Summary

A significant part of the adaptive control literature has –justifiably– been devoted to the stability/robustness/performance issue, under conditions of insufficient excitation. This practically important problem originates in the desire to allow the adaptation to adjust the estimated parameters at all times. Such adaptive controllers exhibit good RMS behavior and they are able to track slowly varying parameters. However, their main drawback is arguably the possibility of parameter drifts in low SNR situations (insufficient excitation). These drifts may, in turn, induce error bursts that limit the practical usefulness of adaptation.

Analytical studies of adaptation bursting have often concentrated on the nonlinear dynamics aspect of the phenomenon. In this work, we approach the problem from an optimization point of view. Parameter drifts can then be interpreted as non-robustness of ill-posed optimization problems. In this framework, error bursts are an immediate consequence of Lipschitz continuity of the parameter estimates (finite adaptation gains).

Guided by these observations, we construct simple but general bursting scenaria that expose the fundamental performance limitations of adaptive algorithms. We show that for a wide class of adaptive algorithms, error burst magnitudes proportional to the diameter of the parametric uncertainty set can be obtained when perturbations are present but there is lack of sufficient excitation. Algorithms in this class are essentially characterized by fading memory and finite adaptation gains. An interesting by-product of this analysis is a proof that dead-zone algorithms (a common remedy in the time-invariant case) are not immune to bursting when estimating time-varying parameters.

In fact, the estimation of time-varying parameters in the absence of sufficient excitation poses a challenging theoretical problem. In this case, achieving *limsup* performance guarantees is as hard as achieving L_∞ performance from initial conditions that are zero in the state errors but arbitrary in the parameters. In a somewhat loose interpretation of the results, burst suppression in the general case seems to require estimators with "infinite" adaptation gains or injection of excitation.

Main Result

Under seemingly benign assumptions, adaptive algorithms satisfy the following performance *lower* bound:

For any $\delta > 0$, there exist bounded, piecewise continuous regressor (or input) and a disturbance d with $\|d\|_\infty \leq \delta$ and such that the worst-case, lim-sup estimation error proportional to the diameter of the parametric uncertainty set.

Example

Bursting scenaria can be constructed systematically to expose this fundamental performance limitation of adaptive algorithms. In particular, the severity of the bursts in adaptive control may be significantly amplified by the temporary "destabilization" of the closed loop.

Consider the plant with input disturbance d

$$y_p = \frac{b}{s+a}[u_p + d]$$

with nominal parameters $a = 0, b = 1$. Let the control input u_p be designed so that the nominal plant output tracks the output of the reference model

$$y_m = \frac{1}{s+1}[r]$$

for any bounded reference input r. For example, $u_p = [r, y_p]\theta$ where θ is updated by a gradient-based adaptive law with projection. For simulation convenience, we update θ in discrete-time with sampling interval $T_s = 0.2$, according to the following equations:

$$\theta(k+1) = \mathbf{\Pi}_{\mathcal{M}}\left[-\frac{2\epsilon_1(k)\zeta(k)}{m(k) + 2\zeta^\top(k)\zeta(k)}\right]$$

The estimation error and regressor vector are taken as the sampled versions of their continuous time counterparts

$$\begin{aligned} \epsilon_1 = \theta^\top\zeta - \frac{1}{s+1}[u_p] \quad &; \quad \zeta = [y_p, \frac{1}{s+1}[y_p]]^\top \\ m = 1 + m_2 \quad &; \quad \dot{m}_2 = -0.75m_2 + u_p^2 + y_p^2 \end{aligned}$$

The projection set $\mathcal{M}$ is selected to contain the nominal controller parameter vector ($\theta_* = [1, -1]^\top$); here $\mathcal{M} = [0.3, 3] \times [-4, 4]$.

Further, define the reference input and disturbance signals as

$$\begin{aligned} r &= R_1[\sin(4t) + \sin(t)] + R_2 \\ d &= \mathrm{sat}_{0.5}[-Ky_p] \end{aligned}$$

where $\mathrm{sat}_{0.5}$ denotes a saturation nonlinearity with linear region $[-0.5, 0.5]$ (clearly, $\|d\|_\infty \leq 0.5$).

It follows that whenever R_1, R_2 are sufficiently small so that Ky_p is in the linear region of the saturation, and r is PE, the adaptation algorithm drives the parameter estimates towards the point $[1, K-1]$. Thus, if $K - 1 > 0$, the nominal unperturbed closed-loop is unstable, something that becomes evident in the form of a burst as soon as the disturbance is removed and/or the magnitude of the reference input is increased. (The burst magnitude is essentially independent of the disturbance bound.)

Remedies

Several techniques can be used to remedy this situation. One is the use of some form of a dead-zone that modifies the optimization objective so that error convergence to zero is not required. The other was introduced in [IJACSP] and uses a "standard" adaptive scheme but employs set-membership concepts to estimate and reduce the parametric uncertainty set on-line.

