2009 American Control Conference
Hyatt Regency Riverfront, St. Louis, MO, USA
June 10-12, 2009

Stability Overlay for Adaptive Control Laws Applied to Linear Time-Invariant Systems

Paulo Rosa, Jeff S. Shamma, Carlos Silvestre and Michael Athans

***Abstract*— Two broad classes of adaptive control algorithms can be found in the literature: i) stability based, with minimal assumptions on the plant; ii) performance based, with relatively more stringent assumptions on the plant. This paper proposes a solution, referred to as Stability Overlay (SO), to enable stability guarantees in performance based algorithms. In our methodology, the performance based adaptive control laws are only responsible for designating the controller that should be selected; the SO decides whether this controller should or not be used, based upon its most recent history of utilization. We argue that using two algorithms in parallel – the SO for stability purposes and any other suitable for the performance requirements – leads to higher levels of performance while guaranteeing stability of the adaptive closed-loop for bounded (but unknown) disturbances. The SO methodology is applicable to both time-invariant and time-varying, nonlinear and linear systems. However, due to space limitations, we only consider linear time-invariant (LTI) plants in this paper. The theory is illustrated with an example.**

I. INTRODUCTION

Adaptive control laws are needed in many practical applications, where a single (non-adaptive) controller is not able to achieve the stability and/or performance requirements. However, many adaptive control laws can lead to unstable closed-loop systems when connected to a plant with even the slightest discrepancies from the family of admissible plant models. This issue was first described in Ref. [1] in the so-called Rohrs *et al.* counterexample. Very small disturbances can be responsible for destabilizing the closed-loop because of the unavoidable unmodeled high frequency dynamics, present in every physical system.

This paper presents a solution to the stability problem common to many closed-loop linear time-invariant (LTI) systems with performance based adaptive control laws. The strategy developed herein, referred to as Stability Overlay (SO), takes into account both stability objectives – often robust to a very wide class of disturbances and model uncertainty – and performance requirements – that, in general, assume a stronger knowledge about the plant to be controlled. The algorithm presented in the sequel is based upon [2], and assesses the "rewards" received by each controller after its most recent utilization, without any prior information on the bounds of the exogenous disturbances and sensors noise. A controller is then disqualified or not, based upon its rewards, in a similar way to what is done in Refs. [3], [4], [5] and in the references therein. However, in our approach, we suggest that the SO should only be responsible for the stability of the plant, and thus another algorithm should run in parallel in order to accomplish the performance requirements. Therefore, our methodology differs from Refs. [3], [4] in the sense that the controllers rewards are not used to decide which controller leads to the highest closed-loop performance, but rather to guarantee that a controller which is not able to stabilize the plant is not persistently selected. The Lyapunov-based solution presented in [5] relies on the model of the plant and hence requires stronger assumptions than the ones presented in the sequel.

For the proposed SO methodology, it is not required to know the plant model to be controlled nor the disturbance properties. Still, it is clear that the performance of the closed-loop can be severely affected if no knowledge is available about the plant. Nonetheless, the model-free characteristic of the present method ensures robustness to several types of model uncertainty. In a sense, if the actual plant is *close* to a plant model in the family used to design the adaptive control law, then the adaptation runs as usual without (or with minor) intervention of the SO. If, however, the actual plant or disturbance properties do not match the ones used during the design, the closed-loop system may become unstable. Therefore, instead of blindly continuing to use the adaptation law, we assess the norm of the inputs and outputs of the system, and eventually switch to a controller that is able to stabilize the plant, as long as such controller belongs to the set of *legal* controllers that the SO is allowed to use.

Therefore, the SO can be seen as a *safety device* that can be used with many adaptive algorithms, achieving high levels of performance while providing robust stability guarantees for several different types of modeling errors.

The applicability of the SO is illustrated by the so-called Rohrs *et al.* counterexample [1]. The main result presented in the sequel is able to provide stability guarantees for the integration of this adaptive law with the SO, with a few changes in the original algorithm.

This paper is organized as follows. We start by introducing the notation and formally posing the problem in Section II.

This work was partially supported by Fundação para a Ciência e a Tecnologia (ISR/IST pluriannual funding) through the POS_Conhecimento Program that includes FEDER funds, by the PTDC/EEA-ACR/72853/2006 OBSERVFLY Project, and by the NSF project #ECS-0501394. The work of P. Rosa was supported by a PhD Student Scholarship, SFRH/BD/30470/2006, from the FCT.

