

Flight testing of low-order anti-windup compensators for improved handling and PIO suppression

Oliver Brieger, Murray Kerr, Ian Postlethwaite, Jorge Sofrony, Matthew C. Turner

Abstract—This paper presents the results of recent flight tests of several anti-windup (AW) compensators on the German Aerospace Centre’s (DLR) Advanced Technologies Testing Aircraft (ATTAS). The objectives of the tests were twofold: to demonstrate the potential for rigorously designed low order AW compensators to reduce the pilot-involved-oscillation (PIO) proneness of the aircraft and improve the handling qualities; and to compare a variety of low-order AW compensators to determine the importance of different design parameters. The AW compensators were assessed based on pilot handling qualities ratings (HQRs) and PIO ratings (PIORs). These ratings, and supporting pilot comments and flight data, demonstrate that the AW compensators improved the handling qualities and reduced the PIO proneness of the aircraft, albeit to different degrees. The results also provide a basic understanding of the relationship between design parameters and the response of the piloted aircraft during periods of rate saturation.

I. INTRODUCTION

The problem of actuator saturation is well documented in the aeronautical field. In particular, actuator rate constraints have long been linked with performance and stability degradation and are known to be a leading cause of pilot-involved-oscillations (PIOs) in fly-by-wire aircraft [13], [5]. Subsequently, the issue of PIO proneness of an aircraft, including the specific effect of rate saturation, has now become an important design consideration.

This research aims to determine the ability of additional saturation compensators to reduce the effects of rate saturation, in particular those designed via the anti-windup (AW) philosophy. Previous work applying AW to aircraft systems includes [15], [10], [14], [17], [8], [4], [3], [2]. The work presented herein builds on work of [17] and the recent “SAIFE” (Saturation Alleviation In-Flight Experiment) flight tests reported in [4], [3]. The SAIFE tests demonstrated that AW control theory could be employed to design compensators systematically to reduce PIO tendencies and improve handling qualities. In some cases, pilot handling qualities rating (HQR) and PIO rating (PIOR) reductions of up to two relative to no AW compensation were achieved in these tests.

This paper describes a second set of flight tests (“SAIFE II”) conducted in August 2007. The objective of the SAIFE II tests was to demonstrate the potential for low order, robustly designed AW compensators for reducing PIO proneness of aircraft and for improving handling qualities. A further objective was to compare a variety of low-order AW design

O. Brieger is with German Aerospace Centre, Institute of Flight Systems, Braunschweig, Germany. Email oliver.brieger@dlr.de

M. Kerr is with Deimos Space S.L., Madrid, Spain. Email murray.kerr@deimos-space.com

I. Postlethwaite and M.C. Turner are with the Department of Engineering, University of Leicester, UK. Email {ixp, mct6}@le.ac.uk

J. Sofrony is with the Department of Mechatronics, Universidad Nacional de Colombia, Colombia. Email jsofronye@unal.edu.co

Research supported by the UK EPSRC.

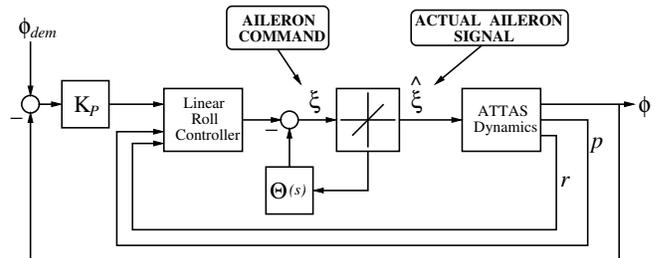


Fig. 1. Structure of the PVS.

methods and compensators to determine the importance of different design parameters. As in SAIFE I [4], [3], the flight tests were conducted on the Advanced Technologies Testing Aircraft (ATTAS) at the German Aerospace Centre (DLR), Braunschweig, Germany, and the AW compensators were assessed using HQRs and PIORs.