Unfortunately, none of these works in the time-varying case.

A dead-zone can be introduced in the adaptive law by simply replacing the estimation error ϵ_1 with

$$\epsilon_{dz} = \begin{cases} \epsilon_1 - d_0 & \text{if } \epsilon_1 > d_0 \\ 0 & \text{if } |\epsilon_1| \le d_0 \\ \epsilon_1 + d_0 & \text{if } \epsilon_1 < -d_0 \end{cases}$$

where d_0 is the dead-zone threshold (here, $d_0 = 0.6$).

A different approach to burst suppression is to decrease the size of the effective parametric uncertainty using set-membership (SM) estimation ideas [IJACSP]. In this approach, the adaptive controller employs two estimators: One is responsible for the updating on the controller parameters via a simple gradient scheme with projection. The other is an SM etimator, responsible for updating the parametric uncertainty set used for projecting the parameters of the first estimator. That is, the parameter updates are performed according to

$$\theta(k+1) = \mathbf{\Pi}_{E_k \cap \mathcal{M}} \left[-\frac{2\epsilon_1(k)\zeta(k)}{m(k) + 2\zeta^\top(k)\zeta(k)} \right]$$

The sets E_k are updated by the auxiliary set-membership estimator operating in hybrid mode. The SM updates of the parametric uncertainty set E_k (i.e., R_k and c_k) are performed every 0.5 time units, with a threshold $\mu(t) = 0.6/\sqrt{m} + 5e^{-0.25t}$ which is consistent with dead-zone case. In addition, for comparison purposes, we also simulate the closed-loop response where the controller parameters are updated by an SM estimator alone.

Simulation Results

The figures show the closed-loop responses (tracking error and parameter estimates) with the four adaptive controllers. These are fixed set projection, dead-zone, set membership, and updated projection sets; the error plots are shifted and clipped to emphasize steady-state details.

In all cases, we begin with zero initial conditions while the reference input and the perturbation alternate between the following two phases:

Phase 1 : *(Drift phase)* $K = 5, R_1 = 0.1, R_2 = 0$ (85 time units).

Phase 2 : *(Burst phase)* $K = 0, R_1 = 0, R_2 = 1$ (15 time units).

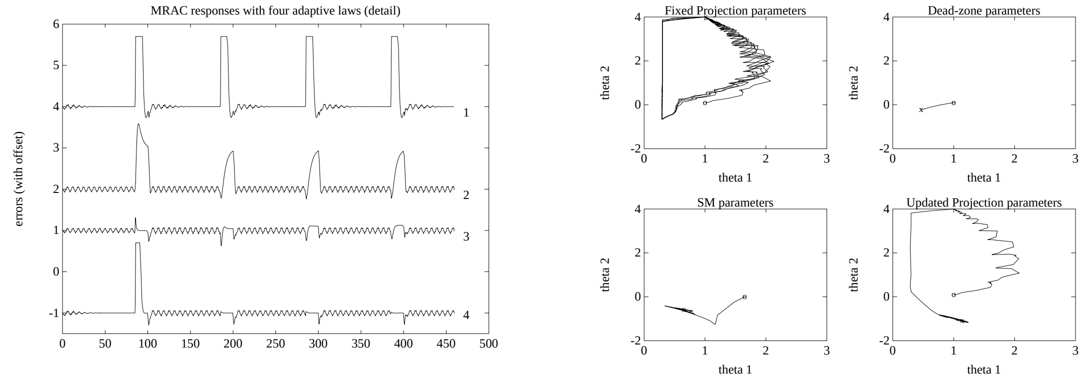

Figure 1:
a. Tracking error response of MRAC with the four adaptive laws. Reference input and disturbance selection: Phase-1 followed by Phase-2. Key: 1: fixed projection set, 2: dead-zone, 3: SM, 4: updated projection set.
b. Parameter trajectories of MRAC with the four adaptive laws. Reference input and disturbance selection: Phase-1 followed by Phase-2.

As expected, the tracking error bursts are considerably reduced with dead-zone adaptation. This improvement, however, is obtained at the expense of a deterioration of the RMS performance. The latter is manifested by a significant increase of the error in the Phase-2 intervals where the reference input is large (the peak error is close to two, its estimated worst case value.) Similar conclusions can be drawn for the MRAC with the SM estimator, although in this case the deterioration of the RMS performance is considerably smaller. This can be attributed to the more efficient utilization of information by the SM estimator.

On the other hand, the MRAC with updated projection sets exhibits good asymptotic performance in both the RMS and *lim sup* sense. Its initial transient is similar to the MRAC with fixed projection set since, at that point, the available data offer very little information about the parametric uncertainty set. During the rest of the cycles, however, the reduction of the latter does not permit any significant parameter drift towards the destabilizing region and, consequently, limits the size of the error bursts. In contrast to the dead-zone and SM adaptation, this algorithm ensures the "convergence" of the tracking error to zero during the Phase-2 intervals where the plant disturbance is absent.