Proceedings of the 17th World Congress
The International Federation of Automatic Control
Seoul, Korea, July 6-11, 2008

Practical approaches to low-order anti-windup compensator design: a flight control comparison

Murray L. Kerr [*,1,2] Matthew C. Turner [*,2] Ian Postlethwaite [*,2]

[*] *Control and Instrumentation Research Group, Dept. of Engineering, Univ. of Leicester, Leicester, LE1 7RH, UK.*
Email: {mlk9,mct6,ixp} @le.ac.uk

Abstract: This paper considers three different methods for low-order anti-windup (AW) compensator design and compares their application to a realistic flight control problem, where the dominant actuator nonlinearity is aileron rate saturation. The compensator design methods all rigorously enforce at least local exponential stability via absolute stability results, but differ in both the construction of the AW compensator itself and the performance requirements used in the design. In particular, the paper compares a low order AW "optimal" design method in which the compensator poles and zeros are chosen by the designer and the accompanying gains synthesised optimally; a new low order method in which the optimisation procedure optimally chooses both the zeros and the gains; and a recently introduced classical design method where loopshaping is used to completely determine the gains and dynamics of the compensator. The methods are compared using a high-order nonlinear model of the lateral dynamics of an experimental aircraft, on which similar compensators have recently been flight tested.

Keywords: Anti-windup; Asymptotic stabilization; LMIs; QFT; Aerospace applications.

1. INTRODUCTION

Modern anti-windup (AW) compensation techniques are numerous and diverse, owing perhaps to the different branches of control theory to which they can be linked, e.g. nonlinear control, robust control, and PID control. Although the problem of control system saturation ("windup") was initially linked to integral type controllers, it was soon evident that saturation problems occurred in many types of control systems. The subsequent "anti-windup" methods which were developed initially were somewhat ad hoc and were rarely accompanied by meaningful performance or stability guarantees. However, developments in the 1980s and 1990s saw the marriage of these anti-saturation ideas with techniques from robust and nonlinear control, thereby giving AW a solid theoretical foundation.

From this basis, many AW techniques which enforce global or local stability, $\mathscr{L}_2$ performance, and similar design metrics have been proposed (Mulder et al. (2001); Turner and Postlethwaite (2004); Grimm et al. (2004)) over recent years. Despite the attractive features of many of these "advanced" schemes, few of them have found their way into application and none are commonplace in practical control engineering. The reason for this is that the apparent advantages of these theoretically strong compensation techniques have to be tempered with the practical drawbacks inherent in many of them. For instance, many of these modern schemes produce compensators of high order making them problematic to implement in systems where computational resources are at a premium; many also do not explicitly control pole magnitudes or locations which are a vital consideration for discrete implementation. Moreover, the proving ground of these modern AW techniques tends to be the realm of low order linear plant models which are gross simplifications of real physical systems. Although successful application to these simple models is a necessary part of the validation of such techniques, many important practical considerations are not present in these simplified models.

Apart from the validation of individual design methods, the literature discussing objective comparisons of various AW techniques is scarce. Many papers (Turner and Postlethwaite (2004); Grimm et al. (2004); Queinnec et al. (2006)) demonstrate performance improvement using AW compensation over using no AW compensation, providing some justification to the proposed methods, but relatively few provide a fair and critical comparison between different AW methods.

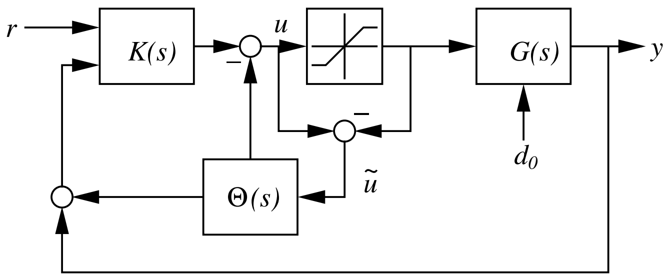

Fig. 1. Standard anti-windup configuration

The aim of this paper, therefore, is to compare three techniques for AW compensator design for a complex, realistic flight control example. To ensure they are feasible for practical implementation, the methods are chosen to produce low order compensators (all compensators produced are 2nd order) and to contain no direct-feedthrough terms which may impede implementation. The techniques we discuss are; (i) an "optimal" low order synthesis method discussed in (Turner and Postlethwaite (2004); Turner et al. (2007)) in which the designer chooses the AW compensator dynamics and LMIs optimally construct the AW compensator gains; (ii) a classical AW method proposed in (Kerr et al. (2007)) where bounds in the Nichols chart (NC) are used to guide the designer in the construction of the AW compensator; and (iii) a new low order zero-optimisation method, similar to (i) but where the LMI optimisation also has the freedom to place optimally the zeros of the AW compensator and the designer simply chooses the poles.

2. ANTI-WINDUP DESIGN

A reasonably generic illustration of the AW problem is shown in Figure 1 where $G(s) = [G_1(s), G_2(s)]$ is the nominal linear plant, $K(s) = [K_1(s), K_2(s)]$ is the nominal linear controller,

[1] Corresponding author.
[2] Work supported by the UK EPSRC.