II. PROBLEM STATEMENT

The problem considered is that of alleviating the effects of rate saturation on the ATTAS aircraft depicted in Figure 1. This figure shows the structure of the roll-axis pilot-vehicle-system (PVS). The ATTAS dynamics capture the aircraft roll response to rate limited aileron actuator commands $\hat{\xi}$, with the linear roll controller being a rate feedback controller, acting on the roll and yaw rates (p and r respectively). The pilot is assumed to react mainly to attitude errors between a high-level attitude demand, ϕ_{dem} , and the real roll attitude, ϕ . The input constraint in the PVS comes from the rate limiter, which acts on the signal ξ , limiting its rate to be below a threshold value. The problem considered here is that of designing an additional compensator Θ , the AW compensator, that acts on the effect of the rate limiter, to modify favourably the aileron command from the linear roll controller, which is assumed fixed.

A. AW Philosophy

The standard problem in AW control is that of magnitude saturation, with extensive theory and design tools developed for this problem - see [18], [19], [21], [7], [9] and references therein. At first, due to the presence of rate limiting, the problem depicted in Figure 1 appears significantly different to that in AW control. However, by making suitable definitions as indicated in Figure 2, the problem depicted in Figure 1 can be formalised in a way that casts it as a magnitude saturation problem and hence in the standard form for AW control. Figure 2 again shows the nominal PVS under consideration. Here $G_i(s)$ denotes the aircraft dynamics and $K_i(s)$ denotes the nominal controller (linear roll controller). Between these two elements is the rate-limit

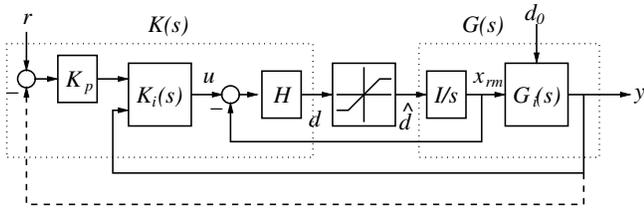


Fig. 2. Structure of the nominal PVS



Fig. 3. ATTAS testbed

nonlinearity which has been modelled as a first order system interconnected with a saturation nonlinearity, which models the limits on the actuator rates. In Figure 2, y denotes the vector of measurements, including those which are observed by the pilot, K_p , who forms an outer control loop. The linear control elements can then be captured by a redefined effective controller $K(s)$ and plant $G(s)$ as follows

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

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

In this configuration, the AW compensator, $\Theta(s)$ is driven by the signal $\hat{d} = Dz(d)$, which is the difference between the ideal and actual control signal rates - a signal internal to the rate-limit. Although access to this signal may appear unrealistic, it is possible, since software rate-limits are normally placed before physical rate-limits. See [16], [4], [11] for more details.

B. ATTAS Aircraft

The ATTAS aircraft used for the flight testing of the AW controllers is a highly modified VFW 614 aircraft operated by DLR (see Figure 3). It features various customized systems, including an adaptive fly-by-wire flight control system capable of hosting different controller designs and a mechanical back-up control system monitored by a safety pilot, which enables the realistic assessment of a flight control system by an evaluation pilot without having to meet extensive certification requirements. The AW compensators were designed based on the ATTAS VFW 614 dynamics and scheduled control laws.

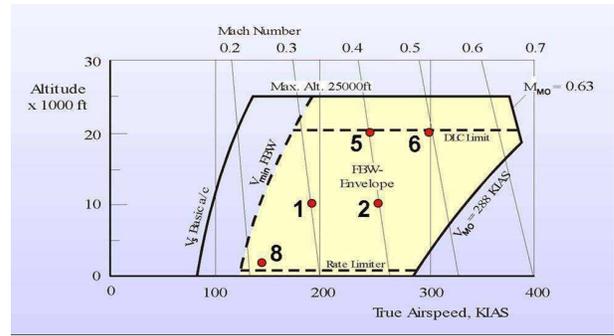


Fig. 4. Test points within the ATTAS envelope: (1) 10000 ft, Ma 0.3 / (2) 10000 ft, Ma 0.4 / (5) 20000 ft, Ma 0.4 / (6) 20000 ft, Ma 0.5 / (8) pattern altitude, 135 KEAS.

