Taking robust LPV control into flight on the VAAC Harrier[*]

George Papageorgiou Keith Glover[†]
Department of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, England
tel: +44-1223-332752 fax: +44-1223-332662

Glenn D'Mello Yoge Patel
Flight Management and Control Department, DERA, Bedford MK41 6AE, England

Abstract

This paper describes the application of a new multivariable control system design procedure, linear parameter-varying (LPV) $\mathcal{H}_\infty$ loop shaping, to design a flight controller for the pitch dynamics of the Vectored thrust Aircraft Advanced flight Control (VAAC) Harrier. The particular LPV design procedure used was developed by Papageorgiou and Glover and is an advancement of the $\mathcal{H}_\infty$ design method successfully flight tested in 1993 on the VAAC Harrier in that the control laws are automatically scheduled over the flight envelope. Automatically scheduled control laws are a strong and very likely candidate for the next generation of flight control systems (FCS). The flight test of the LPV controller described here is the first of its kind in the international aerospace community and is recognised as a key milestone in the general development of next generation FCS design methods.

Keywords: Flight control, LPV models, $\mathcal{H}_\infty$ loop shaping, Robust gain scheduling, VAAC Harrier.

1 Introduction

The aim of this paper is to demonstrate that the design procedure proposed by Papageorgiou and Glover in [13] is suitable for designing the flight controllers of the next generation fighter aircraft. This is done by using the proposed design procedure, that is an advancement of the $\mathcal{H}_\infty$ design method successfully flight tested in 1993 (see [7]) in that the control laws are automatically scheduled over the flight envelope, to design a flight controller for the short period dynamics of the VAAC Harrier. The performance of the resulting control law is evaluated using piloted simulations and flight testing.

There are a number of papers that deal with the design of an automatically gain scheduled controller for an aircraft, e.g. see [8, 1, 3, 2, 10]. The report on which this paper is based though (see [12]) is by far the most comprehensive case study to date. Furthermore, flight testing an automatically gain scheduled controller is a world first.

A design procedure is not only judged by the performance of the controllers that it generates. Whether a design procedure is suitable or not for designing flight controllers depends on a number of factors. For example, the structure of the flight controller must have visibility in terms of its functionality. That is the controller is not a black box, it is made up of several blocks and that the functionality of each block is clear, i.e. it is easy to work out which part of the controller does what. The level of visibility is important to the engineers and managers who need to work with the control law at the later stages of the total design process. It also makes qualification and certification of the control law an easier task.

It is noted that Section 5 (the flight evaluation) is identical to Section 6 of [14], the other conference paper based on [12], and Section 6 (the post-flight analysis) is a more detailed version of [14, Section 7]. The flight evaluation and post-flight analysis details were included in this paper because it is our desire that this paper is a summary of [12].

2 Motivation for the LPV approach

2.1 Robust control background

In this section we give a brief commentary on the more complete background material given in [12, Appendices A and B] which should be referred to for more precise assertions and mathematical statements.

Robust control is becoming a mature field in theoretical developments and emerging applications. One of the effective design procedures is termed $\mathcal{H}_\infty$ loop shaping [9] and combines the traditional notions of bandwidth and loop gain together with modern ideas of robustness into a single framework. Particular advantages of the method are: 1) the ease with which the controller weighting functions can be tuned to give an appropriate system bandwidth; and 2) the use of a generalised stability margin that supersedes gain and phase margins (this is a dimensionless quantity where a margin of greater than 0.3 is considered good and less than 0.2 unsatisfactory).

The approach has now been used quite widely and a considerable body of experience has been accumulated. Particular successes have been in Harrier control by Hyde [7] and in helicopter control [15]. In [7] the plant dynamics varied significantly over the flight envelope and the controller needed to be scheduled with operating point. However the underlying theory is based on a nominal linear model with the controller synthesised to accommodate substantial uncertainty. This controller scheduling has been reasonably effective but is not supported by any theoretical guarantees.

A general theory for the robust control of nonlinear systems is not computationally tractable and useful progress requires an intermediate level of complexity that, for example, incorporates scheduling requirements whilst remaining computationally tractable. One such recent advance has been the use of LPV plant models. In this framework the system dynamics are written as a linear state-space model with

---

[†] correspondence e-mail: `kg@eng.cam.ac.uk`

the coefficient matrices functions of external scheduling variables. Assuming that these scheduling variables remain in some given range then analytical results can guarantee the level of closed loop performance and robustness. The computational techniques rely on the solution of a set of linear matrix inequalities (LMIs) which, although challenging, are convex problems with reliable solvers available.

