Unfalsified Control for Slowly Varying Plants using Fading Memory and Windowing

Shin-Young Cheong and Michael G. Safonov
Department of Electrical Engineering - Systems
University of Southern California, Los Angeles, CA 90089-2563 USA
E-mail: sycheong@usc.edu, msafonov@usc.edu

A new cost function with fading memory and a finite-duration time-window is introduced in order to limit the effect of old data in unfalsified adaptive control applications where the plant varies slowly or infrequently with time. The effectiveness of the approach is demonstrated via a simple simulation in which a plant with a gain, which switches periodically, is stabilized by a time-windowed unfalsified adaptive control law that switches between two candidate PID controllers, neither of which alone would be able to stabilize the time-varying plant without switching. The result demonstrates that time-windowed/fading-memory unfalsified methods can be effective for adaptive control of varying plants, even when the plant fails to satisfy the usual 'feasibility' requirement of unfalsified control that it must be stabilizable by one of the candidate controllers.

I. Introduction

WHEN no one fixed controller is adequate to control an unknown or highly uncertain plant, adaptive control methods can theoretically be used to iteratively identify a suitable controller from a given pool of candidate controllers. Early adaptive control methods were unreliable because they required excessive assumptions such as minimum phase plant, known upper bound of plant order, or no measurement noise.[7] Because of this, most engineers had been reluctant until recently to use adaptive control methods for safety critical applications. To overcome the limitations of the earlier adaptive methods, various new paradigms for adaptive control have been proposed.[1-2] Unfalsified control, one of the newest paradigms, introduced by Safonov and Tsao[1] is attractive because it provides a unified framework which explains the behavior of adaptive controllers in terms of the minimization of a certain data-driven cost function whose value for each controller can be computed at each time from measured plant data. The basic idea is that the adaptive control supervisor unit chooses a controller from the pool that either minimizes the cost function or at least maintains the cost at or below a prescribed cost level. The unfalsified control paradigm has facilitated the discovery of new classes adaptive control laws that reliably stabilize unknown plants under only the very weakest of assumptions; viz., that there exists at least one fixed but *a priori* unknown controller in the candidate controller pool that can stabilize the plant.[3]

At the heart of every adaptive control system is a unit called the *supervisor* that selects the currently active controller from a pool of candidate controllers. In the unfalsified control paradigm, the supervisor is modeled as a device that evaluates and compares the unfalsified performance levels of candidate controllers using a data-driven cost function. Then, the supervisor tags candidates that achieve a prescribed unfalsified cost level as unfalsified controllers. Other controllers that fail the test at a given unfalsified cost level are not used unless the unfalsified cost level of the currently active controller increases above another controller by at least some small amount ε called the hysteresis constant. The basic elements of unfalsified control theory are described in Reference 1, and the use of the theory for reliable adaptive PID controller gain tuning was described in Reference 4. The relationship with Morse-Mayne-Goodwin convergence lemma[2] and importance of using *cost-detectable* cost-functions whose behavior accurately reflects stability was addressed in Reference 3. Moreover, Reference 3 provides the proof of stability.

In the unfalsified control paradigm, the choice of cost function plays the key role in determining the performance and behavior of the adaptive control system. Previous unfalsified control works have developed several cost functions for achieving performance goals and maximal stability robustness using the concept of cost detectability.[3] Nevertheless, these did not handle the possibility of the falsification of all controllers, which can happen when no one controller in the candidate pool can robustly stabilize the time-varying plant — even when the plant varies only slowly or infrequently over a set of plants each of which is stabilizable by one on the controllers in the pool. In this paper, we address this problem by introducing fading memory and time-windowing modifications into the unfalsified control cost function. Angeli and Mosca (Reference 5) suggest a similar approach for handling time-varying plants using fading memory.

This paper is organized as follows. Basic concepts of the unfalsified control are reviewed in Section 2. Section 3 formally introduces the procedure for adding time-windowed fading memory to unfalsified control cost functions. Section 4 contains an example that demonstrates the advantages of the approach via a simple simulation that shows that the time-windowed fading-memory modification to the cost-function allows an unfalsified adaptive system to work without the usual feasibility assumption of unfalsified control. Conclusions are in Section 5.