P. Rosa, C. Silvestre and M. Athans are with Institute for Systems and Robotics - Instituto Superior Tecnico, Av. Rovisco Pais, 1, 1049-001 Lisboa, Portugal prosa@isr.ist.utl.pt, cjs@isr.ist.utl.pt, athans@isr.ist.utl.pt

M. Athans is also Professor of EECS (emeritus), M.I.T., USA

J. S. Shamma is with Georgia Institute of Technology, School of Electrical and Computer Engineering, Atlanta, Georgia, United States of America shamma@gatech.edu

In Section III, some properties of LTI systems are derived. The main result for LTI plants is presented in Section IV. In Section V, simulation results of the Rohrs *et al.* counterexample integrated with the SO are shown. Finally, in Section VI some conclusions about the SO are presented.

II. Preliminaries and Notation

We define $|x|$ as the euclidian norm of $x \in \mathbb{R}^n$, and $\|A\|$ as the induced norm of the matrix A, i.e.,

$$\|A\| = \sup_{x\neq 0} |Ax|/|x|.$$

We further define, for any $\sigma > 0$,

$$\left\| z|^{\sigma}_{[t_1,t_2]} \right\| = \sup_{\tau\in[t_1,\,t_2]} e^{-\sigma(t_2-\tau)}|z(\tau)|,$$

and

$$\left\| z|_{[t_1,t_2]} \right\| = \sup_{\tau\in[t_1,\,t_2]} |z(\tau)|.$$

Throughout this paper, we consider an LTI plant described by

$$\dot{x} = Ax + Bu + F\xi, \quad x(0) = x_0 \tag{1a}$$
$$y = Cx + G\theta \tag{1b}$$
$$z = \begin{pmatrix} y \\ u \end{pmatrix} = \begin{pmatrix} Cx + G\theta \\ u \end{pmatrix} \tag{1c}$$
$$u = K_{\alpha(t)} y = \begin{pmatrix} K_{\alpha(t)} & 0 \end{pmatrix} z \tag{1d}$$
$$K_{\alpha(t)} \in S_0 := \{K_1, K_2, \ldots, K_{N_c}\}. \tag{1e}$$

The output variables $z(.)$ and $y(.)$ can include performance outputs such as the ones obtained by filtering the plant output and the control input with the weights $W_y(s)$ and $W_u(s)$, respectively – see [6] for more details on using performance weights. Furthermore, $x_0 \in \mathbb{R}^n$ is a fixed (but unknown) initial condition, $\xi(.) \in \mathcal{L}_\infty$ is a bounded (but unknown) exogenous disturbance, $\theta(.) \in \mathcal{L}_\infty$ is the bounded (but unknown) measurement noise and $u(.)$ is the control input. S_0 is the set of *eligible* controllers which are considered to be, without loss of generality, constant matrix gains. N_c is the number of legal control laws (and thus the number of elements in S_0), and K_i, for $i \in \{1, 2, ..., N_c\}$, represents a controller. We argue that the control laws of any adaptive or nonadaptive system should be robust to model uncertainty.

Define a *finitely switching* control input as

$$u_{\mathrm{fs}}(t) = \begin{cases} K_{\alpha(t)}(y(t)), & 0 \le t < t_0; \\ K_i(y(t)), & t \ge t_0. \end{cases} \tag{2}$$

Figure 1 depicts the output feedback interconnection between the plant and the controllers K_i, selected through signal $\alpha(t)$.

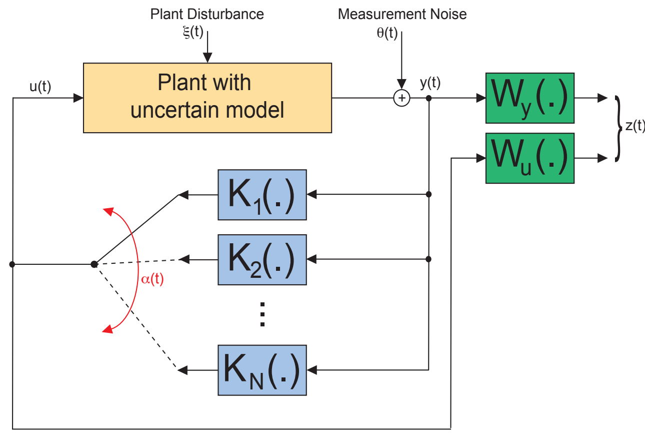

Fig. 1. Feedback interconnection between the plant (1) and the controllers K_i, selected through signal $\alpha(t)$.

