Thesis submitted for the degree of

Doctor of Philosophy

at the University of Leicester

by

Sarah Lee Gatley B.Eng., M.Sc. (Leicester)

Department of Engineering

University of Leicester

July 2001

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Abstract

A robust integrated flight and propulsion control (IFPC) system is designed and analysed for an experimental vertical/short take-off and landing (V/STOL) aircraft configuration, using multivariable design techniques. The model used for the design is based on the DERA VAAC Harrier wide envelope model (WEM) airframe with a Rolls-Royce Spey engine. This provides a detailed, large-scale, interacting, full-envelope, nonlinear simulation model, thought likely to be representative of the next generation of V/STOL aircraft. The centralised IFPC system is evaluated in piloted simulation trials on DERA's real-time all-vehicle simulator (RTAVS), the results of which indicate that level 1 or 2 handling qualities are achieved over the low-speed powered-lift region of the flight envelope. The application of the structured singular value, μ, and the ν-gap metric to the problem of evaluating the robustness properties of this multivariable IFPC system is presented. The centralised controller is subsequently partitioned into decentralised lower-order airframe and engine subcontrollers, in order to address implementation and testing issues. The partitioned system is seen to retain largely both the performance and robustness properties of the centralised system. Due to the particular implementation structure used for the centralised $\mathcal{H}_\infty$ loop shaping controller, the partitioning procedure described can be applied to general two-degree-of-freedom control systems. A scheme to guarantee maximum limits on safety-critical engine variables is developed and applied, to preserve the structural integrity of the engine during extreme manoeuvres. As the system frequently operates on the position limits of the engine actuators, an anti-windup scheme is implemented to maintain the performance of the system during actuator saturation.

Acknowledgements

I would like to thank my supervisor Professor Ian Postlethwaite and co-supervisor Dr. Declan Bates for their help and advice throughout the course of this work. I would also like to express my thanks to Dr. Martin Hayes for his advice on robustness analysis.

I would like to thank DERA Bedford, for the use of the nonlinear aircraft model, and for the useful comments and suggestions during the piloted simulation. Thanks also to DERA Pyestock for their help in understanding the problems of engine control. Thanks to the EPSRC for the three years of funding which they provided, without which none of this would have been possible.

I would also like to thank all of the control group for making this time at Leicester both enjoyable and interesting. Most especially, my thanks go to Simon and to my family for their continued support and encouragement.

Contents

List of Figures

List of Tables

Nomenclature

ACT	active control technology
C_l	lift coefficient
c.g.	centre of gravity
DERA	Defence and Evaluation Research Agency
DPF	diagonal perturbation formulation
$ENOZA$	exit nozzle area, (0.8307 to 0.1602 sine petal angle)
EPR	engine pressure ratio
$ETAD$	elevator position, ($-15°$ to $+15°$)
$ETASTK$	pitch reaction control system position, ($-15°$ to $+15°$)
$FNOZ$	front nozzle position, ($-5°$ to $+120°$)
$HPCPC$	high pressure compressor aerodynamic, (%)
$HPSM$	high pressure compressor surge margin, (%)
IAS	indicated air speed
IFPC	integrated flight and propulsion control
IGV	inlet guide vane angle, ($-8°$ to $+35°$)
LFT	linear fractional transformation
LMI	linear matrix inequality
$LPCPC$	low pressure compressor aerodynamic, (%)
$LPSM$	low pressure compressor surge margin, (%)
LPV	linear parameter varying
MFF	main fuel flow, (0 to 1.2 kgs^{-1})
$NHPC$	high pressure spool speed, (%)
$NLPC$	low pressure compressor spool speed, (%)
P	x body axis angular rate ($rads^{-1}$)
$P6$	combustion chamber pressure, (Pa)
Q	y body axis angular rate ($rads^{-1}$)
R	z body axis angular rate ($rads^{-1}$)

RCS	reaction control system
RCV	reaction control valves
$RNOZ$	rear nozzle position, ($-5°$ to $+120°$)
RTAVS	real-time all-vehicle simulator
SISO	single-input, single-output
SINPUT	input scaling matrix
SOUTPUT	output scaling matrix
$SPLIT$	engine thrust split, (0 to 1, where 0 = all to front, 1 = all to rear)
STOVL	short take-off and vertical landing
$T6$	high pressure compressor outlet temperature, (K)
$T10$	high pressure turbine stator outlet temperature,(K)
$T30$	jet pipe temperature, (K)
THM	thrust magnitude, (kilo Newtons)
U	x body axis velocity (ms^{-1})
V	y body axis velocity (ms^{-1})
VAAC	vectored thrust aircraft advanced flight control
V/STOL	vertical/short take-off and landing
VD	vertical speed
VH	horizontal speed
VT	velocity along the flight path (airspeed), (ms^{-1})
W	z body axis velocity (ms^{-1})
WEM	wide envelope model
X	total x body axis force
Z	total z body axis force
α	aircraft angle of incidence (attack), (deg)
γ	flight path angle, (deg)
$\dot{\gamma}$	flight path rate, ($degs^{-1}$)
θ	pitch angle, (deg)
χ	% of complex uncertainty added to enable a complex μ analysis
KA	airframe subcontroller
KE	engine subcontroller
E	error signals
z	controlled variables
u	control inputs

$\cdot_a$	airframe variables
$\cdot_e$	engine variables
$\cdot_{ea}$	interface variables
$\cdot_c$	commands

Chapter 1

Introduction

This chapter will introduce the aims and benefits of integrated flight and propulsion control (IFPC). Research into the design of IFPC systems is motivated by the desire to exploit potentially significant gains, both in terms of improved flying qualities and reduced pilot workload, which may be obtained by the use of propulsive system generated forces and moments for aircraft manoeuvring in the low speed region of the flight envelope. While the most obvious current application of this new 'extra' control power is in the area of super-manoeuvrable V/STOL fighter aircraft, these technologies may also offer significant benefits to future civil commercial applications. Research into control technology requirements for future supersonic transport aircraft and hypersonic aero-spacecraft have also indicated the necessity of fully integrating the airframe and engine control systems [64].

The use of the propulsion system to generate forces and moments for aircraft manoeuvring results in significantly increased coupling between the airframe and engine subsystems. In its simplest form, this coupling may be unidirectional (propulsive forces and moments affecting airframe states) but in general the use of novel effectors such as reaction control systems will also affect the engine operating point. This coupling necessitates an integrated approach to the overall flight control system design problem, in order to ensure that: optimal use is made of the various propulsion effectors for aircraft manoeuvring control, and that limitations due to engine safety considerations are taken account of in the overall design.

§ 1.1 describes some of the future challenges of IFPC system design. The use of advanced actuators is motivated, providing a brief description of how they work and the performance gains they offer. The opportunities associated with having an increased number of actuators will also be considered. The high workload of the Harrier pilot and how this may be reduced by adopting an integrated approach will be discussed. Possible modes of operation that the aircraft may fly in are also introduced. A description is given of some of the control power requirements in the low speed region of the flight envelope. Consideration will be given to the

general approaches to airframe and engine control, along with details of advanced engine control schemes. Some of the proven advantages of using an integrated approach to the problems of flight and propulsion control will be highlighted. The use of active control technology (ACT) is also motivated. § 1.2 summarises the advantages described in § 1.1. § 1.3 provides an overview of the main contributions of this thesis. Finally, § 1.4 will give the details of the structure of the thesis.

1.1 Future challenges in the control of highly manoeuvrable and V/STOL aircraft

Traditionally, tactical aircraft have been designed with the philosophy that flight control and propulsion control are inherently independent functions. For those particular missions where coordinated flight control and propulsion system commands are required, such as terrain following, the pilot effectively acts as the controls integrator [71]. For high performance aircraft, this task has become so difficult that autothrottles have been developed to ease pilot workload and improve aircraft performance. Future aircraft will be designed as multi-mission, multi-function aircraft. These aircraft will rely more and more on the propulsion system to provide forces and moments to enhance the flight control function. The engine is therefore becoming an important element in the attitude control system of the vehicle with the result that the dynamic coupling between the airframe and propulsion subsystems must be now accounted for in the flight control system design [60]. The combined requirements of supersonic flight and powered lift regimes provide very difficult constraints on the design of the aircraft, and as a result the dynamics are inherently much harder to control than for standard aircraft. In particular, a high degree of cross-coupling is expected when effecting changes in thrust magnitude and thrust direction. Such aircraft will undoubtedly be statically unstable and hence will not be controllable by the pilot without assistance from the flight control computer. The Harrier is an example of a current aircraft that does in fact use the propulsion system to augment the forces and moments generated by the airframe. On the Harrier however, the cross-coupling is not that high in many modes of operation as the thrust vector acts close to the centre of gravity. This enables the pilot to fly the aircraft without recourse to complex multi-axis feedback systems. However, even this moderate amount of cross-coupling impacts significantly on pilot workload. With increased coupling in future advanced V/STOL aircraft, feedback around all axes will be a requirement. The integration of propulsion control systems and flight control systems has been shown to significantly improve airplane performance parameters such as thrust, range and rate of climb [72, 53].

1.1.1 Powered lift flight control

The following section is a brief description of why aircraft have a limited angle of attack at which they may be flown, and how this may be increased through the use of advanced actuators [57]. When considering standard aerofoils at low angles of attack, the flow separation occurs at or close to the trailing edge. As the angle of attack, α, increases, the separation point gradually moves towards the leading edge. In this process, the lift coefficient, C_l, continues to increase and, at some point, attains a maximum value. Beyond this value of α, the lift coefficient drops, and the aerofoil is said to have stalled. This type of stall generally occurs on thick aerofoils and is characterised by a gradual loss of lift beyond stall, as shown by curve (a) in Figure 1.1. For thin aerofoils that have sharply curved leading edges, a slightly different type of stall occurs. From the leading edge and up to the minimum pressure point, a large favourable pressure gradient exists that tends to promote the existence of a laminar boundary layer. Beyond the minimum pressure point, an adverse pressure gradient exists toward the trailing edge. Because of the sharp curvature of the leading edge of a thin aerofoil, the adverse pressure gradient is sufficiently strong to cause flow separation. If the Reynolds number is low, this type of flow separation is of a permanent nature, so there is no subsequent reattachment, and, as a result, the aerofoil stalls. Because the flow separation occurs close to the leading edge, this type of stall is usually abrupt and is marked by a sudden loss of lift beyond the maximum lift. This is shown by curve (b) in Figure 1.1.

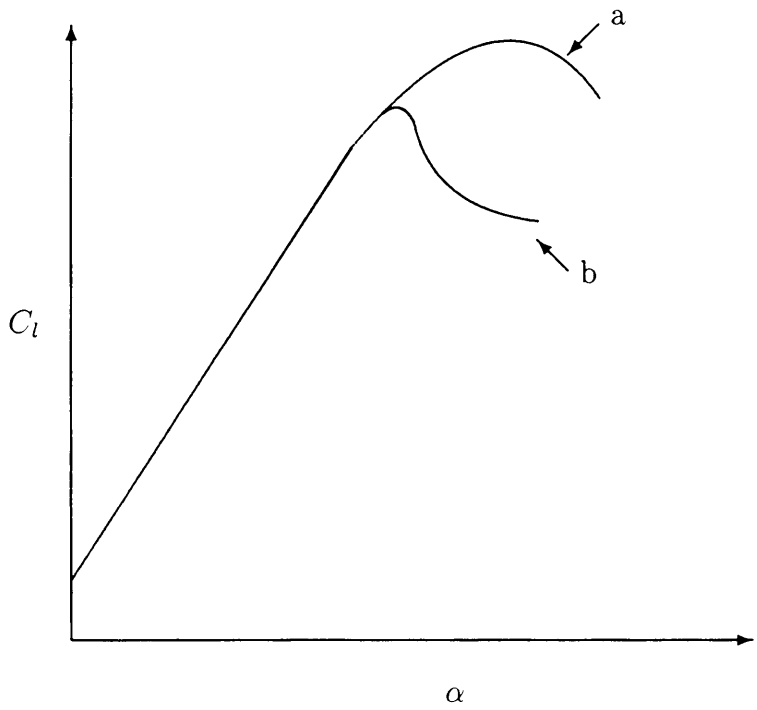

Figure 1.1: Schematic illustration of stall of airfoil sections.

If the freestream Reynolds number is sufficiently high, then there is an increasing tendency for the transition to occur in the separated boundary layer and cause the separated boundary-layer flows to reattach back to the upper surface, forming a bubble in a similar fashion to that for a circular cylinder. The reattached turbulent boundary layer is more capable of resisting the adverse pressure gradient than the laminar boundary layer. The reattached turbulent flow eventually separates from a point closer to the trailing edge.

The following gives details of some of the enabling technologies that are emerging. They provide merely the capability to achieve the required improvements. The realisation of all these potential benefits depends on the development of new V/STOL flight and propulsion control concepts, which effectively integrate many different control sub-problems, while maintaining pilot workload within acceptable limits. Recent developments in aerospace propulsion have resulted in advanced actuators that use air from the propulsion system to augment forces and moments produced by the airframe actuators [60]. This may be through the use of a reaction control system, where bleed air from the engine is diverted to jets that may be used for pitch, roll and yaw control. Other possible uses are the ducting of engine air through slots in the wing to entrain free stream air, where high-energy air is blown tangentially on the upper surface of the wing and the flap. This addition of energy helps to delay boundary-layer flow separation on the upper surface. As a result, the maximum lift coefficient increases. Remote-augmentation-lift-systems may also be used, in which engine air is heated with an augmentor as it is exhausted towards the ground. Alternatively, engine air may be blown over a lifting surface to modify its lifting characteristics. An example of this is a blown flap, where high-pressure air is blown over the upper surface of the flap. In this way, considerable energy is added to stabilise the upper surface boundary layer that delays flow separation, making it possible to achieve an 80% increase of maximum lift [45]. In one boundary-layer suction method, low energy fluid from the upper surface is removed by the application of suction. This process helps to delay the flow separation to higher angles of attack. Significant increases in maximum lift coefficients have been obtained by this method. Another of its advantages is that the application of suction also stabilizes the laminar boundary layer by delaying the flow transition. This results in a significant reduction in the skin-friction drag and, as skin-friction drag can account for half of the total aircraft drag at cruise, any reduction is a significant advantage. An example of the application of this method is the C-140 JetStar aircraft.

Alternatively, thrust vectoring can be used to control an aircraft in the low-speed, low dynamic pressure conditions experienced at high angles of attack when the traditional aerodynamic control surfaces will have lost most of their effectiveness because of their immersion in low-energy stalled flow. The application of multiaxis thrust vectoring provides increasing control

authority at high angles of attack, leading to improved low-speed agility. The use of thrust vectoring can eliminate the limitations on the maximum attainable roll rate capability imposed because of insufficient aerodynamic control authority at high angles of attack. The primary benefit of thrust vectoring is the availability of control power at high angles of attack, low dynamic pressure conditions where the conventional aerodynamic control will have lost most of its effectiveness.

1.1.2 Controls integration

An integrated approach to airframe/propulsion system configurations through the use of advanced actuators could be used to satisfy more demanding operational requirements, such as increased manoeuvrability, to abort and recover from departure, to reduce approach speeds, to enable supercruise, and to allow high angle of attack capability, stealth, and carefree handling throughout the flight envelope [26, 69]. While the most significant activity is in the area of military applications, benefits are also available to commercial applications. These benefits include lower approach speeds, resulting in reduced runway length requirements, while redundant attitude controls may be used to improve flight safety. The resulting steep climbout and approach paths due to V/STOL capability would also reduce the noise exposure to surrounding communities, and increase airport capability [80]. Furthermore, due to the decreased time to reach the required altitude, there will be a reduction in the use of fuel. This will have the direct benefit of reducing emissions.

The control system designers may have to contend with many more primary motivators than on standard aircraft. When considering longitudinal control, typical configurations may have both canards and elevator, a reaction control system (RCS), independent front and rear nozzle angle control and split front/rear thrust control. This increased flexibility can be used to produce a redundancy of actuation or to replace existing actuators, such as eliminating aerodynamic surfaces to reduce aircraft observables to aid stealth. This would also have the effect of reducing the drag of the aircraft. Alternatively, if the actuators supply redundancy, then the option opens for scheduling their use to provide maximum efficiency as a function of flight condition (or failure state of the system). A simple example is the use of thrust vectoring for pitch control in flight regions of low dynamic pressure and the use of the elevator in regions of high dynamic pressure. This approach can be used to increase the flight envelope, improve manoeuvrability, and allow for greater reconfiguration options following actuator failure. Alternatively, it may allow the use of smaller aerodynamic control surfaces. Traditionally, the size of an aerodynamic control surface is determined from control effectiveness at low-speed considerations. The fact that the thrust vectoring can provide the necessary control authority

at low-speed or low dynamic pressure conditions makes it unnecessary to size the aerodynamic surfaces from low-speed considerations. However, at high-speed or high dynamic pressure conditions, the aerodynamic control surfaces become effective, and much smaller surfaces can be used compared to those determined from low-speed considerations. The use of small-size aerodynamic control surfaces leads to lower structural weight, reduced trim drag and improved overall performance.

1.1.3 Reducing pilot workload

The Harrier has already been mentioned as an example of an existing aircraft that utilises thrust generated by the engine to directly affect the airframe attitude. Its unique, two-inceptor left-hand arrangement (throttle *and* nozzle lever) prevents simultaneous modulation of the thrust magnitude and direction. Instead, the pilot is required to use his or her experience to judge the left-hand control combination required for the task; specifically to apportion the total thrust vector into the appropriate longitudinal and normal components, hence resulting in a high pilot workload [11]. The right hand centre stick commands the aircraft pitch. An example of the high pilot workload is that during the final approach to the hover. The decelerating transition is usually performed by selecting increasing nozzle angles in a series of discrete steps, each at a suitable speed or range from the landing point. During the final stages of the deceleration, the pilot maintains a nominally constant pitch attitude, using the centrestick, while regulating the aircraft's flight path angle with the throttle. That is, it is the pilot that effectively integrates the engine and airframe subsystems. Owing to the basic aircraft being statically unstable at low speed, the pilot is required to monitor and control the pitch attitude continuously. As the aerodynamic lift reduces, the pilot is required to progressively increase engine thrust to maintain the necessary total normal force. Additionally, the pilot must also make fine nozzle angle and pitch attitude adjustments to control the aircraft's speed and longitudinal position, allowing for local wind conditions, during the landing sequence. A further aim of an IFPC approach is to reduce the pilot workload. Adopting such a scheme would allow the use of just two inceptors. The controller will use the actuators regardless of whether they are airframe or engine related. Any coupling would automatically be compensated for, so any change in the nozzle angles would not require a correction in the aircraft pitch, and vice versa.

1.1.4 Control power reconfiguration

The benefits of integrating the flight and the propulsion control systems can be grouped loosely into two categories. The first category includes the benefits of increased flight and manoeuvre envelopes. The second category includes the benefits of reconfiguring to accommodate various

modes of operation. Possible modes and situations in which they might be useful include a 'maximum thrust-specific-fuel-consumption' mode in which the goal is to optimise the installed performance of the engine taking into account deterioration and engine-to-engine build differences [54, 60]. A second mode is 'maximum thrust' in which total thrust is more important than thrust response rate or fuel consumption. A third mode is 'maximum thrust response rate' for precision landing or attitude control. One means of implementing this mode would be to spool-up the engine and then spoil its increased thrust by using internal geometry to reduce the cycle efficiency. The result would be a system where thrust can be recovered/dumped by changing areas rather than by spooling up/down the engine. One penalty for operating in this mode would be an increase in fuel consumption.

1.1.5 Control power for advanced manoeuvres

V/STOL operation in confined spaces and near obstacles in a steady wind may expose the aircraft to strong wind gradients and turbulence [25]. The major difficulties are encountered during descent onto the landing pad. For the transition to cruise, or from cruise to powered-lift flight, it is important to ensure a sufficient performance envelope for manoeuvring with satisfactory angle of attack or stall margin. This performance envelope must provide the ability to accelerate or decelerate in level flight, and to climb or descend at constant airspeed, or an appropriate combination of both. Control authority in addition to that for flightpath and airspeed must be sufficient to stabilise the aircraft.

It is well recognised that the total control authority of a V/STOL aircraft must be sufficient to trim the aircraft at a stabilised flight condition, to suppress external disturbances imposed on the aircraft, and to provide for manoeuvring the aircraft as dictated by the mission and the particular phases of flight. Each of these demands on control authority are important to the designer and must be dealt with in the course of an aircraft's configuration definition and refinement. Control necessary to trim throughout the flight envelope is associated with specific details of a particular configuration and must be considered on an individual basis. The appropriate generalisation to be made concerning trim control is that considerable attention should be given in the design process to minimise this requirement. Control authority frequently comes at considerable cost to the propulsion system, airframe weight and overall aircraft performance, and the greatest fraction possible should be available for executing manoeuvres associated with the mission tasks and for operating in the mission environment.

To meet with military specifications, the aircraft's control system must be capable of trimming the aircraft in a steady hover in the presence of a wind from the most critical direction and hold sufficient control in reserve to meet the pitch, roll, yaw and thrust control demands

simultaneously. The wind magnitude most frequently cited for this requirement is 35 knots. This combined requirement for trim and manoeuvre is assumed to account for stabilisation in the presence of gusts and does not require an additional control increment to be added to counteract gust upsets. Since the entire control capability in hover is derived from the propulsion system, either through direct use of thrust or through bleed air from the compressor, this combined control requirement is of serious concern to the designer. An example of the impact a Level 1 flying qualities requirement has for a conceptual design of a four-power lift-cruise fan aircraft results in a combined requirement for trim and manoeuvre control for attitude and height that is over 30% of thrust in excess of that required for an aircraft of the same weight without the trim and manoeuvre control. To meet this requirement for vertical take-off would seriously penalise the aircrafts range/payload performance. The thrust necessary to meet Level 2 flying qualities is not reduced substantially. Not all designs present a concern of such magnitude; however, any thrust held in reserve for control often reduces the aircraft's mission effectiveness, and if excessive, could prevent it from proceeding beyond the design stage.

Other options may be considered as a means for alleviating this substantial demand on the propulsion system for control. Short take-off is unlikely to impose as stringent a demand for control reserve as does sustained hover in ground effect. If heavyweight take-offs are performed with some ground roll, the vertical take-off demand need no longer be met at those weights. Meeting the vertical landing requirement will be substantially easier on the propulsion design. Another option is to prioritise the axes for control demand and meet their individual needs accordingly. It is generally accepted that the yaw axis is less critical for control than roll or pitch. Roll is considered to have first priority, followed by pitch. The Harrier is a good example of a design that has successfully adopted the approach of shared reaction control demand. In that aircraft, an individual control axis has full authority available when the aircraft is operating below maximum thrust and no other controls are activated. When the aircraft is at maximum thrust and more than one control is demanded, the demand is shared, with the least proportional reduction applied to roll, then to pitch and yaw in that order.

1.1.6 Airframe control

This section will give a brief description of what is required for traditional full envelope airframe control. The primary function of the flight control system is to control the position, orientation and rates of the vehicle in space. Furthermore, the control system of an aircraft must contribute to its safe and economic operation, such that the intended flight mission can be accomplished and unexpected events can be handled. The capability to change the equilibrium or trim-lift

coefficient, in order that the aircraft is capable of flying at any desired angle of attack within its aerodynamic limits, is also required. The overall aim is to smooth out the variations in flight dynamics, to rectify the dynamic deficiencies and to provide acceptable levels of flying and handling qualities. This matter may be complicated by the fact that an aeroplane may be stable under some conditions and unstable under other conditions. The position requirements for the control system design are often expressed in terms of desirable paths and tracking accuracy required about that path, e.g. in terrain following/terrain avoidance. Other requirements can include maximum accelerations and turn rates for such manoeuvres as air-to-air combat. Some may also specify performance objectives in terms of rise time, overshoot and steady-state error. The most important feature is safety, and it must be demonstrated to be safe in all known operating conditions and in conditions following various degrees of component malfunctions or failure.

Flight control system designs are typically regulators with input shaping where the design criteria are often expressed in terms of equivalent system models [61]. These designs can be complex, as the plant to be controlled can be marginally stable or unstable, and the plant parameters can be strong functions of the operating point. One factor which simplifies the design is that the states that are to be controlled are typically reconstructed easily from measured outputs. The solutions also generally avoid operation on limits, mainly for safety reasons, as this may lead to instability. Flight control specialists have developed procedures for designing regulators and shaping filters that accomplish their goals. The open-loop instability, which is related to agility of the aircraft, is utilised to obtain better performance and manoeuvrability of the closed-loop system. Caution is advised in [16] when applying modern multivariable control techniques to aircraft to ensure due consideration is given to the aircraft dynamics and transparency of the applied control scheme.

The aircraft equations of motion can usually be defined adequately by a six degree-of-freedom model comprising two longitudinal modes (short period and phugoid) and three lateral-directional modes (spiral, roll, and dutch roll) [57]. Often, these modes are sufficiently decoupled to permit the longitudinal and lateral-directional degrees-of-freedom to be treated separately in the design and implementation of a control logic. The phugoid mode is a lightly damped, long period oscillatory mode, otherwise known as the long-period mode. As such, it can be quite annoying if left to die by itself. The pilot can easily operate the longitudinal control and kill the phugoid mode. Physically during the phugoid motion, the pitch angle increases and the aircraft gains altitude and loses forward speed. The angle of attack remains constant so that a drop in forward speed amounts to a decrease in lift and flattening of the pitch attitude. As a result, at the top of the cycle, the pitch angle goes to zero. Beyond this point, the airplane

begins to lose altitude, the pitch angle goes negative and the airspeed increases. Once the pitch attitude is nearly level, the air speed is at its maximum and the cycle repeats once more. Conversely, the short period mode is fast - $\leq$ 10 seconds - and heavily damped, so it is just felt as a bump by the pilot or passengers, and the pilot does not need to take any action to kill it.

1.1.7 Engine control

Propulsion control designs are fundamentally different to flight control designs. Their primary aim is protection. The objective is to operate the engine at desired levels of thrust and airflow in the presence of disturbances, failures and plant variations without violating constraints on component stability or physical limits [61]. However, the problem is multi-objective, as the objectives change as different limits are encountered. The control therefore is dominated by operating limits. That is, during a transient, the engine will spend much of its time on, for example, the minimum stall boundary or the maximum temperature boundary. A very important effect of this is that an equivalent system model of a closed-loop engine depends upon the transient performed in addition to the flight conditions. For example, at a given condition, an engine may be able to accelerate twice as quickly as it can decelerate. Acceleration is limited by a surge margin constraint, whereas deceleration is limited by a burner blowout constraint. Secondary considerations include maximising response rates, minimising fuel consumption and maximising lift. For IFPC design, the task also involves interpreting the aircraft control commands in terms of engine process commands and then using the engine control effectors to achieve them while respecting these restrictions. One factor which simplifies the design is that the open-loop plant is stable, often with all the poles occurring on the negative real axis. The overall problem is one of steady-state control rather than transient control. Consequently, engine controls are typically a series of integral loops with transients controlled by tracking limits. This is the nature of the classical accel/decel schedule approach to propulsion control in which the ratio of fuel to burner pressure is set during most of a transient by limits programmed into the control that protect, for example, against over-temperature, over-speed, burner blowout and surge [53, 72].

There are several other factors which affect propulsion control design and make it different from flight control design. These include the fact that the critical parameters to control, such as surge margin and thrust, are not easily sensed. The importance of engine thrust is obvious since it is the principal output of the engine. Some systems to measure thrust are under development and show promise, but the typical engine control system relies on controlling measurable quantities that have a high correlation with thrust. Common examples are engine pressure ratio, turbine exit temperature and the rotor speeds. The choice of variable is dependent on

the characteristics of the particular engine. Turbine temperature is an important variable to control since it is one of the most critical variables in determining the thermodynamic efficiency of the propulsion system. In general, the higher the turbine inlet temperature, the higher the cycle efficiency, resulting in more thrust for less fuel. Methods for raising the maximum allowable turbine inlet temperature are the object of continuing research by engine manufacturers. This maximum temperature is dictated by the ability of the material used to withstand the elevated temperature. The control problem is therefore primarily to guarantee that the turbine inlet temperature remains below a safe threshold. However, since sensors cannot withstand the temperatures either, variables that are highly correlated to the turbine inlet temperature are controlled instead. The lower turbine interstage or turbine exit temperatures are common choices.