II. Review of Unfalsified Control

Elementary concepts and terms of the unfalsified control are briefly provided in Section A below and further detail may be found in Reference 1. Section B explains the role and use of fictitious signals in adaptive PID control.

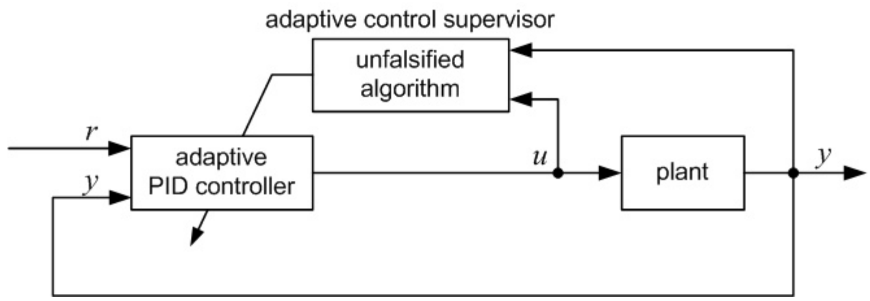

Figure 1. Configuration of adaptive PID controller

A. Theory

The unfalsified control theory begins from the definition of controller falsification based on the measured data and a family of target sets of acceptable triples of signals (r, y, u)

$$\boldsymbol{T}_{spec}(\gamma,\tau) = \{(r, y, u) \mid T_{spec}(r, y, u, \tau) \ \le \gamma\}$$

where $T_{spec}(r, y, u, \tau)$ a scalar-valued *cost function*, and γ and τ are real numbers that represent cost and time respectively. Details of the specification will be addressed in Section IV. The definition of unfalsification and the main theorem are as follows.

Definition 1.[1] A controller $K \in \mathbf{K}$ is said to be *falsified* at *cost level* γ at time τ by plant data (y, u) if this data is sufficient to deduce that the performance specification $(r, y, u) \in \boldsymbol{T}_{spec}(\gamma,\tau)\ \forall r$ would be violated if that controller were in the feedback loop. Otherwise, the controller K is said to be *unfalsified*. The least value of γ for which a controller *K* is unfalsified is called the *unfalsified cost-level* of *K* at time τ. ♦

Theorem 1.[1] A controller $K \in \mathbf{K}$ is *falsified* at *cost level* γ at time τ by plant data (y_0, u_0) if, and only if, for each triple $(r_0, y_0, u_0) \in \boldsymbol{P}_{data}(\tau) \cap \boldsymbol{K}$, there exists at least one pair (y_1, u_1) such that

$$(r_0, y_1, u_1) \in \boldsymbol{P}_{data}(\tau) \cap \boldsymbol{K} \cap \boldsymbol{T}_{spec}(\gamma,\tau)$$

where $\boldsymbol{P}_{data}(\tau) = \{(r, y, u) \mid y(t) = y_0(t) \text{ and } u(t) = u_0(t) \text{ for all } t \le \tau\}$ and $\boldsymbol{K} = \{(r, y, u) \mid u = K\begin{bmatrix} r \\ y \end{bmatrix}\}$. ♦

In the unfalsified control paradigm the adaptive control supervisor selects the active online controller from among the currently unfalsified controllers in the candidate pool set **K**, switching to a new controller when new data falsifies the current online *controller.* It is known[3] that if a few very mild assumptions hold and if additionally the $T_{spec}(r, y, u, \tau)$ is monotone in time t and satisfy certain other 'cost detectability' conditions, then the following feasibility assumption is sufficient to guarantee that the adaptive system is robustly stable.

Assumption 1. (Feasibility[3]) The candidate controller set **K** contains at least one robustly stabilizing and performing controller. ♦

In Section III, the unfalsified algorithm is described in further detail, along with a new modification for handling slowly time-varying plants that may violate this feasibility assumption.

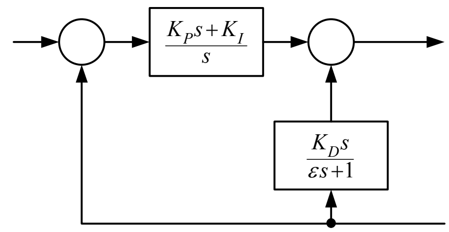

Figure 2. PID controller

B. Fictitious Reference signals in PID controller

In this paper, an adaptive PID controller is applied as shown in Figure 1. In particular, the PID controller has a structure of Figure 2 and this is the same structure used in Reference 4.