C. Flight Test Programme

To demonstrate the effectiveness of AW compensation, the SAIFE II campaign considered two up-and-away flight condition (FCs). However, for reasons of space this paper will concentrate on FC6, corresponding to a trim point of Mach 0.5 and 20,000 feet (see Figure 4). The testing focused purely on the roll axis, since the ATTAS is comparatively agile in roll, with the AW compensators designed to reduce the effects of rate limiting on this axis only. Two evaluation tasks were conducted: the Handling Qualities During Tracking (HQDT) task and target tracking (birdy) task. The HQDT technique is designed to evaluate PIO susceptibility and involves successive bank angle captures where the pilot flies in a manner designed to expose PIO tendencies [3], [1]. The tracking task requires the pilot to track closely a generic birdy target (aircraft symbol) projected onto the main head down display. The birdy performed a predefined sequence of ramp and step-type roll attitude changes (dashed line in Figure 6), requiring the pilot to perform and assess gross acquisition and fine tracking. To highlight the effects of rate saturation during these evaluation tasks, the software imposed aileron rate limits were reduced to 50% of the full authority values. More details can be found in [3].

The testing at FC6 involved the evaluation of six AW compensators. Five of these compensators were low-order compensators that were designed to be robust and work over the whole up-and-away set of flight conditions (FCs 1, 2, 5 and 6); see Sections III and IV for details of their design. The remaining compensator was one of those tested in the previous SAIFE I flight test campaign, and was specifically designed for the aircraft trimmed at FC6. The testing of the AW compensators at FC6 was performed in three sets. Each set included two AW compensators and the case of no AW as reference. In each set, the pilots always evaluated the no AW case first and were aware of this. This was requested by the pilots, and, as the primary objective was to compare compensators rather than assess them against no AW, this was not considered to be a problem.

III. LOW-ORDER AW DESIGN METHODS

The theory and design tools for AW control have improved greatly in the last decade and now AW controllers can be rigorously and systematically designed using modern tools, such as LMIs. Despite these advances, few of these modern designs have made it into application. This may be partly

due to their theoretical advantages being accompanied by certain practical drawbacks. One of these drawbacks is that the resulting AW compensators tend to have large state-dimensions. Given a plant $G(s)$ with state dimension n_p , a compensator with the same state dimension is often produced. It is normally undesirable to implement a control law which requires an extra n_p states just to handle saturation.

Another disadvantage with these modern “optimal” compensators is that their dynamics may be unnecessarily complex and could actually induce poor time domain behaviour. Consequently there has been a keen interest in low order compensators (those with order lower than n_p) which retain the advantages of modern full order compensators, but which are also low order and hence readily implementable. Unfortunately, the AW literature contains few systematic low-order AW compensator design procedures. However, recently this has changed, and several design methods for have been developed ([7], [20], [2], [11]). Furthermore, earlier heuristic methods, notably those based on the Quantitative Feedback Theory (QFT) ([12], [22]) have also been improved.

In this flight test programme, three low-order design methods were considered; two LMI based methods, [20] and its extension in [11], and one classical heuristic method [12]. However, based on ground simulator evaluations, for the flight tests only compensators designed using [20] and [12] were chosen for flight testing. It is these two design methods which are discussed in further detail here. The overall goal of each of these AW design methods can be stated as

Goal 1: To design stable $\Theta(s)$ such that

- 1) The system in Figure 2 is globally asymptotically stable for the deadzone in the Sector $[0, \varepsilon I]$, $\varepsilon \in (0, 1)$.
- 2) $\deg(\Theta(s)) = 1$ or 2.
- 3) During periods of saturation, the deviation of the system response from that of the ideal unsaturated linear response is small.

The first condition ensures the system is stable for $|u| \leq \bar{u}/(1 - \varepsilon)$, where \bar{u} is the actuator saturation level. The second condition ensures that all the compensators are of first or second order, which is sufficiently low to be practically implementable. The third condition captures the performance objective which has become standard in AW control (See [20],[18],[21] for further details).

A. Standard LMI Based Low Order Design

For the control problem considered herein, the low-order method proposed in [20] involves the partitioning of the AW controller into two parts

$$\Theta(s) = F(s)\tilde{\Theta}, \quad (3)$$

where $F(s)$ is a stable transfer function matrix chosen by the designer and $\tilde{\Theta}$ is a gain matrix optimally synthesised by an LMI optimisation procedure. Optimality is with respect to the minimisation of an upper bound γ on the deviation of the saturated response from the ideal linear response. For a full discussion of this method, see [20], [19]. 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 $F(s)$, including the poles and zeros.