The final variables that are critical to control are the surge margins on the fan and the compressor. Surge margin is a measure of the difference in pressure ratio at the stability limit of the compression component and the current operating point (at a constant rotor speed). There are several factors that are important in selecting the design values of surge margin that the control must maintain. The trade is essentially between engine safety and engine performance [60]. For engine performance it is typically desirable to keep surge margin very low. The operating line, which is given by the desirable steady-state operating points of the compression element, that has maximum cycle efficiency is typically very close to the stall line of the compressor. Therefore, a control law designed to control the engine to maximum efficiency will have a small surge margin. Furthermore, the desirable operating points generally get closer to the stall line as altitude increases. Conversely, for engine protection it is desirable to keep surge margin as large as possible, in order to ensure stable engine operation in the presence of all disturbances. A typical buildup of the surge margin consists of entries that account for engine transient operation, construction differences and modelling uncertainty. The surge margin limit is a characteristic of the engine [39]. Figure 1.2 shows a simple steady speed characteristic for a multi-stage compressor in terms of pressure ratio against mass flow. If the compressor is in the area of stable operation (point A) and the flow were to drop slightly, the operating point would move up to a higher pressure ratio to return to the initial value. However, if it is in the unstable region (point B) a slight drop in flow would result in a slight loss in outlet pressure, leading to a further loss in flow and a further pressure loss. The result of this is that the flow actually becomes negative and comes out of the compressor inlet. This is only a momentary effect, as when the back pressure has cleared itself the positive flow is re-established. However, if the outlet remains too restricted the flow will reverse again when the pressure exceeds the compressor capability. Thus, the flow surges back and forth, and the limit of stability is called

the surge point. Surge results from flow breakdown or stall within the compressor.

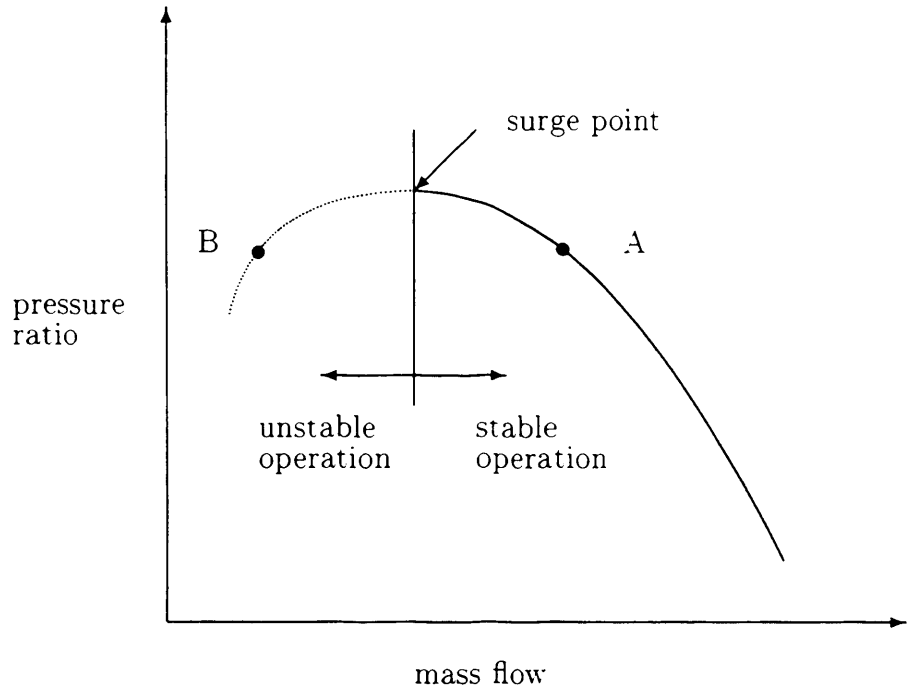

Figure 1.2: Simplified constant speed characteristic demonstrating regions of stable and unstable operation.

The surge margin may be traded in an active control scheme in areas in the flight envelope where the standard surge margin is conservative and limiting performance and economy [72]. The surge margin is the difference between the surge line and the working line, and is shown in Figure 1.3.

1.1.8 Multivariable engine control

This section will describe some of the advanced control techniques that have been applied to the problem of engine control. In [38], multivariable robust control techniques are applied to the RM12 jet engine for the Swedish fighter aircraft JAS 39 Gripen. The RM12 engine is presently controlled by a mixed electronic (DEC) and hydromechanical (MFC) system. The control principles traditionally used for this engine are relatively simple, with several SISO control loops and open-loop scheduling. In normal mode the control function is split between hydromechanical calculations, and calculations done by software in DEC. The main function of the basic control is to control the compressor rotational speed, NH, using the fuel flow. The MFC also schedules the settings of the compressor's variable vanes, CVG, as a function of NH and the compressor inlet temperature, $T25$. The fuel flow calculated by the MFC is then trimmed by DEC. There are basically three control loops in DEC that can effect the fuel flow. These are fan rotational speed (NL), dNH/dt, and compressor outlet static pressure ($PS3$)

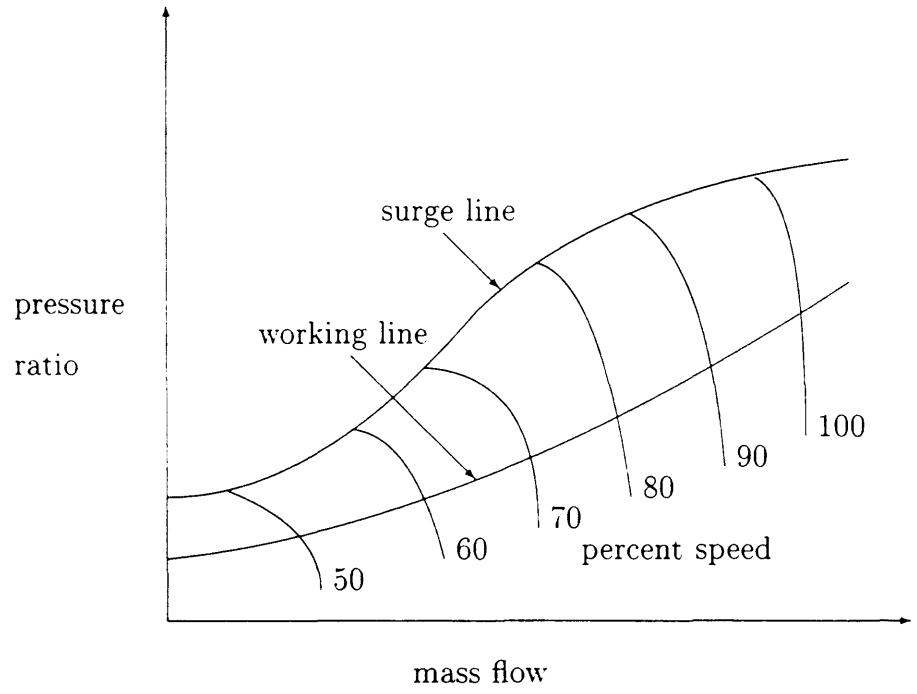

Figure 1.3: Basic over-all compressor characteristic, covering the operational speed range.

control loops. These three control loops calculate their respective fuel requests, and then the actual MFF request is calculated, based on the minimum-wins-selection principle. Another control loop in DEC is the turbine downstream temperature ($T5$) limiter. The maximum $T5$ is compared to a lead-filtered measured $T5$, and an integral regulator opens the nozzle area if the temperature is too high. The nozzle area is, however, mainly open-loop scheduled as a function of power lever angle, PLA. Other control signals from DEC are the fan variable vanes (FVG), and the afterburner fuel flow (WFR). Both are open-loop scheduled, FVG as a function of the intake temperature ($T1$) and fan rotational speed (NL), and WFR as a function of PLA, $T1$, and $PS3$. This is a very simplified description of the control function today. Since the jet engine is inherently multivariable, and the control loops obviously have been treated as single, it has been necessary to introduce a number of cross-couplings to stabilise the system. There are also a large number of 'fixes' and special cases that are necessary because the majority of the control signals are mainly open-loop scheduled.

The main objective of the engine control system is to control the engine to a desired thrust level. Since the thrust is not measured in flight, it is necessary to transform the thrust reference into references to measured engine variables. This can be done in a number of ways. In the concept in use today for the RM12, NL and $T5$ are used to control the thrust. This way of controlling the thrust is often referred to as the 'max temperature' concept, since the engine is always controlled to the maximum allowed turbine temperature. A consequence of this is that a new engine produces more thrust than specified. Another way of controlling the engine is called the

'constant thrust' concept. With this concept the fan rotor speed, NL, and the engine pressure ratio, EPR, are controlled. These variables together give a good measure of the air flow through the engine, and from that also the thrust. The variation in thrust between different engines can then be made small. The thrust will not, however, be constant, as the name suggests, since the outlet temperature also influences the thrust equation. In [38], the NL/EPR control strategy is used as the normal control mode. The other engine measurements of interest, NH, $T5$, and $PS3$, are also used by the feedback controller to avoid large deviations from the desired operating points. The four engine control signals for the core engine operation are all used by the multivariable controller, but only two of them, MFF and the variable exhaust nozzle area, $A8$, are free to define the steady-state engine operating point. FVG and CVG are used only transiently by the controller, and in steady state they are held at their reference values. Three controllers were designed to cover the operating range of the engine, these being scheduled by the conventional high gain feedback method. This scheme results in successful tracking of large PLA demands, ranging from idle to maximum speeds. When tested in nonlinear simulations, the responses using the advanced control method are significantly quicker and smoother than those achieved using the present conventional controllers.

In [49, 47], a fuel flow acceleration and deceleration (accel/decel) limit and a fan surge margin (DP/P) limit are implemented as a part of the engine control scheme. The accel/decel schedule is a set of variable bounds on the fuel flow as a function of the engine core rotor speed ($N25$), high pressure compressor inlet temperature ($T25$), and high pressure compressor discharge static pressure ($PS3$). This limit schedule is determined a priori using an open-loop nonlinear simulation of the engine. The schedule imposes a rate limit on core rotor speed as a function of the minimum and maximum air/fuel ratio (lean and rich blowout), the maximum turbine temperature and the compressor surge margin. The air/fuel ratio is a function of the inverse of $MFF/PS3$, since $PS3$ is indicative of the combustor air mass flow rate. These bounds are implemented as minimum and maximum bounds on the $MFF/PS3$ ratio.

A typical limited thrust response involves the $MFF/PS3$ ratio increasing to the acceleration limit at nearly constant corrected engine core rotor speed $N25R$ ($= N25\sqrt{(T_{std}/T25)}$, where T_{std} is the standard temperature, 245.52°C) due to the rotor speed lag. Then, as the rotor accelerates, the $MFF/PS3$ ratio tracks the limit value until the desired thrust setting is reached. At that point, the $MFF/PS3$ ratio decreases to the steady state scheduled value. Note that during the transient, when $MFF/PS3$ is determined by the acceleration limit, the engine input MFF is determined as a function of the plant outputs $PS3$, $N25$ and $T25$ (using the definition of the $N25R$). The effect of this feedback is that there is a rate limit on the fuel flow when the fuel flow limit is encountered. Other limits, such as the maximum fan speed and

the minimum burner pressure, also affect the fuel flow. MFF is decreased if the maximum fan speed is exceeded, and increased if the burner pressure falls below the minimum value. The fan surge margin also needs to be limited. During a gross thrust decrease the surge margin decreases quickly as the fan pressure ratio increases. Then, with the surge margin limited, the much slower fan rotor speed follows along the surge limit line until the desired operating point is reached. The ability to maintain the fan surge margin greatly depends on the available control effectors.

Overall, in [49, 47] two controller designs were reported, the first being the $N2$ control loop design. This control design strives for independent control of the fan speed and the three thrusts: aft nozzle thrust ($FG9$), ejector thrust (FGE), and ventral nozzle thrust (FGV). The three thrust commands from the airframe control system are used to generate a total gross thrust command, FGT_c. The FGT calculation is relative to the aft nozzle (i.e., the thrust that would be generated if all the mass flow from all the nozzles passed through the aft nozzle). FGT_c is used to schedule a command value for the fan speed. The second control design strives for independent control of the pressure difference/pressure (DP/P, related to the fan surge margin), and the three thrusts. The steady state value for $N2$ is determined indirectly, since the three thrusts determine the total gross thrust, and FGT and DP/P establish $N2$. If a steady-state optimal criterion is used to design the fan speed schedule (specific fuel consumption, for example), then a DP/P schedule that is generated using the trim values for DP/P at steady state operating points determined by the $N2$ schedule will satisfy the same optimal criterion in steady state. If the DP/P limit is reached, then all of the nozzle areas are opened as a percentage of their contribution to the total nozzle area. Thus, the fan surge margin is shared over all three nozzles and all three thrusts are affected when the fan surge margin is reached. These schemes produced good results when tested in nonlinear performance.

[59] adopts a multi-mode structure for the controller design in order to preserve structural integrity of a Rolls-Royce Spey Mk202 engine by limiting engine variables to specified safe values. The efficiency of the engine, and the thrust produced, depends on the pressure ratios generated by the two compressors. If the pressure ratio across a compressor exceeds a certain maximum, it may no longer be able to hold the pressure head generated and the flow will tend to reverse its direction. Since the principle of operation of an aero-engine is based upon compression, combustion and expansion of air, pressures and temperatures at certain stations inside the engine body can, in some conditions, rise to dangerously high levels. This can have an adverse effect on the life span of the engine, and in the extreme case can cause catastrophic failure or engine blowup. The control scheme aims to:

- control engine thrust by controlling the ratio of the high pressure compressor's outlet

pressure to engine inlet pressure ($PS6PS1$),

- control the low pressure compressor's surge margin by controlling the Mach number indicator (ΔP/P) at its bypass exit (DP/P),
- control high pressure compressor's non-dimensionalised spool speed ($NHPCSL$),
- limit low pressure compressor's spool speed (NL), and
- limit the total temperature at the hot duct entry ($TT15$).

With three inputs one can independently control only three outputs. Three controllers are thus designed - a primary thrust controller and two limiters. These three controllers may then be switched between, depending upon which controlled output is more significant at any given time. The rationale behind this strategy is that it is the large reference demands on the thrust output which cause the NL and/or $TT15$ limits to be exceeded. Higher thrust levels require greater fuel flow, and hence higher temperatures and speeds. The reference demands for NL and $TT15$ are therefore set to their maximum values for the two limiters, and then use lowest-wins logic to select the controller that demands the lowest fuel. Thus while the thrust controller asks for the least fuel, as would normally be the case, it would be on-line, but when its fuel demand exceeds that of either one of the limiters, the implication is that the particular limit is about to be violated, and hence that limiter will be selected on-line. The limiter will hold the engine constant at the maximum value until the thrust demand is relaxed, and the thrust controller resumes control. The controller will thus have three modes of operation: a primary or thrust control mode, a NL limiting mode, and a $TT15$ limiting mode. The individual controllers are designed using the method of two-degree-of-freedom $\mathcal{H}_\infty$ loop shaping. These have been tested on the actual Spey engine at the engine test facility at DERA Pyestock. The results were very good across the range of engine operation, with the controllers designed at the nominal point being sufficiently robust to ensure no additional controllers were required, as opposed to the six or more controllers normally scheduled between the idle and full power range of operation.

1.1.9 Recent progress using advanced control methods

When systems are not integrated, each subsystem must be able to operate in a worst-case combination with the other subsystems, meaning large operating margins are required. Integration allows these margins to be reduced when the full margins are not required, resulting in higher thrust, lower fuel consumption, greater manoeuvrability, increased range, or better safety and reliability. In a conventional engine control system, the engine stall margin is large enough to

accommodate the worst case combination of engine and airplane induced disturbances. In the advanced digital engine control system (ADECS), carried out by the NASA Ames Research Center's Dryden Flight Research facility, the use of an integrated approach allowed the stall margin to be modulated in real time based on the current engine and aircraft requirements and flight conditions [53]. This permits the unneeded stall margin to be traded for increased engine performance by increasing thrust, reducing fuel flow, or lowering operating temperatures. The exchange between unneeded engine stall margin and increased engine performance is implemented by uptrimming the engine pressure ratio (EPR). Engine thrust was increased by as much as 10.5% at subsonic flight conditions by uptrimming the EPR. The additional thrust significantly improved the aircraft performance. Rate of climb was increased by 14% at 40,000 ft and the time to climb from 10,000 to 40,000 ft was reduced by 13%. A 14% increase in acceleration at 40,000 ft and 24% increase at 50,000 ft was obtained at intermediate and maximum power, respectively. Other benefits that may be realised are improved inlet stability, and reduced engine temperatures, propulsion system drag, trim drag, and control surface size. [72] also showed that using the propulsion system to augment the flight controls reduced the take-off gross weight of a supersonic cruise aircraft by up to 7 percent. In a further ADECS study, more precise inlet control was obtained while inlet stability margins were reduced. The overall result of the flight research was that range was increased by 5 percent. Altitude control capability was improved by an order of magnitude as compared with manual control. These improvements were achieved through individual applications of integration between the airframe and engine, such as uptrimming the engine pressure ratio, or by varying the inlet geometry. The aim of a complete integrated flight and propulsion control would be to consider all of the benefits in a single system.

In order to change, for example, the angle of attack of an aircraft, the pilot would traditionally adjust the elevator by applying a forward push or backward pull on the centre stick. As the flight speed increases, stick forces increased significantly and exceeded normal pilot capabilities [57]. This led to the development of powered controls. Initially, power was supplied only to assist the pilot by means of hydraulic boosters. With such a system, a proportion of the stick force was still needed to be applied by the pilot. If a booster failed, the system would revert to full manual control. Aircraft using such semimanual systems include the Boeing 707 series and the B-52 Stratofortress Bomber. As flight speed increased further, it became more and more difficult to keep the stick forces within the capability of an average human pilot, and this led to the development of fully powered control systems. With such systems, the pilot hardly does anything more than operate a small hydraulic valve, requiring a small nominal force. The stick force virtually remains constant irrespective of the speed or the deflection of the control surface.

However, to give the pilot some idea of the stick forces involved, an artificial 'feel' is introduced to the system. This may take the form of a simple spring providing an increased resistance with increase of control column movement, or it can be a complex simulator providing forces proportional to the airspeed or Mach number. In the same way, these 'artificial' control systems may be used to change the characteristics of how the individual aircraft handle, resulting in a standard 'feel' to all aircraft. It also gives additional freedom in how the actuators are used. Rather than simply using the centre stick to deflect the elevator to achieve some change in the aircraft pitch, a force applied to the centre stick may be used to command changes in both the elevator angle and the pitch reaction control system.

Active control technology (ACT) refers to aircraft that have flight control systems that do not use a mechanical link between the control surfaces and the cockpit controls. Such aircraft usually incorporate flight control computers that allow modification of the response to control inputs. Another generic term for this type of flight control system is 'fly-by-wire', but this implies that the signals are electrical and are transmitted through wires, whereas fibre optics may in fact be used [43]. ACT has become the basis for the flight control system design on many new commercial and military aircraft. Ideally ACT eliminates the compromise between good handing qualities, which require a large tail and forward centre of gravity (c.g.) and good performance, which requires a small tail and aft c.g.. Such ideal handling qualities may require the use of very high gains to make the aircraft respond 'naturally' regardless of the configuration. It can be shown that there are three factors that prevent this ideal situation: (1) what constitutes a 'natural' response is not well understood; (2) filtering necessitated by flexible modes, noise, controller characteristics and digital processing limits the use of very high gains and (3) high gains can only be achieved with adequate control power (large tail). For conventional aircraft, the handling qualities are established primarily by the configuration. This results in a compromise with performance (primarily tail size). With ACT, the shape of the response to a control input can be drastically modified from a conventional aircraft. It is important to realise that historically the response of aircraft was not something that could be significantly altered. A great deal of effort was expended to extend the centre-of-gravity range, or to decrease the required tail size just a few percent. In the end, the flying qualities resulted from a compromise with performance, and by today's standards were not good. At forward c.g., the stick forces tended to be very heavy and the gust response was excessive. At aft c.g., the stability was often near neutral. Hence, the answer to the question of what is the best response is not simply to make the aircraft fly 'naturally'. In fact the best response shape has been found to be dependent on task. By using the high throughput and large memory of current on-board computers, along with high-speed communication data buses, it will be

possible to implement a wide range of integrated flight and propulsion control modes onboard future aircraft. Such modes will increase total vehicle effectiveness without significant weight or cost penalties. Apart from providing basic stabilisation and integrated management of the flight and powerplant control systems, the introduction of ACT will also provide the opportunity to add control features to reduce the pilot workload further during low-speed flight.

1.2 Summary of advances leading to IFPC

The use of advanced aircraft actuators has been motivated by future multi-mission, multi-function aircraft which will rely increasingly on propulsion system generated forces and moments to augment those generated by the airframe control surfaces. It has been demonstrated how this will enable the aircraft to operate at higher angles of attack and in conditions of low dynamic pressure. The idea that the redundancy resulting from an increased number of actuators could be used to aid reconfiguration following actuator failure, or to the scheduling of actuators according to flight conditions, has been introduced. Alternatively, it has been seen that the extra actuators may be used to reduce the size of, or even eliminate altogether, some of the conventional airframe control surfaces. Further advantages of the use of the advanced actuators such as improved manoeuvrability, a reduction in approach speeds and reduced fuel consumption and emissions have been discussed. The possibility of using an integrated approach to reduce the pilot workload has been considered. It has been shown that although the current Harrier already uses engine generated forces, its thrust vectoring acts through the centre of gravity of the aircraft, minimising the coupling of controlled outputs. However, this still results in a high pilot workload, particularly due to its 3-inceptor control strategy. It has been argued that adopting an integrated approach to the control design, using active control technology, would enable this to be simplified to a 2-inceptor strategy with no coupling into uncommanded outputs. Some control power requirements that would be necessary for advanced manoeuvring at low speeds have been considered. The main objectives in aircraft and engine control have been discussed. It is shown that airframe control is concerned with the position, rates and orientation of the aircraft, achieved using control systems that are typically regulators with input shaping. In contrast, the main aim of engine control is multi-objective protection, with transient control governed by operating limits. Some examples of advanced control techniques applied to engine control have been presented.

1.3 Contributions of the thesis

The main contributions of this thesis are:

1. Design of a 2-inceptor longitudinal centralised integrated flight and propulsion controller using the Spey-WEM aircraft concept model.

2. Piloted simulation of this centralised IFPC system at DERA Bedford using the real-time all-vehicle simulator (RTAVS).

3. Design of a thrust vectoring controller that does not use either the elevator or the reaction control system.

4. Use of μ-analysis to determine at which points in the flight envelope stability may be guaranteed, where uncertainties may be due either to aircraft parametric uncertainties, or to dynamic uncertainties determined by the point in the flight envelope.

5. The use of the ν-gap metric to investigate the limits of stable operation of a single controller within the flight envelope.

6. Partitioning of the centralised IFPC controller into airframe and engine subcontrollers, in order to address practical implementation issues.

7. Implementation of an engine limiting scheme to ensure safe operation of the overall IFPC system.

The results of (1) have already appeared in [6]. The design of a centralised controller allows the interactions between the airframe and engine subsystems to be considered at the design stage. Using a multivariable robust design method allows the use of a two-inceptor strategy, rather than the normal three-inceptor strategy which results in a high pilot workload. The results of (2) have already appeared in [7]. The piloted trials were carried out using the DERA Bedford real-time all-vehicle simulator (RTAVS). The comments by the pilots were favourable, awarding level 1 and level 2 handling qualities ratings. The findings of (3) have appeared in [33]. This work involved the design of a centralised controller that did not use the elevator or the reaction control system. This configuration may be used to simplify the scheduling task, or it may be used in the event of an actuator failure. The nonlinear simulation results, when compared with the full centralised system, compared well, with little degradation in performance. The results of (4) and (5) has been submitted to [30] and are to appear in [31] as applied to the centralised IFPC system. It is not expected that it will be possible to design a single controller for use across the whole flight envelope. This means some form of scheduling or switching of controllers will be required. An important part of this application is identifying the regions in the flight envelope where the nominal controller can no longer guarantee stability. Use of μ-analysis and the ν-gap metric allows these points to be identified. The outcomes of (6) are

to appear in [35], and have been submitted to [5]. Although a centralised IFPC system has been successfully designed, it will not be practical to implement such a high order controller on an actual aircraft. In addition, the lack of clarity associated with a centralised controller will hinder the process of flight clearance. Further, the airframe and engine would in general be designed and manufactured by different companies, each of which will be responsible for their particular subsystem. This means that the engine manufacturer will require a separate engine controller in order to allow the required extensive validation and testing procedures to take place. In order to address these issues, the centralised controller is partitioned into airframe and engine subcontrollers, connected by a specified structure. The application of (4) is also to appear in [35] as applied to the partitioned IFPC system. The results of (7) have already appeared in [32], and are to appear in [34]. Although four of the engine internal variables have been included as controlled outputs of the system, it is known this will not be sufficient to guarantee all of the variables will remain within their limits for all possible pilot demands. Indeed, it is not intended to act as such, since this would limit the achievable performance of the airframe controlled varaibles. Therefore, some additional nonlinear control is required in order to protect the engine from dangerous over stress and over temperature. Overall, the partitioned system with engine protection scheme implemented was seen to provide good nonlinear simulation performance results.

Previous work done by the author in this area includes the design and partitioning of a 3-inceptor centralised controller designed for the Spey-WEM aircraft [8]. This control strategy did not have the advantage of significantly reducing the pilot workload, as it followed the same strategy as is used for current Harrier control, whereby the right-hand centre stick is used to demand pitch rate, and the two left-hand inceptors are used to demand thrust magnitude and thrust direction. μ-analysis was used to find at which point in the flight envelope stability could no longer be guaranteed by the nominal controller [40] and at which point performance could no longer be guaranteed [41].

1.4 Structure of the thesis

This thesis concerns the design and analysis of a robust centralised integrated flight and propulsion controller, its subsequent partitioning into airframe and engine subcontrollers, piloted simulation and the investigation of an engine limiting scheme. After this introductory chapter, the thesis is structured as follows:

Chapter 2 : Integrated flight and propulsion control

The advantages and disadvantages of centralised, decentralised, and direct decentralised IFPC

designs are discussed. It is concluded that a combination of the centralised and decentralised approaches, exploiting the positive aspects of each, would produce the most useful approach to IFPC. Therefore, it is decided to design a centralised controller, to ensure the coupling between the subsystems is accounted for, and to partition this centralised controller into airframe and engine subcontrollers. The partitioning allows the benefits of a decentralised approach to be replicated, which include a more transparent structure and it allows the engine manufacturer a subcontroller with which to carry out the necessary validation and testing of the engine to ensure its safe operation.

Chapter 3 : The Spey-WEM V/STOL aircraft concept and model

The Defence Evaluation and Research Agency (DERA) Bedford aircraft simulation model, comprising of the Harrier based airframe and the Rolls-Royce Spey engine is introduced. This is expected to be representative of future V/STOL aircraft. Details of the pitch reaction control system model, the intake drag and thrust modelling, and the integration of the overall system are presented. In addition, the chapter gives a description of how the actuators are modelled.

Chapter 4 : Mathematical background

This chapter introduces some of the mathematical background that may be useful in the understanding of the $\mathcal{H}_\infty$ loop shaping controller design, μ-analysis and ν-gap metric analysis. The $\mathcal{H}_\infty$ loop shaping design procedure is shown to follow classical notions of open-loop shaping to achieve closed-loop performance, while guaranteeing stability to coprime factor uncertainty. μ-analysis is shown to detect left-half plane poles crossing the imaginary axis, so indicating the onset of instability. This may be used to detect perturbations to the nominal system that will cause instability. These perturbations may be due to changes in the flight envelope or to parameter variations. The ν-gap metric is seen to be a natural companion to the design method of $\mathcal{H}_\infty$ loop shaping. A significant advantage of this approach is that it allows the calculation of a frequency dependent result, unlike the gap-metric whereby unnecessarily conservative conclusions may be reached.

Chapter 5 : Design and piloted simulation of a centralised IFPC system

The performance specifications provided by DERA Bedford are presented. These include the maximum output coupling allowable between airframe controlled outputs and safety limits which internal engine variables must remain within. A centralised controller design is presented, with a discussion of the choice of scaling matrices and weighting functions. Details are given of the real-time all-vehicle simulator (RTAVS) at DERA Bedford on which a piloted simulation was carried out. The piloted simulation results are presented, in terms of time

history plots, pilot comments and pilot handling qualities ratings. In addition, details of a controller designed for an alternative aircraft configuration that does not use the pitch reaction control system or the elevator is presented. Such an alternative configuration may arise from an actuator failure for example. The approach may also be used to simplify the scheduling problem.

Chapter 6 : Partitioning of the centralised IFPC system

Although a centralised controller is designed, it will not be practical in terms of implementation and flight certification. Furthermore, the engine and airframe are manufactured by separate manufacturers, and so a separate engine subcontroller is required to ensure the engine manufacturer is able to carry out the required extensive validation and testing. The method used to partition the centralised system is presented in this chapter. It is shown that while the partitioned system results in greater transparency with lower order subcontrollers, the performance matches well with that of the centralised system.

Chapter 7 : Robustness analysis

In this chapter, μ- and ν-analyses are applied to both the centralised system and the partitioned system in order to explore the robust stability in the face of parameter uncertainty and varying flight dynamics. When applied to the centralised system, the results of the μ-analysis, ν-gap analysis and the nonlinear simulation results give reasonably consistent results. Although the partitioning appears to be successful when considering the performance of the nonlinear simulation, when a μ-analysis is carried out it becomes clear that the system has lost the robustness properties associated with the centralised system. An important feature of the partitioning method is the ease with which the subcontrollers may be redesigned. Once the engine subcontroller structure is modified, the final partitioned system is shown to preserve both the time domain performance and the robust stability characteristics of the centralised system.