A set of candidate controllers is **K** and the i^{th} controller in **K** is $K_i = \{k_{P_i}, k_{I_i}, k_{D_i}\}$. Fictitious reference signals and error signals associated with the i^{th} candidate controller can be computed by the equations

$$\tilde{r}(K_i) := \tilde{r}_i = y + \frac{s}{k_{P_i}s + k_{I_i}}\left(u + \frac{k_{D_i}s}{\varepsilon s + 1}y\right) \quad \text{and} \quad \tilde{e}(K_i) := \tilde{e}_i = \tilde{r}_i - y \tag{1}$$

which are obtained by inverting Figure 2 to solve for the r and e in terms of (u, y). We assume that for all controllers in the candidate pool the three $(k_{P_i}, k_{I_i}, k_{D_i})$ PID gains are either all positive or all negative, which ensures that the fictitious reference signal generator systems defined by (1) are well-posed and stable.[1]

In the implementation of unfalsified adaptive control, the adaptive control supervisor unit uses the fictitious reference signal associated with each candidate controller to iteratively update that controller's current unfalsified cost level. This is possible because, as shown in Reference 1, if the system defined by (1) is well-posed and stable then the i^{th} controller K_i is unfalsified at *cost level* γ at time τ by plant data (y,u) if, and only if,

$$J(K_i, u, y, \tau) := T_{spec}(\tilde{r}(K_i), u, y, \tau) \le \gamma\ . \tag{2}$$

III. Problem Formulation

Selection of a suitable cost function is a significant task in building an unfalsified adaptive control law. In this paper, the cost-detectable monotone function $T_{spec}(\tilde{r}(K_i), u, y, t)$ normally used to evaluate unfalsified cost level for the candidate controller K_i at each time t is replaced by the following *non-monotone* modified cost function

$$J(K_i, u, y, t) = \begin{cases} \max\limits_{l \in (t-\tau_0, t)} \left\{ e^{-\lambda(t-l)} T_{spec}(\tilde{r}_i, y, u, l) \right\}, \ t \ge \tau_0 \\ \max\limits_{l \in (0,t)} \left\{ e^{-\lambda(t-l)} T_{spec}(\tilde{r}_i, y, u, l) \right\}, \ t < \tau_0 \end{cases} \tag{3}$$

where $\tilde{r}_i = \tilde{r}(K_i)$. As shown in Reference 1, the i^{th} controller K_i is unfalsified at cost level γ at time t by plant data (y,u) if, and only if,

$$J(K_i, u, y, t) \le \gamma\ . \tag{4}$$

The modified cost (3) has two differences as compared to the original unmodified cost (2).

1) **Time-Windowing**. The max operator term ignores old data outside the recent-time window $(t - \tau_0, t)$.

2) **Fading-Memory**. The term $e^{-\lambda(t-l)}$ exponentially reduces the effect of the older data.

We note that a similar exponential factor arises in Reference 6 in conjunction with scale-independent hysteresis switching algorithms.

These differences make it possible for older data to be de-weighted and very old data to be ignored. This can be useful when the plant changes slowly or infrequently. In practice, slow changes of the environment (temperature, air pressure, etc.) or sudden but infrequent component failures are possible causes of such changes in the plant. In case of a cost function with monotone non-decreasing property as in Reference 3 without fading memory or windowing of data, a controller falsified once would not be recycled even though it would be the best candidate controller at sometime after being falsified. But, with the modified cost function (3) having the fading memory and time-windowing features of 1) and 2) above, Assumption 1 can be relaxed, as we will demonstrate via an example.