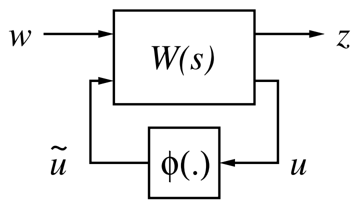

Fig. 2. Equivalent AW problem

$y \in \mathbb{R}^p$ is the measured output and $u \in \mathbb{R}^m$ the control input. The AW compensator, $\Theta(s) \in \mathscr{RH}_\infty^{(m+p)\times m}$, is driven by the signal $\tilde{u} = u - \mathrm{Sat}(u) = \mathrm{Dz}(u)$. Upon activation, the AW compensator $\Theta(s)$ injects two signals into the system, one at the controller input and one at the controller output. Following (Turner and Postlethwaite (2004); Turner et al. (2007)), the nominal plant $G(s)$ and the controller $K(s)$ are assigned the following state-space representations

$$G(s) \sim \left[\begin{array}{c|cc} A_p & B_{pd} & B_p \\ \hline C_p & D_{pd} & D_p \end{array}\right], \tag{1}$$

$$K(s) \sim \left[\begin{array}{c|cc} A_c & B_{cr} & B_c \\ \hline C_c & D_{cr} & D_c \end{array}\right]. \tag{2}$$

There are several different ways to frame the goals of the AW problem but, in essence, most procedures seek to establish stability of the system in Figure 1 and also to ensure swift return from saturated nonlinear to nominal linear behaviour. The AW problem is thus frequently re-cast as that depicted by the Figure 2 in which $W(s)$ represents the overall linear system, containing the dynamics of $G(s)$, $K(s)$ and $\Theta(s)$, and $\phi(.) = \mathrm{Dz}(\cdot)$, the deadzone nonlinearity. The exogenous input vector $w(t)$ and output $z(t)$ are used to capture, as discussed later, the performance of the system. The AW problem is then to choose $\Theta(s)$ to ensure: (i) asymptotic stability of the origin of the system when $w(t) \equiv 0$; and (ii) that the output $z(t)$ remains small, in some sense, when $w(t)$ is present.

Figure 2 depicts a general AW problem in which $w(t)$ and $z(t)$ are not specified. However, as argued in Teel and Kapoor (1997); Weston and Postlethwaite (2000); Turner and Postlethwaite (2004), it is desirable to design $\Theta(s)$ to ensure "optimal" recovery of linear performance. In this case, Figure 2 can be seen as a "mismatch" system representing the deviation of the nonlinear (saturated) system's behaviour from linear behaviour. In this case $w = u_{lin}$, the linear control signal from the ideal system without saturation, and $z = z_d$, the deviation of the performance outputs from those of the ideal system without saturation, which one would like to minimise (Weston and Postlethwaite (2000); Grimm et al. (2004); Villota et al. (2006)). In the subsequent work an objective of the design methods is to minimise the $\mathscr{L}_2$ gain of the mismatch system in some way.

2.1 Local stability

When $G(s) \in \mathscr{RH}_\infty$ (i.e. $G(s)$ is stable), it is well known that there always exists an AW compensator, $\Theta(s)$ which can guarantee global stability of the system in Figure 1 (Teel and Kapoor (1997); Weston and Postlethwaite (2000)). When $G(s) \notin \mathscr{RH}_\infty$ (i.e. $G(s)$ is not stable), establishing stability is more difficult and it often suffices to guarantee stability locally (Gomes da Silva Jr. et al. (2003); Hindi and Boyd (1998)). In this case, it is assumed that $\phi(\cdot)$ is no longer the standard deadzone which inhabits the Sector$[0, I]$, but is instead confined to some narrower sector, i.e.

$$\phi(\cdot) \in \mathrm{Sector}[0, \varepsilon I] \quad \varepsilon \in (0,1).$$

This is equivalent to assuming that $|u_i| \leq \bar{u}_i/(1-\varepsilon)\ \forall i$, with $\bar{u}_i$ the saturation level in the i'th channel. It then follows that, regardless of the locations of the poles of $G(s)$, there always exists a sufficiently small ε such that local stability is guaranteed. As ε approaches unity, stability is closer to being provided globally. Note that by choosing $\varepsilon < 1$, one can establish both local stability for unstable systems and $\mathscr{L}_2$ performance properties in small regions of the signal space of u, making it easier to obtain better local performance, typically at the expense of large-signal stability. In the remainder of the paper, each of the design methods treats this "local" stability/performance objective.

3. LOW ORDER ANTI-WINDUP

It is well known in the AW literature that, given $G(s) \in \mathscr{RH}_\infty$ with state dimension n_p, a full order (n_p'th order) AW compensator guaranteeing global exponential stability always exists (Teel and Kapoor (1997); Weston and Postlethwaite (2000); Grimm et al. (2003)). It is also well known that the lower bound on the $\mathscr{L}_2$ gain, from either exogenous inputs to arbitrary performance outputs, or of the so-called mismatch system, is achieved by certain *full order* AW compensators (Grimm et al. (2003, 2004)). These properties roughly carry over to unstable systems, which can be handled by appropriately modifying the sector condition used in AW synthesis.

These compelling arguments supporting the use of full order (n_p'th order) compensators are usually offset by other considerations. Perhaps the most obvious of these is that, in practice, it is undesirable to implement a control law which requires an extra n_p states just to handle saturation constraints; most AW compensators favoured by practitioners are static or low order (Rundquist and Stahl-Gunnarsson (1996); Hanus et al. (1987)). Another possible disadvantage with full order compensators is that the dynamics required to push the $\mathscr{L}_2$ gain close to optimal may be unnecessary and may actually give rise to behaviour which may hinder performance observed in the time domain due to excessively slow poles or approximate pole-zero cancellations. Thus there has been an understandable interest in low order compensators (compensators with order lower than n_p) that combine the theoretical advantages of full order compensators (feasibility, performance) but also retain the practical advantages of the low order compensators (ease of implementation, predictable dynamic behaviour).