Chapter 8 : Limiting of safety-critical engine variables

The technique used to ensure the internal engine variables remain within safe limits is presented. The specifications comprise of maximum or minimum limits for 10 internal engine variables, which ensure safe operation of the engine. The technique presented is based on the theory that there are many combinations of engine actuator positions that will result in, for example, the same internal engine temperatures. If an internal engine variable approaches a limit, the actuator positions just prior to the limit being violated are noted. If the engine variable continues towards its limit, the noted actuator positions serve as temporary, artificial

position limits. This approach enables variations within the engine, for example due to manufacturing tolerances or engine wear, to be accounted for on-line. The nonlinear simulation results of this approach implemented on the partitioned system are presented, and it is shown to be effective.

Chapter 9 : Conclusions and further work

The main conclusions of this thesis are stated. Some suggestions of further work in order to improve on the current designs and to increase the flight envelope are presented.

Chapter 2

Integrated Flight and Propulsion Control

This chapter will motivate the Integrated Flight and Propulsion Control (IFPC) approach used. § 2.1 will consider why recent actuator advances have resulted in the need for an integrated flight and propulsion control approach, while § 2.2 will compare the advantages and disadvantages of a centralised approach with decentralised approaches. § 2.3 provides a summary of this chapter.

2.1 Approaches to Integrated Flight and Propulsion Control

Chapter 1 introduced some of the advances in actuator technology that have resulted in an increase in the generation of forces and moments by the propulsion system used to directly affect the airframe. In this particular study, both thrust vectoring nozzles and reaction control system jets are used. These may be used in order to achieve significant performance gains by increasing the manoeuvrability and the flight envelope of the aircraft. However, they will also result in greatly increased interactions between the airframe and engine that will need to be considered during the design stage. This will be particularly important at low speeds where the effects of the propulsion system generated forces and moments will be most significant. Therefore, if the benefits are to be realised, a different approach to control design will be necessary. It has traditionally been possible to design the airframe and engine controllers separately, with the resulting independent subcontrollers implemented together for the first time on the aircraft itself. It is the pilot that effectively integrates the independent systems by judging, for example, what change in thrust will be necessary in order to facilitate a desired change in flight path angle. Engine design specifications and performance requirements have traditionally been driven by the vehicles in which they were to be installed, as an engine's thrust-to-weight ratio and airflow requirements are major factors which affect an aircraft's configuration [60]. In addition, vehicle configurations and mission capabilities have been affected by the propulsion system's achievable performance. Any interactions that did occur between the airframe and

engine subsystems, such as the effects of altitude on intake flow, were treated as disturbances to be rejected. Independent optimisation of each control system is therefore usually compromised by worst case assumption of the other systems, which may be rather conservative and limiting in terms of achievable performance [53]. This has, however, been a valid approach, as the interactions between the two have not been significant enough to warrant further effort. This is true even for the traditional Harrier, in which propulsion system generated forces and moments are applied through the use of thrust vectoring, as the thrust vectoring acts close to the centre of gravity. Although the configuration results in a high pilot workload, it is still possible for the aircraft to be flown without stability augmentation provided it is kept within limited conditions [68]. There are two main routes that may be taken in order to address the issue of increased coupling between the airframe and engine: the use of either a centralised or a decentralised design approach. Details of the two methods are considered in the next section.

2.2 Centralised vs Decentralised approaches to IFPC design

The first approach we shall examine is the design of a centralised controller that considers the overall integrated system as the design plant, as shown in Figure 2.1 [29]. The centralised design plant consists of the airframe, engine and all interactions between the two. This is expected to yield an optimal solution in that it will include all of the interactions within the systems which would therefore be accounted for in the centralised design. This controller may be considered to be 'optimal', in the sense of revealing the limits of performance for the overall system. One of the arguments against a centralised approach is that aerodynamic and engine control requirements are fundamentally different and that it is therefore difficult to capture all the design requirements, for example, in a single $\mathcal{H}_\infty$ cost function. Therefore, a major issue in the centralised design approach is the choice of design method. Not only should the synthesis technique allow a formulation of centralised control design criteria that adequately reflects the performance specifications of the total system - the airframe integrated with the propulsion system - but it should result in controllers of reasonable order with guaranteed performance and robustness characteristics. Robustness is of special importance because there are many modelling uncertainties and errors associated with the design plant due to neglected/unknown dynamics, nonlinearities, parameter uncertainties and actuator rate and position limits. This problem is largely overcome by the use of an $\mathcal{H}_\infty$ loop shaping design methodology. This technique has been successfully applied to industrial aerodynamic and aeroengine control problems. However, this will result in a centralised controller of very high order, which may contain unrealistic feedback paths and be complex to implement. Complex centralised controllers are difficult to certify for flight-clearance as they inherently suffer from a lack of transparency with

regard to subsystem functionality. In addition, the actual airframe and engine are still to be produced by manufacturers working independently of one another, each of which will be responsible for their respective part. It is therefore essential that the engine manufacturer has an engine controller with which to carry out the necessary extensive testing and validation in order to ensure adequate performance and integrity of the propulsion system in the presence of operational and safety limits.

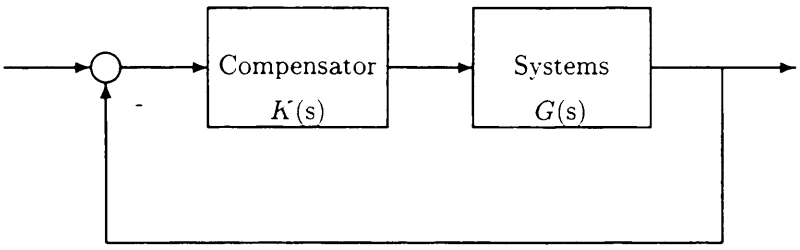

Figure 2.1: Centralised IFPC system.

An alternative approach to IFPC system design is a decentralised approach whereby the integrated plant is partitioned into loosely interconnected subsystems where each subsystem is driven by the outputs of the other subsystems [21]. Significant effort is required for this, as it may be a non-trivial task which includes modal, controllability, observability and steady-state gain analyses in order to obtain the measures necessary to group the system variables. Low order controllers can then be designed for each of the subsystems by specialists in that particular area. This will result in a clear overall structure with easy to implement subcontrollers. This decentralised structure was considered to be important in the case of ship control, in order to increase the chances of the approach being accepted quickly and to reduce difficulties with training and commissioning procedures [36]. The structure of decentralised control generally consists of a high level controller that generates specifications that the (low level) subsystems must meet in order to guarantee the overall system performance requirements are met. In other words, the command flow is downward, from the most general level, which addresses overall mission goals and requirements, to the most specific level of actuator commands to the implementation of those commands. Figure 2.2 shows this structure, whereby the compensator $K_m(s)$ is the mission level compensator that addresses only the control of the airframe outputs. $K(s)$ is a subsystem level compensator (or group of compensators) that addresses simultaneously the control of the engine outputs and the generation of interface variables that are those outputs of the subsystems which are treated as inputs or controls in the mission level problem. Having such a structure allows mission level and subsystem level problems to be addressed separately. Figure 2.3 shows how the mission level problem can be identified. Its goal is to design a compensator and a set of required closed-loop equivalent actuator characteristics such that mission level goals are met. These actuator characteristics can simply be assumed

achievable performance of the subsystems, or they may be required filters for loop shaping to achieve robustness.

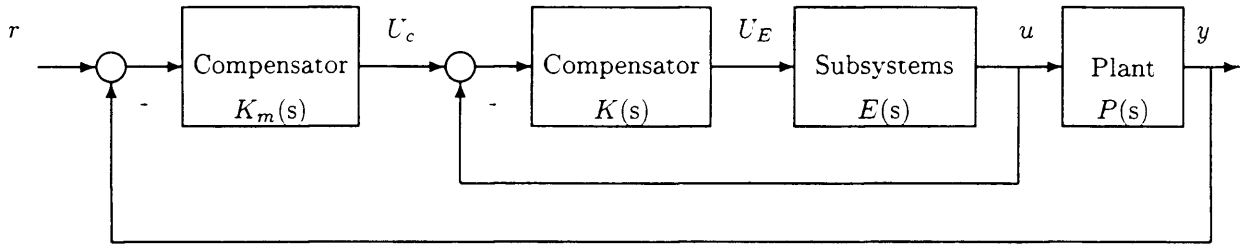

Figure 2.2: Hierarchical IFPC system.

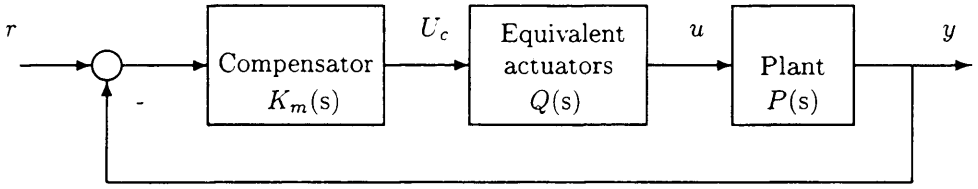

Figure 2.3: Mission level problem.

One approach to the higher level controller design is to exploit the fact that there may be several combinations of actuator positions that will produce the same overall effect on the airframe, and that at each flight condition there may be a most efficient combination of actuator positions that will produce the desired change in aircraft orientation [73]. This would result in a control selector that has knowledge of how effective each actuator is at producing, for example, pitch moments as a function of flight condition and thrust level. With this information, the control selector could use a combination of all effectors to achieve the aircraft's commanded movement, as no distinction is made between propulsive and airframe effectors. This structure is of particular use in the case of actuator saturation and actuator failure, as the control selector will simply redistribute the control across the remaining actuators.

In other words, the high-level requirements may be achieved in terms of generalised actuators [21, 61, 62, 48]. The high level regulator generates inputs in terms of generalised control commands (Figures 2.2 and 2.3). These generalised control commands are then translated to commands to individual effectors via the control selector. This method is achieved in several stages. A high level controller is designed that produces desired plant inputs. The controllers for each of the subsystems must then be designed in order to achieve these required actuator changes to ensure the demand is met. Initially, the subsystems are formulated as generalised controls, so allowing the high level and the low level designs to be considered separately. This allows the initial design of the high level controller. The generalised actuators may be the assumed achievable performance of the subsystems. There is one generalised control for each degree-of-freedom of, for example, a rigid body aircraft. The control requirements are then

mapped into specifications for the subsystems. These specifications should allow independent designs of the subsystem controllers, guarantee mission level performance and be in a form the subsystem designers can use. These subsystem specifications come in four main parts. The first is a definition of the nominal input-output characteristic for each subsystem. The second is a tolerance band within which the subsystem controllers must match the nominal. This is derived by determining the allowable level of uncertainty from stability/performance robustness measures in the mission level design. This analysis of the mission level design provides an upper bound on the magnitude of uncertainties as a function of frequency. In the case of [61, 48], it is assumed that the tolerances are based on uncertainties at the input. This means that all uncertainties are assumed to occur in the subsystems ('equivalent actuators'). This uncertainty bound may then be approximated by a (low order) transfer function matrix to produce a full matrix of fictitious uncertainties, an upper bound for the norm of which is the stability robustness measure. Hence this is a candidate matrix for representing model uncertainties. Even though this representation may not reflect the actual uncertainties present, it does provide an acceptable bound around the desired nominal characteristics. The third is a bound or constraint that must be satisfied in order to accommodate coupling that exists among the subsystems. The fourth is a definition of the maximum ranges and rates for each subsystem response. This last requirement is determined by the large amplitude response characteristics of the subsystems. For good performance, the subsystem response must match the nominal closely. Once the performance specifications for the subsystem are set, the engine compensator can be designed. If the subcontroller design cannot meet the performance specifications, then the specifications need to be relaxed. This would require a redesign of the high level regulator, the generation of a new set of subsystem specifications, and the redesign of the propulsion control. This can create a time consuming, highly iterative process. After evaluating the performance of the engine subsystem controller design, the entire integrated system model response with separate subcontrollers is evaluated. The performance of the integrated system is also compared to that of the separate individual subsystems. The majority of the work to be done using this approach comes with producing specifications for the subsystems to ensure performance, stability and robustness. One of the main drawbacks associated with this method is the fact that the interconnections between the subsystems are not easily accounted for, so the specification of loosely coupled subsystems is crucial. If this is not done properly, this approach is unlikely to produce acceptable results.

An alternative decentralised approach is the direct decentralised method [15]. When no particular structure is imposed on the controller, the resulting centralised design may produce a relatively complicated controller which is difficult to implement and validate. Applying the

direct decentralised method, the controller has an upper triangular structure imposed on it, in order to allow for the separate validation of the engine controller. Therefore, the propulsion controls are generated based only on the engine measured variables, while the flight controls are based on both engine and airframe measurements. This structure is contrary to the general decentralised approach, whereby the airframe controller generates demands for both the airframe and engine subsystems, in order to efficiently use the control power from both the airframe and engine to achieve the required airframe change. This structure is chosen to represent a well-separated propulsion system which may be analysed and tested separately. One of the drawbacks of designing a centralised controller then partitioning it is the trade that must be made between performance matching and complexity. It is this trade the direct decentralised approach is attempting to avoid. The approach of [15] is based on the connection between $\mathcal{H}_\infty$ optimal control and algebraic Riccati-like equations that result from the selected control structure, in order to satisfy the decentralised structure. The controller structure is first fixed to obtain the required closed-loop system states description. The Riccati-like equations are then solved by iterations. While this iteration scheme worked well for small-size problems, typically 4 or 5 state systems, for larger problems the procedure can consume a significant amount of computer time. The structure of the method is shown in Figure 2.4.

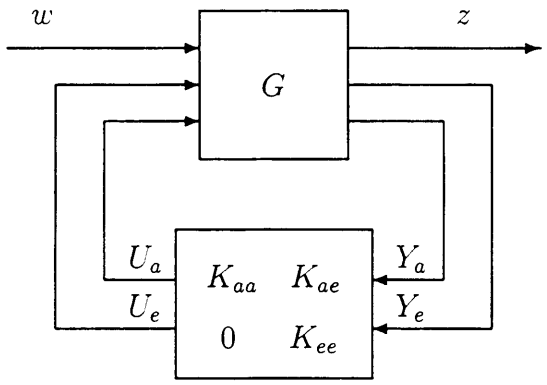

Figure 2.4: Direct decentralised scheme

One of the drawbacks of this method is that it is specific to the particular flight condition at which it is designed. In the given example, no airframe actuators are included to enable any control to be achieved in the higher speed region of the flight envelope. The airframe control is achieved purely through the use of the propulsion system actuators. The linear simulation responses of the decentralised controller are compared with a centralised controller, whereby the centralised controller works quite well and the decentralised controller results in degraded performance.

Both the centralised and decentralised approaches have individual strong points, but each has

significant drawbacks. A combination of the two approaches that exploits the best aspects of each while avoiding the problems associated with each may provide a favourable solution [27]. The decentralised method has the advantages of resulting in lower order, easy to implement subcontrollers which can easily be redesigned as long as they remain within certain specifications, while maintaining a hierarchical structure. Furthermore, it will provide a separate engine subcontroller, so enabling the engine manufacturer to carry out the required validation and testing to ensure safe operation. The centralised method on the other hand has the advantage of accounting for the interactions that occur within the subsystems of the full plant at the design stage. An ideal approach would be to design a centralised controller which may be partitioned in such a way that results in separate subcontrollers with the interactions between the subsystems accounted for. Previous work done using this approach can be found in [8, 9, 10]. Details of partitioning the centralised controller are given in Chapter 6.

2.3 Summary

This chapter has introduced the two main approaches to IFPC system design: centralised and decentralised. It has been argued that the centralised approach will result in a design that has considered all of the interactions between the subsystems. However, this would be complex to implement, difficult to certify for flight clearance and it would not provide the engine subcontroller that the engine manufacturer would need in order to ensure the safety of the engine prior to it being fitted in the airframe. The decentralised approach has been shown to result in easy to implement subcontrollers with a reasonably transparent structure. This would allow experts in the significantly different fields of airframe and engine control to design the individual subcontrollers. However, partitioning the centralised plant into subsystems may be complex and time consuming, and the assumption that the subsystems are only loosely coupled may not be valid. It is concluded that a combination of the two approaches would enable the good points of each to be captured while avoiding the negative aspects of each.

Chapter 3

The Spey-WEM V/STOL aircraft concept and model

This chapter will introduce the DERA Spey-WEM nonlinear simulation model. § 3.1 serves to indicate the physical location of the actuators on the standard Harrier, and introduces some of the airframe measurements used. § 3.2 will give a brief overview of the modelling of the subsystems, including that of the airframe, engine, reaction control system, intake drag and thrust, and how these subsystems are integrated together, and of how the actuators have been modelled. Finally, § 3.3 provides a summary of this chapter.

3.1 Conventional Harrier configuration

Figure 3.1 [74] shows the top, side and front views of the standard Harrier with the positions of the airframe control surfaces, thrust vectoring nozzles and the reaction control system actuators indicated. Although the airframe model used in this work is based on the Harrier, the experimental configuration of the aircraft featured in this thesis is not exactly the same as that of the standard Harrier. One of the most important differences in actuator positions is the movement of the front pair of nozzles forward and down.

Figure 3.2 shows how the airframe variables are measured, where:

U	x body axis velocity	VD	vertical speed
V	y body axis velocity	VH	horizontal ground speed
W	z body axis velocity	VT	velocity along flight path
α	angle of attack	X	total x body axis force
γ	flight path angle	Z	total z body axis force
θ	pitch angle		

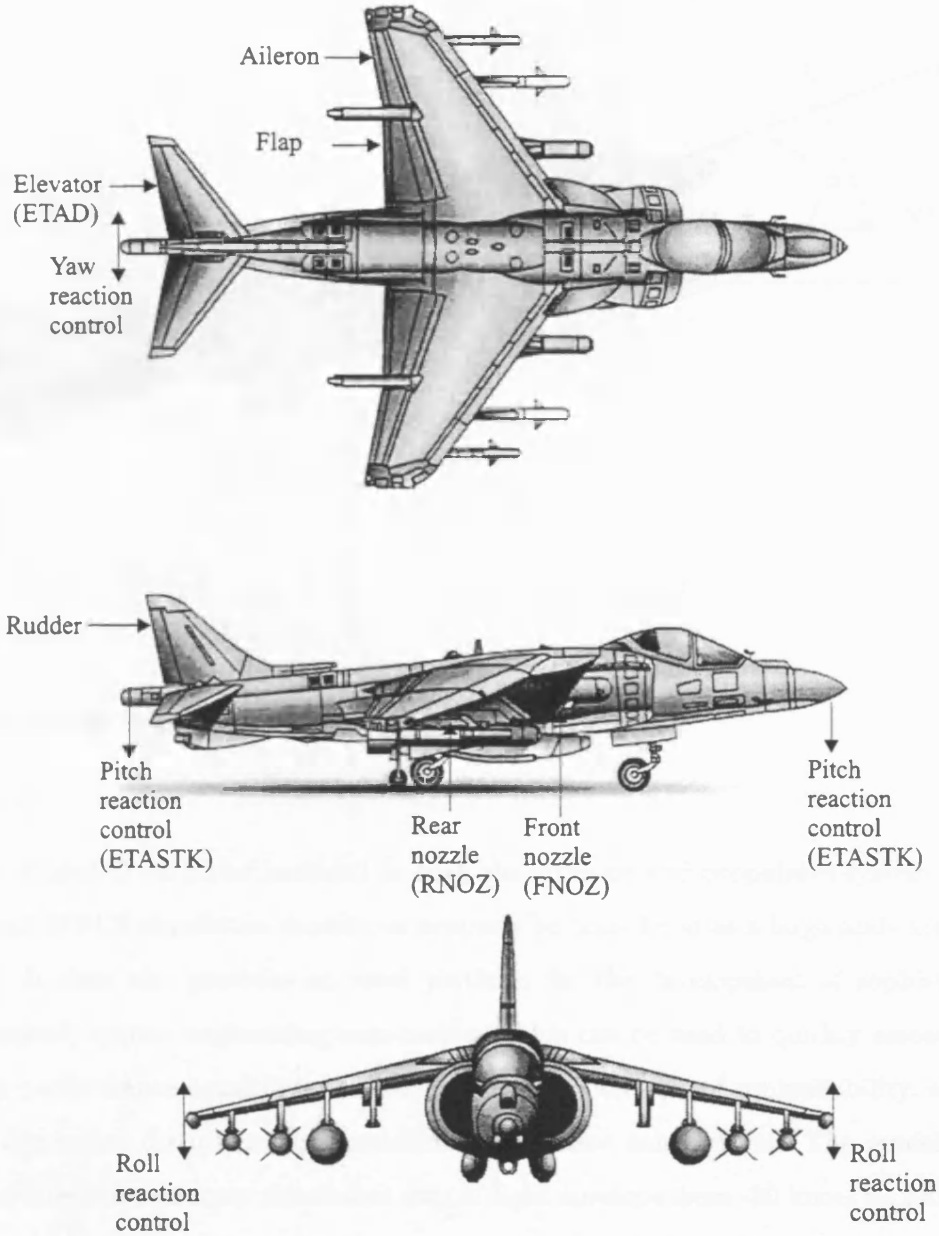

Figure 3.1: Actuator positions on a standard Harrier.

3.2 Spey-WEM model details

The aircraft simulation model used in this study has been developed at DERA Bedford in order to investigate the problems and opportunities associated with the integration of flight and propulsion control systems for V/STOL aircraft. The DERA Bedford IFPCS model provides a high fidelity dynamic representation of the airframe/engine interactions encountered in real STOVL flight configurations, whilst not specifically simulating any single aircraft. It thus provides an excellent platform for evaluating different approaches to integrated control system design, both in synthesis and analysis. Due to the degree of modelling detail and the

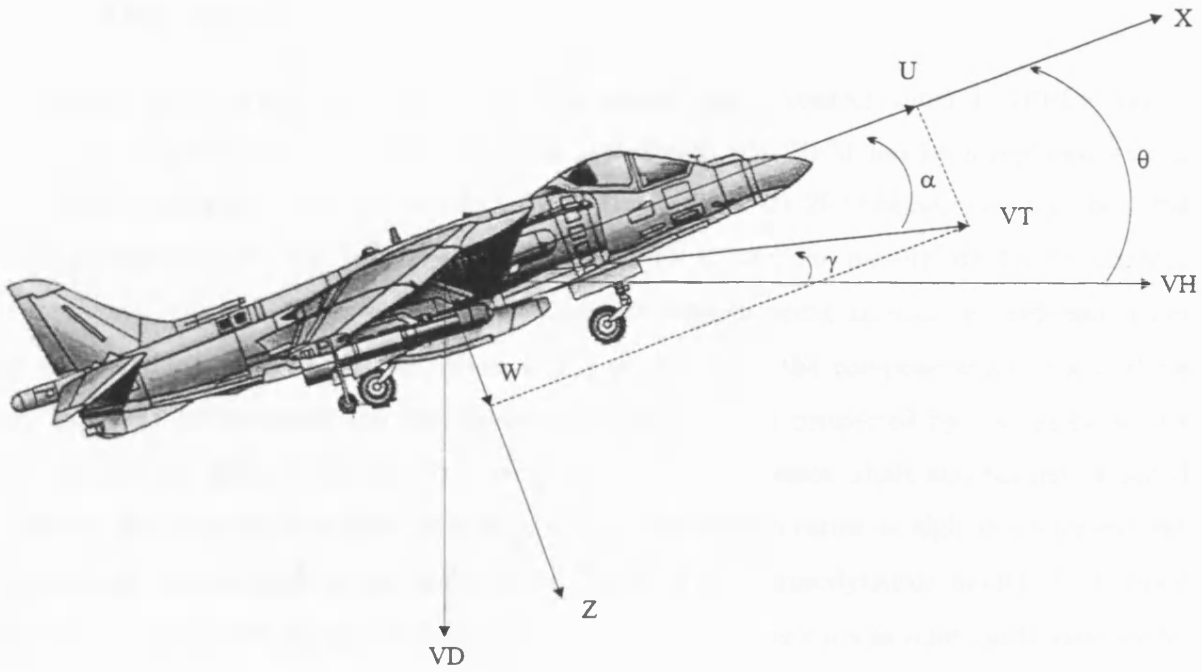

Figure 3.2: Airframe measurements.

number of control variables included in both the airframe and propulsion system simulations, the overall IFPCS simulation model can properly be considered as a large scale interconnected system. It thus also provides an ideal platform for the development of sophisticated (and quite general) system engineering approaches, which can be used to quickly assess a proposed system's performance capabilities and deficiencies, by analysis of controllability, observability, system dynamics, disturbances, interactions and sensor noise effects. The model offers a six degree-of-freedom nonlinear simulation over a flight envelope from -20 knots to 250 knots. Linearised models for controller design purposes can be generated over the full flight envelope. The following subsections will describe the airframe model, the engine model and the integration between the two. The full six degree-of-freedom modelling is very detailed, including details of how conditions differ according to altitude.

3.2.1 Harrier based airframe

The airframe model used is based on the nonlinear DERA Bedford Harrier T.Mk4 wide envelope model (WEM). This model has been used extensively in the vectored thrust aircraft advanced flight control (VAAC) Harrier research programme [58, 9, 10, 69, 68] and has been established through flight trials as being representative of the real aircraft [44]. Atmospheric data is varied according to altitude. The airframe model control surfaces are the elevator, aileron, rudder and flap.

3.2.2 Spey engine

In order to fully explore the possibilities for advanced engine control under an IFPCS framework, the original Pegasus engine previously included in the WEM has been replaced with a high fidelity thermodynamic model of the Rolls-Royce Spey Mk 202 engine. This has the same basic architecture, for the purposes of control, as the EJ200, the powerplant for Eurofighter [18]. Again, this model has been verified through tests as being an accurate representation of the real engine [59, 63, 2]. As shown in Figure 3.3, both the compressor and the turbine are split into low-pressure and high-pressure stages, and are connected by concentric shafts which rotate at different speeds. Each combination of compressor, shaft and turbine is called a spool. The Spey has a by-pass ratio of only 0.6, compared to ratios as high as 4.8 for current commercial engines such as the Rolls-Royce Trent. The thermodynamic model of the Spey allows the control law designer full access to engine parameters such as inlet guide vane angle, fuel flow rate and exit nozzle area. The inlet guide vane angle is used to match airflow to the fan characteristic, as the inlet guide vanes direct the flow to suit the first-stage rotor. The exit nozzle area may be used to optimise the engine cycle over a range of flight conditions with regard to thrust, specific fuel consumption and engine life, assisting in the reduction of life-cycle costs. The engine thrust is vectored through four nozzles similar to the standard Harrier. The effect of front/rear thrust split on engine performance is assumed to be negligible. The Pegasus engine provides up to 9kgs^{-1} of compressed air for the pitch, roll and yaw reaction control system at negative, zero and 'sub-aerodynamic' forward speeds of the aircraft [3]. This represents about 5% of the engine throughput of 200kgs^{-1}, or about 11% of the hot-core throughput. Total thrust and high pressure bleed flow to the reaction control system (RCS) is scaled to match Pegasus performance and no duct losses are modelled in the rotating nozzles. The engine model actuators are main fuel flow, exit nozzle area and inlet guide vane angle. The engine model has a total of 10 inputs and 15 states and produces 38 outputs. Two of the engine model outputs (high pressure compressor exit temperature and the square root of the temperature at the combustor exit) are inputs to the reaction control system model, one engine model output (total thrust) is an input to the thrust subsystem and one engine model output (intake mass flow rate) is an input to the intake drag model.