III. Properties of LTI Closed-Loop Systems

In this section, we derive some properties of LTI closed-loop systems that are going to be useful in the proof of the main result of this paper. Consider the LTI plant described by (1). Without loss of generality, we assume that the N_c controllers in S_0 are static output feedback controllers, i.e., each K_i, for $i \in \{1, \ldots, N_c\}$, is a constant matrix. Notice that, as shown in [2] if the controllers are dynamic and described by

$$\dot{x}_c = A_c x_c + B_c y$$
$$u = C_c x_c + D_c y,$$

then, for each controller, one can rewrite the closed-loop system as

$$\begin{pmatrix} \dot{x} \\ \dot{x}_c \end{pmatrix} = \begin{pmatrix} A & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} x \\ x_c \end{pmatrix} + \begin{pmatrix} B & 0 \\ 0 & I \end{pmatrix} u_{\mathrm{aug}} + \begin{pmatrix} F \\ 0 \end{pmatrix} \xi$$
$$z_{\mathrm{aug}} = \begin{pmatrix} C & 0 \\ 0 & I \end{pmatrix} \begin{pmatrix} x \\ x_c \end{pmatrix} + \begin{pmatrix} G \\ 0 \end{pmatrix} \theta$$
$$u_{\mathrm{aug}} = \begin{pmatrix} D_c & C_c \\ B_c & A_c \end{pmatrix} z_{\mathrm{aug}}.$$

In this form, switching between dynamic controllers means switching the "static output feedback" matrix $\begin{pmatrix} D_c & C_c \\ B_c & A_c \end{pmatrix}$ which only depends on controller matrices. The only restriction is that all dynamic controllers in this setup must be of the same order, which in practice does not represent a shortcoming, since we can solve that issue by having some controllers that use only a subset of the available states, while forcing the remaining ones to go to zero.

We follow closely the steps in [2] to show that, under natural assumptions, the linear system (1) has the following properties:

Property 1: For any finitely switching input (2) and any t_0, $\Delta T > 0$,

$$\left\| z|^{\sigma}_{[0,t_0]} \right\| < \infty \Rightarrow \left\| z|^{\sigma}_{[0,t_0+\Delta T]} \right\| < \infty,$$

i.e., the closed-loop does not have a finite escape time.

***Property** 2:* There exist a control law, K_{i^*}, and positive constants σ and l^*, such that for any $0 < \gamma < 1$, there exists a $\Delta T^* \geq 0$ that satisfies the following condition. For any finitely switching control input (2) with $K_i = K_i^*$,

$$\left\| z|^{\sigma}_{[0,t_0+\Delta T]} \right\| \leq \gamma \left\| z|^{\sigma}_{[0,t_0]} \right\| + l^*,$$

for all $\Delta T \geq \Delta T^*$.

Property 1 ensures there are no controllers in the *legal* set (referred to as *eligible* controllers) that take the output of the plant to infinity in finite time. Finally, Property 2 states that there is at least one eligible controller that satisfies a desired stabilization condition. Parameter l^* accounts for the exogenous disturbances and the initial condition.

***Assumption** 1:* The switched linear system (1) satisfies:

- $\Re\{\lambda_j(A + BK_iC)\} < 0, \forall_j$ for some $i \in \{1, \ldots, N_c\}$.
- The pair $[A, C]$ is observable.
- The exogenous disturbances and the measurement noise are bounded by some (possibly unknown) constants ξ_0 and θ_0, respectively, i.e., $|\xi(.)| \leq \xi_0$ and $|\theta(.)| \leq \theta_0$.

***Proposition** 1:* Under Assumption 1, the linear system (1) has the Properties 1–2.

The proof of Proposition 1 can be found on the extended version of this paper, available on the Internet.

Remark 1 The discounted norm $\left\| z|^{\sigma}_{[0,\, t_0+\Delta T]} \right\|$ can be interpreted as a state-norm estimator (cf. [7], [8]). Although in a different perspective, these norm estimators are first order systems where the inputs are the discounted norms of the inputs and outputs of the plant. Therefore, our decisions of *disqualifying* or not a controller, according to this interpretation, will be based upon the estimate of the norm of the actual state of the closed-loop system. $\diamond$

IV. MAIN RESULT

The reward, $r(n)$, after using controller $K(n)$ during the time interval $t_{n-1} \leq t < t_n$ is defined as

$$r(n) = \begin{cases} 1, & \left\| z|^{\sigma}_{[0,\, t_n]} \right\| \leq \gamma \left\| z|^{\sigma}_{[0,t_{n-1}]} \right\| + l(n) \\ 0, & otherwise, \end{cases} \qquad (3)$$

where γ is a fixed scalar with $0 < \gamma < 1$ and $l(n)$ is going to be specified next.