The above "pure" LPV model is not quite matched to the flight control problem where the scheduling variables are in fact system states (e.g. airspeed and angle of attack), rather than bounded external variables. An approach to this problem is to generate so-called "quasi-LPV" models, which are applicable when the scheduling variables are measured states, the dynamics are linear in the inputs and other states, and there exist inputs to regulate the scheduling variables to arbitrary equilibrium values. In this case the resulting "quasi-LPV" model gives a theoretically exact model of the nominal system, however in practice numerical tests need to be performed on a gridded set of the scheduling variables.

If such a "quasi-LPV" model is obtained, then the LPV techniques can determine whether there exists a controller, which will itself be an LPV system, that meets a given closed loop specification over the stated operating envelope. The resulting controller is then automatically scheduled with the scheduling parameters. The analytical results will however only be strictly valid if it can independently be ensured that the scheduling parameters, and perhaps their rates of change, stay within the assumed limits.

2.2 Deriving a reliable LPV model

The model of the VAAC Harrier available for this study was the WEM (Wide Envelope Model). The WEM does not match the assumptions required for the "quasi-LPV" model generation since: 1) it was embedded in Fortran code which needed to be decoded to determine the underlying analytic model (which was very time-consuming); and 2) the model was not linear in all the inputs and all the states that were not scheduling variables.

On the basis of the required flight envelope (Table 1) it was decided to choose the airspeed, V, and angle of attack, α, as the scheduling variables. The nonlinear model is then approximated by a "quasi-LPV" model which can be nonlinear in V and α but must be linear in all the remaining variables. This is described in detail in [14] and the quality of some approximations is given in [14, Figs. 2 and 3]. The engine model was not studied in as much detail as the rigid body dynamics and was replaced by a set of standard linearisations. As can be appreciated from the description of [14] the derivation of this model is not straightforward and requires a good appreciation of the relative importance and parameter dependence of the terms in the equations of motion to obtain an accurate model. A list of reasonable assumptions for this system is given in [14]. The accuracy of the resulting "quasi-LPV" model can be observed by comparing its transient response with that for the full WEM simulation model and some typical responses are given in [14, Figs. 5 and 6] with generally good agreement.

3 The controller specifications

The Harrier has three distinct flight regimes: jet-borne flight (lift provided by the engines), wing-borne flight (lift provided by the wings) and the transition region (lift from both the wings and engines). The controller must work well over the following flight envelope (wing-borne flight):

True air speed:	80 → 180m/s (155 → 350kts)
Angle of attack:	$-3^\circ \rightarrow 12^\circ$
Height:	500 → 20000ft

Table 1: The flight envelope

It is desired that the control law enables the following four functions: 1) a pitch rate, q, demand system; 2) a pitch attitude, θ, hold; 3) an angle of attack, α, limiter; and 4) good disturbance rejection. Thus, the Harrier was flown in the three inceptor configuration. By three inceptor it is meant that the pilot controls the aircraft (longitudinally) using the stick (closed loop demand), the throttle and the nozzle angle.[1]

The controller will be designed to meet two specifications widely used within the aerospace industry: 1) robust stability and nominal performance requirements expressed as an exclusion region in a Nichols plot [11, Figs. 27.7 and 27.9]; and 2) the Gibson criteria [6] which are guidelines for achieving precision tracking and avoiding pilot induced oscillations (PIOs). It is desired that the transfer function from stick force in pounds to pitch angle in degrees achieves at least level 1 in [11, Fig. 27.12], and that the ratio of dropback (θ_{db}) to steady state pitch rate (q_{ss}) in [11, Fig. 27.15] should be

$$0 < \theta_{db}/q_{ss} < 0.25 \text{ seconds.}$$

The larger the value of θ_{db}/q_{ss}, the faster the flight path angle response to a pitch rate demand.

The tailplane must not be used excessively. Frequency limiting has been implemented on the Harrier so that its structural modes are not excited by the tailplane. When the frequency limiter is active, the tailplane position is held fixed for a minimum of 0.48 seconds. To avoid frequency limiting, the tailplane demand must not contain high frequency components (above 3.125Hz).

To ensure safety, the certification authorities require that the angle of attack remains within certain limits. If these limits are exceeded, then an independent monitor (IM) disengages the experimental control law and the safety pilot takes charge. Angle of attack is one of many variables monitored by the IM [4].