3.2.3 Reaction control system

The reaction control system (RCS) is modelled as a separate subsystem, having a total of 16 inputs, including the pitch, roll and yaw reaction control system actuator positions, the average nozzle angle, and engine model outputs high pressure compressor exit pressure and the square root of the high pressure compressor outlet temperature. One of the RCS model outputs (air

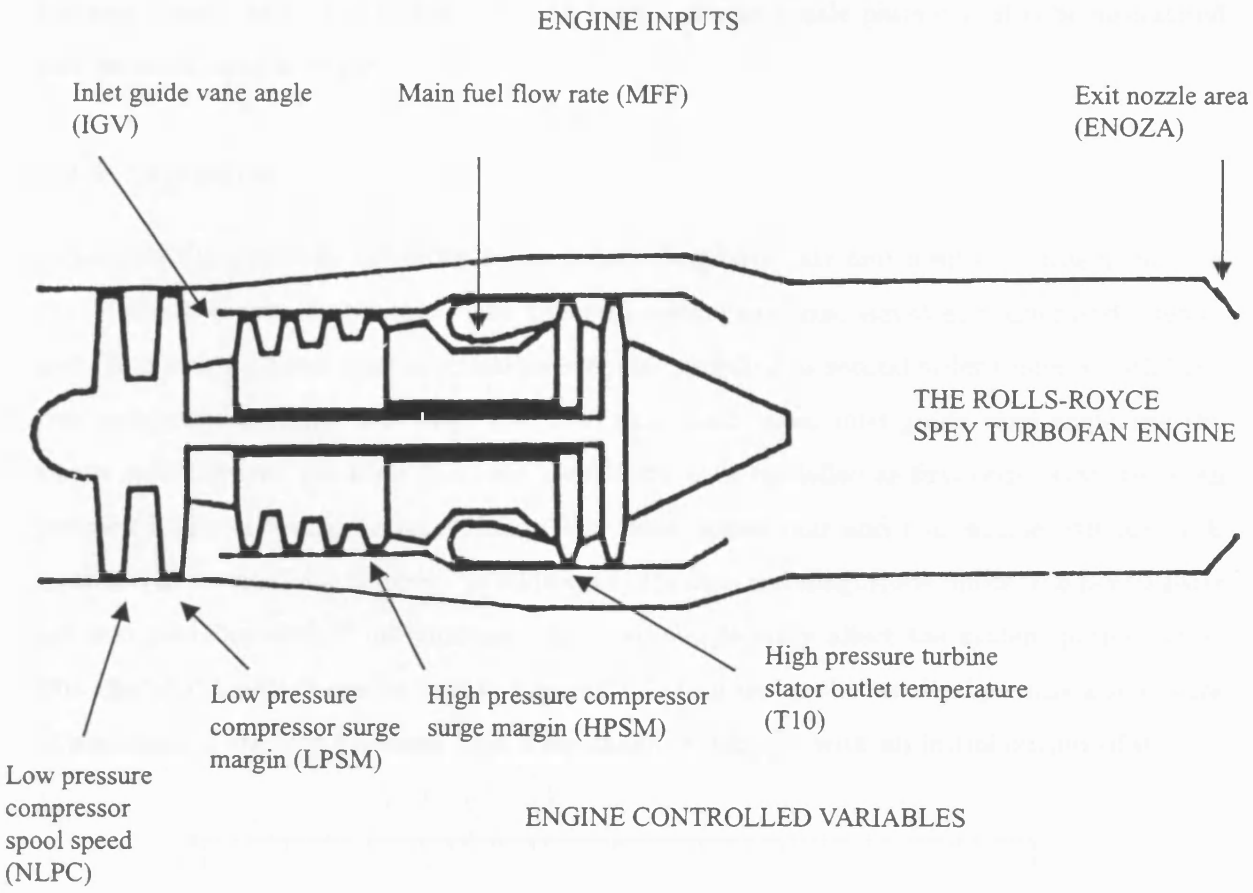

Figure 3.3: Schematic of the Spey engine.

mass flow rate) is fed back as an input to the engine (high pressure bleed flow, in kgs^{-1}, scaled by -1 to give the fraction of bleed lost overboard from exit of high pressure compressor).

3.2.4 Intake drag

The intake drag model inputs include the intake mass flow rate (an engine output), scaled to that of a Pegasus, along with VT and α. The prime requirement of the intake is to minimise the pressure loss up to the compressor face while ensuring that the flow enters the compressor with a uniform pressure and velocity, at all flight conditions. Non-uniform or distorted flow may cause compressor surge which can result either in engine flame out or severe mechanical damage due to blade vibration induced by unsteady aerodynamic effects.

3.2.5 Thrust

Thrust vectoring is also modelled as a separate subsystem. This has the gross thrust output from the engine as an input, along with the front and rear nozzle pair angles and the thrust split. To increase design difficulty, the front pairs of nozzles has been moved forward and downward to displace the centre of thrust from the centre of gravity, which introduces thrust/pitching

moment interactions. The thrust from the front and rear nozzle pairs can also be modulated and vectored independently.

3.2.6 Actuators

Representative nonlinear actuation systems including both rate and position limits have been placed on all control motivators. The airframe control surfaces, elevator, rudder and aileron, and the reaction control system actuators are each modelled as second order systems with both rate and position limits. The main fuel flow, exit nozzle area, inlet guide vane angle and the thrust split between the front and rear nozzles are each modelled as first order systems, with position limits and variable rate limits. The front nozzle pair and rear nozzle pair are each modelled as fourth order systems. In addition to the rate and magnitude limits, the nozzle pairs are also modelled with 6° of backlash which can significantly affect the system performance. The effects of backlash can be seen in Figure 3.4, which shows the result of passing a sine wave of amplitude 2 through backlash with a deadband width of 1 with an initial output of 0.

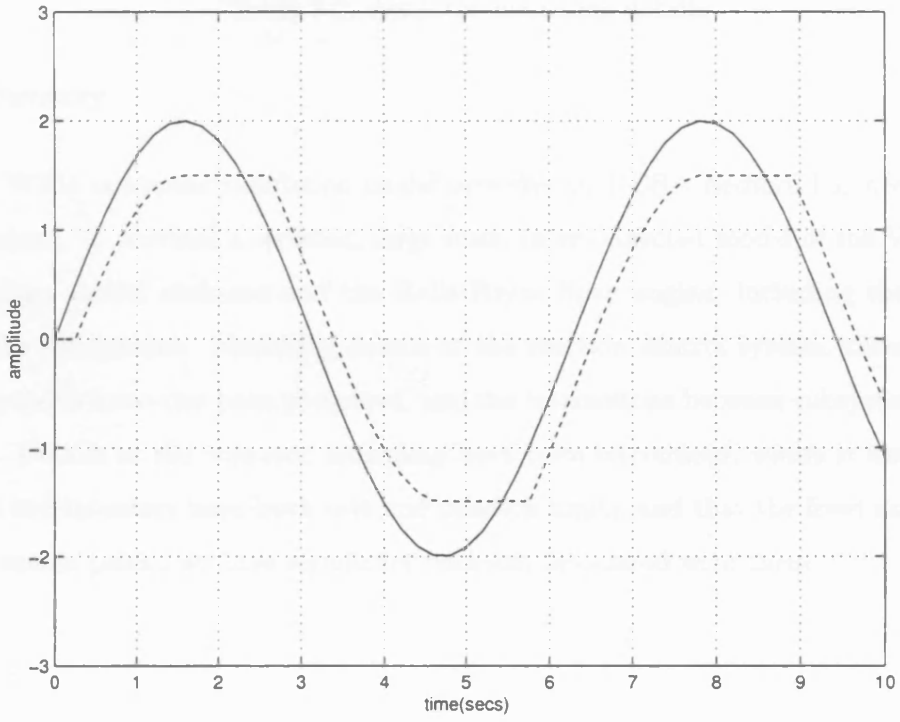

Figure 3.4: Input to backlash (—) and output from backlash (- -).

Additional details of the engine actuators are given in Table 3.1 and of the remaining actuators in Table 3.2.

ACTUATOR	position limit (lower)	position limit (upper)
main fuel flow servo	0 kgs^{-1}	1.2 kgs^{-1}
exit nozzle servo	0.1602 sine petal angle	0.8307 sine petal angle
inlet guide vane servo	-5°	35°

Table 3.1: Engine actuator modelling details

ACTUATOR	natural frequency ω ($rads^{-1}$)	damping factor ζ	rate limit (lower)	rate limit (upper)	position limit (lower)	position limit (upper)
elevator servo	31.6	0.87	-60°s^{-1}	60°s^{-1}	-15°	15°
pitch reaction control servo	31.6	0.87	-60°s^{-1}	60°s^{-1}	-15°	15°
thrust vectoring nozzle servo	11.9	0.96	-90°s^{-1}	90°s^{-1}	-5°	120°
thrust vectoring nozzle air motor	0.5	0.47	-90°s^{-1}	90°s^{-1}	-5°	120°

Table 3.2: Actuator modelling details

3.3 Summary

The Spey-WEM nonlinear simulation model provided by DERA Bedford has been introduced in this chapter. It provides a detailed, large scale, interconnected model of the VAAC Harrier wide envelope model airframe and the Rolls-Royce Spey engine, including the interactions between the subsystems. Modelling details of the reaction control system, thrust and intake drag subsystems have also been presented, and the interactions between subsystems have been discussed. Details of the actuator modelling have been introduced, where it has been shown that all of the actuators have both rate and position limits, and that the front and rear thrust vectoring nozzle pairs also have significant backlash associated with them.

Chapter 4

Mathematical background

This chapter aims to provide an overview of the main mathematical approaches applied in this thesis. § 4.1 introduces the $\mathcal{H}_\infty$ loop shaping design procedure, both in terms of augmenting the open-loop plant with weighting matrices in order to achieve closed-loop performance, and in terms of the robust stabilisation routine. § 4.2 introduces μ-analysis, which is used to detect the onset of instability due to uncertainties within a system. § 4.3 presents the ν-analysis approach, which is used to determine how close two plant are in terms of their closed-loop behaviour, and so determine if a controller will guarantee stability for both plants. Finally, § 4.4 gives a summary of this chapter, and of how these methods will be useful.

4.1 $\mathcal{H}_\infty$ loop shaping controller design

The $\mathcal{H}_\infty$ loop shaping design procedure is well documented, and was first introduced by McFarlane and Glover [50]. The plant used for the controller design is first scaled at the input and the output. The plant input scaling is necessary in order to reflect the relative actuator capabilities. The plant output scaling is used to specify the coupling requirements. If one unit of cross coupling is equally undesirable in each output, the identity matrix should be used. The design method is essentially a two stage process, details of which are now presented.

4.1.1 Open-loop shaping

The scaled open-loop plant is augmented with (generally diagonal) weighting matrices. The choice of these follows classical notions of loop shaping, whereby the singular values of the open-loop scaled plant are shaped in order to achieve the closed-loop requirements of nominal performance with good robustness, or good disturbance and noise rejection properties. In order to achieve these requirements, the open-loop plant should be shaped such that it has a high gain at low frequencies, a low gain at high frequencies and a moderate roll-off rate in gain at the desired bandwidths (generally about -20dB/decade).

The weighting matrix $W_1(s)$ is chosen to add integral action and ensure reasonable roll-off rates for the open-loop singular values around the desired crossover frequencies. The constant (non-dynamic) weighting matrix k is used to adjust control actuation requirements to respect the various actuator rate and magnitude limits. The constant matrix W_2 is used to prioritise outputs. The second stage of the $\mathcal{H}_\infty$ loop shaping design method involves the use of $\mathcal{H}_\infty$ optimisation to compute a controller block K_∞ which robustly stabilises the shaped plant, P_s, against a particular type of uncertainty description, based on stable perturbations to each of the factors in a coprime factorisation of the plant. The order of the resulting controller K_∞ is equal to the order of the weighted plant. Some trade-off is generally required between the performance and robustness requirements. The way in which the weighting functions are absorbed into the final controller K is shown in Figure 4.1.

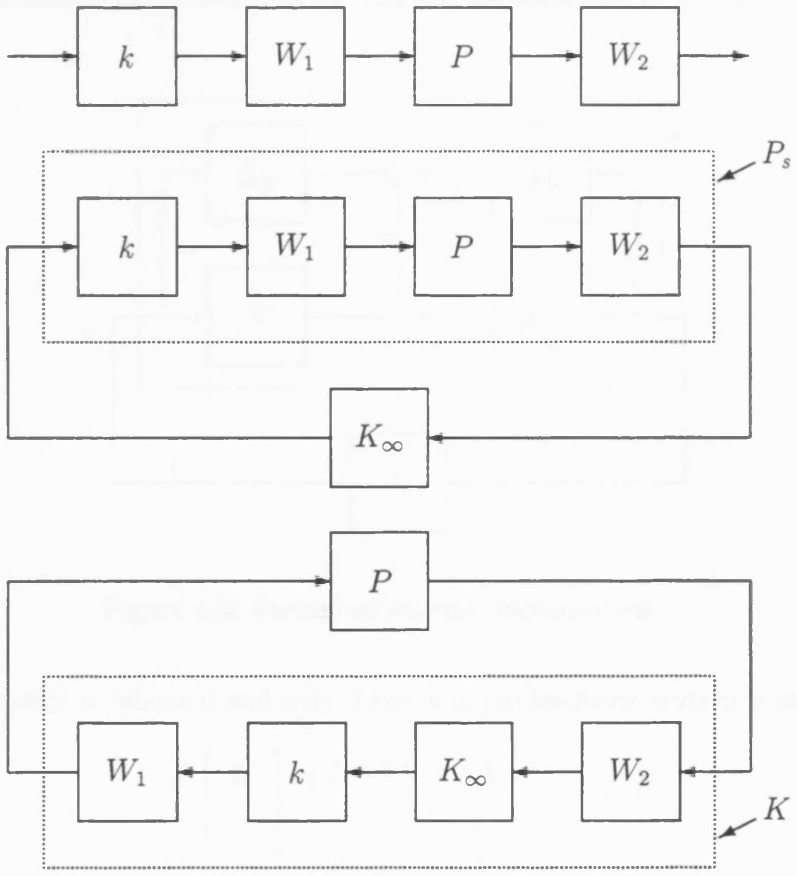

Figure 4.1: $\mathcal{H}_\infty$ loop shaping design procedure.

4.1.2 Robust stabilisation

The resulting shaped plant is robustly stabilised with respect to coprime factor uncertainty using $\mathcal{H}_\infty$ optimisation. Modelling the uncertain plant as a normalised coprime factorisation has the advantage of not requiring the perturbed plant to have the same number of unstable poles as for the nominal plant. In addition, all of the components (factors and perturbations)

are stable, even if the plant is unstable.

The normalised left coprime factorisation of the plant is given by

$$P = \tilde{M}^{-1}\tilde{N}$$

where

$$\tilde{M}\tilde{M}^\star + \tilde{N}\tilde{N}^\star = I$$

The family of perturbed plants is given by:

$$P_\Delta = \left\{(\tilde{M} + \tilde{\Delta}_M)^{-1}(\tilde{N} + \tilde{\Delta}_N) : \| [\tilde{\Delta}_N, \tilde{\Delta}_M] \|_\infty < \epsilon\right\}$$

as shown in Figure 4.2, where ϵ is the stability margin which is to be maximised. The problem is to find the largest class of perturbed systems that may be stabilised by a single fixed controller, i.e. maximise ϵ.

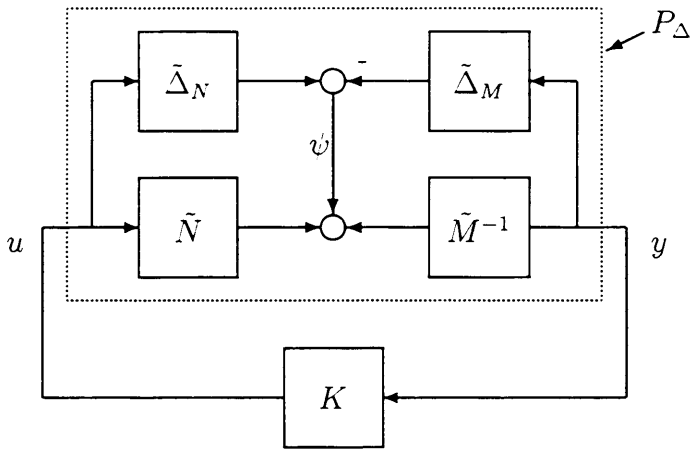

Figure 4.2: Perturbed coprime factorisation.

The stability property is robust if and only if the nominal feedback system is stable and

$$\gamma := \left\| \begin{bmatrix} K \\ I \end{bmatrix} (I - PK)^{-1} \tilde{M}^{-1} \right\|_\infty \leq \frac{1}{\epsilon}$$

where γ is the $\mathcal{H}_\infty$ norm from ψ to $\begin{bmatrix} u \\ y \end{bmatrix}$ and $(I - PK)^{-1}$ is the sensitivity function for this positive feedback arrangement. One advantage of this approach is that no γ iteration is required: the lowest achievable value of γ and the corresponding maximum stability margin ϵ are given explicitly as

$$\gamma_{min} = \epsilon_{max}^{-1} = \left(1 - \left\|[\tilde{M}, \tilde{N}]\right\|_H^2\right)^{-\frac{1}{2}} = (1 + \rho(XZ))^{\frac{1}{2}}$$

where Z and X are the unique positive definite solutions to the following algebraic Riccati equations

$$(A - BS^{-1}D^TC)Z + Z(A - BS^{-1}D^TC)^T - ZC^TR^{-1}CZ + BS^{-1}B^T = 0$$

$$(A - BS^{-1}D^TC)^TX + X(A - BS^{-1}D^TC) - XBS^{-1}B^TX + C^TR^{-1}C = 0$$

where

$R = I + DD^T$

$S = I + D^TD$

The controller is then given by

$$K_\infty = \left[\begin{array}{c|c} A + BF + \gamma^2(L^T)^{-1}ZC^T(C + DF) & \gamma^2(L^T)^{-1}ZC^T \\ \hline B^TX & -D^T \end{array}\right]$$

where

$F = -S^{-1}\ (D^TC + B^TX)$

$L =(1 - \gamma^2)I + XZ$

which guarantees

$$\left\| \begin{bmatrix} K \\ I \end{bmatrix} (I - PK)^{-1}M^{-1} \right\|_\infty \leq \gamma$$

for a value of $\gamma < \gamma_{min}$. This has the advantage of avoiding the γ iteration generally associated with $\mathcal{H}_\infty$ control.

If the resulting value of γ is small (ideally ≤ 4), this indicates that the chosen loop shapes are compatible with the design robustness requirements. This also indicates that the shaped open-loop is not changed much by the inclusion of K_∞ in the open-loop transfer function, as the degradation is limited to frequencies where the specified loop shape is sufficiently large or sufficiently small to be insignificant. A value of $\gamma = 4$ corresponds to the plant being robust to a 25% variation in coprime factor uncertainty.

A particular implementation structure, for $\mathcal{H}_\infty$ loop shaping controllers is shown in Figure 4.3 [44]. The constant prefilter $K_\infty(0)W_2$ is formed to ensure zero steady state tracking error, assuming integral action in W_1. This configuration is preferred, as it is not prone to producing overshoot to step responses, as the references do not directly excite the dynamics of K_∞. Note that K_∞ has been designed in the frequency domain for robustness objectives rather than time-domain performance.

4.2 Robustness analysis using the structured singular value, μ

The structured singular value is a matrix function denoted by μ. Its use relies strongly on the class of general linear feedback loops called linear fractional transformations (LFTs) [55]. It is

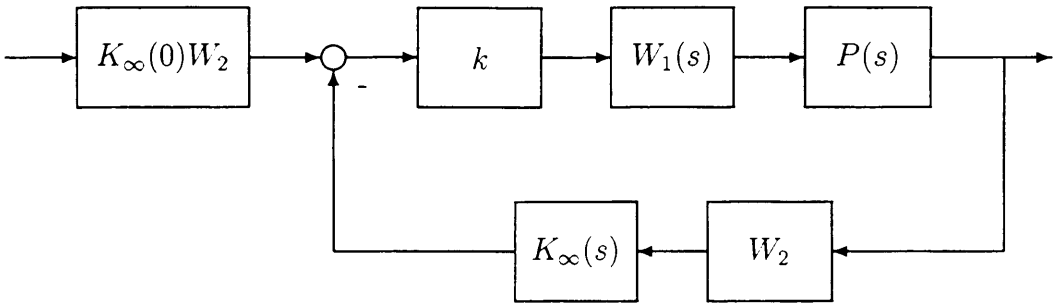

Figure 4.3: $\mathcal{H}_\infty$ loop shaping implementation structure - centralised IFPC system.

through the use of LFTs that any interconnected system can be rearranged to fit the general framework shown in Figure 4.4, which in turn may be rearranged as shown in Figure 4.5.

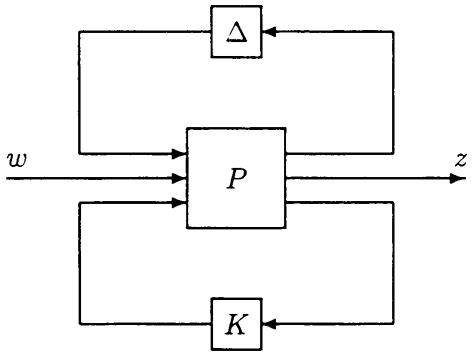

Figure 4.4: General control configuration.

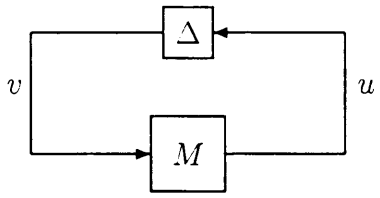

Figure 4.5: Standard block diagram for μ-analysis.

Consider the loop shown in Figure 4.5. This represents the loop equations

$$u = Mv$$

$$v = \Delta u.$$

As long as $I - M\Delta$ is nonsingular, the only solutions u, v to the loop equations are $u = v = 0$ [4]. However, if $I - M\Delta$ is singular, then there are infinitely many solutions, and the norms $||u||, ||v||$ of the solutions can be arbitrarily large. This constant matrix feedback system will be described as 'unstable'. Likewise, the term 'stable' will describe the situation when the only solutions are identically zero. In this context then, we will be looking for a measure of the smallest structured Δ that causes 'instability' of the constant matrix feedback loop shown in Figure 4.5, assuming stability in the first place.

The structured singular value is a function which provides a generalisation of the singular value, $\overline{\sigma}$, and the spectral radius, ρ [78, 70]. μ can be used to get necessary and sufficient conditions for robust stability. In the definition of μ, there is an underlying structure Δ (a prescribed set of block diagonal matrices), on which everything depends. This structure may be defined differently for each problem depending on the uncertainty and performance objectives of the problem. Defining the structure involves specifying three things: the total number of blocks, the type of each block, and their dimensions. The problem shall be examined with Δ structurally restricted. In particular, the block diagonal matrix Δ will be considered, with two types of blocks: repeated scalar and full blocks. Let S and F represent the number of repeated scalar blocks and the number of full blocks respectively. To bookkeep their dimensions, the positive integers $r_1, \ldots, r_S; m_1, \ldots, m_F$ are introduced. The i_{th} repeated scalar block is $r_i \times r_i$, while the j_{th} full block is $m_j \times m_j$. With those integers given, $\mathbf{\Delta} \subset \mathbb{C}^{n \times n}$ is defined as

$$\mathbf{\Delta} = \{diag\,[\delta_1 I_{r1}, \ldots, \delta_s I_{rS}, \Delta_1, \ldots, \Delta_F] : \delta_i \in \mathbb{C},\ \Delta_j \in \mathbb{C}^{m_j \times m_j}\}.$$

For consistency among all the dimensions, we must have

$$\sum_{i=1}^{S} r_i + \sum_{j=1}^{F} m_j = n.$$

Often, norm-bounded subsets of $\mathbf{\Delta}$ are needed, so the following notation is introduced:

$$\mathbf{B\Delta} = \{\Delta \in \mathbf{\Delta} : \overline{\sigma}(\Delta) \leq 1\}$$

To keep the notation as simple as possible in the first equation, all of the repeated scalar blocks are placed first; in actuality, they can come in any order. Also, the full blocks do not have to be square, but restricting them as such saves a great deal in terms of notation.

The definition of the structured singular value shall be motivated by asking the following question: Given a matrix $M \in \mathbb{C}^{p \times q}$, what is the smallest perturbation matrix $\Delta \in \mathbf{\Delta}$ in the sense of $\overline{\sigma}(\Delta)$ such that

$$\det(I - M\Delta) = 0\ ?$$

The principle is thus to detect the crossing of one of the closed-loop poles through the imaginary axis. That is, we are interested in finding

$$k_{min} := \inf\{\overline{\sigma}(\Delta)\ :\ \det(I - M\Delta) = 0,\ \Delta \in \mathbf{\Delta}\}.$$

It is easy to see that

$$k_{min} = \inf\{k_m\ :\ \det(I - k_m M\Delta) = 0,\ \Delta \in \mathbf{B\Delta}\} = \frac{1}{\max_{\Delta \in \mathbf{B\Delta}}\ \rho(M\Delta)}.$$

$1/k_{min}$ is called the structured singular value and it is denoted by $\mu(M)$.

Definition 4.1 *For $M \in \mathbb{C}^{n\times n}$, $\mu(M)$ is defined as*

$$\mu(M) := \frac{1}{\min\{\overline{\sigma}(\Delta) \ : \ \Delta \in \mathbf{\Delta}, \ \det(I - M\Delta) = 0\}}$$

unless no $\Delta \in \mathbf{\Delta}$ makes $I - M\Delta$ singular, in which case $\mu(M) := 0$.

Clearly, $\mu(M)$ depends not only on M but also on the allowed structure for Δ. This is sometimes shown explicitly by using the notation $\mu_{\Delta}(M)$. A value of $\mu = 1$ means that there exists a perturbation with $\overline{\sigma}(\Delta) = 1$ which is just large enough to make $I - M\Delta$ singular. A larger value of μ is 'bad' as it means that a smaller perturbation makes $I - M\Delta$ singular, whereas a smaller value of μ is 'good'. The computation of the exact value of μ has been shown to be NP hard (it may not be computable in a polynomial time), and so in practice it is not possible to compute the exact value of μ, except for very low order systems.

Several reasons exist for handling μ rather than its inverse, $1/\mu(M)$, the multiloop stability margin [22, 55]. As a first point, the structured singular value $\mu(M)$ can not take an infinite value, since the nominal closed-loop is asymptotically stable, whereas the multiloop stability margin may be infinite (if no structured model perturbation exists, which destabilises the closed-loop). Furthermore, the structured singular value can be considered as an extension of classical algebraic notions, namely the spectral radius and the maximal singular value.

The most well-known use of μ as a robustness analysis tool is in the frequency domain. Suppose $G(s)$ is a stable, real rational, multi-input, multi-output transfer function of a linear system. For clarity, assume G has q_1 inputs and p_1 outputs. Let $\mathbf{\Delta}$ be a block structure and assume that the dimensions are such that $\mathbf{\Delta} \subset \mathbb{C}^{q_1 \times p_1}$. Feedback perturbations to G that are themselves dynamical systems with the block diagonal structure of the set $\mathbf{\Delta}$ will be considered.

Let $\mathcal{M}(\mathbf{\Delta})$ denote the set of all block diagonal and stable rational transfer functions that have block structures such as $\mathbf{\Delta}$.

$$\mathcal{M}(\mathbf{\Delta}) := \{\Delta(\cdot) \in \mathcal{RH}_{\infty} \ : \ \Delta(s_0) \in \mathbf{\Delta} \text{ for all } s_0 \in \overline{\mathbb{C}}_+\}$$

where $\overline{\mathbb{C}}_+$ denotes the closed right-half plane.

Theorem 4.1 *Let $\beta > 0$. The loop shown in Figure 4.6 is well-posed and internally stable for all $\Delta(\cdot) \in \mathcal{M}(\mathbf{\Delta})$ with $||\Delta||_{\infty} < \frac{1}{\beta}$ if and only if*

$$\sup_{\omega\in\mathbb{R}} \mu_{\mathbf{\Delta}}(G(j\omega)) \leq \beta.$$

Hence, the peak value on the μ plot of frequency response determines the size of perturbations that the loop is robustly stable against.

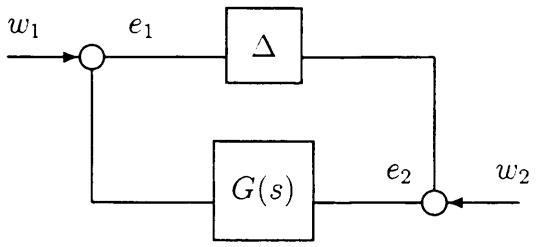

Figure 4.6: Determining loop stability.

It is this use of μ-analysis that is going to be of benefit in the design of IFPC controllers across the flight envelope. It is not expected that a single controller will be capable of delivering the required performance over all of the flight envelope, and that some form of scheduling/switching will be required. μ-analysis may be used to determine at which points the nominal controller will no longer guarantee stability. This will enable the designer to ensure stability across the flight envelope while scheduling/switching between a minimum number of controllers.

4.3 Robustness analysis using the ν-gap metric

Frequently applied open-loop methods used to quantify the differences between two plants may produce highly inaccurate results for closed-loop behaviour. An often cited example is the comparison of the following plants:

$$P_1(s) = \frac{1}{s},\ P_2(s) = \frac{1}{s+0.1}.$$

The closed-loop complementary sensitivity functions corresponding to P_1 and P_2 with unit feedback are relatively close and their difference is

$$||P_1(I+P_1)^{-1} - P_2(I+P_2)^{-1}||_\infty = 0.0909,$$

but the difference between P_1 and P_2 is

$$||P_1 - P_2||_\infty = \infty$$

This shows that the closed-loop behavior of two systems can be very close even though the norm of the difference between the two open-loop systems can be arbitrarily large.

The shortfall of the gap metric, as introduced by [82], is that is not easily related to the frequency response of the system [83]. On the other hand, the ν-gap metric introduced by Vinnicombe [76], has a clear frequency domain interpolation and can, in general, be computed from frequency responses.

First, a generalised stability margin is defined as:

$$B_{P,K} := \left\| \begin{bmatrix} P \\ I \end{bmatrix} (I-KP)^{-1} \begin{bmatrix} -K & I \end{bmatrix} \right\|_\infty^{-1}.$$

This stability margin is that found with the $\mathcal{H}_\infty$ loop shaping design procedure ($= \epsilon = \gamma^{-1}$).