Remark 2 We recall that $z(t)$ may be an augmented and/or filtered output. For instance, if the nonadaptive controllers were designed using some kind of dynamic weights, then those same weights can be used to filter the output of the plant and hence create a new vector to compute the rewards of the controllers. $\diamond$

Figure 2 describes the Stability Overlay (SO) algorithm. The notation $S = S\backslash K(n)$ means "the exclusion of element $K(n)$ from set S". The initial set of eligible control laws is denoted S_0, while K_0 is the first control law selected, $\Delta T(n)$ is the period control law $K(n)$ is used, l_0 is the initial value of $l(n)$ in (3), and l_{inc} and ΔT_{inc} are the increments for $l(n)$ and $\Delta T(n)$, respectively, whenever all the control laws have *failed* in their most recent utilization.

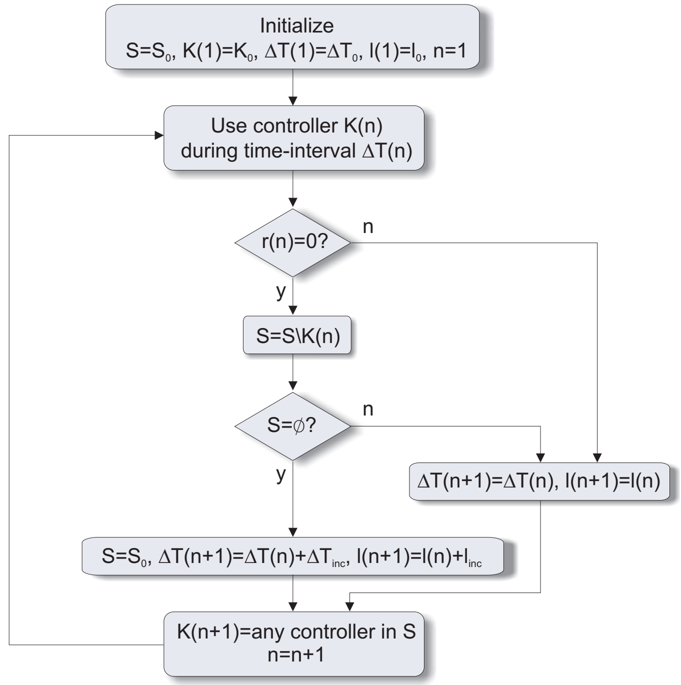

Fig. 2. Stability Overlay (SO) algorithm

In reference to the SO algorithm and Section III, we can summarize the calculations suggested for LTI plants:

- choose arbitrary $\sigma > 0$;
- choose arbitrary positive $\gamma < 1$;
- use proof of Proposition 1 to derive values for ΔT_0, l_0.

Remark 3 It is important to stress that smaller values of ΔT_0 and l_0 can actually lead to better transients, since we only derived (very conservative) upper bounds for those parameters. Nevertheless, these calculations should point out *reasonable* magnitudes for the parameters whenever both the plant and the controllers are linear and time-invariant. Moreover, the stability guarantees are not affected by the selection of the initial values for these parameters. $\diamond$

***Theorem** 1:* If Assumption 1 is satisfied, the Stability Overlay Algorithm results in $\left\| z|^{\sigma}_{[0,t]} \right\|$ bounded.

Proof of Theorem 1

We first recall Claim 1 of Ref. [2]:

***Claim** 1:* The parameters $\Delta T(n)$ and $l(n)$ are uniformly bounded, i.e.,

$$\begin{aligned} \Delta T_{\text{max}} &= \lim_{n\to\infty} \Delta T(n) < \infty \\ l_{\text{max}} &= \lim_{n\to\infty} l(n) < \infty \end{aligned} \qquad (4)$$

Proof: The parameters $\Delta T(n)$ and $l(n)$ are increased whenever every control law, K_i, in its most recent utilization resulted in a zero reward, i.e.,

$$\left\| z|^{\sigma}_{[0,\, t_n]} \right\| > \gamma \left\| z|^{\sigma}_{[0,t_{n-1}]} \right\| + l(n). \qquad (5)$$