As DERA did not specify desired closed loop time domain properties, the design strategy followed was to try and obtain the best possible closed loop performance (i.e. a fast and robust closed loop) subject to: 1) the limitations imposed by the plant, measurement noise and frequency limiting; and 2) meeting the Nichols plot exclusion region specifications and the Gibson criteria (guidelines for acceptable closed loop performance were provided by [7, 5]).

4 The controller architecture

To ensure that the functionality of each block within the controller is clear, the controller architecture is fixed before any designing takes place. This is done by careful examination

---

[1] The controller designed by Hyde in [7] was flown in the two inceptor configuration, i.e. the pilot controls the aircraft using only the stick and a left-hand inceptor (both closed loop demands).

of the closed loop requirements, past experience and good knowledge of the performance and accuracy of the available actuators and sensors. Of course, by restricting the structure of the controller, the controller will not be optimal (optimal with respect to the closed loop requirements); there is an inherent trade-off between visibility and optimality. As the control law must enable a pitch rate demand, a pitch attitude hold and an angle of attack limiter (see previous section), the architecture illustrated in Fig. 1 was chosen.

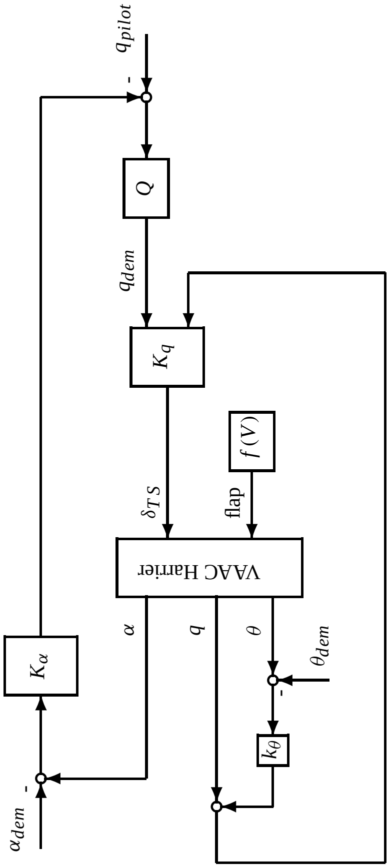

Figure 1: The controller architecture; δ_{TS} denotes the tailplane deflection

The function of each block within the controller is described below:

1. The inner loop controller K_q will be designed so that the closed loop transfer function from q_{dem} to q does not change much due to modelling uncertainty. Thus, the designer has a robust platform on which to build the control system.

2. The pitch attitude hold will be implemented using the gain k_θ and

$$\theta_{dem} = \begin{cases} \theta_{hold}, & |q_{dem}| \leq 0.1^\circ/\text{s}, \ |q| \leq 1^\circ/\text{s and} \\ & |\phi| \leq 6^\circ \\ \theta, & \text{otherwise} \end{cases}$$

   Assuming that the θ-hold is disengaged, i.e. $\theta_{dem} = \theta$, then θ_{hold} is the value of θ at the time instant when the conditions $|q_{dem}| \leq 0.1^\circ/\text{s}$, $|q| \leq 1^\circ/\text{s}$ and $|\phi| \leq 6^\circ$ become true (note that ϕ denotes the Euler roll angle).

3. The command pre-filter Q is a lead-lag filter that will be designed so that the transfer function from q_{pilot} to q, with no α limiting, meets the Gibson criteria.

4. The α-limiter will be implemented using the outer loop controller K_α and

$$\alpha_{dem} = \begin{cases} f_\alpha \alpha_l, & \alpha < f_\alpha \alpha_l \\ \alpha, & f_\alpha \alpha_l \leq \alpha \leq f_\alpha \alpha_u \\ f_\alpha \alpha_u, & \alpha > f_\alpha \alpha_u \end{cases}$$

   where α_l and α_u are the lower and upper IM incidence limits respectively, and f_α is a real number that takes values in $(0, 1]$.

5. The flap schedule, $f(V)$, has been chosen so that during an acceleration/deceleration, if θ is kept fixed, then the flight path angle does not decrease/increase.

Having fixed the controller architecture, the designer then chooses the controller's "gains" using a suitable design procedure. Our approach to controller design using LPV methods is given in detail in [12, Appendix B]. The basic idea is to: 1) establish loop shaping weights at a set of points in the operating envelope and then interpolate these weights over the whole envelope; and 2) use established LPV computational techniques to determine whether a satisfactory controller can meet the specification set by the weights over the whole envelope. Iteration is required between the problem set-up, the synthesis tools and simulation.