- 3) Choose weighting matrices $W_p > 0, W_r > 0$ to trade-off performance and robustness.
- 4) Minimise γ subject to the LMI in eq. (14) of [20].
- 5) Form $\tilde{\Theta} = [\tilde{\Theta}'_1 \quad \tilde{\Theta}'_2] = LU^{-1}$, with L and U given in step 4.
- 6) Form $\Theta(s)$ according to equation (3).

Note that this low order technique requires the designer to *specify fully* the dynamics, $F(s)$, of the compensator; only the gain matrix, $\tilde{\Theta}$, is synthesised 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 ([20], [19])

B. Classical Loopshaping Based Low Order Design

For the control problem considered herein, the low-order method proposed in [12] involves the loopshaping of the AW compensator $\Theta(s)$ to satisfy stability and performance constraints that are captured as exclusion regions in the Nichols Chart (NC). This can also be done robustly, by enforcing the constraints for a finite set of discrete plant cases $\{G\}$, as in QFT. The design procedure can enforce a number of constraints and here we enforce three. The first is absolute stability via the Popov Criterion, which gives rise to standard exclusion regions in the NC [12]. The second is the enforcement of the OLOP criterion [6]. As the OLOP criterion is given in the NC, this can be done in a straightforward manner. The third is the enforcement of a lower bound γ_l (see [12]) on the deviation of the saturated response from the ideal linear response. The design procedure is summarised below.

Procedure 2: 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; the OLOP criterion; and enforcement of γ_l .
- 4) Find the union of the exclusion regions at each $\omega \in \Omega$.
- 5) Loopshape $\Theta(s)$ to satisfy the regions at each $\omega \in \Omega$.

IV. AW DESIGNS FOR THE ATTAS

A. Pilot model

For control system design and system analysis, the pilot in Figure 2 can be described by a variety of models. In contrast to the work reported in [4], the compensators discussed in this paper (except AWC 7) are designed with the pilot modelled by a linear transfer function K_p . To capture PIO proneness of the PVS, here this is taken as either a simple gain ([6]) or a gain plus time delay with lead-lag characteristics (i.e. Neale Smith model). Although a gross approximation of real human behaviour, such models do provide a useful indication of pilot behaviour during a closed-loop task. As PIO behaviour, by definition, involves “closed-loop” pilot behaviour, it was reasoned that resistance to PIO behaviour would be delivered as a result.

B. Choice of design parameters

As mentioned above, the compensators discussed herein were designed either using a low-order LMI-based method, which requires the fixing of the compensator dynamics

and the optimisation of compensator gains, or a classical loopshaping design procedure. Key parameters which the designer is required to choose in both methods are sector sizes for which the stability and performance results hold, dictated by the parameter ϵ ; and the poles and zeros of the compensators. In addition to affecting the stability and robustness of the compensator designs, the poles and zeros are important in the transient response of the AW compensators. In each case the dynamics were chosen in order to satisfy stability requirements and to be relatively fast, so that the pilot would not perceive the effects of AW compensation on the aircraft response during open-loop tasks. The chosen parameters are listed in Table I. AWC 7 is not shown because it was a full order compensator from SAIFE I [4].

Compensator	Method	ϵ	Poles	Zeros
AWC 1	LMI	0.97	$s = \{-2.9, -3\}$	none
AWC 3	LMI	0.97	$s = \{-0.8\}$	none
AWC 9	Classical	0.999	$s = \{-1.5, -1.6\}$	none
AWC 10	Classical	0.999	$s = \{-18, -1\}$	$s = \{-1.2\}$
AWC 6	Classical	0.999	$s = \{-2.5, -2.55\}$	none

TABLE I
AW COMPENSATOR PARAMETERS

C. Analysis

The compensators were selected on the basis of their performance in time-domain nonlinear simulations (as in [17]), \mathcal{L}_2 gains (either upper bounds or lower bounds - see [11]) and their OLOP plots. In contrast to previous work, more weight was given to the OLOP plots, shown in Figure 5 for all compensators tested. As mentioned above, classical controllers can be explicitly designed to OLOP specifications, while the LMI designs had their OLOP points checked *a posteriori*. Not all compensators had OLOP points below the boundary, although they were all sufficiently far from the inverse describing function to be considered not susceptible to limit-cycle behaviour.