However, by Property 2, there exists a $\Delta T^* \geq 0$, a positive constant l^*, and at least one control law which satisfies the condition

$$\left\| z|^{\sigma}_{[0,t_0+\Delta T]} \right\| \leq \gamma \left\| z|^{\sigma}_{[0,t_0]} \right\| + l^*, \tag{6}$$

provided that $\Delta T(n) \geq \Delta T^*$. This implies that the condition (5) cannot be satisfied infinitely often with $\Delta T(n)$ and $l(n)$ increasing without bound. ■

According to Claim 1, there is at least one control law that is going to be used infinitely many times. For this control law, $r = 1$. Thus, all other control laws are going to be used at most a finite number of times. According to Proposition 1, the output is going to remain bounded during that (bounded) time interval where $r = 0$. For some t_0, the rewards obtained for $t \geq t_0$ are positive. Since the output is bounded at $t = t_0$, it will remain bounded for $t > t_0$. This means that the closed-loop system is stable, which concludes the proof. □

It is important to stress that we do not describe how to choose the controller to be put in the loop. In fact, we allow any controller in set S to be selected. The choice of the controller is responsible for the performance of the closed-loop and should be taken care of by some adaptive control law that (probably) takes into account the model of the plant and the disturbances properties. One example of the applicability of the SO with a model reference adaptive control architecture is presented in the following section.

However, the applicability of the SO is much wider. Many types of adaptive laws are eligible to be integrated with the SO, such as the schemes based on the identification of the plant parameters (see, for instance, [9], [10] and references therein), or the estimator-based methodologies in [11]. Preliminary results of the integration of the SO with the Robust Multiple-Model Adaptive Control (RMMAC) methodology, introduced in Refs. [12], [13], [14], [15] and references therein, are also available but are not presented in this paper.

Remark 4 It should be noticed that the choice of the parameters for the algorithm may be very sensitive in some cases, depending upon the plant dynamics and the disturbances intensity. In fact, if the norm of the output of the closed-loop system grows *very fast* whenever a destabilizing controller is picked, and if the time required to disqualify a controller is very large, one may not get "practical stability". This means that, although a stabilizing controller is eventually selected, the transients may not be reasonable from a practical point of view. ◇

Remark 5 The reason for increasing ΔT can be explained in a very intuitive manner, that relates it to the classical performance/robustness tradeoffs. If ΔT is large enough, stabilizing controllers are not going to be disqualified. However, we may also be using destabilizing controllers for a long time, since we only switch to another one after at least a time interval ΔT. In case ΔT is very small, we may find the *right* controller faster, but we may also disqualify stabilizing controllers just because they were not used long enough. Therefore, large values of ΔT guarantee stability at the cost of large transients. ◇

Remark 6 We stress that the observability (or detectability) assumption does not have to be necessarily satisfied. However, if indeed there are unobservable states, then the fact of $\left\| z|^{\sigma}_{[0,t]} \right\|$ being bounded does not imply that the closed-loop system is stable, but only that it is input/output stable. ◇

V. ROHRS ET AL. COUNTEREXAMPLE

To illustrate the usefulness of the SO, we use the so-called Rohrs *et al.* counterexample – see Ref. [1]. We use the reference model adaptation law referred to as Continuous-Time Algorithm 1 and the same terminology as in Ref. [1]. Figure 3 depicts the architecture of this methodology.

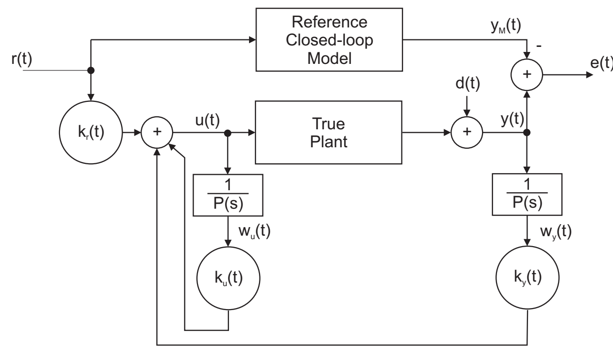

Fig. 3. Continuous-Time Algorithm 1

The reference input is denoted by $r(t)$, the control input by $u(t)$, the disturbances by $d(t)$, and the output by $y(t)$. Signal $y_M(t)$ is the output of the reference model and $P(s)$ represents a polynomial. The adaptive gains are denoted by $k_r(t)$, $k_u(t)$ and $k_y(t)$. The adaptation law evolves according to