In [12, Appendix D] it is described in detail how the design was achieved for a single-input/single-output controller of the short period dynamics with the tailplane as input. The performance of the control law designed has been analysed with powerful handling quality metrics in [12, Section 5.1] and is shown to meet the required closed loop specifications listed in the previous section.

The controller synthesis step is computationally challenging and with the current state of these computational tools care needs to be taken in the choice of grid size and basis functions for parameter-dependent solutions of LMIs. The optimisation problem that generated the controller that was flight tested took about 28 hours to solve on a Pentium II 400MHz (57 grid points and 540 basis function variables).

The implemented controller is itself an LPV system with the state-space matrices dependent on the scheduling variables and their rates of change, which will need to be measured. Details of this together with the choice of sampling period are given in [12, Appendix E].

5 Flight evaluation

5.1 The VAAC Harrier

The VAAC Harrier, a T.Mk4 (two-seat trainer) illustrated in Fig. 2, has been specially equipped with a digital fly-by-wire FCS to form part of a rapid prototyping environment for research. The FCS, which includes the flight control computers, head-up and head-down displays (HUD and HDD) and an inertial navigation system, has full authority control of the tailplane, ailerons, rudder, flaps, nozzles and engine as well as the reaction control system which provides attitude and heading control at low speeds.

When the FCS is engaged, the subject pilot controls the aircraft from the rear cockpit and the IM software continuously

Figure 2: A photograph of the VAAC Harrier

checks the flight condition to ensure that the aircraft remains within it's safe flight envelope. If the IM detects a potential breach of this envelope, it will cause the FCS to disengage and control will then revert to the safety pilot in the front cockpit who then flies the aircraft using the standard T.Mk4 mechanical flight controls. This combination of mechanical reversion mode, assisted by the IM, allows flight control software to be flown without the need for a rigorous clearance process. In principle, software turnarounds can be completed in hours.

5.2 Preparations for flight

Use of the simulation facilities based at DERA[2] Bedford helped to ensure the limited flight time of one hour was used efficiently. Initially, the controller was exercised in the VAAC flight lab, a hardware-in-the-loop facility designed to allow functional checks. It incorporates actual FCS hardware elements, such as the VAAC Harrier's flight control computers and displays, and permits real-time operation of control software with the aircraft aerodynamics, engine and actuators being represented by mathematical models. Testing in the flight lab provided confidence that the controller software was compatible with the VAAC Harrier's operating environment and that the FCS could be engaged successfully.

A preliminary evaluation of the handling qualities of the controller was conducted using the fixed-base Real-Time All Vehicle Simulator (RTAVS). This provided valuable feedback and enabled some gross tuning of controller performance prior to flight test. However, due to known areas of uncertainty in the mathematical model of the VAAC Harrier, and slightly different inceptor characteristics, no attempt was made to fine tune handling qualities. Instead the simulator exercise was used to identify parameters in the controller that affected handling qualities and these were implemented in the software as "variable gains" for in-flight tuning. Examples of variable gains included stick gearing and time constants for stick pre-filters and sensor noise filters.

Portability of software between the flight lab, RTAVS and the VAAC Harrier was achieved with minimal effort due to the use of automatic code generation. This permits the controller, implemented in Simulink block diagram form, to be ported without any modification from a desktop PC (used for initial design and analysis), through to simulation facilities (flight lab and RTAVS), and onto the VAAC Harrier for flight test. A consistent data-logging interface is also provided, allowing time histories of selected controller variables to be recorded directly in Matlab-readable files for post-flight analysis.

5.3 Flight test

The preparations described above provided sufficient confidence in the controller and so no ground testing was conducted prior to flight. During the flight, on-board telemetry allowed trials personnel to monitor progress in real-time. Video cameras in the aircraft transmitted images of the rear cockpit outside world (viewed through the HUD) and rear cockpit.

The flight itself was divided into a number of specific test points, beginning with a familiarisation/tuning phase. This was to allow the test pilot to familiarise himself with the controller characteristics over the specified speed range and to address any major deficiencies, via the in-flight "gain" tuning feature, prior to the evaluation phase. The first test point comprised a series of pitch attitude capture tasks, with specified task performance criteria, conducted at a range of airspeeds (150, 200 and 250 knots). During this test point, the incidence limiting and pitch attitude hold functions were deselected. Next, the performance of the pitch attitude hold was evaluated by making throttle and nozzle angle changes to induce pitch disturbances. To assess the incidence limiting function, the pitch attitude hold was engaged before the pilot throttled back. This resulted in increasing incidence, as airspeed reduced, until the incidence limit was encountered, at which time the limiting function lowered the nose of the aircraft to maintain incidence. Finally, the robustness of the controller was investigated by performing a deceleration to the hover, well below the minimum design speed (80 m/s) of the controller.

