in red while the successful pairs are labeled in green.


Fig. 3. Full region-of-convergence for the adaptive system from Figure 2. The conservative estimate of the ROC in blue is overlaid on top of the full ROC R_0. The analytical estimate η_{analytic} is shown in cyan.

maximum upper bound β_{UB}. The convergence (or lack of it) of the trajectories changes the search parameter β_{sim} used to generate subsequent samples. This ends up clustering the data points near the boundary of R_0 at a ball with its radius set to the current β_{UB}. The inadvertent clustering is apparent in Figure 2 by the large density of data points clumped together near $4 \leq ||W^*|| \leq 8$. The emergent clustering also explains why the successful region down near $0 \leq W_1^* \leq 5$ and $W_2^* < -9$ is not explored further: the procedure limited its β_{sim} to search around $\beta_{UB} = 1.767$. After 3000 successful pairs were simulated, the procedure returned the size of the η-sublevel set as $\gamma_l^* = \eta = 1.414$. Note that this discrepancy between β_{UB} and η results from the conservativeness necessary in computing a certifiable bound. However, previous analytical methods [7] can only provide verified estimate $\eta_{\text{analytic}} = 1.015$.

For comparison, Figure 3 displays the actual region-of-convergence for the system. The invariant set $\Omega_{V,\eta}$ calculated by the estimation procedure is also overlaid on top of the full R_0. This figure highlights a number of important considerations. First, the region-of-convergence may be highly non-convex as the boundary of R_0 for $W^* < 0$ values is oddly shaped and jagged. Additionally, this figure also highlights the aforementioned conservativeness of ROC estimation procedures. The invariant set $\Omega_{V,\eta}$ is not able to capture values on the outskirts of R_0 due to the limitations of the process. While a tighter estimate of the ROC is desirable, the verified barrier certificate is still extremely useful as a proven safe set and larger than previous analytical methods.

VI. CONCLUSION

This paper presents a method for computing invariant sets to produce verifiable estimates of the region-of-convergence of learning-based adaptive controllers. These level sets can subsequently be used to create barrier certificates for safety. This procedure provides an automated offline tool to supplement existing analytical methods that can become increasingly unwieldy with the complexity of learning-based adaptive controllers. This work extends recent work in region-of-attraction analysis, a similar but separate problem for nonlinear systems without global attraction to the origin. The paper presents results for a CL-MRAC example problem to demonstrate the ability of the procedure to find verifiable estimates of the region-of-convergence. While these estimates may be conservative, they do provide theoretical guarantees of safety with their associated barrier certificates. Ultimately, the procedure is intended as a step in the process to address the larger problem of formally verifying adaptive systems.

REFERENCES

[1] Gang Tao. *Adaptive Control Design and Analysis*. Adaptive and Learning Systems for Signal Processing, Communications, and Control. Wiley-Interscience, Hoboken, NJ, 2003.
[2] Kumpati S. Narendra and Anuradha M. Annaswamy. *Stable Adaptive Systems*. Dover, Mineola, NY, 1989.
[3] Eugene Lavretsky and Kevin A. Wise. *Robust and Adaptive Control*. Springer, 2013.
[4] Girish V. Chowdhary. *Concurrent Learning for Convergence in Adaptive Control Without Persistency of Excitation*. PhD thesis, Georgia Institute of Technology, December 2010.
[5] Girish Chowdhary, Hassan A. Kingravi, Jonathan How, and Patricio A. Vela. Bayesian nonparametric adaptive control using gaussian processes. *IEEE Transactions on Neural Networks and Learning Systems*, 26(3):537–550, March 2015.
[6] Tansel Yucelen. *Advances in Adaptive Control Theory: Gradient- and Derivative-Free Approaches*. PhD thesis, Georgia Institute of Technology, May 2012.
[7] Girish Chowdhary, Tansel Yucelen, Maximilian Muhlegg, and Eric Johnson. Concurrent learning adaptive control of linear systems with exponentially convergent bounds. *International Journal of Adaptive Control and Signal Processing*, 27(4):280–301, April 2013.
[8] David Henriques, Joao G. Martins, Paolo Zuliani, Andre Platzer, and Edmund M. Clarke. Statistical model checking for markov decision processes. In *International Conference on Quantitative Evaluation of Systems*, 2012.
[9] P. Zuliani, A. Platzer, and E. M. Clarke. Bayesian statistical model checking with application to simulink/stateflow verification. In *Hybrid Systems: Computation and Control*, 2010.
[10] Ufuk Topcu, Andrew Packard, and Peter Seiler. Local stability analysis using ssimulation and sum-of-squares programming. *Automatica*, 44(10):2669–2675, 2008.
[11] Ufuk Topcu, Andrew K. Packard, Peter Seiler, and Gary J. Balas. Robust region-of-attraction estimation. *IEEE Transactions on Automatic Control*, 55(1):137–142, January 2010.
[12] James Kapinski, Jyotirmoy Deshmukh, Xiaoqing Jin, Hisahiro Ito, and Ken Butts. Simulation-guided approaches for verification of automotive powertrain control systems. In *American Control Conference*, 2015.
[13] James Kapinski, Jyotirmoy Deshmukh, Sriram Sankaranarayanan, and Nikos Arechiga. Simulation-guided lyapunov analysis for hybrid dynamical systems. In *Hybrid Systems: Computation and Control*, 2014.
[14] Naira Hovakimyan and Chengyu Cao. *L1 Adaptive Control Theory*. SIAM, 2010.
[15] Ufuk Topcu. *Quantitative Local Analysis of Nonlinear Systems*. PhD thesis, University of California, Berkeley, 2008.