$$\dot{k}(t) = -\Gamma w(t) e(t),$$

where $\Gamma = \Gamma^T > 0$ and

$$e(t) = y(t) - y_M(t)$$

We use the example in Ref. [1] where the plant model is given by

$$Y(s) = \left(\frac{2}{s+1}\right) \frac{229}{s^2 + 30s + 229} U(s), \tag{7}$$

the reference model is described by

$$Y_M(s) = \frac{3}{s+3} R(s), \tag{8}$$

and

$$k_y(0) = -0.65 \quad \text{and} \quad k_r(0) = 1.14. \tag{9}$$

Notice that, according to the design procedure,

$$P(s) = 1 \quad \text{and} \quad k_u(t) = 0 \,\forall_t.$$

Let the reference be given by

$$r(t) = 2$$

and the output additive disturbance by

$$d(t) = \frac{1}{2}\sin(8t)\,.$$

The first step in order to apply the SO is to discretize and bound the gains k_y and k_r, so that we have a finite set of controllers. Each pair (k_y, k_r) defines a controller. Therefore, if k_y is divided into n_y bins and k_r is divided into n_r bins, the set S_0 will have $n_r n_y$ controllers.

Another design decision has to be made regarding what to do when the adaptive control law choses a controller that was previously disqualified, because it received a zero reward. A simple approach is to use the controller *closest* to that one, in a geometric sense, although it may not be very effective in terms of performance. Another strategy is to decrease the gain if it was increasing before a disqualified controller was obtained, and vice-versa. *We stress that a specific strategy need not be used to guarantee stability. The boundedness of the output $y(t)$ is guaranteed a priori by the use of the SO, no matter how we schedule the controllers in the SO algorithm, so the only aspect the control engineer has to account for is the performance.*

A. Simulation of the Rohrs et al. Counterexample without the Stability Overlay

In this subsection, we replicate the results in Ref. [1], just for comparison purposes. Figures 4 and 5 illustrate that the closed-loop system becomes unstable as time goes by, due to the unmodeled dynamics, excited by the disturbances, that were not accounted for during the design of the adaptive control law. The infinite gain operators [1] from $e(t)$ to $u(t)$ and from $e(t)$ to $k(t)$, inherently present in the Continuous-Time Algorithm (CA) 1, are responsible for the instability of the closed-loop.

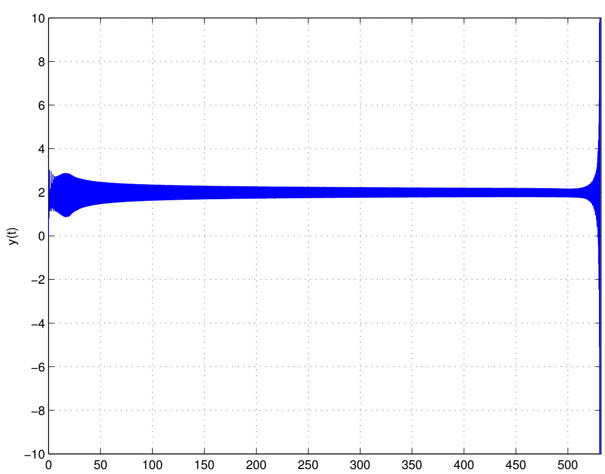

Fig. 4. Output $y(t)$ of the closed-loop system using the continuous-time algorithm 1, without the stability overlay.

B. Simulation of the Rohrs et al. Counterexample with the Stability Overlay

We now analyze the continuous-time algorithm supervised by the SO. The continuous gains k_r and k_y are discretized

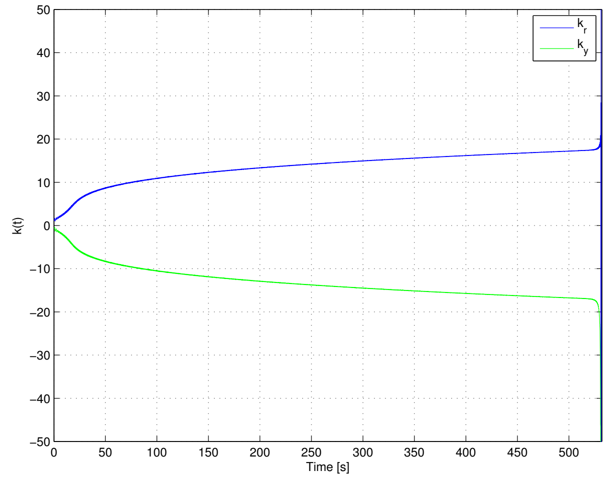

Fig. 5. Adaptive gains time evolution for continuous-time algorithm 1, without the stability overlay.