Unfortunately, the AW literature has seen relatively few systematic design procedures for low order AW compensators. Perhaps the main reason is that the procedure for synthesising low order compensators does not appear, in general, to be convex. In order to circumvent this problem, a similar procedure to that used in low-order $\mathscr{H}_\infty$ control can be used (Galeani et al. (2006)), where an alternating projection method is used to "convexify" the nonconvex optimisation procedure; it is noted that this is not guaranteed to work. Similarly, one could use model reduction methods to obtain a low order compensator with similar input-output characteristics to a good full-order design. An alternative is to fix the dynamics of the AW compensator and then to optimise the associated gains and/or zeros (Turner and Postlethwaite (2004); Biannic et al. (2007)). Although again no feasibility is guaranteed, a useful guide to the dynamics required can be found by examining the dominant modes of a full order design. A third method, which can be particularly useful in single-input-single-output (SISO) problems, is to design AW compensators using classical loopshaping methods in which the dynamics are chosen based on the frequency response of the linear part of the system and exclusion regions in the NC (Kerr et al. (2007)). It is these last two methods which we discuss and compare.

The overall goal of all the AW design methods discussed herein can therefore be stated as

Goal 1: To design $\Theta(s) \in \mathscr{RH}_\infty$ such that

(1) The system in Figure 1 is globally asymptotically for all $\phi(.) \in \mathrm{Sector}[0, \varepsilon I]$, $\varepsilon \in (0,1)$.
(2) $\deg(\Theta(s)) < n_p$.

(3) The "gain" of the mismatch system is sufficiently small, i.e. $\|z_d\|_2 < \gamma\|u_{lin}\|_2$ for some $\gamma > 0$.

3.1 Standard Low order design

The simple static/low-order method proposed in Turner and Postlethwaite (2004) involves the partitioning of the AW into two parts

$$\Theta(s) = \begin{bmatrix} \Theta_1(s) \\ \Theta_2(s) \end{bmatrix} = \begin{bmatrix} F_1(s)\tilde{\Theta}_1 \\ F_2(s)\tilde{\Theta}_2 \end{bmatrix} \quad (3)$$

where the $F_1(s)$ and $F_2(s)$ are two stable transfer function matrices chosen by the designer and $\tilde{\Theta}_1$ and $\tilde{\Theta}_2$ are two gain matrices produced by an LMI optimisation procedure. The objective of the LMI optimisation is to minimise the $\mathscr{L}_2$ gain of the nonlinear map from w to z in Figure 2, where the system is the mismatch system and hence $w = u_{lin}$ and $z = z_d$. For a full discussion of this method, see Turner and Postlethwaite (2004); Turner et al. (2007). A brief design procedure is as follows:

Procedure 1: Gain optimisation

(1) Choose $\varepsilon \in (0,1)$. This dictates the size of sector for which the system is stable.
(2) Choose $F_1(s)$ and $F_2(s)$, including the poles and zeros and their directions.
(3) Choose weighting matrices $W_p > 0, W_r > 0$ to trade-off performance and robustness.
(4) Minimise γ subject to the LMI in equation (4).
(5) Form $\tilde{\Theta} = [\tilde{\Theta}_1' \quad \tilde{\Theta}_2']' = LU^{-1}$.
(6) Form $\Theta(s)$ according to equation (3).

Remark 1: This low order technique requires the designer to *specify fully* the dynamics of the compensator and only synthesises the gains in an optimal fashion. Although this may seem restrictive, the dynamics of a full order compensator can be a useful guide in choosing these dynamics and it is often possible to synthesise a low-order compensator with very similar performance to full-order compensators (Turner and Postlethwaite (2004); Turner et al. (2007)). ◦

Using this design method, a compensator $\Theta(s)$ is synthesised which guarantees (i) asymptotic stability for all $\phi(.) \in \text{Sector}[0, \varepsilon I]$; and (ii) the local $\mathscr{L}_2$ gain bound γ of the mismatch system weighted by W_p is minimised, and some robustness to additive plant perturbations is provided (Turner et al. (2007)).

3.2 Low order design with zero optimisation

Although the foregoing method represents a simple approach to low order AW design, it suffers from several drawbacks. Firstly, the optimisation may not be feasible due to fixed poles/zeros; and, secondly, even if feasible, the optimisation is performed over a restricted set of low order compensators. Consequently, some sort of iteration in pole/zero choice is typically required in order to achieve a satisfactory design. To alleviate some of these problems, a new design approach is proposed below that provides both gain and zero optimisation.

It is well known in the control literature that optimisation of the residual gains of a transfer function, subject to quadratic constraints and objectives, results in a convex optimal control problem. Exploiting this fact, a design method is proposed based on representing the filters in their dyadic partial fraction expanded form. This idea has in fact already been used in AW control in the recent independent work of Biannic et al. (2007), although the problem solved in that work differs from that presented here in terms of the performance objective, the lack of explicit use of the mismatch system and in that the partial fraction expansion used is not in dyadic form. The latter is important, as it highlights some of the remaining limitations of the method. Interestingly, when the filters are placed in the dyadic partial fraction expansion form, they are actually in an identical form to that considered in the previous section and given in Turner and Postlethwaite (2004). Hence the existing design theory in Turner and Postlethwaite (2004), and its later development to basic robust AW control in Turner et al. (2007), can be exploited and reused.