The ν-gap metric can be used as an alternative method for quantifying both stability and performance robustness, particularly for controllers designed using the $\mathcal{H}_\infty$ loop shaping methodology. Given any two plants with normalised right and left coprime factorisations, $P_1 = N_1 M_1^{-1} = \tilde{M}_1^{-1}\tilde{N}_1$ and $P_2 = N_2 M_2^{-1} = \tilde{M}_2^{-1}\tilde{N}_2$, then the ν-gap metric is defined follows:

Definition 4.2

$$\delta_\nu(P_1, P_2) = \begin{cases} \|\Psi(P_1, P_2)\|_\infty, & \textit{if } \det\ \Theta(j\omega) \neq 0\ \forall\ \omega \textit{ and} \\ & \textit{wno } \det\ \Theta(s) = 0, \\ 1 & \textit{otherwise} \end{cases}$$

where $\Theta(s) := N_2^{\sim} N_1 + M_2^{\sim} M_1$, $\Psi(P_1, P_2) := -\tilde{N}_2 M_1 + \tilde{M}_2 N_1$ *and wno* $g(s)$ *denotes the winding number of* $g(s)$ *evaluated around the standard Nyquist contour.*

In effect the ν-gap metric measures the 'distance' or 'difference' between two plants from a closed-loop, as opposed to open-loop perspective. The ν-gap metric can also be computed directly from the system transfer matrices without first finding the normalised coprime factorisations [77, 78].

Theorem 4.2 *The ν-gap metric can be defined as*

$$\delta_\nu(P_1, P_2) = \begin{cases} ||\Psi(P_1, P_2)||_\infty, & \textit{if } \det(I + P_2^{\sim} P_1) \neq 0\ \forall\omega \textit{ and} \\ & \textit{wno} \det(I + P_2^{\sim} P_1) + \eta(P_1) - \eta(P_2) - \eta_0(P_2) = 0, \\ 1, & \textit{otherwise} \end{cases}$$

where $\Psi(P_1, P_2)$ *can be written as*

$$\Psi(P_1, P_2) = (I + P_2 P_2^{\sim})^{-1/2}(P_1 - P_2)(I + P_1^{\sim} P_1)^{-1/2}.$$

where $\eta(G)$ *is the number of open right-half plane poles of G(s) and* $\eta_0(G)$ *is the number of imaginary axis poles of G(s).*

This alternative formula is useful when doing a hand calculation or when computing from the frequency response of the plants since it is not necessary to compute the normalised coprime factorisations.

The following theorem states that robust stability can be checked using a frequency-by-frequency test.

Theorem 4.3 *Suppose* (P_0, K) *is stable and* $\delta_\nu(P_0, P_1) < 1$. *Then* (P_1, K) *is stable if*

$$b_{P_0,K}(\omega) > \psi(P_0(j\omega), P_1(j\omega)),\ \forall\omega.$$

where

$$\psi(P_0(j\omega), P_1(j\omega)) = \overline{\sigma}(\Psi(P_0(j\omega), P_1(j\omega))).$$

and

$$b_{P_0,K}(\omega) = \overline{\sigma}\left(\begin{bmatrix} P(j\omega) \\ I \end{bmatrix} (I - K(j\omega)P(j\omega))^{-1} \begin{bmatrix} -K(j\omega) & I \end{bmatrix} \right)$$

Moreover,

$$\arcsin b_{P_1,K}(\omega) \geq \arcsin b_{P_0,K}(\omega) - \arcsin \psi(P_0(j\omega), P_1(j\omega)), \ \forall\omega$$

and

$$\arcsin B_{P_1,K} \geq \arcsin B_{P_0,K} - \arcsin \delta_\nu(P_0, P_1).$$

The significance of the preceding theorem can be illustrated using Figure 4.7.

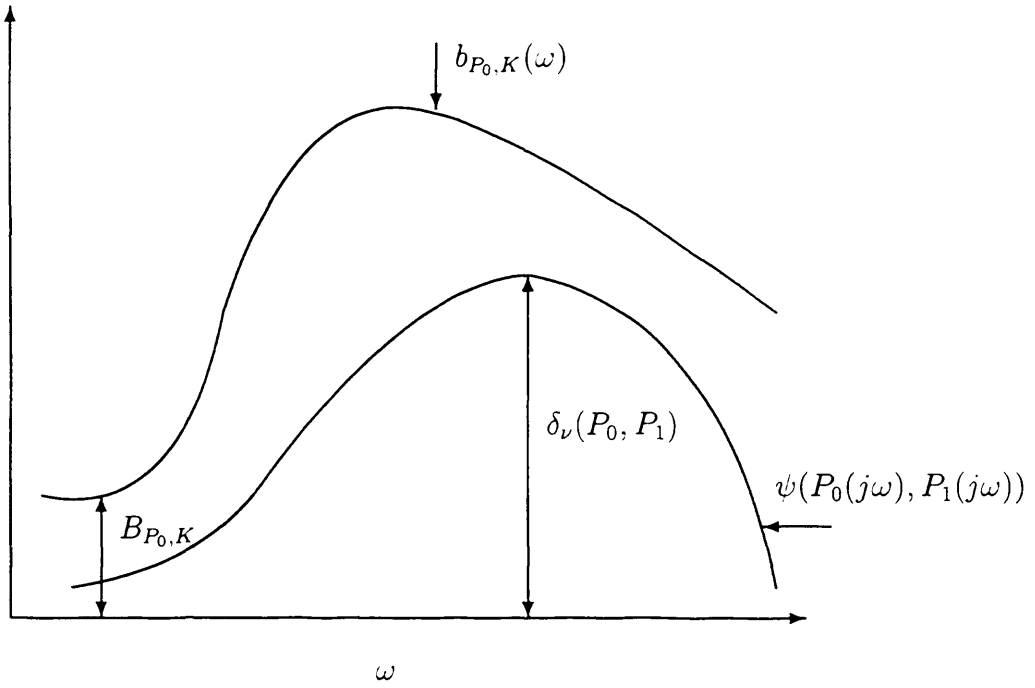

Figure 4.7: K stabilises both P_0 and P_1 since $b_{P_0,K}(\omega) > \psi(P_0, P_1)$ for all ω.

It is clear from the figure that $\delta_\nu(P_0, P_1) > B_{P_0,K}$. Thus a frequency-independent stability test cannot conclude that a stabilising controller K for P_0 will stabilise P_1. However, the frequency-dependent test in the preceding theorem shows that K stabilises both P_0 and P_1 since $b_{P_0,K}(\omega) > \psi(P_0(j\omega), P_1(j\omega))$ for all ω. Furthermore,

$$B_{P_1,K} \geq \inf_\omega \sin(\arcsin\ b_{P_0,K}(\omega) - \arcsin\ \psi(P_0, P_1)) > 0.$$

The following theorem is one of the main results on the ν-gap metric.

Theorem 4.4 *Let P_0 be a nominal plant and $\beta \leq \alpha < B_{opt}(P_0)$.*
(i) For a given controller K,

$$\arcsin\ B_{P,K} > \arcsin\ \alpha - \arcsin\ \beta$$

for all P satisfying $\delta_\nu(P_0, P) \leq \beta$ if and only if $B_{P_0,K} > \alpha$.
(ii) For a given plant P,

$$\arcsin\ B_{P,K} > \arcsin\ \alpha - \arcsin\ \beta$$

for all K satisfying $B_{P_0,K} > \alpha$ if and only if $\delta_\nu(P_0, P) \leq \beta$.

The preceding theorem shows that any plant at a distance less than β from the nominal will be stabilised by any controller stabilising the nominal with a stability margin of β. Furthermore, any plant at a distance greater than β from the nominal will be *destabilised* by *some* controller that stabilises the nominal with a stability margin of at least β.

Similarly, one can consider the system robust performance with simultaneous perturbations on the plant and controller.

Theorem 4.5 *Suppose the feedback system with the pair (P_0, K_0) is stable. Then*

$$\arcsin\ B_{P,K} \geq \arcsin\ B_{P_0,K_0} - \arcsin\ \delta_\nu(P_0, P) - \arcsin\ \delta_\nu(K_0, K)$$

for any P and K.

To summarise, the ν-gap metric is a natural dual to the generic performance measure $B_{P,C}$ [14]. In fact, it is the least conservative measure of the difference between systems in the following *feedback sense*: Two systems are close in the ν-gap metric if any reasonably good (in terms of $B_{P,C}$) controller for one system exhibits similar closed-loop performance with the other. This approach may be used as an alternative to the μ-analysis when assessing at which points in the flight envelope the nominal controller may no longer guarantee stability.

4.4 Summary

$\mathcal{H}_\infty$ loop shaping, μ-analysis and ν-gap analysis have been introduced. $\mathcal{H}_\infty$ loop shaping has been chosen as the design method for the centralised controller. This nominal controller will be designed using a plant model linearised at a chosen point in the flight envelope. It is known that there will be significant variations in this model, both due to velocity changes, whereby the forces and moments are generated differently, or due to parameter variations, whereby parameters such as weight and inertias will change. It is important to know the effects of these

variations in order to determine where in the flight envelope stability may be guaranteed, and so at which points controller switching/scheduling may be required. These variations may be considered to be perturbations to the nominal model, allowing μ-analysis to be used to asses the robustness of the nominal system to these perturbations at different points in the flight envelope. The nominal $\mathcal{H}_\infty$ loop shaping controller will have a particular robustness margin associated with it. An alternative method of considering stability of the nominally designed controller at different points in the flight envelope is through the use of the ν-gap metric. If the difference in two linearised plant models at different points in the flight envelope, as measured by the ν-gap metric, remains below this stability margin over all frequency, then any controller that robustly stabilises one of the plants will also guarantee stability for the other plant.

Chapter 5

Design and piloted simulation of a centralised IFPC system

This chapter will give details of the design of the centralised controller, the piloted simulation and the performance analysis of the system. § 5.1 will detail the system specifications as given by DERA. These include the pitch control law specification, the aircraft and engine limits and test procedures. § 5.2 gives details of the nominal 80 knot plant, and § 5.3 describes the multivariable nature of the plant. The centralised design used in the piloted simulation is introduced in § 5.4, along with a justification of the control structure. Details of the simulator are given in § 5.5, along with the results of the piloted simulation. § 5.6 gives details of the anti-windup strategy applied to the centralised system. § 5.7 presents the design of a pure thrust vectoring controller that does not use either the elevator or the pitch reaction control system. Finally, § 5.8 provides a summary of this chapter.

5.1 IFPC system specifications

The following specifications are those supplied by DERA [12]. The aircraft model is based on the VAAC Harrier with a Spey engine. The front nozzle positions are significantly different from those on the standard Harrier. In order to introduce thrust and thrust vector cross-coupling into the pitch axis, the front pair of nozzles have been moved forward and down. Aircraft control will be achieved via:

- elevator
- reaction control valves (RCV).

Engine control will be obtained via:

- front and rear nozzle deflections
- front/rear thrust split

- inlet guide vanes
- fuel flow
- nozzle exit area.

Note at this point that the control inputs have been separated into airframe and engine. The front and rear nozzle deflections are considered to be a part of the engine control as they are physically a part of the propulsion system. The control law will be tested from 50 knots to 150 knots indicated air speed (IAS) and will have a set of structural, aerodynamic and engine limits to adhere to. The following specification will list the mission level features to be incorporated. At the end of each section there are *'advisory'* statements provided by DERA. There are intended as a guide to design features that would be required in any flight-worthy control law and which may be incorporated at a later date, but are not essential within the scope of this research.

5.1.1 Pitch control law functional specification

The pitch axis control law will follow a two-inceptor strategy. In this scheme fore/aft displacement of the centre stick will produce a change in flight path angle, and displacement of the left-hand inceptor will demand aircraft velocity parallel to this flight path. The following paragraphs describe the functional specifications for these responses. All aircraft responses should aim to be fast and well damped throughout the flight regime.

The following functions will be provided:

- **Flight path manoeuvre demand**. The right hand pitch control law will command flight path rate, $\dot{\gamma}$, and should actively hold flight path with the stick centred. Stick displacement will produce a flight path rate demand up to a maximum of $\pm 3.0°s^{-1}$. Flight path demands should aim to be decoupled from axial manoeuvres, with a maximum demand of ± 2 knots transient speed change during any flight path manoeuvre.
- **Velocity control**. Left-hand inceptor displacement will demand velocity parallel to flight path, VT. Velocity demands should be decoupled from flight path manoeuvres. A maximum transient deviation of $\pm 0.3°$ in flight path angle is acceptable during any velocity manoeuvre.

(***Advisory:*** In a performance limited situation (e.g. maximum engine thrust is reached) there will be a conflict between satisfying the flight path or the velocity demands. In this situation, the control law should always prioritise flight path demands over velocity.)

5.1.2 Aircraft and engine limits

The following limits are to be respected when governing engine and aircraft dynamics. It should be noted that in some areas of the flight envelope the performance requirements given in the functional specification might lie outside the aircraft capability. In this case, the control law should govern aircraft motion to give the maximum performance obtainable within all the limits.

- **Incidence boundary control.** To protect against extreme incidence angles, which can lead to lateral/directional instability, an incidence boundary is necessary. The aircraft angle of incidence, α, should be kept within $+12.0°$ and $-6.0°$.

(***Advisory:*** For reference, maximum use of wing lift should be used at all times in order to conserve fuel and engine life. Harrier training specifies a nominal 6.0° in steady flight with transients within the boundaries acceptable during manoeuvring.)

- **Engine limits.** To protect engine components from dangerous over stress and over temperature, the following set of engine limits need to be respected. Four main limits are given, with the rest of the limits listed as advisory. These are denoted by $-\cdot$ lines on the time history plots of the engine controlled outputs.
  - Low-pressure spool speed, $NLPC$ <102%
  - High-pressure turbine stator outlet temperature, $T10$ <1430K
  - High-pressure compressor surge margin, $HPSM$ >10%
  - Low-pressure compressor surge margin, $LPSM$ >10%

(***Advisory:*** Also limit

- Low-pressure compressor aerodynamic, $LPCPC$ <104%
- High-pressure spool speed, $NHPC$ <101%
- High-pressure compressor aerodynamic, $HPCPC$ <102%
- High-pressure compressor outlet temperature, $T6$ <810K
- Combustion chamber pressure, $P6$ <2300kPa
- Jet pipe temperature, $T30$ <2200K)

5.1.3 Test procedures - off-line testing

The following assumptions can be made about the piloted aircraft in order to aid off-line testing:

- Aircraft flight will be wings level and with zero sideslip.

- The longitudinal stick can be moved at a maximum rate of $20°s^{-1}$.
- Given the above stick rate limit, it is suggested that tests to ascertain system performance at ± 1, ± 2 and $\pm 3°s^{-1}$ flight path rate should be performed.
- The left-hand inceptor can be moved at a maximum rate of 2 units s^{-1}.
- Given the above left-hand inceptor rate limit, it is suggested that tests to ascertain system performance at ± 5, ± 15 and ± 30 knots velocity change from the design point should be performed.
- Demanded flight path angles will be between $+30°$ and $-20°$.
- Mission level inputs will be applied independently.

(***Advisory:*** Mission level inputs tested simultaneously, with the priority given to maintaining flight path demands in limiting situations.).

5.1.4 Test procedures - piloted simulation

The resulting control laws will be tested under piloted simulation on the DERA Bedford real-time all-vehicle simulator (RTAVS) simulators. No lateral and directional control laws will be present, rudder and aileron angles being governed directly by the pilot (as in the standard Harrier). Although the simulation will have six degrees-of-freedom, all assessed manoeuvres will be performed wings level and with zero sideslip. The simulation trial will aim to assess the control law at its design condition, and test robustness and scheduling issues by then repeating the assessment at off-condition points. Piloted simulation will endeavour to quantify:

- stability (fast, predictable well damped aircraft response).
- robustness (stability at off-design conditions, indicating the amount of scheduling needed).
- performance (meeting the design specifications).

5.2 Nominal 80 knot plant

A nominal linear representation of the Spey-WEM model was generated at the 80 knot trim point of the V/STOL flight envelope. At this flight condition coming down to hover, the aircraft is longitudinally unstable, and propulsion system generated forces and moments have largely taken over control of the aircraft from the aerodynamic effectors. This is the transition area in the flight envelope. The resulting state space model of the integrated airframe and engine

system has 35 states - 15 engine states, 16 actuator states and 4 airframe states, and is of the form

$$\dot{x} = Ax + Bu$$

$$y = Cx + Du,$$

The control inputs are given by

$$u = [ETAD, ETASTK, FNOZ, RNOZ, SPLIT, MFF, ENOZA, IGV]$$

The vector of outputs y includes 7 airframe and 21 engine variables. Based on the performance requirements detailed in the previous section, the vector of controlled variables z was chosen as

$$z = [\alpha, VT, \dot{\gamma}, NLPC, T10, HPSM, LPSM]$$

The angle of incidence α was included in z in order to explicitly minimise deviations from its trim point during manoeuvres. The 4 engine variables $NLPC$, $T10$, $HPSM$ and $LPSM$ were chosen as these were the main variables specified to be limited. Overall, this leaves a plant having more controlled outputs than control inputs. This resulting non-square plant increases the interest of the partitioning process that is discussed in the following chapter (Chapter 6).

5.3 Open-loop analysis

Due to the multivariable nature of the system, there is no sole direct correlation between any one actuator and a single output - each actuator will generally affect several of the outputs, and each output may be affected by several of the actuators, as there are significant interactions between subsystems. It would be difficult to apply SISO methods to controller design in this case, as it would not be a simple task to evaluate which inputs and outputs to pair together, or how to minimise the coupling between outputs. While inputs and outputs may be paired for the nominal design, these would be valid only at this point in the flight envelope. For example, while changes in flight path rate, $\dot{\gamma}$, depend on engine generated forces and moments at lower speeds and through the transition region, at higher speeds it will increasingly be changed through the use of airframe generated forces and moments. In addition, as it is a non-square plant, it would need the number of controlled outputs to be reduced.

The linear model is scaled both at the input and at the output. These scalings are important for the specification of relative actuator usage and output cross-coupling trade-offs, and for analysis in the frequency domain. The inputs should be scaled by their expected ranges of operation and, as such, the input scaling, $SINPUT$, is given by:

$$SINPUT = diag\,(15,\ 15,\ 50,\ 50,\ 0.5,\ 0.6,\ 0.5,\ 25)\,.$$

The outputs should be scaled such that equal magnitudes of cross-coupling into each of the outputs is equally undesirable. The output scaling, $SOUTPUT$ is given by:

$$SOUTPUT = diag\left(1,\ 1,\ \frac{180}{\pi},\ \frac{1}{200},\ \frac{1}{800},\ \frac{1}{200},\ \frac{1}{500}\right).$$

The constant matrix $SOUTPUT$ is used to prioritise airframe controlled variables (which must achieve specific handling qualities characteristics) over engine quantities (which have simply to be limited within certain values). This step is crucial to this integrated control strategy, since the centralised controller must allow the engine variables to deviate from their nominal values during flight path manoeuvres. Tight control of all engine variables corresponds to demanding complete decoupling of these variables from large dynamic thrust changes. The resulting controllers would inevitably produce seriously degraded thrust response characteristics in nonlinear simulation. The trade-off between satisfaction of flying qualities performance specifications and respecting limits on safety critical engine variables is thus explicitly built into the design process, via selection of the output scaling. The choice of $SOUTPUT$ was therefore made in order to address the two different levels of control required for the airframe controlled variables and the engine controlled variables. The airframe controlled variables have to meet fairly stringent handling qualities, whereas the engine variables must be kept within specified limits in order to maximise engine life and to avoid operation in dangerous regions where failures may occur. In order to manage these two different requirements, it is possible to select the corresponding weight in the output scaling matrix $SOUTPUT$. This allows the designer to specify for the three airframe controlled variables, one unit of coupling is equally undesirable in each. In other words, coupling of 1 deg in α, 1 ms^{-1} in VT and 1 $degs^{-1}$ in $\dot{\gamma}$ are all equally undesirable. For the engine variables, however, such tight decoupling is not required. Indeed, if such decoupling was demanded, it would seriously degrade the performance of the airframe controlled variables. This is due to the fact that a change in the airframe variables generally requires a change in thrust. Any change in thrust would have the effect of changing the temperatures, pressures and speeds within the engine. If a change in an airframe variable such as $\dot{\gamma}$ is demanded while requesting no change in the trim values of the engine variables, the airframe specified performance will not be achieved. Additionally, if a further reduction in decoupling of the engine variables were allowed, so they were less constrained to remain at their trim values, there is no gain in the achieved airframe variables performance. It simply results in the engine variables exceeding their limits by a greater amount and for a longer time but with no corresponding gain in airframe performance. This is shown in § 5.4.1.

Figure 5.1 shows the open-loop transfer functions from the main fuel flow to each of the 7 controlled outputs (α, VT, $\dot{\gamma}$, $NLPC$, $T10$, $HPSM$, $LPSM$) at 20, 50, 80, 120, and 150 knots for the scaled plant. It can be seen from this that a change in main fuel flow will

result in a change in most of the controlled outputs. Figure 5.2 shows the open-loop transfer functions from each of the 8 actuator inputs to the output $\dot{\gamma}$ for the scaled plant. Here it can be seen that a change in $\dot{\gamma}$ may be facilitated through changes in several actuators. While these dependencies are true at the 80 knot design point, it can be seen that they change throughout the flight envelope. In order to cope effectively with the interactions, it is necessary to adopt a multivariable approach to control design. Figure 5.3 shows the open-loop transfer functions from each of the 8 actuator inputs to the output VT for the scaled plant. Again, it can be seen that a change in VT can be brought about through changes in several actuators.

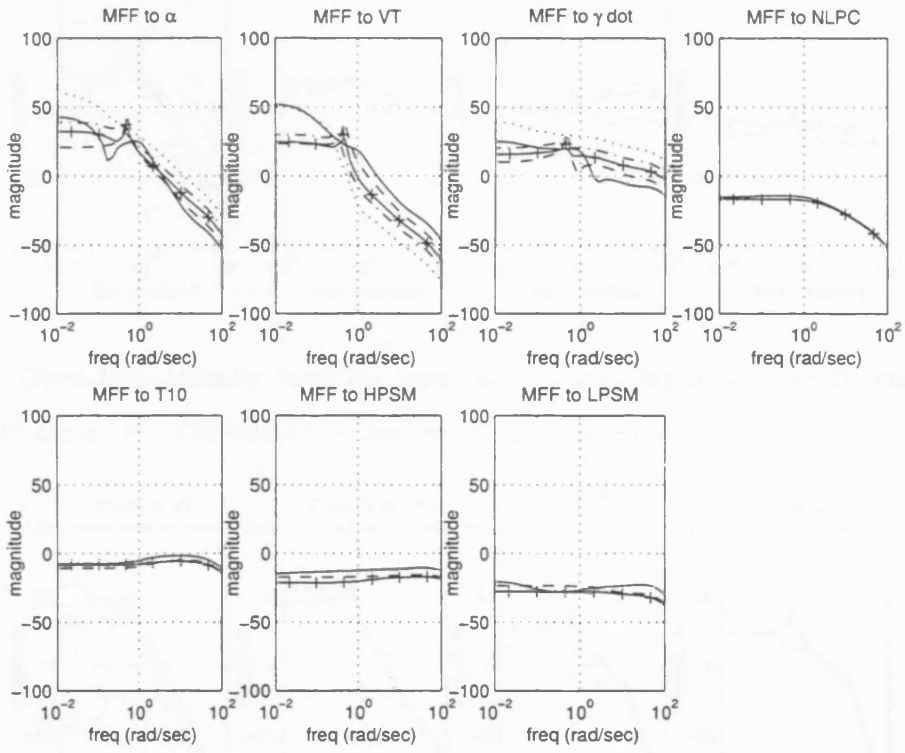

Figure 5.1: Open-loop transfer function from main fuel flow to each controlled output at 20 knots (··), 50 knots (-·), 80 knots (+), 120 knots (- -) and at 150 knots (—).

The scaled linear 80 knot model is shown in Figure 5.4. The peak value around 0.5 rads^{-1} is that associated with the phugoid frequency, described in 1.1.6.

5.4 Centralised IFPC system design

For the centralised controller, the weighting matrices were chosen as:

$$k = diag\,(0.1,\ 0.1,\ 0.24,\ 0.24,\ 0.24,\ 0.2,\ 0.2,\ 0.2)\,,$$

$$W_1 = \frac{s+2}{s} \times I_{8\times 8}, \quad W_2 = I_{7\times 7}.$$

The shaped plant is shown in Figure 5.5.

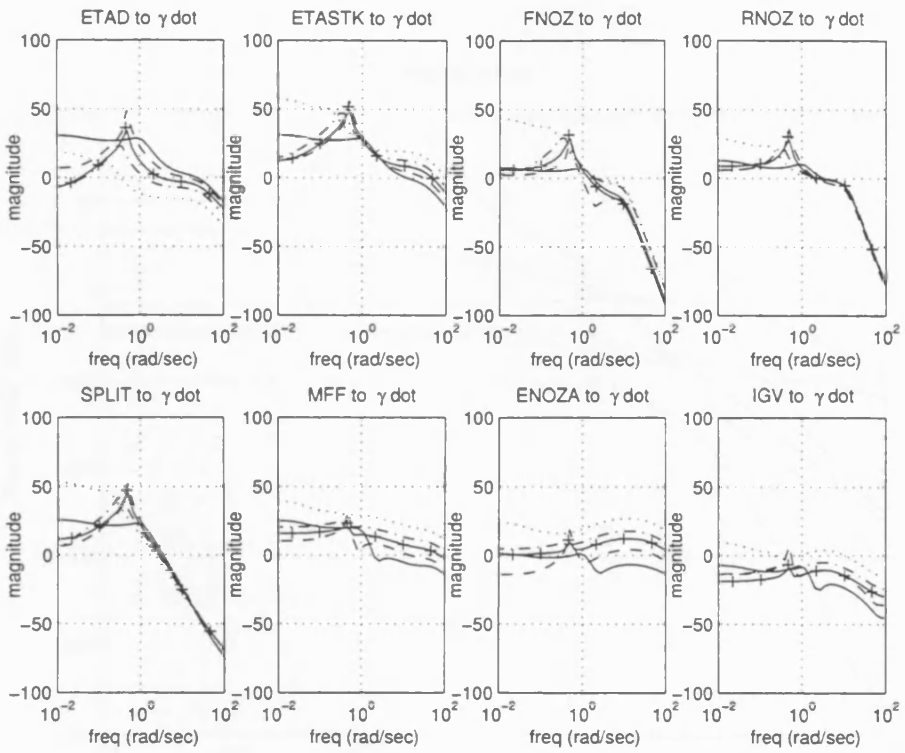

Figure 5.2: Open-loop transfer function from each control input to $\dot{\gamma}$ at 20 knots (··), 50 knots (-·), 80 knots (+), 120 knots (- -) and at 150 knots (—).

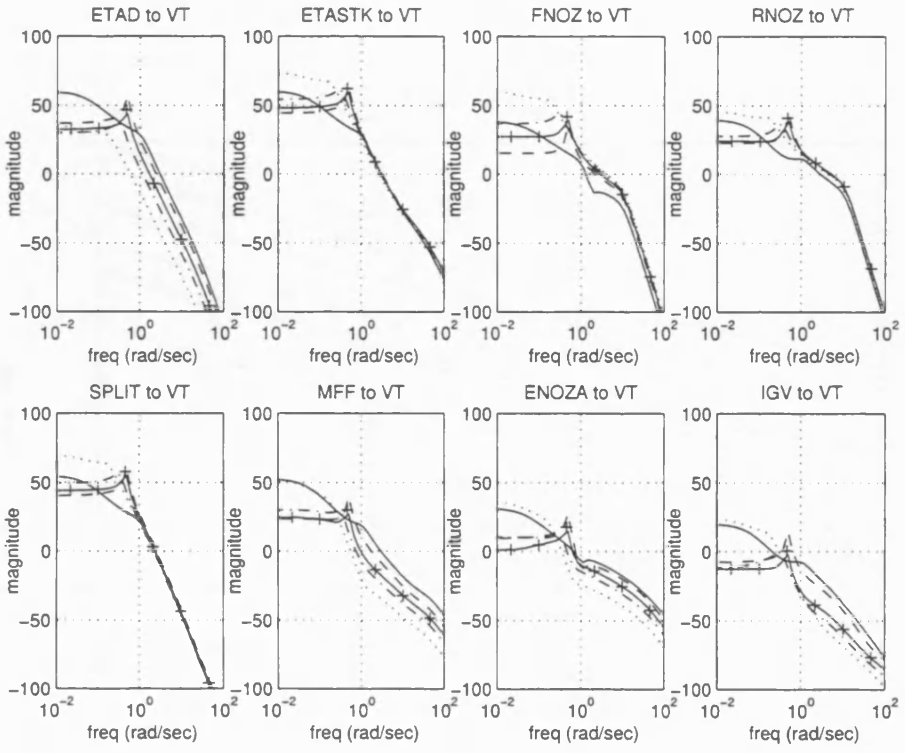

Figure 5.3: Open-loop transfer function from each control input to VT at 20 knots (··), 50 knots (-·), 80 knots (+), 120 knots (- -) and at 150 knots (—).