and bounded, in order to have a finite set of eligible controllers. For this simulation, we use bins of width 2 and the limits are ± 50. Figure 6 shows that the closed-loop is now stable, although some transients are also experienced. The time instants when a controller is disqualified are also represented. If a disqualified controller is selected by the adaptive algorithm, the values of the adaptive gains k_r and k_y are updated to the *legal* ones closest to those obtained by the adaptive law.

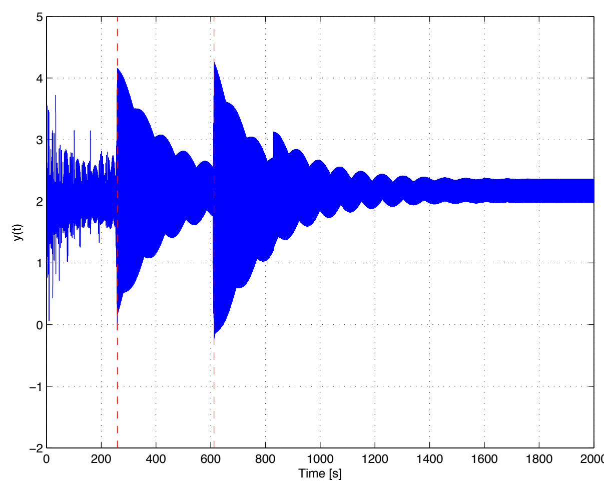

Fig. 6. Output $y(t)$ of the closed-loop system using the continuous-time algorithm 1, with the stability overlay. The red dashed lines indicate the time instants when the currently scheduled controller fails.

Figure 7 depicts the time-evolution of the adaptive gains. It should be noticed that, as soon as the instability is detected, the SO disqualifies the currently used controller and switches to another one. After only two switches, a controller is selected that stabilizes the plant.

We stress that the adaptive gains in Fig. 7 are the ones effectively used in the feedback loop, and that they can differ from those obtained with the adaptive law. The later are only reset by the SO after a failure on the adaptive control law.

Remark 7 Figure 7 indicates that one could simply saturate the adaptive gains and still get the stability result. However, if we do that, and use very large limits for the saturation, we can get an unstable system. On the other hand, if the limits are very small, we can get poor performance. The SO obtains those limits in a natural way, as shown in the simulations. $\diamond$

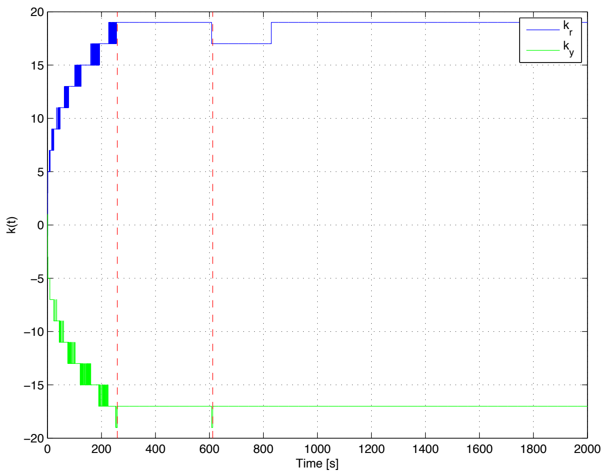

Fig. 7. Adaptive gains time evolution for continuous-time algorithm 1, with the stability overlay. The red dashed lines indicate the time instants when the currently scheduled controller fails.

VI. CONCLUSIONS

This paper proposed a solution, referred to as Stability Overlay (SO), for the stability problem that appears in many adaptive control laws for linear time-invariant (LTI) systems. We take advantage of the on-line evaluation of the selected controller to disqualify those that do not comply with the stability requirements. Unlike other adaptive control strategies, we take into account both stability objectives – often robust to a very wide class of disturbances and model uncertainty – and performance requirements – that, in general, assume a stronger knowledge about the plant to be controlled.

This approach only requires that at least one of the *eligible* controllers is able to stabilize the plant, and can be extended to nonlinear systems.