6 Post-flight analysis

The post-flight analysis was done by processing the test flight data and video, and the pilot report.

The test flight took place in good weather conditions (temperature: 20°C, wind: 10 kts) and the control law engaged without difficulty.

Shortly after the control law was engaged for the first time, the test pilot reported that the coaming lamp of the pitch frequency limiter was on for about 50% of the time. The result of the frequency limiting was a $\pm 2^\circ$ error in pitch attitude. Frequency limiting was not expected, and could not be explained at the time because: 1) the pilot was not in the loop; 2) there was very little turbulence; and 3) the level of measurement noise encountered on the VAAC was expected to be less than that encountered in the flight lab.[3] In an attempt to stop the frequency limiting, the time constant of the pitch rate low-pass filter was increased from 0.05s to 0.1s. The frequency limiting then stopped[4], but there was greater susceptibility to PIOs as the low-pass filter adds about 27° of phase-lag at the loop gain cross-over frequency. Just before the pilot performed a pitch attitude capture at 200 knots (see Fig. 3) he said: "...there is a tendency, when you're in the loop

---

[2] DERA is the UK Defence Evaluation and Research Agency.

[3] It turned out that the frequency limiting was triggered by a high level of measurement noise; the signal conditioning unit that filters the pitch rate measurement had failed.

[4] Frequency limiting was not encountered again until the end of the test flight.

and if you bring your gains up to try and chase an accurate attitude, to get into a small PIO, very small, but it’s quite easy to ease off and stop that.” After the maneuver he said: “...one overshoot and a couple of bobbles because I’m in the loop; I’m now out of loop and it’s holding quite well.” According to the pilot, the PIO pitch bobbling is at 2 to 3Hz with a 0.25° amplitude.

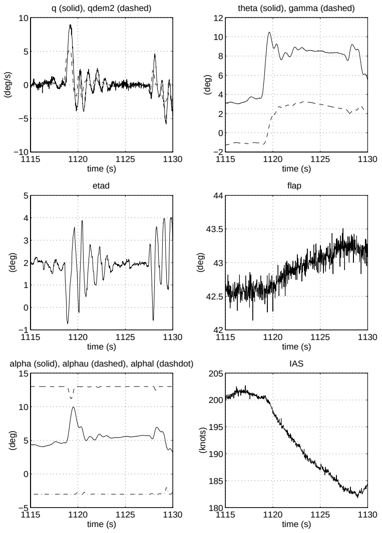

Figure 3: A +5° pitch attitude capture at $h_{init} = 5400$ft, nozzle = 36° and thtdem = 55%; h_{init} is the height at the beginning of the maneuver and thtdem is the thrust demand

The handling of the aircraft improved when the stick sensitivity was reduced by 25%, i.e. full-scale deflection of the stick demands ±7.5°/s instead of ±10°/s.[5] Pitch attitude captures with the new stick scaling are shown in Fig. 4. After the maneuver the pilot said: “...bobbles but capturing within a degree, possibly just outside of a degree, and holds for 5 seconds.” The fact that handling is better is reflected in the Cooper-Harper ratings; the maneuver in Fig. 3 scored a 5 whereas the maneuver in Fig. 4 scored a 4.[6]

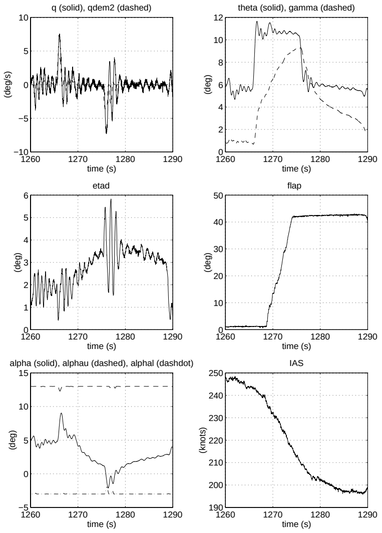

Figure 4: A +5° pitch attitude capture followed by a −5° pitch attitude capture at $h_{init} = 3000$ft, nozzle = 51° and thtdem = 63%

The closed loop simulation model was validated (strictly speaking not invalidated) by driving it with the pilot’s pitch rate demand (`qdem2`) from Fig. 4 and comparing the resulting time histories (see Fig. 5) with the ones illustrated in Fig. 4. It is easy to see, by comparing Figs. 4 and 5, that the pitch attitude responses match up quite nicely. The reason for the flight path angle ballooning and the large speed drop in Fig. 4 (a total of about 50 knots as opposed to only 10 knots), is that the flap deploys during the maneuver. As the flap deploys, the aerodynamic lift and drag increase.