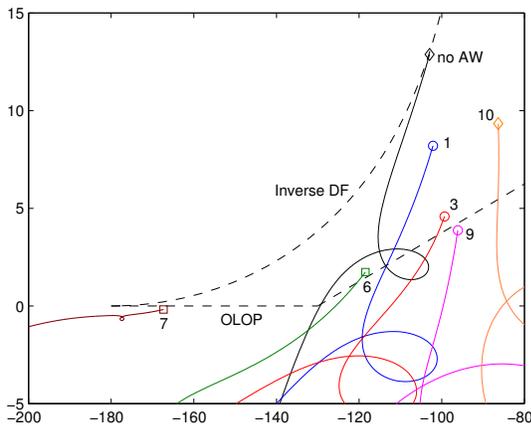


Fig. 5. Frequency Response of the loop transmission around the rate limiter from the onset frequency (0.64 rad/s), and OLOP points. Pilot gain of 1.2. AW compensator number as indicated.

V. FLIGHT TEST RESULTS AND DISCUSSION

The flight test results are presented in three different ways: (i) time-histories depicting the roll axis birdy tracking task;

(ii) RMS measures of roll attitude error and stick activity during birdy the tracking task; and (iii) PIORs and HQRs given by pilots for the HQDT and birdy tracking tasks. While the time histories and RMS measures give some indication of how the AW compensators behaved, they must be treated with some caution because they do not cater for the manner in which the pilot gain (i.e. aggressiveness) changes during flight. Hence the pilot ratings should be given the greatest weight in the analysis of the results.

A. Time Domain Birdy

Figure 6 shows a sample of the aircraft's behaviour during the birdy tracking task, as performed by Pilot 2. The roll attitude response and pilot stick activity are shown for three cases: when no AW is engaged and when AW compensators AWC9 and AWC10 are engaged. Observe that the roll attitude is significantly less oscillatory when either AW compensator is engaged. Also note that with no AW, there are high levels of stick activity which frequently result in the stick being displaced to its limits. When AW is engaged, the stick activity is markedly lower, although it is difficult to identify which compensator performs better.

Note that Figure 6 shows the "best" set of data obtained in the flight tests; i.e. it depicts the set in which the difference between the unaugmented system and the AW-compensated system is most clearly exposed. However, the other birdy tasks conducted by Pilot 2 contained broadly the same features as those presented here, with the AW compensators typically delivering superior tracking and lower stick activity to when no AW was engaged. On the other hand, when Pilot 1 was evaluating the system response, it was more difficult to extract such conclusions from the time domain data alone.

B. RMS errors/stick activity

A rough measure of the success of AW compensation can be extracted from the time domain data by computing (a) the RMS tracking errors for the birdy task and (b) the associated RMS stick activity. These results are tabulated in Table II. The first entry in the table gives the average recorded values for the "no AW" case, as these were recorded in three separate sets. Observe that, when Pilot 2 is evaluating, the aircraft benefits from AW being engaged (with the exception of AWC7): the RMS error is always around 3-4 degrees lower than without AW. There is also less stick activity, with compensators AWC1, AWC9 and AWC10 resulting in about half the stick activity to when no AW is used. It is again more difficult to see an improvement with Pilot I and in fact, these measures appear to indicate that the RMS tracking error is sometimes worse with AW engaged. Note that this data does not account for the pilot's evolving flying technique in which he continually adjusts his "gain" to evaluate the system.

C. PIO and HQ ratings

Table IV shows the PIORs and HQRs awarded by the pilots for the HQDT and birdy tracking tasks. Columns 5-12 contain the ratings awarded for the tasks, as defined by Table III. The last two columns of the table show the total improvement obtained by using a given AW compensator, calculated by summing the points improvement of the four ratings awarded. The table is divided into three sections, each of which represents a test set.