As in the method of Turner and Postlethwaite (2004), the gain and zero optimization low-order method proposed here involves the partitioning of the AW compensator into two parts. Here $\Theta_1(s)$ and $\Theta_2(s)$ are given in dyadic partial fraction form:

$$\Theta_i = D_i + \sum_{j=1}^{k_i} \frac{y_{1i(j)}u_{1i(j)}^H}{s + a_{i(j)}} + \sum_{j=1}^{l_i} \frac{y_{2i(j)}u_{2i(j)}^H s + y_{3i(j)}u_{3i(j)}^H}{s^2 + b_{i(j)}s + c_{i(j)}^2} \quad (5)$$

where k_i and l_i indicate the number of first and second order terms desired in the dynamics of $\Theta_i(s)$. Denoting the output dimension (either m or p) of $\Theta_i(s)$ as q_i, it is convenient to define the following matrices:

$$\begin{aligned}
\mathbf{y_{1i}} &= [\, y_{1i(1)} \;\; y_{1i(2)} \;\; \cdots \;\; y_{1i(k_i)} \,] \in \mathbb{R}^{q_i \times k_i} \\
\mathbf{y_{2i}} &= [\, y_{2i(1)} \;\; y_{2i(2)} \;\; \cdots \;\; y_{2i(l_i)} \,] \in \mathbb{R}^{q_i \times l_i} \\
\mathbf{y_{3i}} &= [\, y_{3i(1)} \;\; y_{3i(2)} \;\; \cdots \;\; y_{3i(l_i)} \,] \in \mathbb{R}^{q_i \times l_i} \\
\mathbf{y_{4i(j)}} &= [\, y_{2i(j)} \;\; 0_{q_i \times 1} \,] \in \mathbb{R}^{q_i \times 2} \\
\mathbf{y_{4i}} &= [\, y_{4i(1)} \;\; y_{4i(2)} \;\; \cdots \;\; y_{4i(l_i)} \,] \in \mathbb{R}^{q_i \times 2l_i}
\end{aligned}$$

$$\begin{aligned}
\mathbf{u_{1i}} &= [\, u_{1i(1)} \;\; u_{1i(2)} \;\; \cdots \;\; u_{1i(k_i)} \,] \in \mathbb{R}^{m \times k_i} \\
\mathbf{u_{2i}} &= [\, u_{2i(1)} \;\; u_{2i(2)} \;\; \cdots \;\; u_{2i(l_i)} \,] \in \mathbb{R}^{m \times l_i} \\
\mathbf{u_{3i}} &= [\, u_{3i(1)} \;\; u_{3i(2)} \;\; \cdots \;\; u_{3i(k_i)} \,] \in \mathbb{R}^{m \times l_i} \\
\mathbf{u_{4i(j)}} &= [\, u_{2i(j)} \;\; u_{3i(j)} \,] \in \mathbb{R}^{m \times 2} \\
\mathbf{u_{4i}} &= [\, u_{4i(1)} \;\; u_{4i(2)} \;\; \cdots \;\; u_{4i(l_i)} \,] \in \mathbb{R}^{m \times 2l_i}
\end{aligned}$$

In this representation, $\mathbf{u_{1i}}$, $\mathbf{u_{2i}}$ and $\mathbf{u_{3i}}$ are the design parameters chosen by the optimisation algorithm. Considering the *degenerate case* where $\mathbf{y_{2i}} = \mathbf{y_{3i}}$, this can be put in the following equivalent form, where the design parameters are captured in the gain matrices $\tilde{\Theta}_i$:

$$\Theta_i(s) = F_i(s)\tilde{\Theta}_i \quad (6)$$

where $F_i(s) \in \mathbb{R}^{q_i \times q_i + k_i + 2l_i}$ and $\tilde{\Theta}_i \in \mathbb{R}^{q_i + k_i + 2l_i \times m}$. Here $F_i(s)$ and $\tilde{\Theta}_i(s)$ take the form:

$$F_i \sim \left[\begin{array}{cc|ccc} \text{diag}[-a_{i(j)}]_{k_i \times k_i} & \mathbf{0}_{k_i \times 2l_i} & I_{k_i \times k_i} & \mathbf{0}_{k_i \times 2l_i} & \mathbf{0}_{k_i \times q_i} \\ \mathbf{0}_{2l_i \times k_i} & \text{diag}\begin{bmatrix} -b_{i(j)} & 1 \\ -c_{i(j)}^2 & 0 \end{bmatrix}_{2l_i \times 2l_i} & \mathbf{0}_{2l_i \times k_i} & I_{2l_i \times 2l_i} & \mathbf{0}_{2l_i \times q_i} \\ \hline \mathbf{y_{1i}} & \mathbf{y_{4i}} & \mathbf{0}_{q_i \times k_i} & \mathbf{0}_{q_i \times 2l_i} & I_{q_i \times q_i} \end{array}\right] \quad (7)$$

$$\tilde{\Theta}_i = \left[\, \mathbf{u_{1i}} \;\; \mathbf{u_{4i}} \;\; D_i^H \,\right]^H . \quad (8)$$

The AW compensator is parameterised in the same form as that for the low order design method discussed in the previous section (Turner et al. (2004, 2007)). Hence the optimisation problem can be posed in exactly the same way, except for a minor change in the dimensions of the state-space matrices in the LMI (4). Due to the change in the structure of $F_i(s)$ the design procedure is slightly different and is given below.

Procedure 2: Gain and zero optimisation

(1) Choose $\varepsilon \in (0,1)$.
(2) Choose the poles using $\{a_{i(j)}, b_{i(j)}, c_{i(j)}\}$, and output pole directions $\mathbf{y_{1i}}$, $\mathbf{y_{2i}} = \mathbf{y_{3i}}$ of $F_1(s)$ and $F_2(s)$.
(3) Choose weighting matrices $W_p > 0$ and $W_r > 0$.
(4) Minimise γ subject to the LMI in equation (4).
(5) Form $\tilde{\Theta} = [\tilde{\Theta}_1' \quad \tilde{\Theta}_2']' = LU^{-1}$.
(6) Form $\Theta(s)$ according to equations (6)-(8).