Algorithm 1. (Hysteresis algorithm for unfalsified control)[4]

INITIAL SETTING:

- a finite set **K** of m candidate controllers $K_i,\ i \in \mathbf{I} = \{1, 2, \cdots, m\}$
- initial cost $\gamma(0) = 0$ and $J(i, 0) = 0,\ \forall i \in \mathbf{I}$
- sampling time Δt
- effective time of the max operator τ_0
- the values of ε, λ
- initial online controller $K = K_m$ at $t = 0$

PROCEDURE at each time $t = k\Delta t$ and for each candidate controller $K_i,\ i \in \mathbf{I} = \{1, 2, \cdots, m\}$:

1. Measure $u(k\Delta t)$, $y(k\Delta t)$ and set $\gamma(k\Delta t) = J(K(k\Delta t), u(k\Delta t), y(k\Delta t), k\Delta t)$.
2. For each $i \in \mathbf{I}$, calculate $\tilde{r}_i(k\Delta t),\ \tilde{e}_i(k\Delta t),$ and $J(i, k) = J(K_i, u(k\Delta t), y(k\Delta t), k\Delta t)$.
3. If $\min_{K_i \in \mathbf{K}} \{J(i, k\Delta t)\} \le \gamma(k\Delta t) - \varepsilon$,

   then $K((k+1)\Delta t) = \arg\min_{K_i \in \mathbf{K}} \{J(i, k\Delta t)\}$.

   Otherwise, $K((k+1)\Delta t) = K(k\Delta t)$ where $K(k\Delta t)$ is online controller at time t.
4. Set k=k+1 and repeat. ♦

Comment: The addition of non-zero fading memory and finite time window parameters, λ and τ_0 , in Equation (3) is the difference as compared to the cost form (2) used in Reference 3. When these parameters are positive numbers, the modified cost need not necessarily increase monotonically as it does in Reference 3, which means that after some time has elapsed a controller can be recycled and re-added to the unfalsified set without increasing the cost level γ . Therefore, when either or both of the parameters λ and τ_0 is non-zero, the set of unfalsified controllers at each cost level γ computed by Algorithm 1 no longer necessarily shrinks monotonically with time as in the previous works[1, 3-4] where there was no fading memory or windowing in the cost function. ♦

At the beginning of the hysteresis switching[2] Algorithm 1, functions and parameters are initialized. The time-window duration τ_0 would normally be selected to be somewhat less than the time-scale over which significant plant variations may occur that cannot be robustly accommodated by any one candidate controller in the controller pool **K**. After initialization, output data $u(t)$ and $y(t)$ are measured and fictitious signals are computed by (1) and (2). The unfalsified cost level $J(i, k)$ of each candidate controller $K_i,\ i \in \mathbf{I} = \{1, 2, \cdots, m\}$ is iteratively updated in step 2 of Algorithm 1, based on the measured data and the fictitious signals by (3) where T_{spec} will be presented in Section IV. If the unfalsified cost level $J(i, k)$ is larger than the current unfalsified cost level γ , the controller K_i is falsified at this cost level and is not used. The PID controller K_i with the currently minimum cost among the unfalsified controllers is examined, and its cost is compared to the previous cost-minimizing controller's cost γ minus a small positive number ε, called the *hysteresis constant.* If unfalsified cost level $J(i, k)$ of this cost-minimizing K_i is less than $\gamma - \varepsilon$, Algorithm 1 updates the currently active controller, replacing it with K_i .

IV. Computer Simulations

We now describe the results of simulations that were performed to verify the effectiveness of the modified cost function. Details of the simulation are as follows. We consider an unfalsified adaptive PID controller as shown in Figure 2. Its three gains $\{K_P, K_I, K_D\}$ are selected from a finite set by the hysteresis algorithm[2]. We shall compare the result of using the hysteresis algorithm with, and without, the fading memory and time-windowing modification of the cost $J(K_i, u, y, t)$. In both of the two cases considered the hysteresis algorithm is the same. Only the cost function $J(K_i, u, y, t)$ differs.

- The plant switches from $G_1(s)$ to $G_2(s)$ at time $t_1 = 100$ and from $G_2(s)$ back to $G_1(s)$ at time $t_2 = 300$.
- $G_1(s) = \dfrac{0.1}{s+0.1}$ and $G_2(s) = -G_1(s) = -\dfrac{0.1}{s+0.1}$
- $\mathbf{K} = \{K_1, K_2\} = \{\{3, 0.1, 0.1\}, \{-3, -0.1, -0.1\}\}$
- The reference input *r*(*t*) is a step function with amplitude 1.