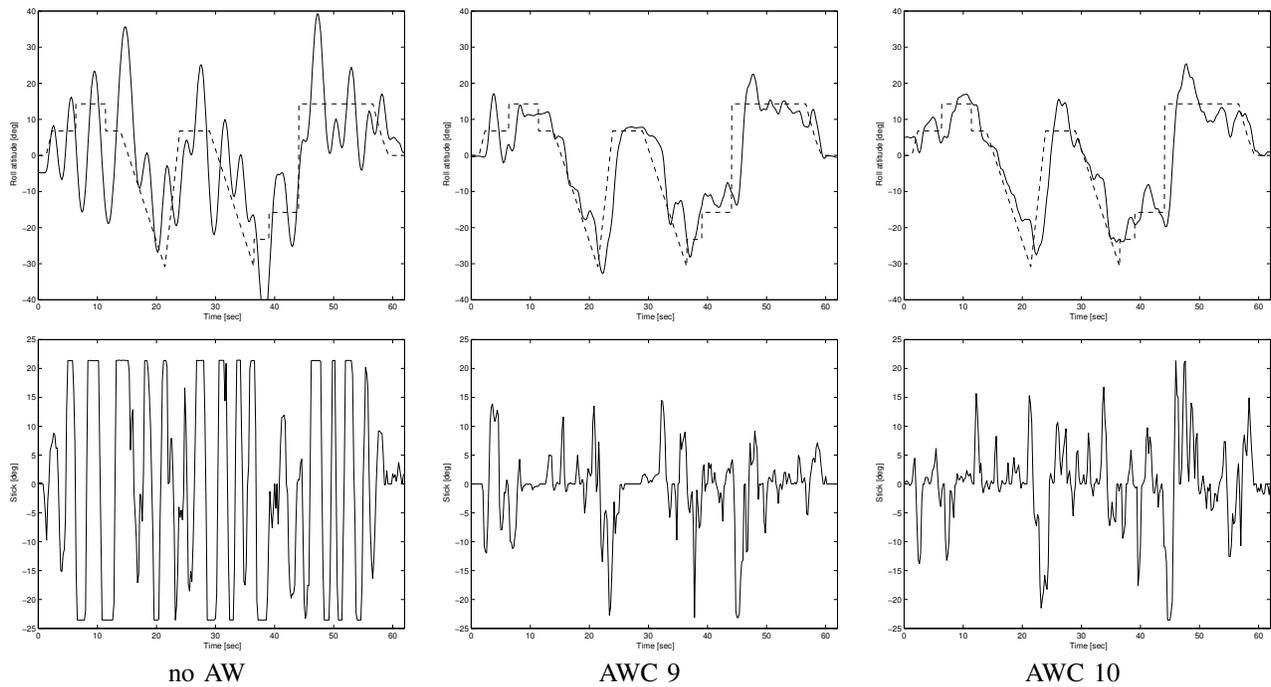


Fig. 6. Birdy tracking task results: Pilot II

Compensator	Pilot I (Q)		Pilot II (Markus)	
	RMS error	RMS stick	RMS error	RMS stick
ave. no AW	7.22	0.19	9.98	0.20
AWC 1	6.21	0.14	5.33	0.08
AWC 3	6.68	0.17	6.85	0.16
AWC 9	8.13	0.21	6.43	0.10
AWC 10	9.45	0.26	6.46	0.11
AWC 6	7.90	0.21	5.76	0.17
AWC 7	7.31	0.21	10.78	0.23

TABLE II
CONDENSED TIME DOMAIN DATA

PIO-c	PIO rating: bank angle capture, HQDT
HQR-g	HQR: gross acquisition, birdy
HQR-f	HQR: fine tracking, birdy
PIO-b	PIO rating: birdy

TABLE III
KEY TO HQR/PIO RATINGS

1) *Set 1 - LMI-based designs*: The first set of compensators consisted of the no AW case, and the LMI designed low-order compensators, AWC1 and AWC3. As with Table II, Table IV suggests that having either compensators AWC1 or AWC3 engaged is preferable to having no AW active. This appears to be the case for both pilots, although it is perhaps most clear for the case of AWC1 and Pilot 2.

2) *Set 2 - Classical-based designs*: The second set of compensators consisted of the no AW case, and the two classically designed compensators, AWC9 and AWC10. Table IV shows strong consistency between the PIORs and HQRs awarded by Pilot 2 and the time domain data depicted in Figure 6 and given in Table II, with substantial performance improvement obtained when either AWC9 or AWC10 was engaged. Comparing the two compensators is difficult with

the ratings indicating a slight preference towards AWC9. In contrast, the time-domain data and the ratings awarded by Pilot 1 are rather inconsistent. For example, Table IV indicates that Pilot 1 preferred AWC9 over no AW, although the RMS values reported in Table II suggest the opposite.