Remark 2: While the gains and zeros of the AW compensator can be chosen optimally, the method still has some limitations.

$$\begin{bmatrix} Q\bar{A}'+\bar{A}Q & B_0U+\bar{B}L-Q\bar{C}_1'\varepsilon & 0 & Q\bar{C}_2' & -Q\bar{C}_1' \\ \star & -2U-\varepsilon(D_{01}U+\bar{D}_1L)+(UD_{01}'+L'\bar{D}_1')\varepsilon & \varepsilon I & UD_{02}'+L'\bar{D}_2' & -U\tilde{\Gamma}'-L'\bar{D}_1' \\ \star & \star & -\gamma I & 0 & I \\ \star & \star & \star & -\gamma W_p^{-1} & 0 \\ \star & \star & \star & \star & -\gamma W_r^{-1} \end{bmatrix} < 0 \tag{4}$$

In particular, the optimisation algorithm can only choose the right singular value directions of the poles (input pole directions) and not the left singular value directions (output pole directions), which are fixed by the choice of $\mathbf{y_{1i}}$ and $\mathbf{y_{2i}}$. Optimisation of the input and output directions is only possible when the multiplicity of the pole is at least equal to the dimension of the output spaces (ie $\max\{m,p\}$). In this case, one can choose the set of output directions for the repeated pole such that they span the output space. Of course in SISO AW controller design with $m = p = 1$, no directional limitations arise. Similar arguments can be made in terms of the zero directions, although the limitations are likely to be less, as the zeros are determined by the choice of the poles, their output and input directions, and the gain matrix, the latter two of which are chosen optimally. $\circ$

3.3 Classical design

A useful AW approach for SISO problems is to design the AW compensator using classical loopshaping methods, in which the AW dynamics are chosen based on the frequency response of the linear part of the system and exclusion regions in the Nichols chart (NC) (Kerr et al. (2007)). These exclusion regions capture stability constraints (absolute stability, describing function) and linear performance constraints, and guide the shaping of the AW compensator. This can also be done robustly, by enforcing the constraints for a finite set of discrete plant cases $\{G\}$, as is the case in Quantitative Feedback Theory (QFT). Here we summarise the design approach in Kerr et al. (2007).

The problem considered is that of the design of only $\Theta_1(s)$ (alternative choices can be made). No structure on $\Theta_1(s)$ is a priori defined. The design procedure can enforce a number of constraints and here we consider two. The first is absolute stability via the Popov Criterion, which gives rise to standard exclusion regions in the NC. The second is the enforcement of a lower bound γ_l on the $\mathscr{L}_2$ gain of the nonlinear map from w to z in Figure 2, where again this system represents the mismatch system and consequently $w = u_{lin}$ and $z = z_d$. The lower bound is calculated by treating the nonlinearity $\phi(\cdot)$ in Sector$[0, \varepsilon I]$ as an uncertain linear gain in $[0, \varepsilon I]$. While this does not ensure that the true gain of the nonlinear map is bounded by some specified level, enforcement of absolute stability does ensure that all maps are bounded. Initial experience has also shown that this optimistic lower bound is a useful indication of the relative performance of AW compensators, which is known to be a potential problem with the upper bound (Grimm et al. (2004)). As the map is linear, it also gives a frequency dependent gain, which aids loopshaping.

The design procedure is summarised below. The algorithms for the calculation of the exclusion regions for absolute stability and performance are slight variations on those standard to QFT and are detailed in Kerr et al. (2007). The bounds in Figure 6 provide an indication of possible region topologies.

Procedure 3: Classical loopshaping

(1) Choose the finite discrete plant family $\{G\}$ and a finite discrete set of design frequencies Ω.
(2) Choose $\varepsilon \in (0,1)$. This dictates the size of sector for which stability and performance are enforced.
(3) Calculate exclusion regions at each $\omega \in \Omega$ and for each $G \in \{G\}$ for absolute stability via the Popov Criterion.
(4) Calculate exclusion regions at each $\omega \in \Omega$ and for each $G \in \{G\}$ for the map from $w = u_{lin}$ and $z = z_d$, with $\phi(\cdot)$ replaced by $[0, \varepsilon I]$, to be less than γ_l.
(5) Find the union of all the exclusion regions at each $\omega \in \Omega$.
(6) Loopshape $\Theta_1(s)$ to satisfy the exclusion regions at each $\omega \in \Omega$ while satisfying any practical design constraints.

Remark 3: This classical loopshaping method does not aim to provide "optimal" AW compensators. Rather, the method aims to provide compensators that balance the need for stability in a reasonably large sector (large ε) and good recovery of linear response (small γ_l), with the desire for low order compensators with dynamics that are suitable for implementation. As there is no constraint on the structure of the AW compensator, feasible compensators of any order are readily designed and if one cannot be found, the exclusion regions provide immediate insight into any conflicting design properties; ie stability versus performance, or the need for higher order compensation. $\circ$

4. THE AIRCRAFT MODEL AND RATE-SATURATION

An important saturation problem in flight control is that of rate-saturation in an aircraft's control surfaces. Over recent years, these rate-saturation problems have become associated with so-called pilot-induced-oscillations (PIOs) which are characterised by undesirable oscillatory behaviour due to aircraft-pilot-controller interactions (NRC (1997)). They have been responsible for several accidents, most notably the SAAB Grippen crashes, and have become the subject of much research both in the control and handling qualities communities.