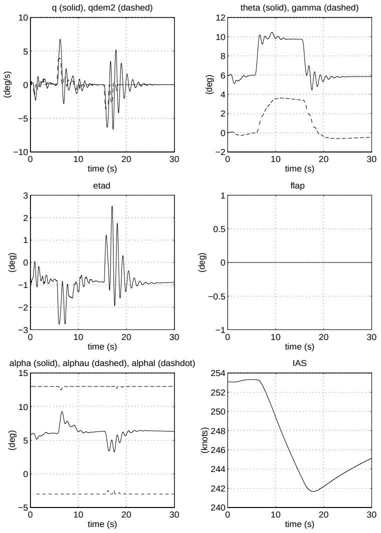

Figure 5: Time histories that result from driving the closed loop simulation model with the pilot’s pitch rate demand (`qdem2`) from Fig. 4

The performance of the θ-hold in the presence of nozzle and throttle changes is shown in Fig. 6. After the maneuver the pilot said: “..., so for about 1.5 to 2 seconds the pitch (θ) was in the wrong place but it readjusted and it reacquired nicely.”

The α-limiter did not perform as desired until f_α was reduced to 78.12% of its nominal value.[7] After performing the maneuver illustrated in Fig. 7 the pilot said: “...it seems to be working (the α-limiter) and doing its job now. It’s hunting (oscillatory behaviour, cycling) the aeroplane to keep the α within limits.” Note that the α-limiter engages because the pilot reduces the throttle and the control law maintains θ approximately constant.

To demonstrate the robustness of the control law in flight, the test pilot was asked to fly the Harrier down to hover (the controller was designed only for wing-borne flight, see Table 1). Unfortunately, the test pilot did not manage to hover the Harrier with the control law engaged; the control law was

---

[5] The stick filter provides linear variation with stick position and a deadzone.

[6] The pilot commented that the handling of the aircraft was easier in the later stages of the test flight as he became more familiar with the behaviour of the control law.

[7] It is probable that the α-limiter did not perform as expected because it uses inertial α whereas the IM uses the α measurement from the vane (inertial α is effectively a filtered version of the vane measurement).

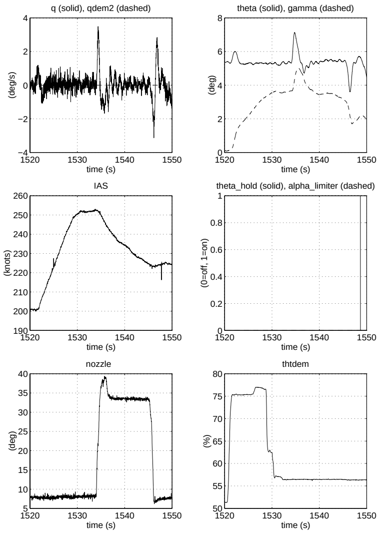

Figure 6: Pitch attitude deviations due to thtdem and nozzle changes at $h_{init} = 5000$ft and flap = 40°

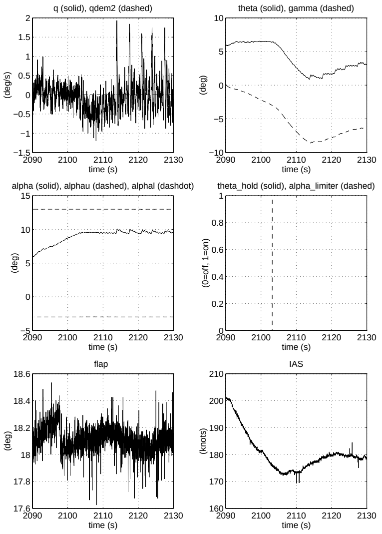

Figure 7: The α-limiter is engaged by reducing thtdem to 37% at $h_{init} = 5300$ft and nozzle = 8°

disengaged by the safety pilot at a height of 150ft and an indicated airspeed of 35kts (approximately 18m/s) because of large pitch attitude oscillations caused by frequency limiting. It is believed that the frequency limiting was triggered by the α-limiter (see [12, Fig. 5.7]). If the α-limiter had been switched off, pre-flight analysis suggests that the pilot should have been able to hover the aircraft. Nevertheless, the control law exhibited very good robustness properties as the pilot was able to fly down to 18m/s without a significant increase in his workload.