3) *Set 3*: This set was the most difficult to interpret. The compensators assessed were the no AW case, a full-order compensator from the SAIFE I tests (AWC7) and a classically designed compensator (AWC6) which had similar OLOP and high gain characteristics to AWC7. As no “no AW” birdy test was performed by Pilot 1, in Table IV, the improvement offered by the compensators is measured against the “no AW” run of Set 2. Thus, it is best to concentrate on Pilot 2, who rated both compensators similarly, although had a preference for AWC6 according to Table IV, which is consistent with Table II where the RMS error in the birdy tracking task is roughly half that of AWC7. The results corresponding to AWC7 were surprising: in the SAIFE I tests, this compensator had shown major improvements over no AW, but this was not the case here. Although the pilot commented that AWC7 imparted minor improvements to the system, this is not reflected in the time domain data. The reason for the disparity is unclear, but one possibility is that during the SAIFE I tests, both pilots appeared to be flying more aggressively (higher gain), which highlights better the differences between the AW and no AW cases.

D. Overall comparison

An overall comparison between the compensators is difficult to make, as it is difficult for pilots to compare more than three compensators. Time-domain data and pilot ratings suggest that AWC1 and AWC9 are the best overall, with both pilots preferring flying with these compensators active to no AW. Interestingly, AWC9 is rated quite differently by the

No.	Flight Condition		Comp	Pilot 1 (Q)				Pilot 2 (Markus)				Improvement	
	Height [ft]	VIAS [kts]		PIO-c	HQR-g	HQR-f	PIO-b	PIO-c	HQR-g	HQR-f	PIO-b	Pilot 1	Pilot 2
6	20000	224	none	5	6	5	3	5	5	5	4	n/a	n/a
6	20000	224	1	3	5	4	3	4	4	4	2	+4	+5
6	20000	224	3	3	5	4	3	3	5	5	4	+4	+2
6	20000	224	none	5	6	5	4	5	6	6	4	n/a	n/a
6	20000	224	9	4	5	5	3	3	5	4	2	+3	+7
6	20000	224	10	5	6	5	5	4	5	4	2	-1	+6
6	20000	224	none	n/a	n/a	n/a	n/a	5	6	6	5	n/a	n/a
6	20000	224	6	4	5	5	4	4	6	5	4	+2	+3
6	20000	224	7	5	6	6	5	4	6	5	5	-1	+2

TABLE IV

SAIFE 2 FLIGHT TEST RESULTS: PIO RATINGS AND HQR'S AWARDED FOR HQDT AND BIRDY TRACKING TASK

two pilots, with Pilot 2 clearly liking this compensator the most and Pilot 1 giving it a more average rating. Another interesting feature was that although AWC1 was perhaps the best of the two LMI designs, in terms of the performance measures given here, the pilots actually opted for AWC3 when asked to choose their favourite in Set 1.

E. OLOP

Comparing the performance of the AW compensators to their OLOP point locations in Figure 5, it is evident that the OLOP point location roughly correlates with the PIO susceptibility in the tests: all the AW compensators (except perhaps AWC7) generally have improved OLOP points relative to no AW and this is also seen in the PIORs, and AWC3 and AWC9 were the pilots' overall favourites and these compensators have perhaps the best OLOP points, being close to the OLOP boundary and far from the inverse describing function. AWC7 appeared to be the least liked by the pilots and, although its OLOP point does lie slightly below the OLOP boundary, its location is perilously close to the inverse describing function, making it vulnerable to small changes in system gain, which may occur in flight and due to changes in piloting technique.

VI. CONCLUSIONS AND FUTURE WORK

The flight tests reported in this paper evaluated several low-order AW compensators on the ATTAS aircraft. Based on the pilot ratings and supporting flight data, the use of low-order AW compensation was found to be advantageous. The compensators reduced the PIO proneness of the aircraft and improved the handling qualities, sometimes dramatically. Moreover, due to their low order, the compensators were readily implementable and appeared to impart quite predictable behaviour. The compensators were explicitly and implicitly designed for improved OLOP points and this roughly correlated with the reductions in the pilot ratings.