4.1 The PIO problem

Figure 3 shows the nominal pilot vehicle system (PVS) under consideration. $G_i(s)$ denotes the aircraft dynamics and $K_i(s)$ denotes the nominal controller. Between these two elements is the rate-limit nonlinearity which, as is often the case, has been modelled as a set of first order systems interconnected with a saturation nonlinearity, which models the limits on the actuator rates. We assume no magnitude saturation nonlinearity is present because the strict rate limits considered typically prevent magnitude limits from being activated and rate limits are the dominant nonlinearity in PIO occurrences. In Figure 3, y denotes the vector of measurements, including those which are observed by the pilot, K_p, who forms an outer control loop. Although at first the configuration in Figure 3 seems somewhat different than the AW problem discussed earlier, by making suitable definitions as indicated in Figure 3, the problem can be cast as a standard AW problem with magnitude saturation by defining $K(s)$ and $G(s)$ as follows

$$d = \underbrace{H[K_{i,1}Kp \quad K_{i,2} - K_{i,1}K_p \quad -I]}_{K(s)} \begin{bmatrix} r \\ y \\ x_{rm} \end{bmatrix}, \tag{9}$$

$$\begin{bmatrix} y \\ x_{rm} \end{bmatrix} = \underbrace{\begin{bmatrix} G_{i,1} & G_{i,2}/s \\ 0 & I/s \end{bmatrix}}_{G(s)} \begin{bmatrix} d_o \\ \hat{d} \end{bmatrix}. \tag{10}$$

In this configuration, the AW compensator, $\Theta(s)$ is driven by the signal $\tilde{d} = \mathrm{Dz}(d)$, which is the difference between the ideal and actual control signal rates - a signal internal to the rate-limit. Although this appears unrealistic, in fact this is quite possible since software rate-limits are normally placed before physical rate-limits to ensure that the latter are never exceeded.

In Figure 3, the pilot is modelled by a linear transfer function K_p, which is normally taken as either a simple gain (Duda

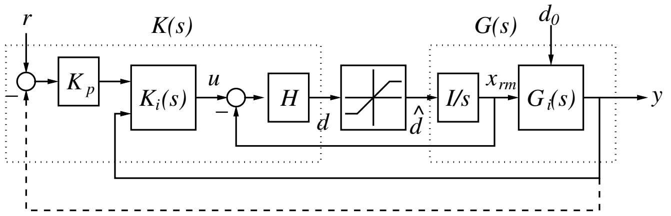

Fig. 3. Structure of the nominal PVS

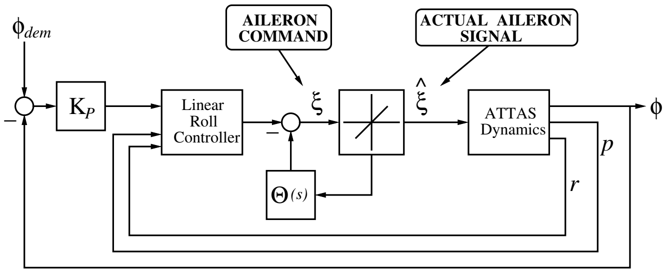

Fig. 4. Structure of the nominal PVS.

(1997)) or a gain plus time delay and with lead-lag characteristics (Neale Smith). Although a gross approximation of real human behaviour, such models do provide a useful indication of approximate pilot behaviour during a closed-loop task.

4.2 The ATTAS aircraft

The three low-order AW design techniques are compared against one another using a high fidelity nonlinear simulation of the ATTAS aircraft. The ATTAS has been the subject of previous PIO research and good summaries of this can be found in Brieger et al. (2007); Sofrony et al. (2006). As in these papers, here the rate-limits are reduced to 50 % (12.5 deg/s) of their nominal value in order to induce higher levels of saturation and hence more oscillatory behaviour; the actuator bandwidth is set at $H = 50$ rad/s. The pilot is assumed to react mainly to attitude errors between a "high-level" attitude demand, ϕ_d and the real roll attitude, ϕ. Based on the OLOP criterion (Duda (1997)), the pilot is modelled as a pure gain; $K_p = 1.2$. Only the roll control loop is considered, partly because this reduces the AW problem to essentially a SISO problem, and also because the flight testing reported in Brieger et al. (2007) only considered roll manoeuvres for structural and safety reasons. The roll control loop structure is shown in Figure 4. The nominal linear controller is effectively a static rate-command-attitude-hold controller with roll rate (p) and yaw rate (r) used for feedback and the controller output being aileron command $\xi \in \mathbb{R}$. Hence, without loss of generality, $\Theta_2(s)$ can be set to zero and design of only $\Theta_1(s)$ considered, i.e. $\Theta(s) = \Theta_1(s)$. In the designs to follow, the ATTAS aircraft was trimmed at Mach 0.5, 20000 feet, which gives a statically and dynamically stable aircraft. The corresponding linear model, including the engine dynamics and rate limit model, is of order 27.

5. RESULTS

5.1 LMI-based designs

The low order LMI designs reported here were designed by first performing a full order Riccati design as described in Brieger et al. (2007); choosing poles close to those considered "dominant" in the full order design; and then using these as described in Procedures 1 or 2. All LMI-based designs used $\varepsilon = 0.97$. Initially, no zeros were included in the design and Procedure 1 produced the following compensator

$$\Theta_1^{(1)}(s) = \frac{0.131}{(s+3)(s+2.9)}, \quad \gamma = 8.07. \tag{11}$$

These same poles were then used in Procedure 2, in which the zeros were optimally assigned (under the constraint of strict properness) which produced the compensator

$$\Theta_1^{(2)}(s) = \frac{-0.045(s-1.008)}{(s+3)(s+2.9)}, \quad \gamma = 3.54. \tag{12}$$

As expected, the $\mathscr{L}_2$ gain bound for $\Theta_1^{(2)}(s)$ is lower than that of $\Theta_1^{(1)}(s)$. Interestingly, the zero in the second compensator is chosen by the LMI-optimisation to be nonminimum phase. This agrees with arguments in Kerr et al. (2007).