In conclusion, it was a successful test flight with respect to engaging the control law, exercising the control law within and well-beyond the design flight envelope limits, and achieving good pilot handling quality ratings. In terms of these points, it was one of the most successful first test flights performed by DERA to date. To summarise:

1. The control law achieved level two flying qualities over the whole flight envelope. The pitch attitude captures with the modified stick scaling (see Fig. 4) scored a 4 on the Cooper-Harper Rating scale, i.e. minor but annoying deficiencies and desired performance requires moderate pilot compensation.

2. Level two flying qualities were achieved even though during most of the flight the nozzles were not fully aft (the controller was designed with the nozzles fully aft), the flaps are soft (in Fig. 3 the flap demand is 50° but the actual angle is less) and the pitch rate low-pass filter introduces a large phase lag at the loop gain cross-over frequency (no post-compensator was used for the controller design). Furthermore, the pilot was able to fly the Harrier near hover.

3. The fact that the test flight went well, and was expected to go well, is because the mathematical model of the VAAC Harrier is sufficiently accurate to enable accurate predictions about the closed loop behaviour (see for example Figs. 4 and 5 which show responses of the real aircraft and the mathematical model for identical pilot inputs), the controller architecture has no major shortcomings and the controller comes with “sensible” closed loop guarantees. A guarantee is “sensible” if it reflects reality; guaranteeing robust stability with respect to uncertainty with a structure that does not reflect reality is not “sensible”. If the guarantees with which a controller comes are “sensible”, then the chances that the controller will work on the real plant are much greater.

4. The pilot thought that pitch PIO susceptibility was evident. This is almost certainly due to the phase lag introduced by the pitch rate low-pass filter. The susceptibility to PIOs might have been less evident had the signal conditioning unit not failed; this is suggested by the RTAVS trials and desktop simulation. When the test done to produce Fig. 5 was repeated with the original pitch rate low-pass filter, the bobbles decreased as seen in Fig. 8.

5. Stick scaling issues must be investigated as stick scaling can significantly improve handling. For this case study the designers had to guess the stiffness of the stick to be able to apply Gibson’s criteria for avoiding PIOs. Motion-based simulation would also have been of value for setting up the control law and for pilot rehearsals.

6. The ability to alter control law parameters in flight proved useful.

7 Conclusions

In this paper, a SISO robust gain scheduled controller has been designed for the short period dynamics of the VAAC Harrier using the design procedure proposed by Papageorgiou and Glover [13]. The performance of the control law has been evaluated using piloted simulations and flight testing.

The proposed design procedure has the following attractive features: 1) it is systematic and it can handle MIMO plants with uncertain nonlinear and/or parameter-dependent dynamics; 2) a detailed description of the plant uncertainty

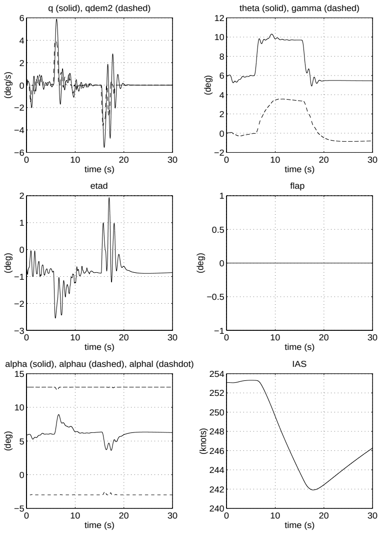

Figure 8: Time histories that result from driving the closed loop simulation model with the pilot’s pitch rate demand (`qdem2`) from Fig. 4, but with the original pitch rate low-pass filter

is not required; 3) it is relatively straightforward to translate the controller specifications into the design parameters, i.e. into weighting functions; and 4) the resulting controllers come with “sensible” closed loop guarantees.

The first (and only) test flight of the designed controller was one of the most successful first test flights performed by DERA to date in that: 1) the control law engaged without difficulty; 2) the control law was exercised within and well-beyond the design flight envelope limits; and 3) level two type responses were achieved throughout the flight envelope even though the signal conditioning unit failed, the stick filter was not tuned properly and the α-limiter used inertial α instead of the α measurement from the vane.