As a caveat, the choice of the parameters for the SO algorithm may be very sensitive if the norm of the output of the closed-loop system grows *very fast* whenever a destabilizing controller is picked, and if the time required to disqualify a controller is very large. In those cases, although a stabilizing controller is eventually selected, the transients may not be reasonable from a practical point of view.

The SO is responsible for disqualifying controllers that are not able to stabilize the plant, while some other adaptive control law, that (possibly) takes into account the plant model and the disturbances properties, is responsible for selecting the *best* controller, i.e., the controller in set S that leads to the highest performance.

In reference to the Rohrs *et al.* counter examples, the stability guarantees were obtained by first discretizing and bounding the adaptive gains, and then by using the same principles as in the previous example. Simulation results were presented, illustrating the benefits of using the SO with the Continuous-time Algorithm 1 of Ref. [1].

The proposed method can be applied to a much wider class of adaptive controllers, with little effort, guaranteeing stability properties otherwise not available. The integration with adaptive architectures, such as the Robust Multiple-Model Adaptive Control (RMMAC), introduced in Refs. [12], [13], [14], [15] and references therein, is currently a topic of research.

VII. ACKNOWLEDGMENTS

We wish to thank our colleagues Antonio Pascoal, Pedro Aguiar, Vahid Hassani and José Vasconcelos for the many discussions on the field of robust adaptive control.

REFERENCES

[1] C. E. Rohrs, L. Valavani, M. Athans, and G. Stein, "Robustness of Continuous-Time Adaptive Control Algorithms in the Presence of Unmodeled Dynamics," *IEEE Transactions on Automatic Control*, vol. AC-30, no. 9, pp. 881–889, September 1985.

[2] I. Al-Shyoukh and J. Shamma, "Switching supervisory control using calibrated forecasts," in *46th IEEE Conference on Decision and Control*, December 2007.

[3] M. Safonov and T.-C. Tsao, "The unfalsified control concept and learning," *IEEE Transactions on Automatic Control*, vol. 42, no. 6, pp. 843–847, June 1997.

[4] D. Wang, A. Paul, M. Stefanovic, and M. Safonov, "Cost-detectability and Stability of Adaptive Control Systems," in *Proceeding of the 44th IEEE International Conference on Decision and Control, and the European Control Conference 2005*, December 2005.

[5] D. Angeli and E. Mosca, "Lyapunov-based switching supervisory control of nonlinear uncertain systems," *IEEE Transactions on Automatic Control*, vol. 47, pp. 500–505, 2002.

[6] K. Zhou and J. Doyle, *Essentials of Robust Control*. Prentice Hall, September 1997.

[7] J. P. Hespanha, D. Liberzon, D. Angeli, and E. D. Sontag, "Nonlinear Norm-Observability Notions and Stability of Switched Systems," *IEEE Transactions on Automatic Control*, vol. 50, no. 2, pp. 154–168, February 2005.

[8] R. A. Garcia and J. L. Mancilla-Aguilar, "State-norm estimator of switched nonlinear systems," in *2002 American Control Conference*, 2002.

[9] P. Ioannou and J. Sun, *Robust Adaptive Control*. Prentice Hall, Inc, 1996.

[10] M. Kuipers and P. Ioannou, "Practical robust adaptive control: Benchmark example," in *Proceedings of the 2008 American Control Conference*, 2008.

[11] J. P. Hespanha, D. Liberzon, A. S. Morse, B. Anderson, T. Brinsmead, and F. de Bruyne, "Multiple model adaptive control, part 2: switching," *Int. J. of Robust and Nonlinear Control,* Special Issue on Hybrid Systems in Control, vol. 11, no. 5, pp. 479–496, April 2001.

[12] M. Athans, S. Fekri, and A. Pascoal, "Issues on robust adaptive feedback control," in *Preprints of the 16th IFAC World Congress, pp. 9-39*, 2005.

[13] S. Fekri, M. Athans, and A. Pascoal, "Issues, progress and new results in robust adaptive control," *Int. Journal of Adaptive Control and Signal Processing*, vol. 20, pp. 519–579, 2006.

[14] ——, "Robust multiple model adaptive control (RMMAC): A case study," *Int. Journal of Adaptive Control and Signal Processing*, vol. 21(1), pp. 1–30, 2006.

[15] S. Fekri, "Robust adaptive mimo control using multiple-model hypothesis testing and mixed-μ synthesis," Ph.D. dissertation, Instituto Superior Tecnico, Lisbon, Portugal, December 2005.