Figure 5 shows the roll attitude response of the aircraft to a 3-2-1-1 reference roll attitude demand. Four different cases are shown: (i) when rate-limits are in place but no AW is used; (ii) when rate-limits are in place and AW compensator $\Theta_1^{(1)}$ is used; (iii) when rate-limits are in place and AW compensator $\Theta_1^{(2)}$ is used; and (iv) the artificial case when the system has no rate-limits. Observe that the system with no AW becomes unstable and enters a limit cycle. With $\Theta_1^{(1)}$ (no zeros) behaviour is greatly improved, but perhaps the best performance is obtained by $\Theta_1^{(2)}$ (optimised zeros) which tracks better the ideal no constraint response and settles down faster than $\Theta_1^{(1)}$.

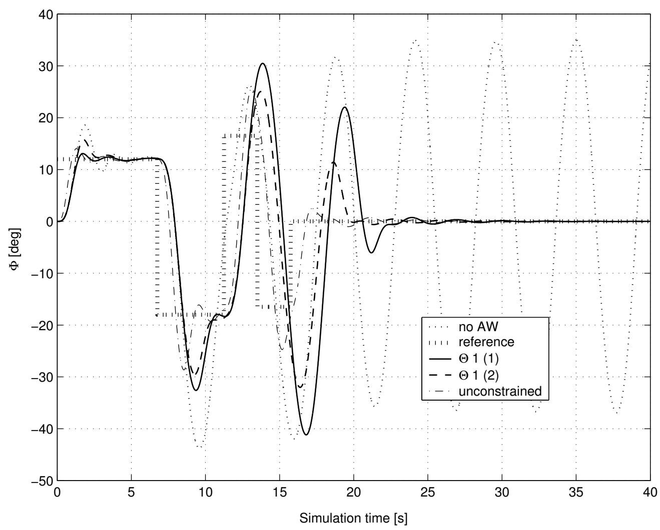

Fig. 5. LMI-based compensator responses

5.2 Classical-based designs

The low order classical design reported here was designed based on Procedure 3. The plant family was the single plant case representing the trimmed aircraft and gain pilot model. The design frequency set Ω contained 20 frequencies, ranging from $1e-5$ to 50 rad/s. For stability enforcement (step (3)) $\varepsilon = 0.999$ was chosen, with the Popov multiplier chosen to be $(1+100s)$. For performance enforcement (step (4)) $\varepsilon = 0.97$ was chosen. The bounds dictating the exclusion regions in the NC are shown at a selection of frequencies in Figure 6, along with the frequency response of the classically designed compensator $\Theta_1^{(3)}(s)$, which has transfer function

$$\Theta_1^{(3)}(s) = \frac{-0.180(s-1.2)}{(s+18)(s+1)}, \quad \gamma_1 = 1.13. \tag{13}$$

The poles of this compensator were then used to initialise a design based on Procedure 2 to optimise the gain and zero location for minimisation of γ. This produced a compensator with transfer function

$$\Theta_1^{(4)}(s) = \frac{-0.186(s-0.656)}{(s+18)(s+1)}, \quad \gamma = 2.76. \tag{14}$$

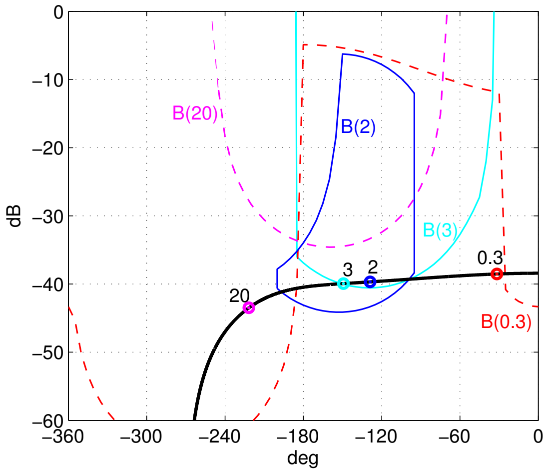

Fig. 6. Classical design bounds and frequency response of $\Theta_1^{(3)}(s)$. (Solid - lower bound; Dashed - upper bound; Solid closed - enclosed feasible region)

Note that the $\mathscr{L}_2$ gain *upper* bound for this 4th compensator is the lowest obtained here (using the Circle multiplier) and the transfer function of this compensator is remarkably similar to the classical design; it appears that the gain and zero location obtained through a logical classical procedure are, given the same poles, mimicked by the zero optimisation procedure.

Figure 7 shows the roll attitude response of these two designs for the same 3-2-1-1 reference demand. As expected, the two AW compensators preserve the stability and performance of the system. The responses obtained using $\Theta_1^{(3)}$ and $\Theta_1^{(4)}$ are very similar, which is not surprising given their similar transfer functions.