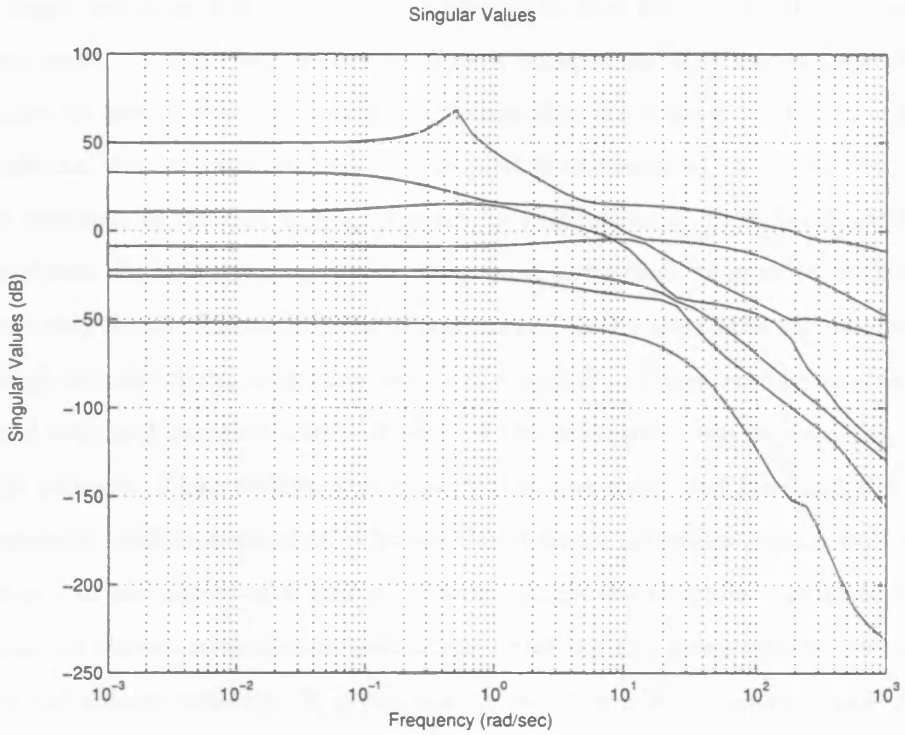

Figure 5.4: Scaled 80 knot linear plant.

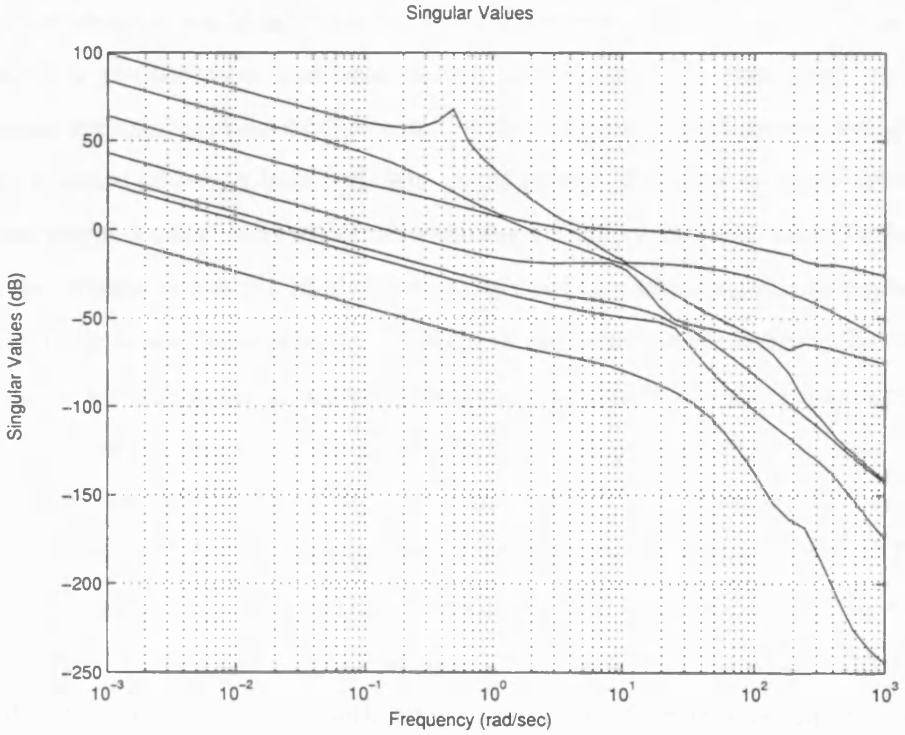

Figure 5.5: Shaped plant.

As indicated in § 1.1, the airframe and engine actuators are used in different ways. When applying a large reference demand on, for example, $\dot{\gamma}$, it is likely that the engine actuators (in particular main fuel flow) will saturate. This is acceptable, as it means that although the demand cannot be met in full, the maximum change that the engine can produce is delivered, subject to internal engine variable limits. The pilot is accustomed to making large demands that are not achieved in full due to engine actuator saturations. The same is not true for the airframe actuators. In this case, operation on a limit is avoided because of possible stability problems that may result. These differing requirements on the use of the various actuators can be met through the use of the weighting matrices k and W_1. These may be selected such that the individual rate and position limits of each of the actuators may be explicitly considered in the design process. The result in this case is that the main fuel flow and the exit nozzle area are frequently used to saturation, whereas the airframe actuators remain well within their position limits. As the choice of k and W_1 have a significant affect on the final value of γ_{opt}, they must also be chosen with this consideration, resulting in a compromise between nominal performance and robust stability. It is possible to select weighting matrices that result in an improved nominal performance, but the robustness stability will deteriorate, as indicated by an increased value of γ_{opt}.

If a greater value had been chosen for the weight k associated with the elevator, it would result in a much greater use in the elevator with no performance benefits. This will simply increase the likelihood of the elevator reaching a position limit unnecessarily in the event of a large demand being made. If a greater value had been chosen for the weight k associated with the pitch reaction control system, it would have resulted in the reaction control system being used to an extent where it would no longer be an efficient use of thrust. If greater or lesser values had been chosen for the weight k associated with the remaining actuators, it would have resulted in either a reduction in robustness margin of the system without any improvement in performance, or it would limit their usage unnecessarily, resulting in degraded performance. The choice of the weighting matrix W_1 was made in order to ensure an acceptable roll-off rate around crossover. The choice of the output weighting matrix W_2 to be the identity matrix was motivated by one unit of cross-coupling (when considering the scaled plant, § 5.3) being equally undesirable in each of the outputs. As this has already been accounted for in the choice of $SOUTPUT$, it is possible to simply use the identity matrix. Overall, the choice of k was limited by the achieved value of γ_{opt}, which, using these weights resulted in a value of 3.99. It is desirable to limit γ_{opt} to a value of ≤ 4 in order to satisfy both robustness and performance requirements. This will guarantee robustness up to 25% uncertainty in the coprime factorisation [6].

Performance properties of the controller described above were initially examined via PC-based

nonlinear time-domain simulations. The following figures show the nonlinear responses of the system to demands on $\dot{\gamma}$ and on VT at 20, 50, 80, 120 and 150 knots. Figures 5.6 to 5.9 show the responses at 20 knots. The demand on $\dot{\gamma}$ results in very good $\dot{\gamma}$ tracking, with all of the controlled outputs remaining within their specified limits. The demand on VT results in a good response. The coupling into $\dot{\gamma}$ exceeds its limit, as do $NLPC$ (just) and $T10$. Figures 5.10 to 5.13 show the responses at 50 knots. Again, it can be seen that $\dot{\gamma}$ is tracked very well, but the $T10$ limit is now exceeded. The responses to a demand on VT are good with slightly less $\dot{\gamma}$ coupling, and $NLPC$ remains just below its maximum limit. Figures 5.14 to 5.17 show the responses at 80 knots. The γ response is seen to be slightly more sluggish, with both the main fuel flow and exit nozzle area now saturating. It can be seen that the coupling into γ is reduced for a demand on VT at this nominal design point. Figures 5.18 to 5.21 show the responses at 120 knots. It can be seen the $\dot{\gamma}$ response is getting rather sluggish, but the coupling into α and VT is still acceptable. Although the transient γ coupling is reduced when a demand is made on VT, there results a significant steady-state error. The coupling into α is just within acceptable limits, but $T10$ exceeds its limit briefly. Figures 5.22 to 5.25 show the responses at 150 knots. The demand on $\dot{\gamma}$ now results in coupling into VT that just exceeds the specifications of 2 knots. The coupling into α for a demand on VT also exceeds 12°. It can be seen that both the front and rear nozzle saturate, resulting in a loss of stability.

At this point, it should be recalled that the controller was designed using a linearised model at the 80 knot point in the flight envelope. The controller was then tested using the linear model before being finally tested on the nonlinear model. During the linear testing of the controller, when making a demand of 30 knots on VT, the coupling into γ remained below 0.15°, well within the maximum allowable specified. All of the actuators remained within their rate and position limits. However, when implemented on the nonlinear model, the coupling for the same demand increased to 2.7°. It is the nonlinear effects of the model that are resulting in the increased coupling, making it difficult to improve the linear nominal controller further.

The level of performance achieved indicated general satisfaction of the specifications set out in § 5.1, and it was thus decided to conduct formal evaluation of the IFPC system, as described in the § 5.5.

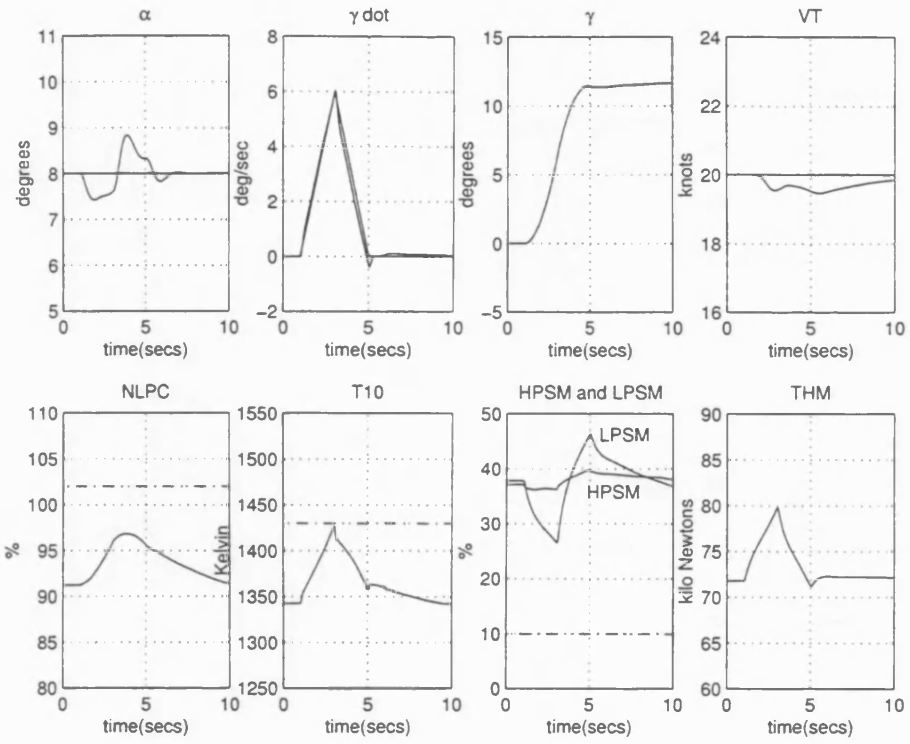

Figure 5.6: Airframe and engine responses to a $\dot{\gamma}$ demand for the centralised controller at 20 knots.

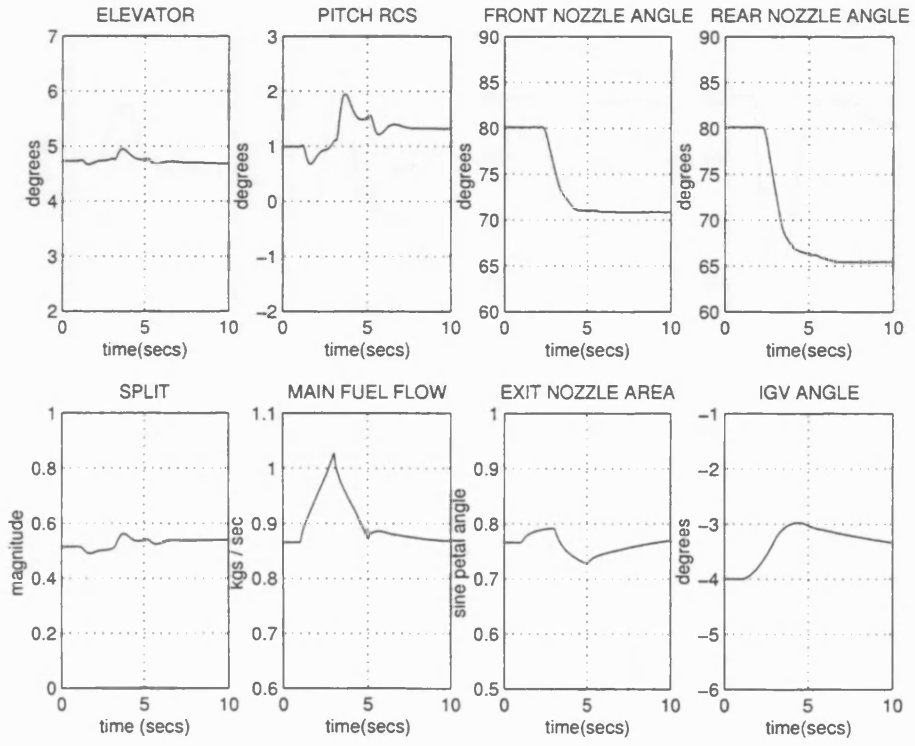

Figure 5.7: Actuator responses to a $\dot{\gamma}$ demand for the centralised controller at 20 knots.

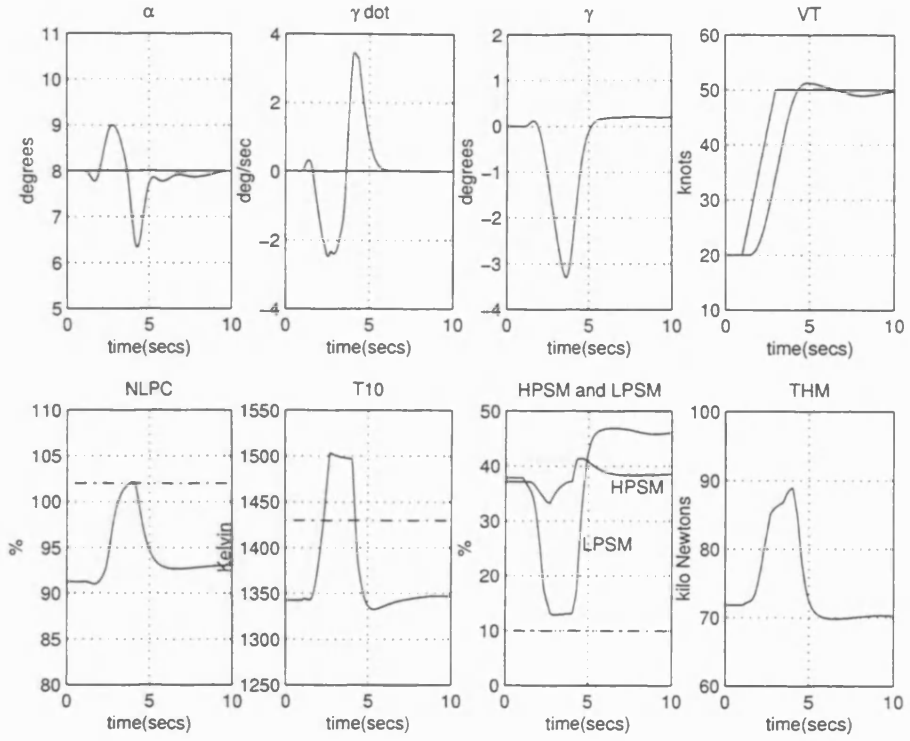

Figure 5.8: Airframe and engine responses to a VT demand for the centralised controller at 20 knots.

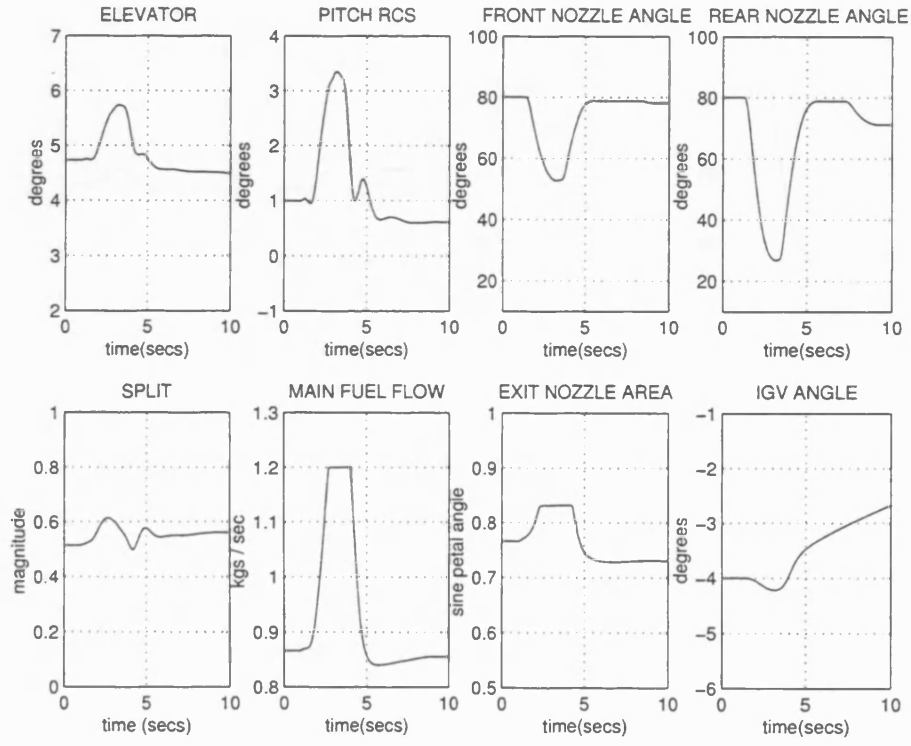

Figure 5.9: Actuator responses to a VT demand for the centralised controller at 20 knots.

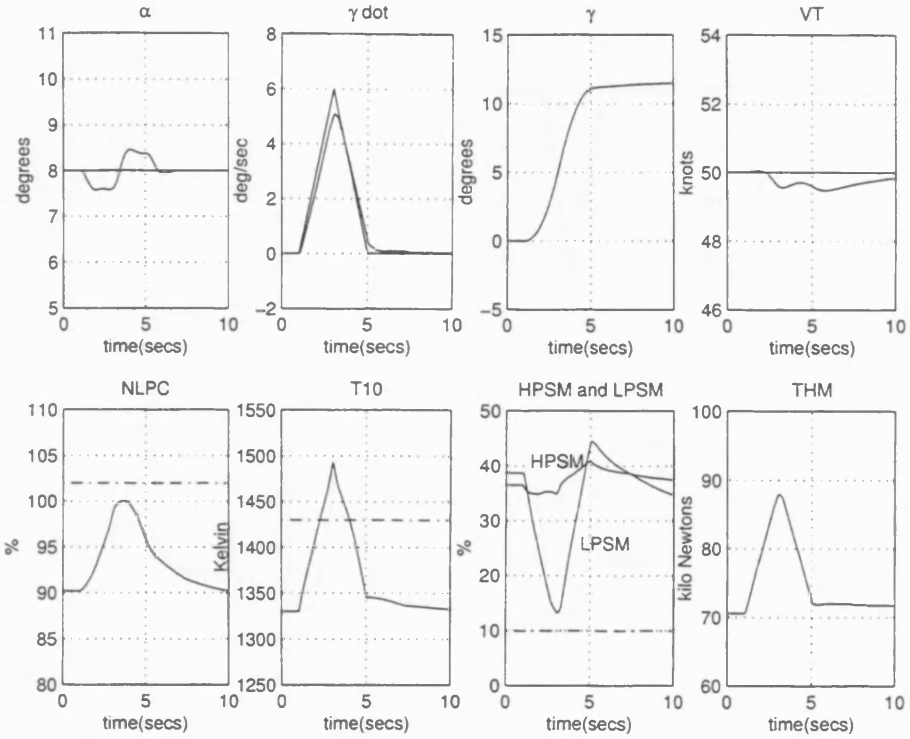

Figure 5.10: Airframe and engine responses to a $\dot{\gamma}$ demand for the centralised controller at 50 knots.

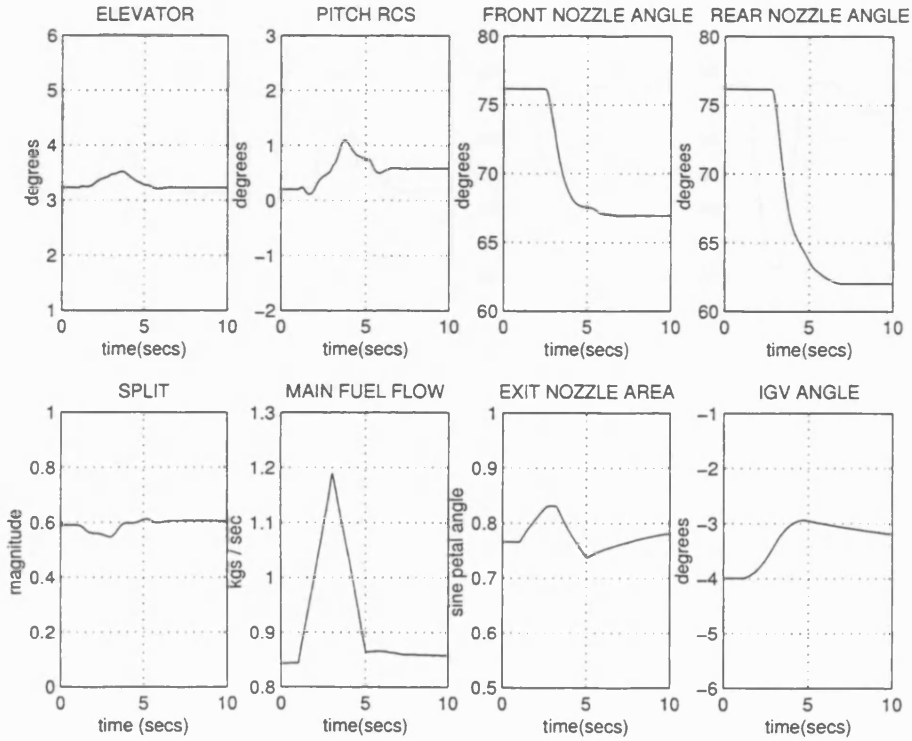

Figure 5.11: Actuator responses to a $\dot{\gamma}$ demand for the centralised controller at 50 knots.

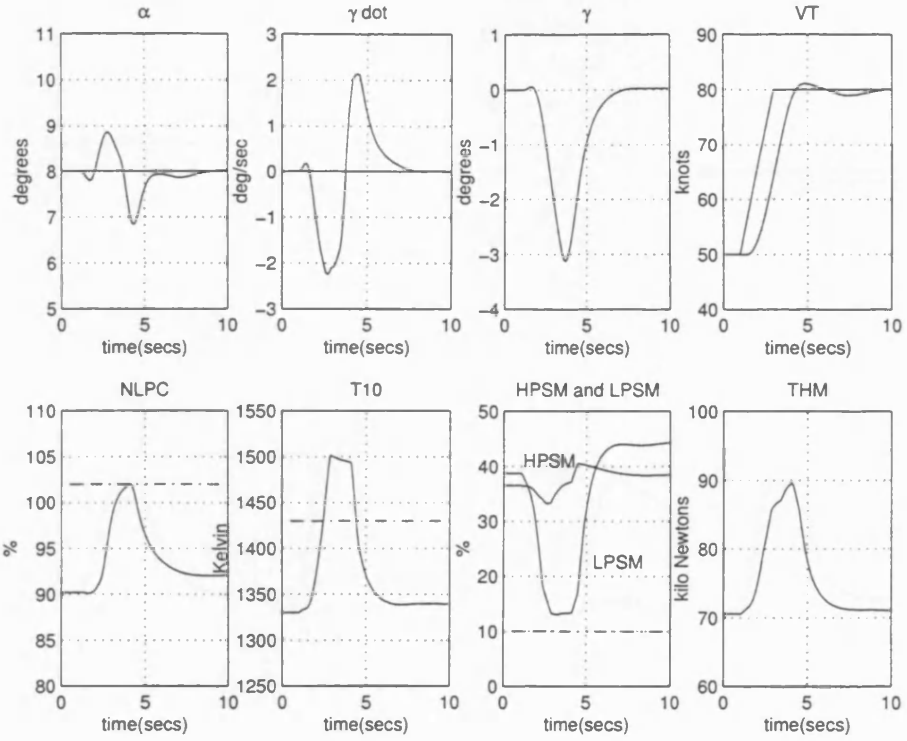

Figure 5.12: Airframe and engine responses to a VT demand for the centralised controller at 50 knots.

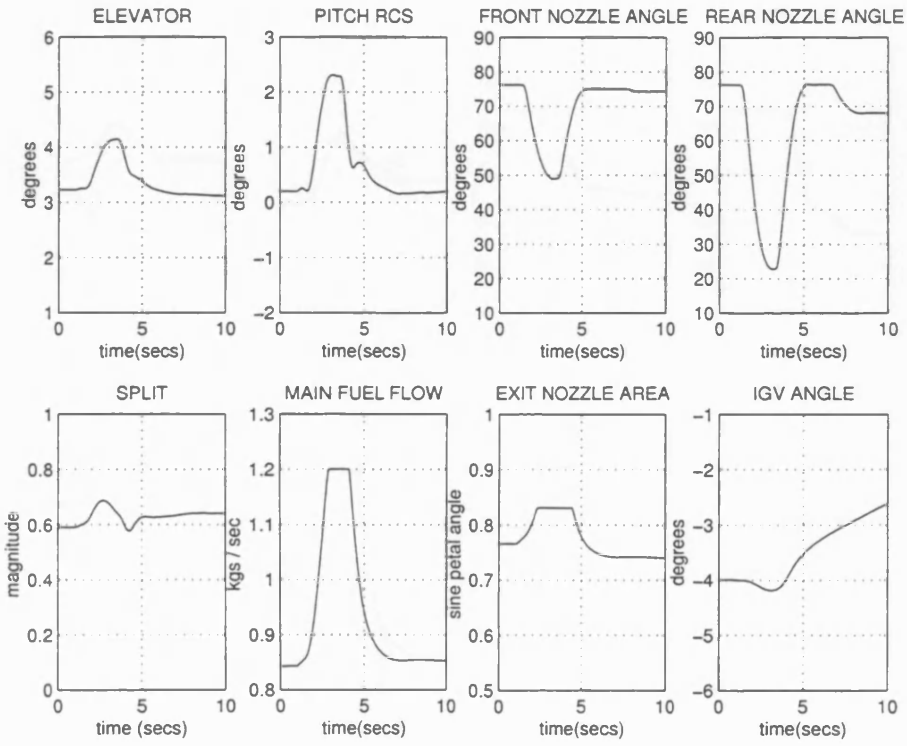

Figure 5.13: Actuator responses to a VT demand for the centralised controller at 50 knots.

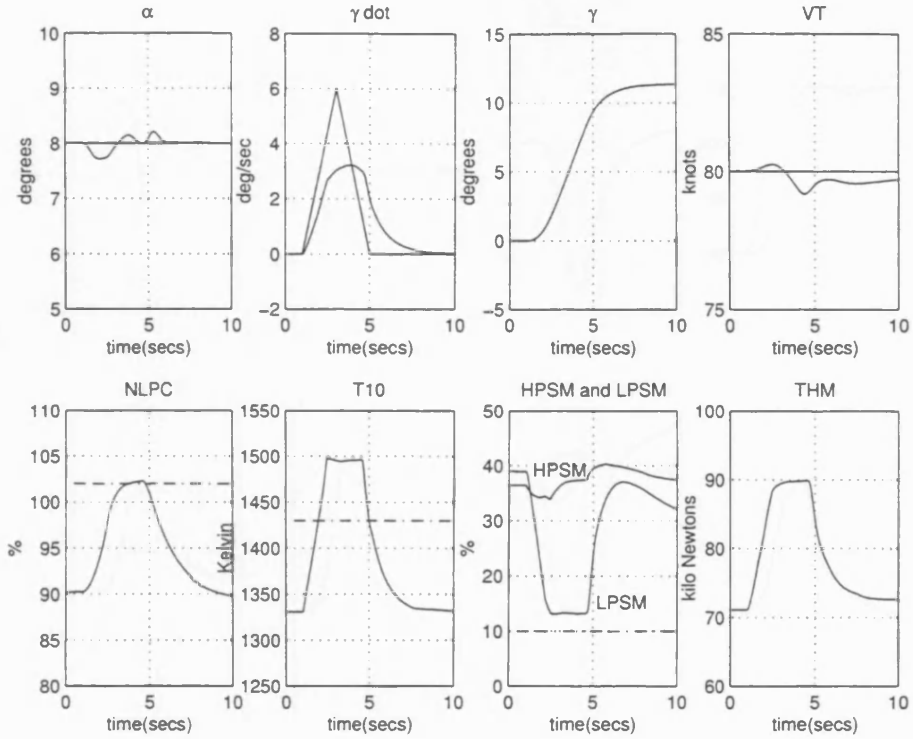

Figure 5.14: Airframe and engine responses to a $\dot{\gamma}$ demand for the centralised controller at 80 knots.