Case 1 (perfect memory):
Cost function (3) with $\tau_0 \to \infty,\ \lambda = 0$;

Case 2 (fading and limited memory):
Cost function (3) with $\tau_0 = 5,\ \lambda = 0.1$.

To select a performance specification T_{spec}, following performance goal is considered

$$\left\| w_1 * \tilde{e}_i \right\|^2_{L_2[t-\tau_0,t]} + \left\| w_2 * u \right\|^2_{L_2[t-\tau_0,t]} \le \left\| \tilde{r}_i \right\|^2_{L_2[t-\tau_0,t]},\ \forall t \ge \tau_0 \tag{6}$$

where $\|\cdot\|$ is a L_2 norm and $*$ denotes convolution. Inequality (6) means that the error and control output should be small compared to the reference signal. The signals w_1 and w_2 are inspired by controller design method of H_∞ mixed-sensitivity loop shaping. Frequency responses of w_1 and w_2 are W_1 and W_2, which satisfy the following H_∞ performance criterion.

$$\left\| \begin{bmatrix} W_1 S \\ W_2 K S \end{bmatrix} \right\|_\infty \le 1 \tag{7}$$

The transfer function S in (7) is a sensitivity function of the system. Based on (6), the performance specification set T_{spec} of i^{th} controller at time *t* is selected to be

$$T_{spec}(r_i(t), y(t), u(t), t) = \frac{\left\| w_1(t) * (y(t) - \tilde{r}_i(t)) \right\|^2_{L_2[t-\tau_0,t]} + \left\| w_2(t) * u(t) \right\|^2_{L_2[t-\tau_0,t]}}{\left\| \tilde{r}_i(t) \right\|^2_{L_2[t-\tau_0,t]}}. \tag{8}$$

<u>*Comment:*</u> When the duration of time-window τ_0 is infinity and fading memory parameter λ is zero as in Case 1 above, the performance specification (8) is *cost-detectable*. As we proved in Reference 3, cost-detectability is sufficient to ensure that the hysteresis switching algorithm robustly converges to a stabilizing controller whenever a stabilizing controller exists in the candidate controller set. Thus, we may reasonably expect for all intents and purposes that the fading memory and time-windowing algorithm will be similarly robust if τ_0 and λ are chosen to be sufficiently large. ♦

There are only two elements in the set of candidate controllers. Controller K_1 stabilizes a plant $G_1(s)$ and controller K_2 stabilizes a plant $G_2(s)$. On the other hand, K_1 cannot stabilize $G_2(s)$ and K_2 cannot stabilize $G_1(s)$. For this simulation, the two plants are the same except for their sign. Likewise, the controllers K_1 and K_2 the same, except for opposite signs of their gains.

The cost level of controller K_2 increases abruptly in the early stage of the simulation because the plant is $G_1(s)$, which is destabilized by K_2. Later, at simulation time t_1, the plant changes into $G_2(s)$. Then, the cost for controller K_1 begins to increase quickly, because K_1 destabilizes the plant $G_2(s)$. In this stage, the traditional Case 1 cost