The tests were designed to evaluate the effect of different design parameters and trade-offs. While this was seen in the tests, with pole location, OLOP point, \mathcal{L}_2 gain and sector size all found to be important parameters, a full discussion of these features was not possible in this paper and will be reported elsewhere. Finally, it should be stated that the pilots commented that, aside from some small undesirable motions in open-loop tasks, they definitely preferred flying the aircraft with the better AW compensators engaged.

REFERENCES

- [1] Anonymous. Flying qualities testing. *Chapter 22, US Air Force Test Pilot School, Edwards Air Force Base*, pages 1–11, 2002.
- [2] J-M. Biannic, C. Roos, and S. Tarbouriech. A practical method for fixed-order anti-windup design. *Proc. NOLCOS.*, 2007.
- [3] O. Brieger, M. Kerr, D. Leissling, I. Postlethwaite, J. Sofrony, and M.C. Turner. Flight testing of a rate-compensation scheme on the ATTAS aircraft. *Proc. DGLR Conference*, 2006.
- [4] 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 Cont. Conf.*, New York, 2007.
- [5] National Research Council Committee on the Effects of Aircraft-Pilot Coupling on Flight Safety. *Aviation Safety and Pilot Control: Understanding and Preventing Unfavorable Pilot-Vehicle Interactions*. The National Academies Press, 1997.
- [6] H. Duda. Prediction of pilot-in-the-loop oscillations due to rate saturation. *J. of Guid., Con. and Dyn.*, 20(3):581–587, 1997.
- [7] S. Galeani, M. Massimetti, A. R. Teel, and L. Zaccarian. Reduced order linear anti-windup augmentation for stable linear systems. *Int. J. of Systems Science*, 37:115–127, 2006.
- [8] S.L. Gatley, M.C. Turner, I. Postlethwaite, and A. Kumar. A comparison of rate-limit compensation schemes for pilot-induced-oscillation avoidance. *Aerospace Science and Technology*, 10:37–47, 2006.
- [9] 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.
- [10] G. Hovmark and G. Duus. Experimental evaluation of phase compensating rate limiters in an aircraft's lateral flight control system. Technical Report TP-120-04, GARTEUR, August 2000.
- [11] M. Kerr, M.C. Turner, and I. Postlethwaite. Practical approaches to low-order anti-windup compensator design: a flight control comparison. *IFAC World Congress*, 2008, to appear.
- [12] 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.
- [13] D.H. Klyde, D.T. McRuer, and T.T. Myers. Pilot-induced oscillation analysis and prediction with actuator rate limiting. *Journal of Guidance, Control and Dynamics*, 20(1):81–89, 1997.
- [14] I. Queinnec, S. Tarbouriech, and G. Garcia. Anti-windup design for aircraft flight control. *IEEE Conf. on Control Applications*, 2006.
- [15] L. Rundquist and K. Stahl-Gunnarsson. Phase compensation of rate-limiters in unstable aircraft. *Proc. IEEE Conf. on Con. App.*, 1996.
- [16] J. Sofrony, M.C. Turner, and I. Postlethwaite. Anti-windup synthesis for systems with rate-limits: a Riccati equation approach. *Proc. SICE Annual Conf.*, 2005.
- [17] J. Sofrony, M.C. Turner, I. Postlethwaite, O.M. Brieger, and D. Leissling. Anti-windup synthesis for PIO avoidance in an experimental aircraft. *IEEE Conf. on Dec. and Con.*, 2006.
- [18] A.R. Teel and N. Kapoor. The \mathcal{L}_2 anti-windup problem: Its definition and solution. *Proc. European Cont. Conf.*, 1997.
- [19] M.C. Turner, G. Herrmann, and I. Postlethwaite. Incorporating robustness requirements into anti-windup design. *IEEE Trans. Aut. Control*, 52(10):1842–1855, 2007.
- [20] 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.
- [21] E. Villota, M. Kerr, and S. Jayasuriya. A study of configurations for anti-windup control of uncertain systems. *IEEE Conf. on Dec. and Con.*, 2006.
- [22] W. Wu and S. Jayasuriya. A new QFT design methodology for feedback systems under input saturation. *J. Dyn. Sys., Meas. and Cont.*, 123:225–23, 2001.