Furthermore, the control law designed for the VAAC Harrier has visibility in terms of its functionality, its performance can be analysed with powerful handling quality metrics and is shown to meet the required closed loop specifications, and it was straightforward to build anti-windup logic into it. Also, it is important to note that the metrics that were used are not specific to $\mathcal{H}_\infty$/LPV controllers.

A more complex design problem than the one tackled in this report, i.e. more channels and a larger flight envelope, might be difficult to solve due to computing power restrictions.

In conclusion, the design procedure proposed by Papageorgiou and Glover should be considered as a strong candidate for designing the next generation flight controllers.

Acknowledgements and Notes

The authors would like to thank Andrew Berry, Dr. Rick Hyde, Dr. Ewan Muir and Fred Scorer for all their help.

This work was funded by the UK Ministry of Defence under DERA Applied Research Project 07b3.6 “Aircraft dynamics and control system integration”. The theory and principles of application of LPV control to combat aircraft were developed under DERA Corporate Research Programme TG3/2.301 “Multivariable control for future combat aircraft” during 1995 to 1998.

This paper was published with the permission of the controller of Her Majesty’s Stationery Office.

References

[1] G.J. Balas, I. Fialho, A. Packard, J. Renfrow, and C. Mullaney. On the design of LPV controllers for the F-14 aircraft lateral-directional axis during powered approach. In *American Control Conference*, 1997.

[2] S. Bennani, D.M.C. Willemsen, and C.W. Scherer. Robust control of linear parametrically varying systems with bounded rates. *Journal of Guidance, Control, and Dynamics*, 21(6):916–922, 1998.

[3] J-M. Biannic, P. Apkarian, and W.L. Garrard. Parameter varying control of a high-performance aircraft. *Journal of Guidance, Control, and Dynamics*, 20(2):225–231, 1997.

[4] DRA. VAAC independent monitor requirements document. Technical Report DRA/AS/FDS/IM/URD/3.0, Defence Research Agency.

[5] J.A. Franklin, M.W. Stortz, P.F. Borchers, and E. Moralez III. Flight evaluation of advanced controls and displays for transition and landing on the NASA V/STOL systems research aircraft. Technical Report 3607, NASA, April 1996.

[6] J.C. Gibson. Piloted handling qualities design criteria for high order flight control systems. In *Criteria for Handling Qualities of Military Aircraft*, volume 333 of *AGARD Conference Proceedings*, June 1982.

[7] R.A. Hyde. $\mathcal{H}_\infty$ *Aerospace Control Design - A VSTOL Flight Application*. Advances in Industrial Control Series. Springer-Verlag, 1995.

[8] L.H. Lee and M. Spillman. Control of slowly varying LPV systems: an application to flight control. In *AIAA Guidance, Navigation, and Control Conference*, 1996.

[9] D. McFarlane and K. Glover. A loop shaping design procedure using $\mathcal{H}_\infty$ synthesis. *IEEE Transactions on Automatic Control*, 37(6):759–769, 1992.

[10] J.B. Mueller and G.J. Balas. Implementation and testing of LPV controllers for the NASA F/A-18 systems research aircraft. In *AIAA Guidance, Navigation, and Control Conference*, number AIAA 2000-4446, August 2000.

[11] E. Muir. *Robust Flight Control: A Design Challenge, GARTEUR*, volume 224 of *Lecture Notes in Control and Information Sciences*, chapter 27, pages 419–443. Springer-Verlag, 1997.

[12] G. Papageorgiou and K. Glover. Design, analysis and flight testing of a robust gain scheduled controller for the VAAC Harrier. Technical Report CUED/F-INFENG/TR.368, University of Cambridge, Department of Engineering, February 2000.

[13] G. Papageorgiou and K. Glover. Design of a robust gain scheduled controller for the High Incidence Research Model. To appear in the AIAA Journal of Guidance, Control, and Dynamics, 2000.

[14] G. Papageorgiou and K. Glover. Development of a “reliable” LPV model for the longitudinal dymanics of DERA’s VAAC Harrier. In *AIAA Guidance, Navigation, and Control Conference*, number AIAA 2000-4459, August 2000.

[15] A. Smerlas, I. Postlethwaite, D.J. Walker, M.E. Strange, J. Howitt, R.I. Horton, A.W. Gubbels, and S.W. Baillie. Design and flight testing of an $\mathcal{H}_\infty$ controller for the NRC Bell 205 experimental fly-by-wire helicopter. In *AIAA Guidance, Navigation, and Control Conference*, number AIAA 1998-4300, August 1998.