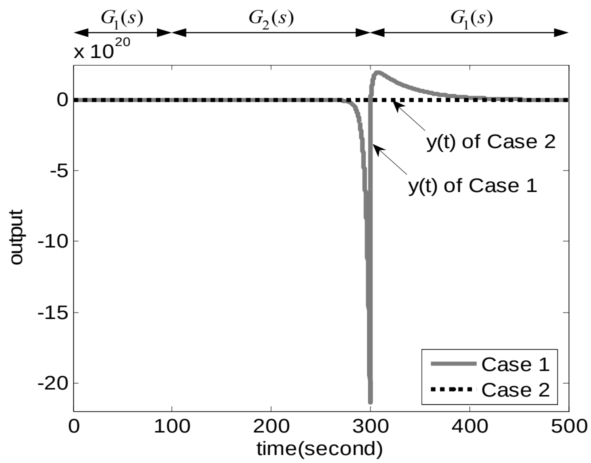

Figure 3. Simulation result of Case 1 and Case 2 in large scale

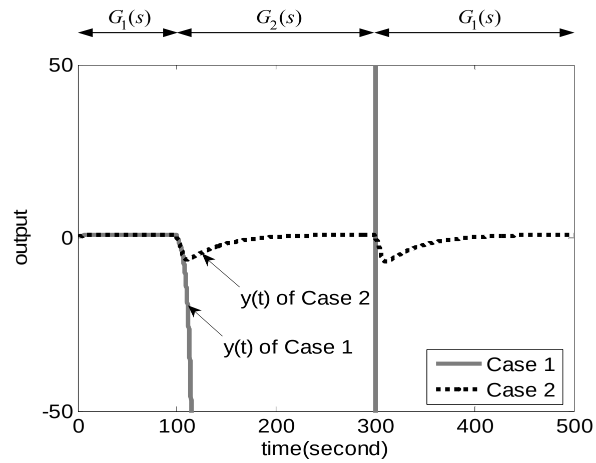

Figure 4. Same as Figure 3, but with vertical axis magnified by factor 0.5×10^{20}.

function of unfalsified control cannot choose a proper controller because the cost level of controller K_1 is still higher than K_2. Therefore, the output $y(t)$ of the Case 1 grows negatively exponentially as shown in Figure 3. At time t_2, the plant comes back to $G_1(s)$ and the output needs long time to become stabilized.

With the Case 2 fading-memory time-windowed cost function, controller K_2, which was initially falsified, is recycled and is available to be selected for stabilization of $G_2(s)$. The stable simulation result in Figure 4 reflects the fact that the currently active controller was successfully switched from K_1 to K_2 when the plant switched from $G_1(s)$ to $G_2(s)$, although several seconds were needed to remove the effect of old data after t_1. Next, after t_2 when the plant switched from $G_2(s)$ back to $G_1(s)$, the Case 2 fading memory cost function allowed the controller to correctly switchback to K_1 again. The results in Figure 3 and Figure 4 show that the fading memory Case 2 cost function controller maintains stability, in contrast to traditional infinite-memory Case 1 cost function, which does not. Case 1 has exponentially increased peak around time t_2, but Case 2 shows good performance at all times.

V. Conclusion

A new unfalsified control cost function with time-windowing and the fading memory was introduced in this paper. Simulations with slowly time-varying plant demonstrate the effectiveness of the modification for stabilizing a plant that violates the usual feasibility assumption of unfalsified control and cannot be robustly stabilized by any single controller in the candidate set. The time-windowing and fading memory of the new cost function works because it allows re-use of controllers that were falsified by very old data that may no longer be relevant in the case of a slowly or infrequently varying plant.

References

[1] Safonov, M. G. and Tsao, T. C., "The unfalsified control concept and learning," *IEEE Trans. Autom. Control*, AC-42(6), June 1997, pp.843-847.

[2] Morse, A. S., Mayne, D. Q., and Goodwin, G. C., "Application of hysteresis switching in parameter adaptive control," *IEEE Transactions on Automatic Control*, 37(9), September 1992, pp.1343–1354.

[3] Stefanovic, M., Wang W., and Safonov, M. G., "Stability and convergence in adaptive systems," *In Proc. American Control Conf.*, Boston, MA, June 30 – July 2, 2004. IEEE Press, New York.

[4] Jun, M. and Safonov, M. G., "Automatic PID tuning: An application of unfalsified control," In *Proc. IEEE CCA/CACSD*, Kohala Coast–Island of Hawaii, HI, August 22-27, 1999, pp.328–333.

[5] Angeli, D. and Mosca, E., "Adaptive switching supervisory control of nonlinear systems with no prior knowledge of noise bounds," *In Proc. IEEE Conf. on Decision and Control*, vol.2, Las Vegas, NV, December 2002, pp.1187 – 1192.

[6] Hespanha, J. P., Liberzon D. and Morse, A. S., "Bounds on the number of switchings with scale-independent hysteresis," *In Proc. IEEE Conf. on Decision and Control,* vol.4, Sydney, NSW, December 2000, pp.3522 – 3627.

[7] Rohrs, C. E., Valavani, L., Athans, M. and Stein, G., "Robustness of adaptive control algorithms in the presence of unmodeled dynamics," *IEEE Trans. Autom. Control*, AC-30(9), September 1985, pp.881-889.