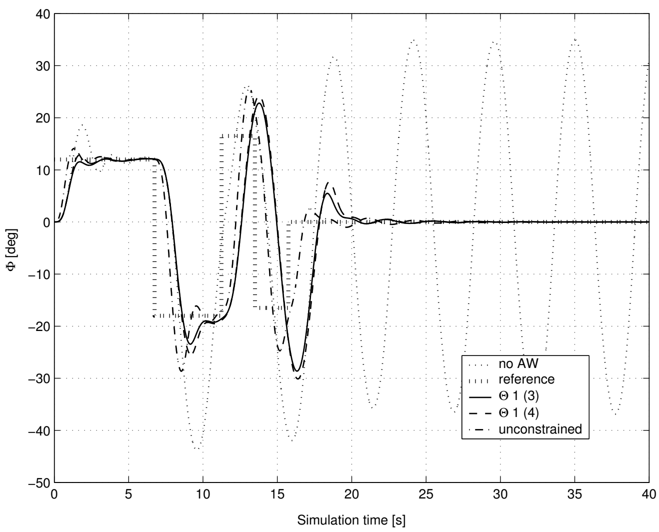

Fig. 7. Classical-based compensator responses

6. CONCLUSIONS

Three low-order AW compensator designs were presented in this work: one based on the optimisation of a given compensator's gains; one based on a classical loopshaping method; and one based on a new method for optimisation of a given compensator's zeros and gains. All three were successfully applied to a realistic flight control problem where aileron rate saturation severely degrades the piloted aircraft system stability and response. As expected, the AW compensator based on the new method for optimisation of zeros and gains performed favourably compared to that based on the existing gain optimisation method, in terms of both the magnitude of the $\mathscr{L}_2$ bound and the time domain response.

Somewhat unexpectedly, the classical loopshaping approach led to the design of a compensator that performed similarly to a compensator designed via the zero and gain optimisation method for the same pole set. This is desirable, as the classical method may provide an indication of good pole and/or zero locations which can serve as a basis for further optimisation. As was seen here, these pole and zero locations can be quite different to those indicated by the full order design and may lead to superior designs. Although not conclusive, this suggests that designs based on ideas from both the optimisation and classical methods may result in the most beneficial AW compensators.

REFERENCES

J-M. Biannic, C. Roos, and S. Tarbouriech. A practial method for fixed-order anti-windup design. *Proc. NOLCOS.*, 2007.

O. Brieger, M. Kerr, D. Leissling, I. Postlethwaite, J. Sofrony, and M.C. Turner. Anti-windup compensation of rate saturation in an experimental aircraft. In *Proc. American Control Conference*, New York, 2007.

H. Duda. Prediction of pilot-in-the-loop osciallations due to rate saturation. *Journal of Guidance, Control and Dynamics*, 20(3):581–587, 1997.

S. Galeani, M. Massimetti, A. R. Teel, and L. Zaccarian. Reduced order linear anti-windup augmentation for stable linear systems. *International Journal of Systems Science*, 37:115–127, 2006.

J.M. Gomes da Silva Jr., S. Tarbouriech, and G. Garcia. Local stabilisation of linear systems under amplitude and rate saturating actuators. *IEEE Transactions on Automatic Control*, 48(5):842–847, 2003.

G. Grimm, J. Hatfield, I. Postlethwaite, A.R. Teel, M.C. Turner, and L. Zaccarian. Anti-windup for stable linear systems with input saturation: an LMI based synthesis. *IEEE Transactions on Automatic Control*, 48(9):1509–1525, 2003.

G. Grimm, A. Teel, and L. Zaccarian. Robust linear anti-windup synthesis for recovery of unconstrained performance. *Int. J. of Robust Nonlinear Control*, 14:1133–1168, 2004.

R. Hanus, M. Kinnaert, and J.L. Henrotte. Conditioning technique, a general anti-windup and bumpless transfer method. *Automatica*, 23(6):729–39, 1987.

H. Hindi and S. Boyd. Analysis of linear systems with saturation using convex optimisation. *Proc. of the American Control Conference*, 1998.

M. Kerr, E. Villota, and S. Jayasuriya. Robust anti-windup design for input constrained SISO systems. *Proc. 8th Int. Symposium on QFT and Robust Frequency Domain Methods.*, 2007.

E.F. Mulder, M.V. Kothare, and M. Morari. Multivariable anti-windup controller synthesis using linear matrix inequalities. *Automatica*, 37(9):1407–1416, 2001.

(National Research Council) NRC. *Aviation Safety And Pilot Control: Understanding and Preventing Unfavorable Pilot-Vehicle Interactions.* National Academy Press, 1997.

I. Queinnec, S. Tarbouriech, and G. Garcia. Anti-windup design for aircraft flight control. *IEEE Conference on Control Applications*, 2006.

L. Rundquist and K. Stahl-Gunnarsson. Phase compensation of rate-limiters in unstable aircraft. *Proc. IEEE Conference on Control Applications*, 1996.

J. Sofrony, M.C. Turner, I. Postlethwaite, O.M. Brieger, and D. Leissling. Anti-windup synthesis for PIO aviodance in an experimental aircraft. *IEEE Conference on Decision and Control*, 2006.

A.R. Teel and N. Kapoor. The $\mathscr{L}_2$ anti-windup problem: Its definition and solution. *Proc. European Cont. Conf.*, 1997.

M.C. Turner, G. Herrmann, and I. Postlethwaite. Accounting for robustness in anti-windup synthesis. In *Proc. American Control Conference*, Boston, 2004.

M.C. Turner, G. Herrmann, and I. Postlethwaite. Incorporating robustness requirements into anti-windup design. *IEEE Transactions on Automatic Control*, 52(10):1842–1855, 2007.

M.C. Turner and I. Postlethwaite. A new perspective on static and low order anti-windup synthesis. *International Journal of Control*, 77(1):27–44, 2004.

E. Villota, M. Kerr, and S. Jayasuriya. A study of configurations for anti-windup control of uncertain systems. *IEEE Conference on Decision and Control*, 2006.

P.F. Weston and I Postlethwaite. Linear conditioning for systems containing saturating actuators. *Automatica*, 36(9):1347–1354, 2000.