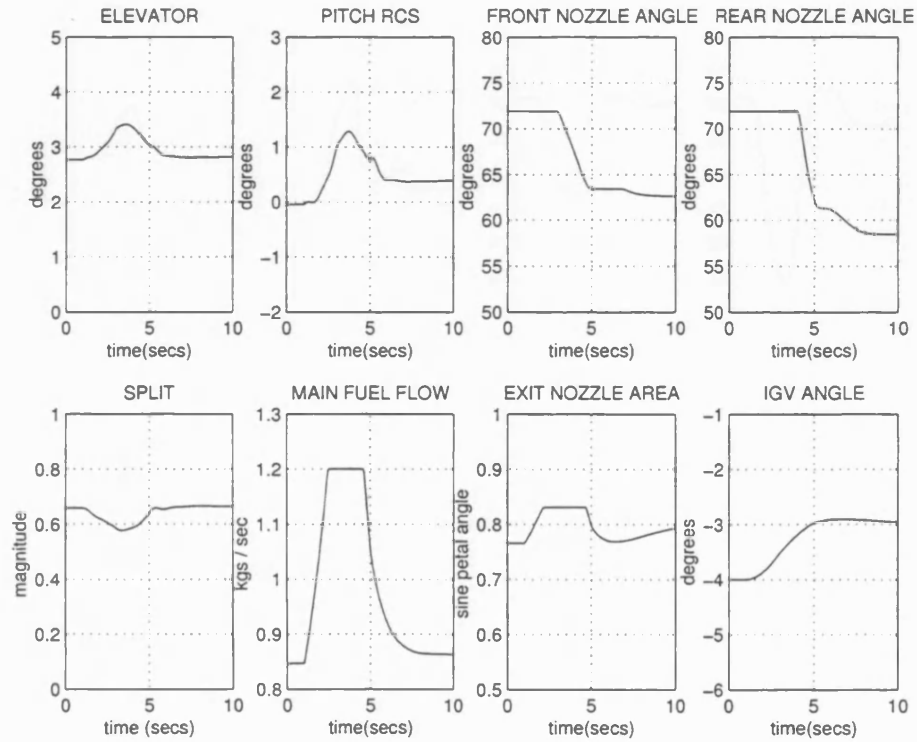

Figure 5.15: Actuator responses to a $\dot{\gamma}$ demand for the centralised controller at 80 knots.

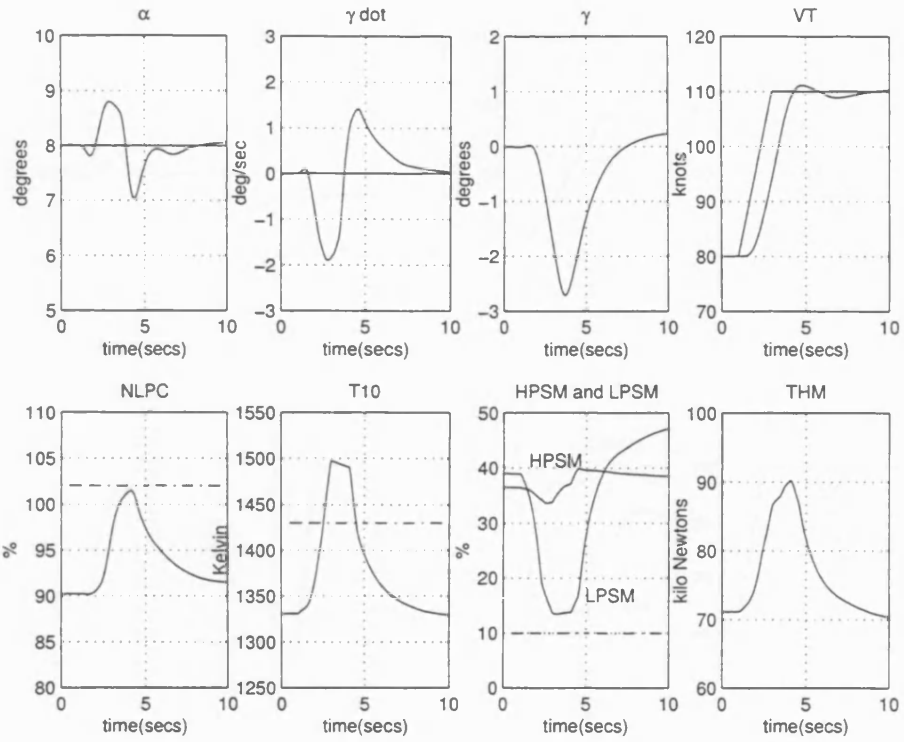

Figure 5.16: Airframe and engine responses to a VT demand for the centralised controller at 80 knots.

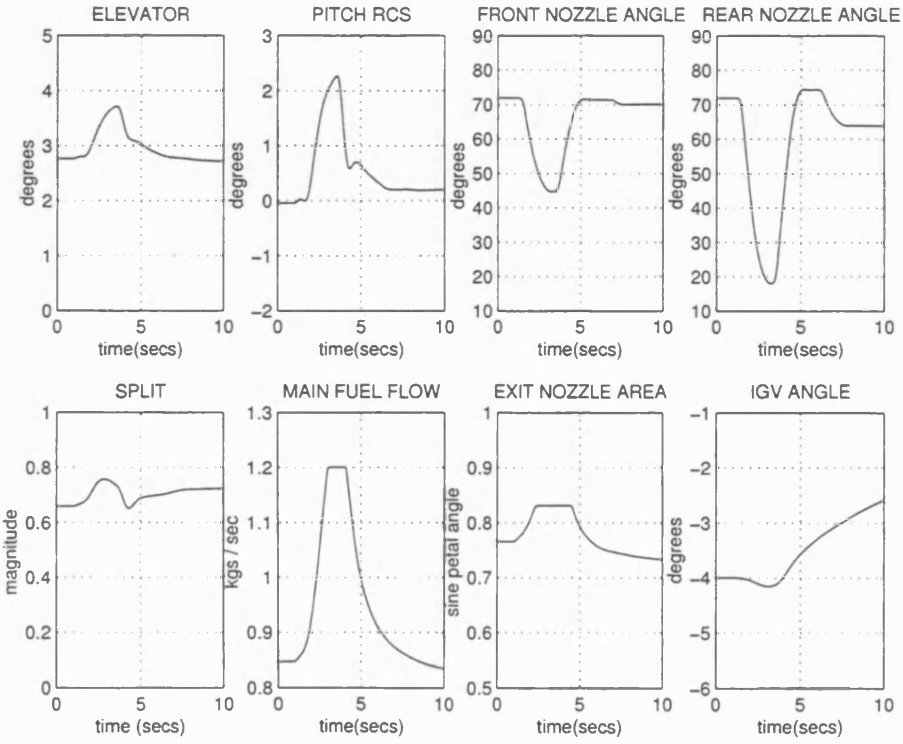

Figure 5.17: Actuator responses to a VT demand for the centralised controller at 80 knots.

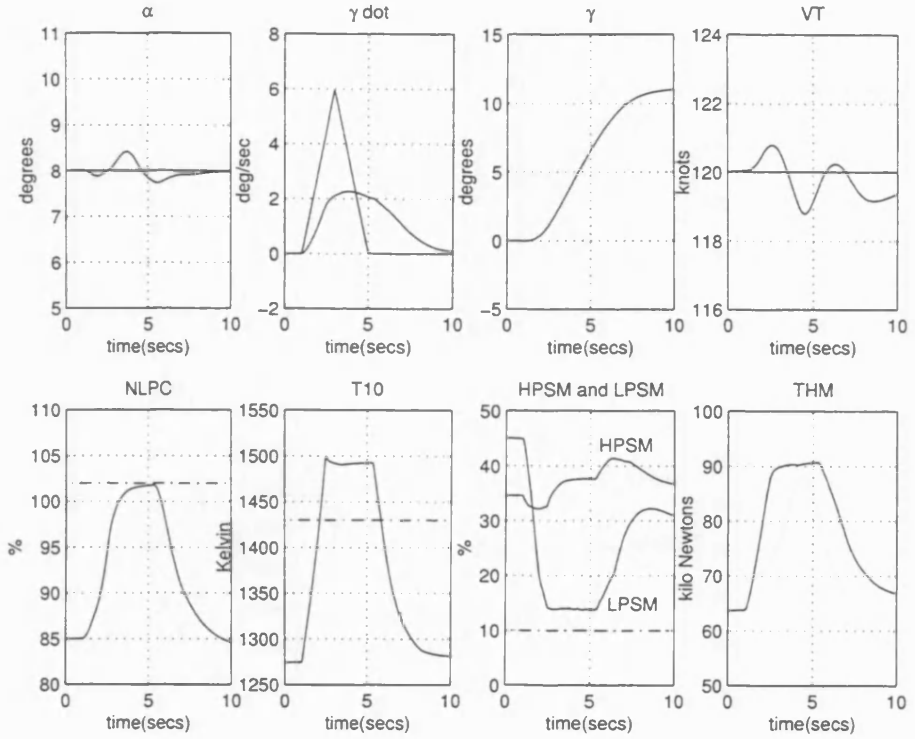

Figure 5.18: Airframe and engine responses to a $\dot{\gamma}$ demand for the centralised controller at 120 knots.

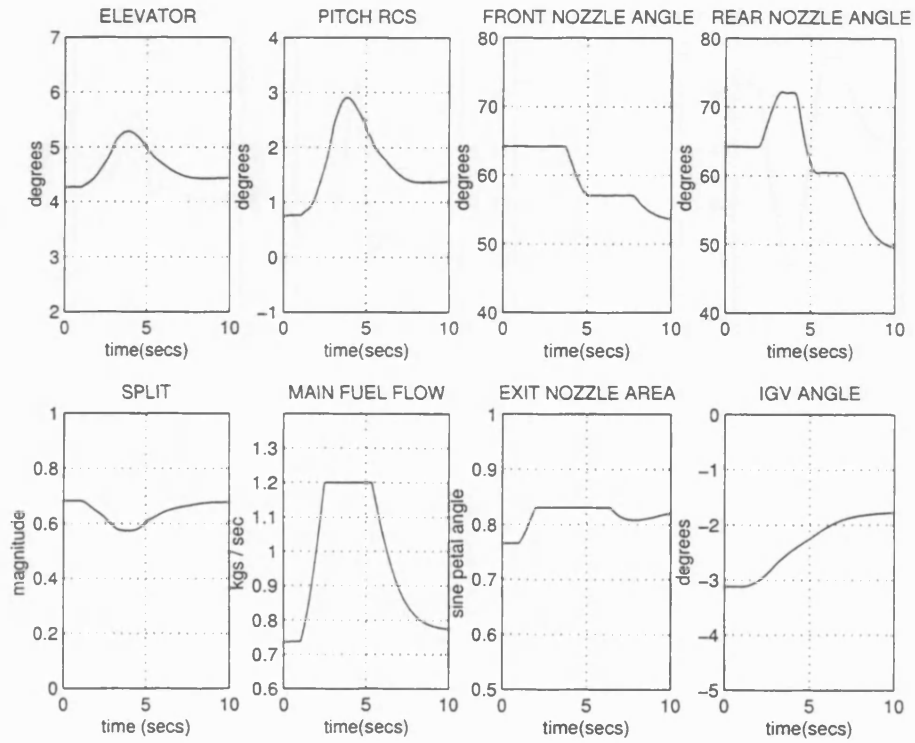

Figure 5.19: Actuator responses to a $\dot{\gamma}$ demand for the centralised controller at 120 knots.

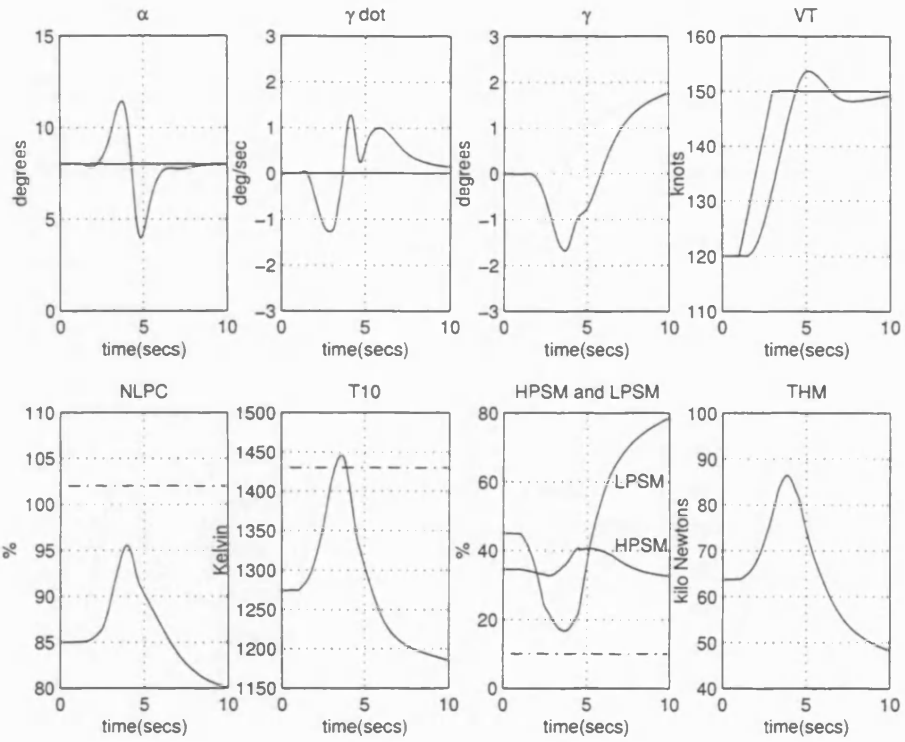

Figure 5.20: Airframe and engine responses to a VT demand for the centralised controller at 120 knots.

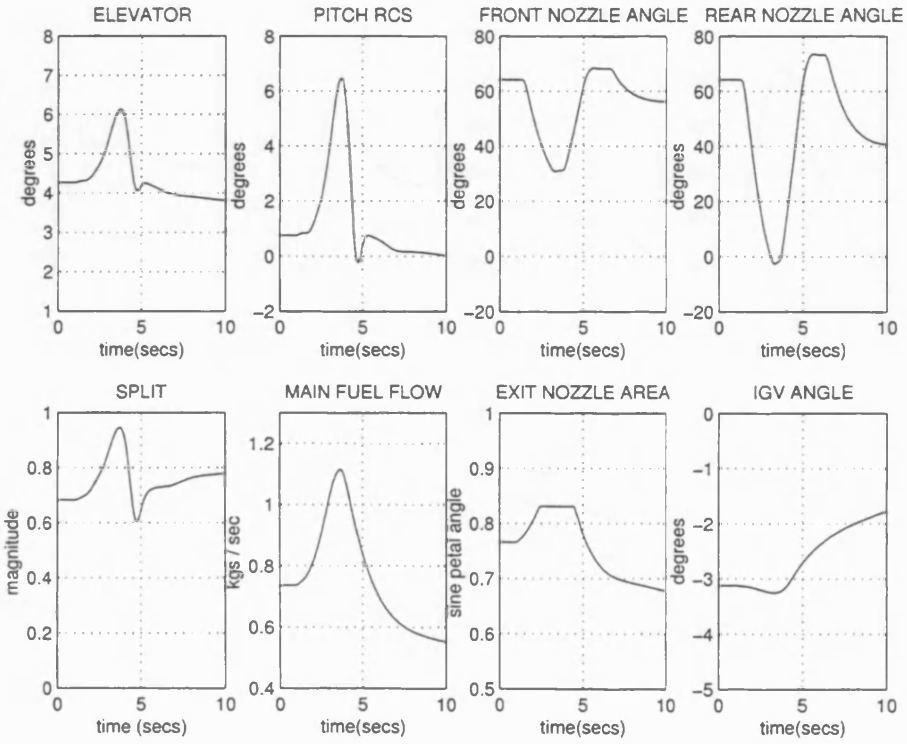

Figure 5.21: Actuator responses to a VT demand for the centralised controller at 120 knots.

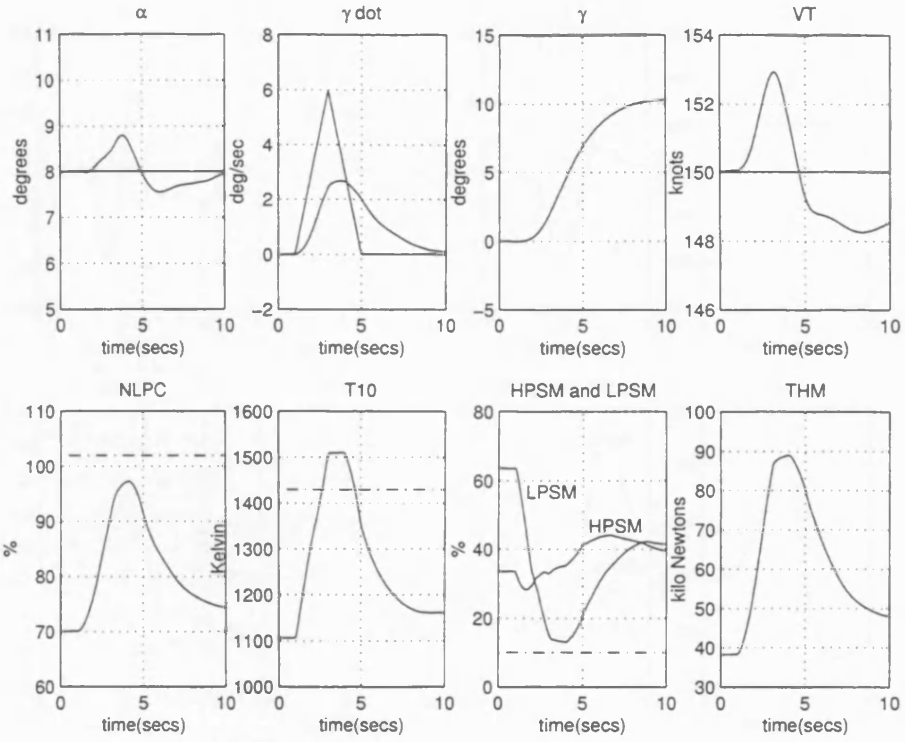

Figure 5.22: Airframe and engine responses to a $\dot{\gamma}$ demand for the centralised controller at 150 knots.

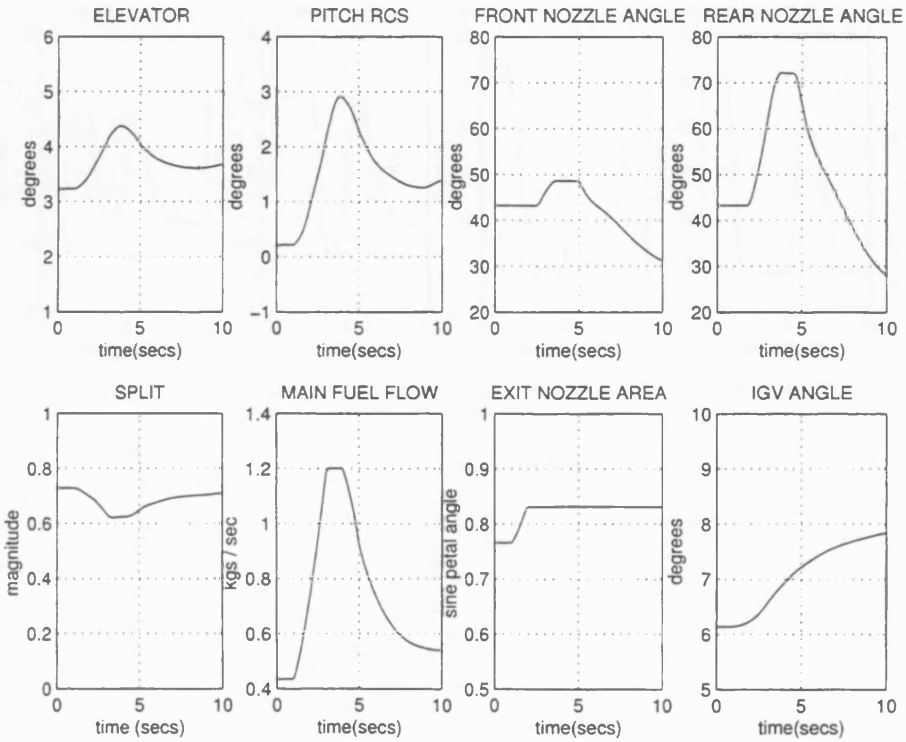

Figure 5.23: Actuator responses to a $\dot{\gamma}$ demand for the centralised controller at 150 knots.

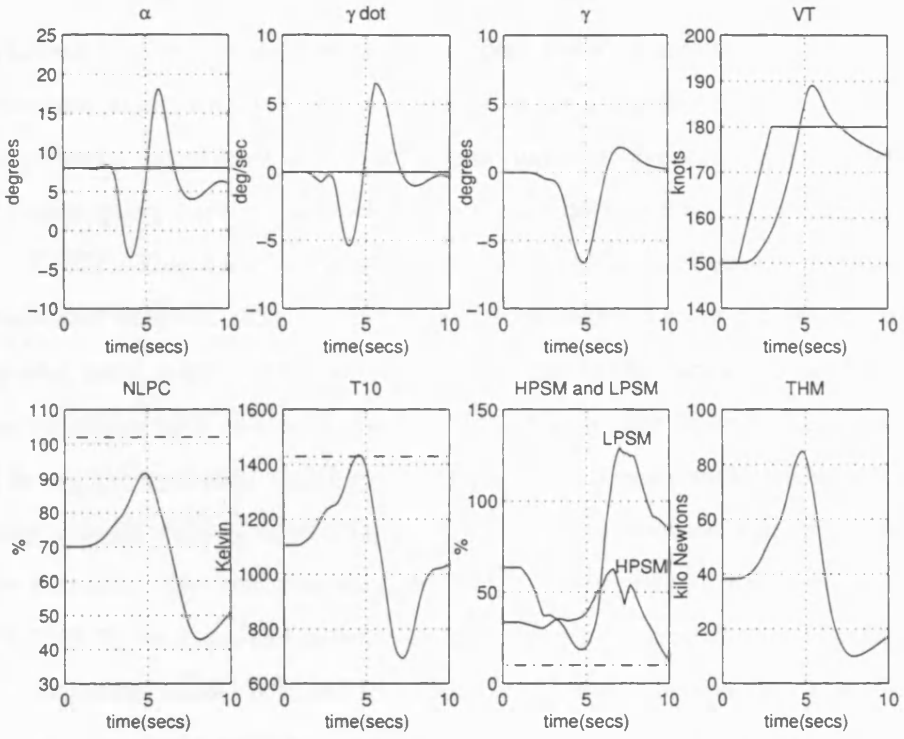

Figure 5.24: Airframe and engine responses to a VT demand for the centralised controller at 150 knots.

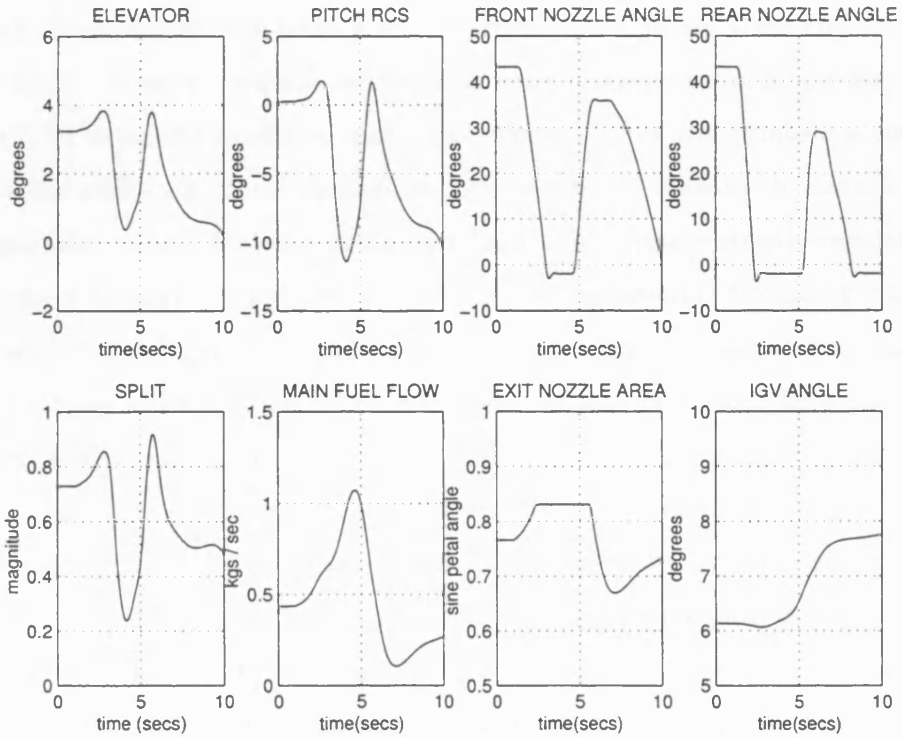

Figure 5.25: Actuator responses to a VT demand for the centralised controller at 150 knots.

5.4.1 Justification of control structure

A novel approach applied to the design of the centralised controller is the inclusion of the 4 engine variables along with the airframe variables as controlled outputs. This is possible despite the different requirements on the engine and airframe controlled variables because the level of decoupling required of each output may be specified using the output scaling matrix $SOUTPUT$. This ensures the engine variables remain sufficiently decoupled from the airframe controlled outputs, enabling the airframe variables to meet reference demands while still maintaining some control over the engine variables. In the case of the airframe controlled outputs, fast responses with no steady-state error in order to meet specific handling qualities are wanted. The engine controlled variables are required to remain within their specified limits, but enforcing this by strictly controlling their transient responses will limit the achievable airframe performance. By including them and through thoughtful use of the output scaling matrix $SOUTPUT$, we have the benefit of having them controlled, but not at the expense of the airframe performance. It is not intended that this will be the only method of limiting the engine variables. Some additional nonlinear scheme will be required in order to guarantee these limits are respected, further details of which are given in § 8. It is possible to prove the usefulness of this approach by comparing the nonlinear simulation results when applying greater or less control on the engine controlled (limited) variables through the use of reduced or increased output scaling respectively for these variables. When reduced output coupling is demanded, the controller will make a strong attempt to regulate these engine variables to their trim values. However, when a reference demand is made, for example on $\dot{\gamma}$, a change in thrust will be required in order to meet that demand. This will produce a change in the engine variables while this thrust demand is met, which will transiently disturb them from their nominal trim values. If the controller demands that they are strictly regulated to their trim values the $\dot{\gamma}$ demand will not be met and an increasing amount of coupling into α and VT will result, as shown in Figures 5.26 and 5.27. Conversely, if greater output coupling is allowed, the engine variables will exceed their limits by a greater magnitude and for a greater time, as shown in Figures 5.28 and 5.29. This does not, however, produce improved performance for $\dot{\gamma}$ or VT. Overall, this shows that these variables are being controlled, but not to an extent that will degrade the airframe performance and not so loosely that they are free to exceed their limits more than is necessary to achieve the demands made for flight path changes. Figures 5.30 and 5.31 show the responses for a demand on VT in the case for tighter control of the engine variables. Although the engine variables do remain within their limits, it is at the expense of increased coupling into γ. Figures 5.32 and 5.33 compare the responses to a demand on VT for the centralised controller and for the controller demanding less control of the engine

variables. Again, this shows no performance improvements, and the engine variables are not being regulated back to their trim values once the airframe demand is achieved.

5.5 Piloted simulation trials of the centralised IFPC system

5.5.1 Implementation on the DERA real-time all-vehicle simulator

Piloted simulation trials were conducted on DERA Bedford's real-time all-vehicle simulator (RTAVS). RTAVS is a fixed base simulator with an immersive, back projected, outside world display with a field-of-view of +135° / -45° vertically by ±135° horizontally. The aircraft simulation models are run on PCs with a networking facility that allows large models to be spread over several processors, or a series of aircraft models to be flown in separate cockpits in the same outside world environment, or both. The cockpit and instrumentation for this trial was representative of a generic fast jet. A dual linear throttle was used with the stick and rudder forces being simulated by springs. The head up display (HUD) used for the trial is shown in Figure 5.34. This display included a pitch attitude indicator, velocity vector diamond, thrust split and front and rear nozzle angle indicators.

The Spey-WEM model and integrated control law were auto-coded from the MATLAB/Simulink [1] environment into C-code and then linked with the RTAVS code to provide executable software for the simulator. Stick shaping and a variable rate limit were added to the $\dot{\gamma}$ and the VT demands respectively. The $\dot{\gamma}$ demand limit was applied by scaling the maximum stick deflections to ±3°. The executable code calls the same FORTRAN modules (aerodynamics, engine and RCS system) as the original Simulink diagram. Finally, the pilot's inputs to the Spey-WEM model were mapped to the cockpit, and the aircraft states were mapped to drive the outside world view and the HUD. The implementation structure of the system is shown in Figure 5.35.

The aim of the trial was to investigate the flying qualities and robustness of the integrated flight and propulsion controller for the Spey-WEM model. The simulator was piloted by three members of DERA Bedford's Flight Management and Control Department. Each of the pilots has extensive experience in the design and testing of advanced control laws for modern STOVL aircraft. To test the flying qualities of the IFPC system a series of manoeuvre tasks were performed at and about the 80 knot design point, as described in the next section. The flying qualities were assessed using the Cooper-Harper Flying Qualities Ratings, as shown in Figure 5.36.

The controller design specifications limited $\dot{\gamma}$ to $\pm 3°s^{-1}$ for full fore and aft stick deflection, but gave no rate limit on VT demands. A rate limit of 10 $knots^{-1}$ was implemented on the

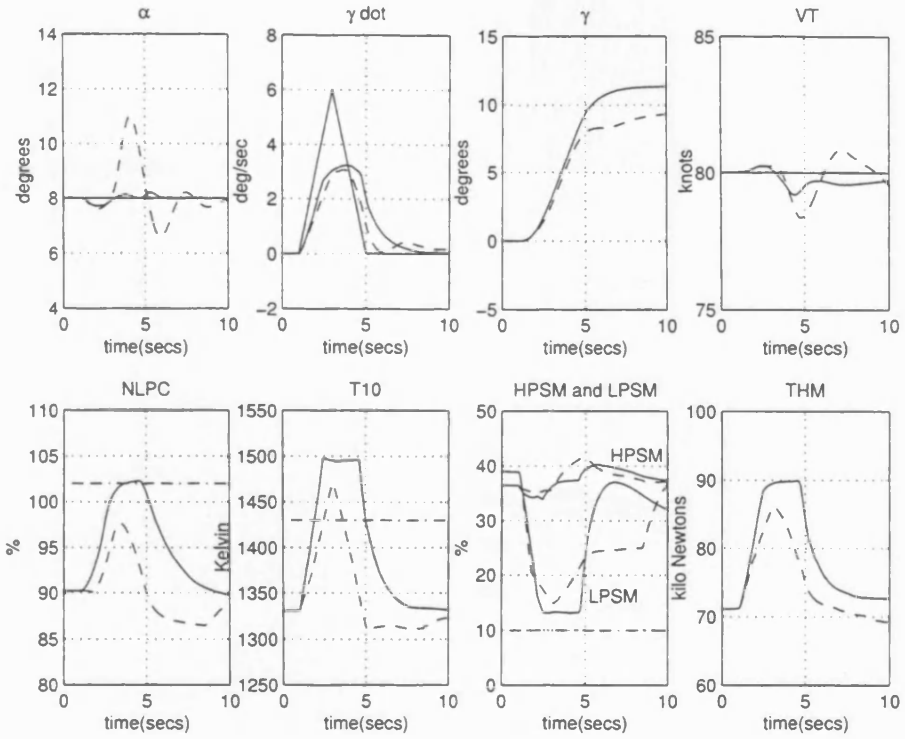

Figure 5.26: Airframe and engine responses to a $\dot{\gamma}$ demand with nominal W_2 (—) and with tighter control of the engine variables (- -) at 80 knot.

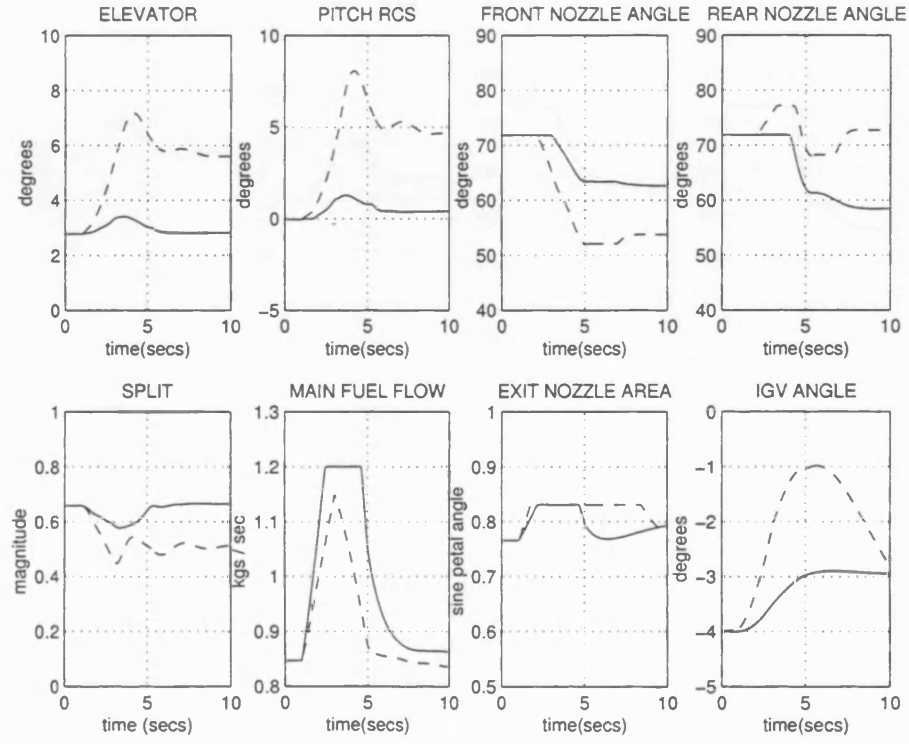

Figure 5.27: Actuator responses to a $\dot{\gamma}$ demand with nominal W_2 (—) and with tighter control of the engine variables (- -) at 80 knots.

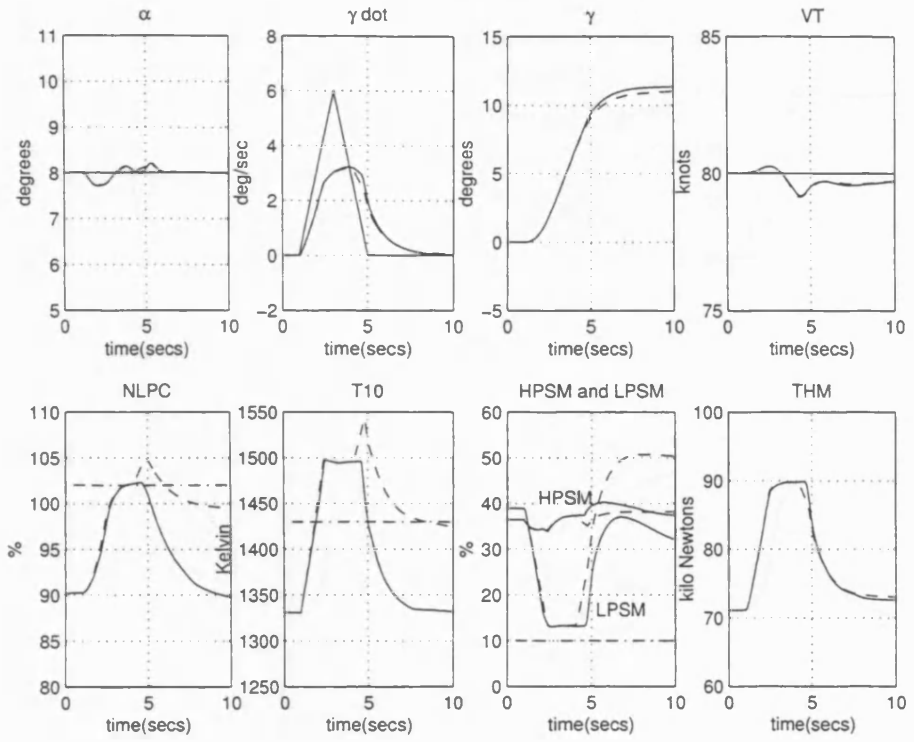

Figure 5.28: Airframe and engine responses to a $\dot{\gamma}$ demand with nominal W_2 (—) and with less control of the engine variables (- -) at 80 knots.

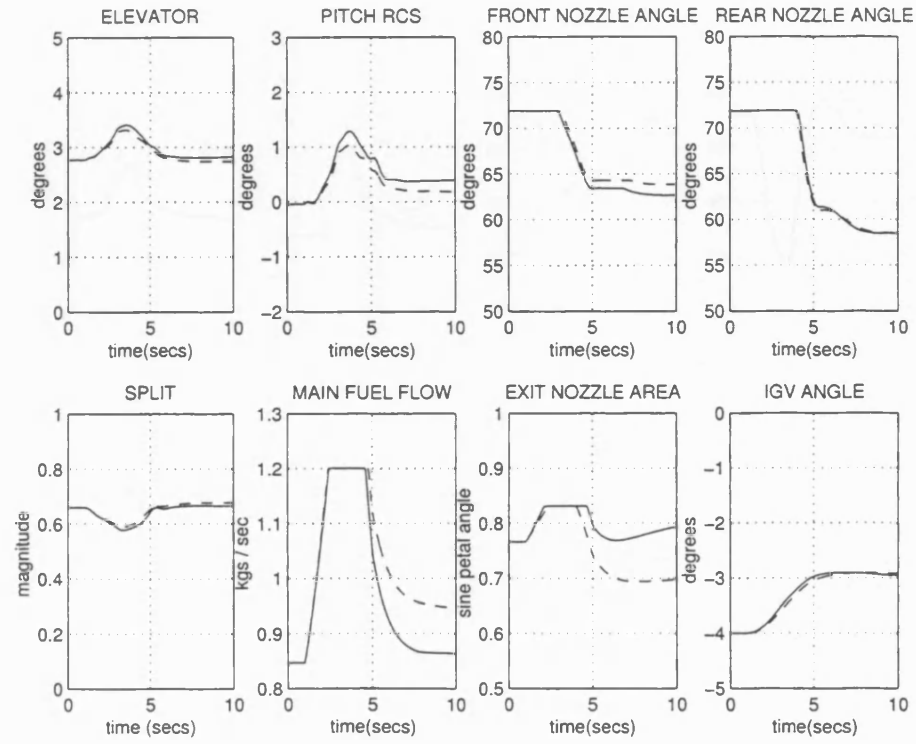

Figure 5.29: Actuator responses to a $\dot{\gamma}$ demand with nominal W_2 (—) and with less control of the engine variables (- -) at 80 knots.

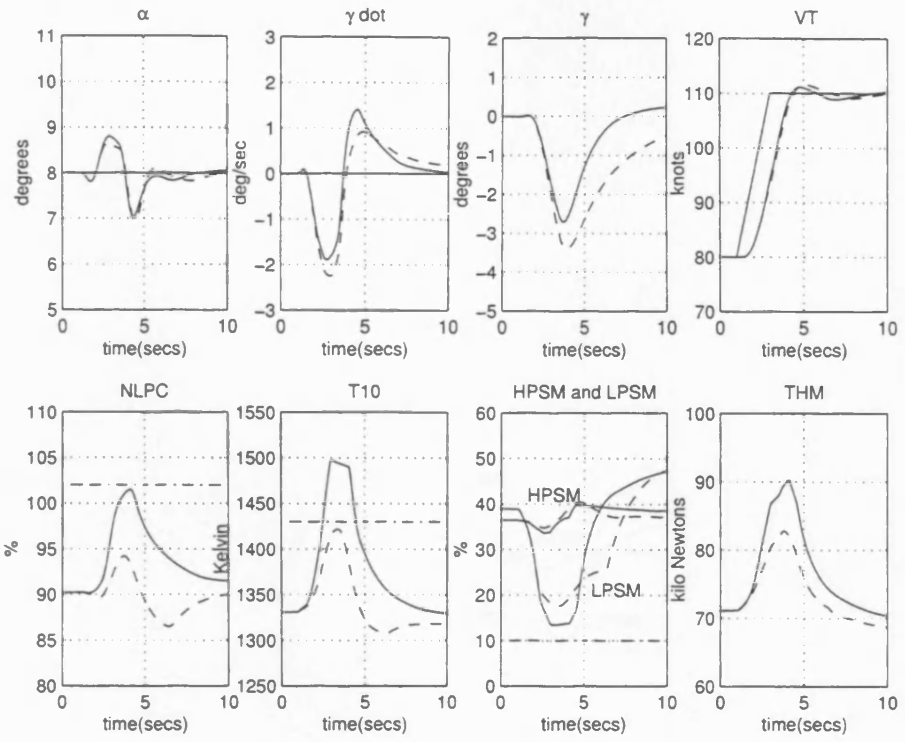

Figure 5.30: Airframe and engine responses to a VT demand with nominal W_2 (—) and with tighter control of the engine variables (- -) at 80 knots.

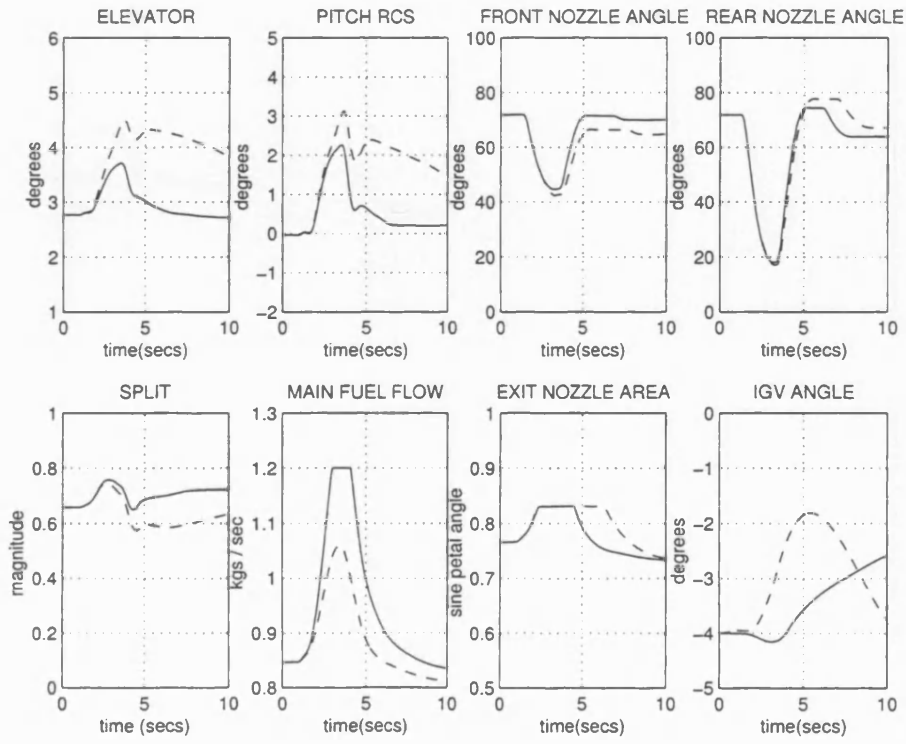

Figure 5.31: Actuator responses to a VT demand with nominal W_2 (—) and with tighter control of the engine variables (- -) at 80 knots.

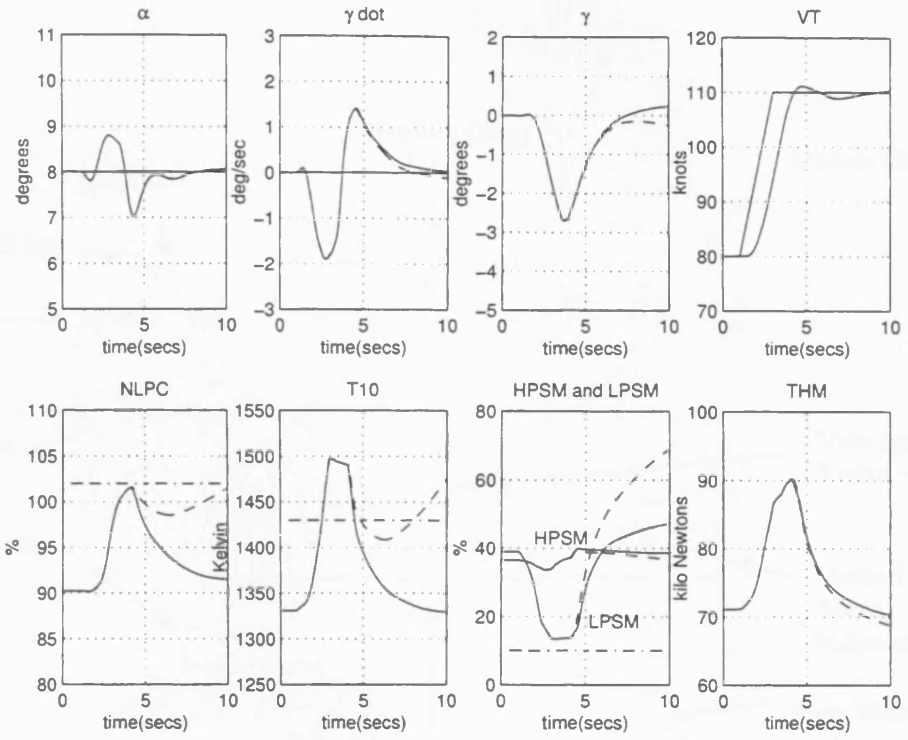

Figure 5.32: Airframe and engine responses to a VT demand with nominal W_2 (—) and with less control of the engine variables (- -) at 80 knots.

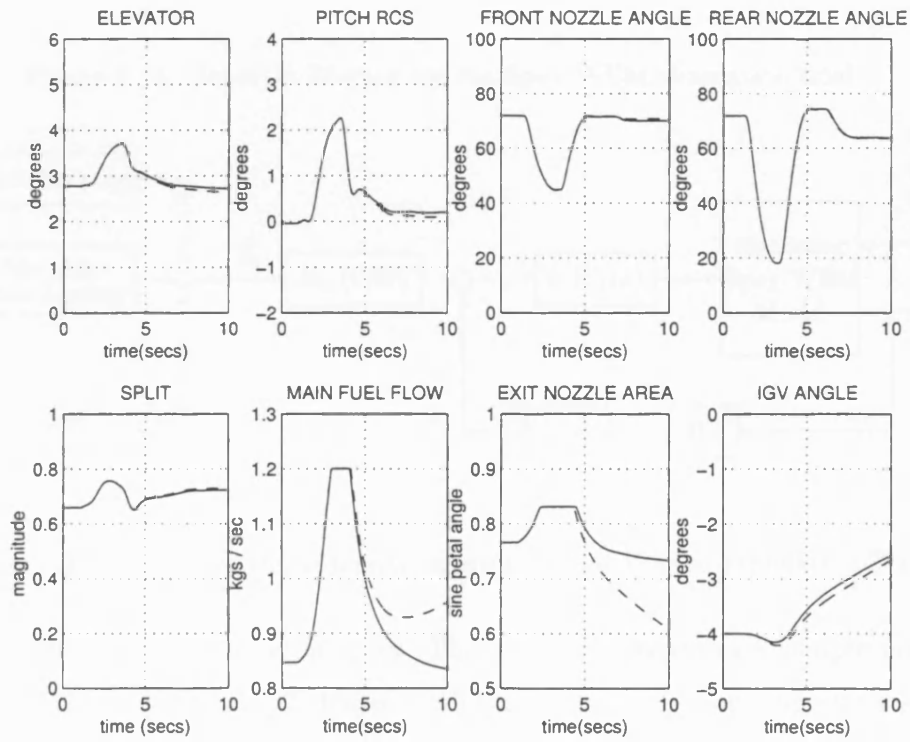

Figure 5.33: Actuator responses to a VT demand with nominal W_2 (—) and with less control of the engine variables (- -) at 80 knots.

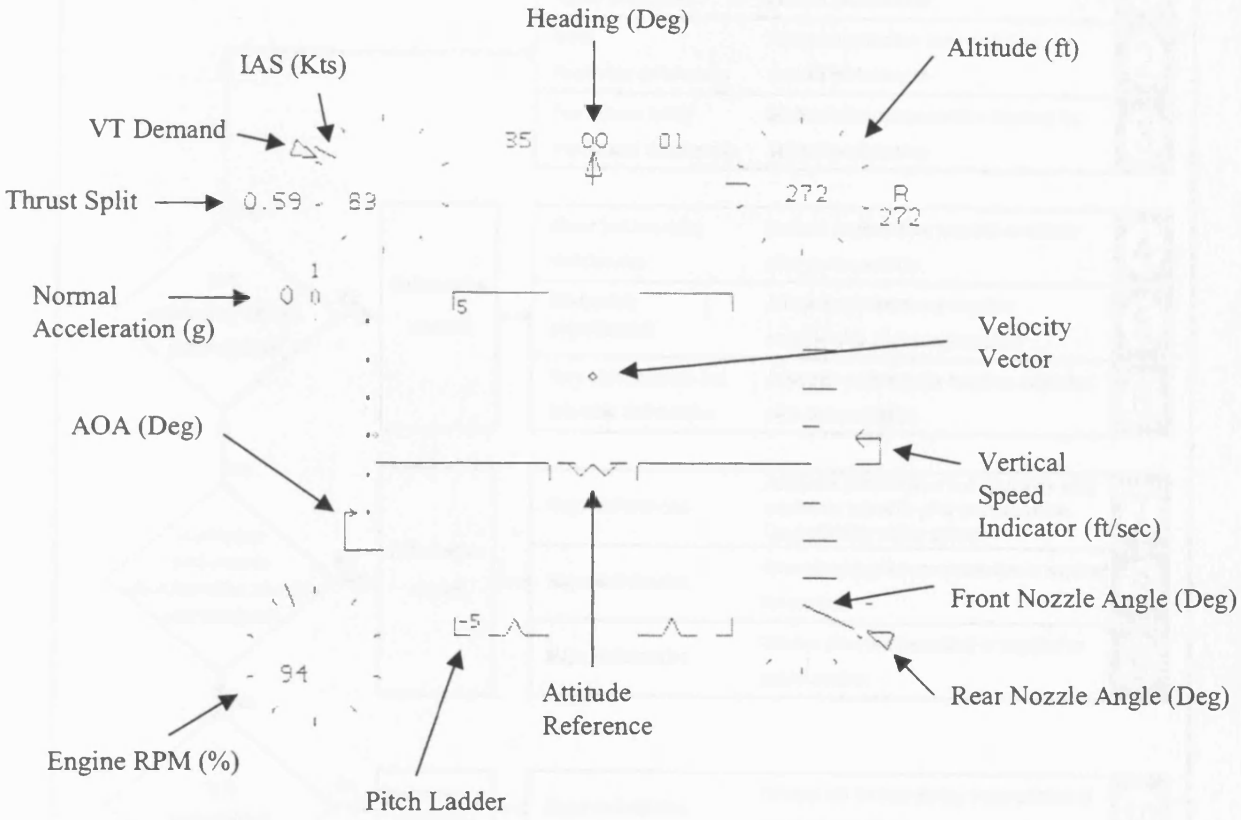

Figure 5.34: Head Up Display for the Spey-WEM simulation trial.

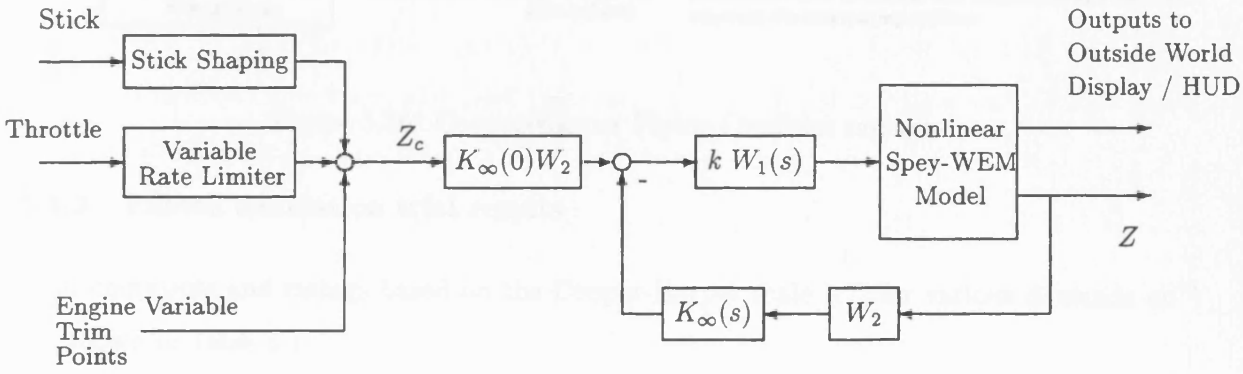

Figure 5.35: Controller Implementation structure for Piloted Simulation Trial.

demand from the pilot's left hand inceptor. This value was chosen as a compromise between providing adequate response to VT demands while minimising the coupling into γ and height. As the controller was designed for longitudinal motion only, the lateral and directional states of the model were fixed at zero. This allowed a direct evaluation of the longitudinal motion of the aircraft without the additional pilot workload involved in lateral/directional control.

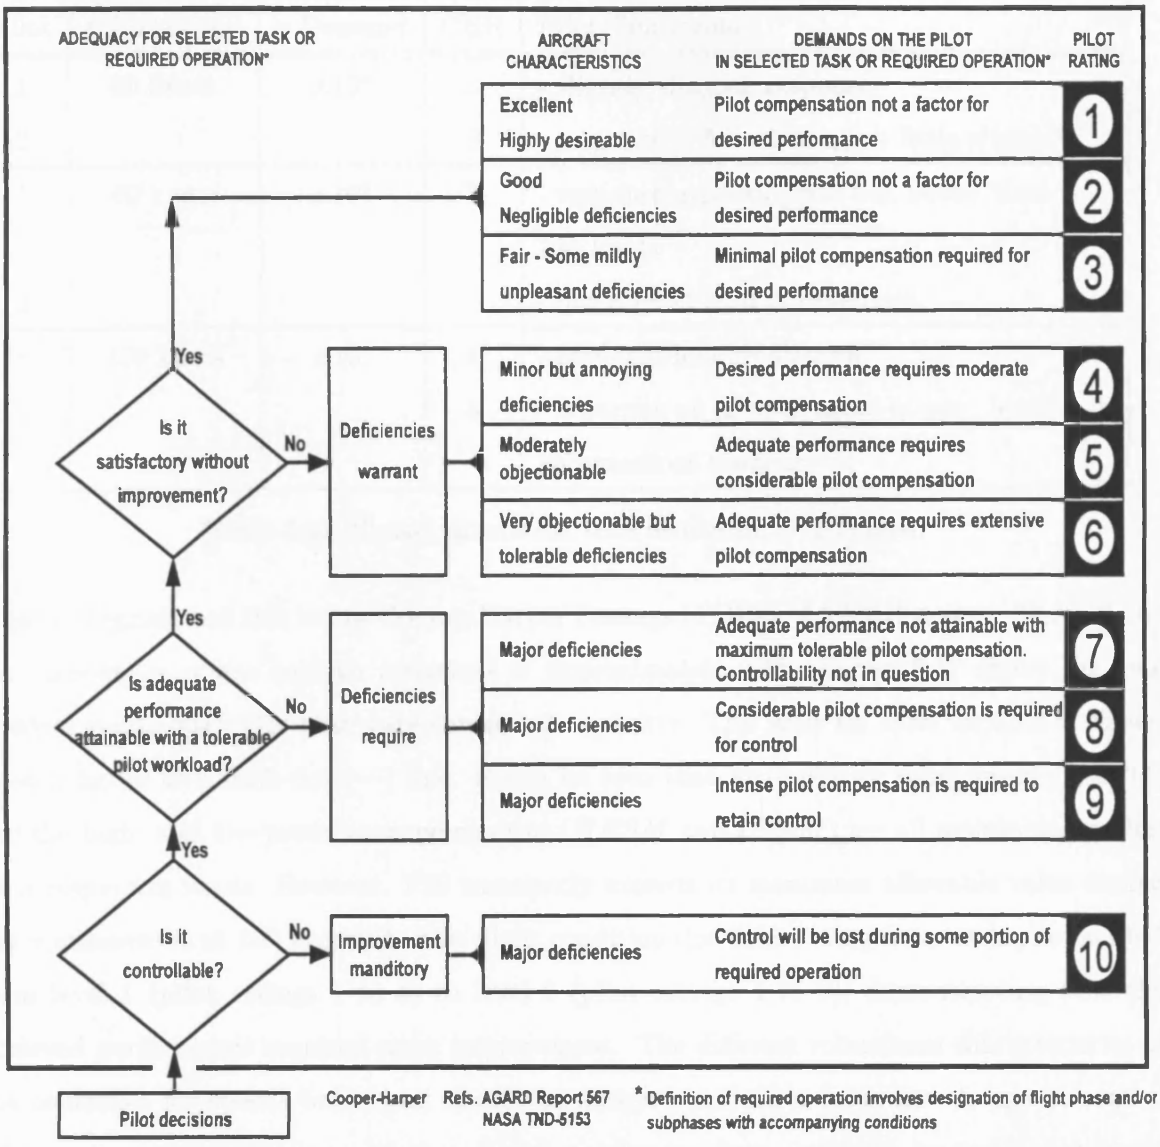

Figure 5.36: Cooper-Harper Flying Qualities ratings.

5.5.2 Piloted simulation trial results

Pilot comments and ratings based on the Cooper-Harper scale [17] for various demands on $\dot{\gamma}$ are shown in Table 5.1.

Time histories of various aircraft states, engine variables and actuators for 10° doublets on γ at both 60 knots and 100 knots are shown in Figures 5.37 and 5.38. $\dot{\gamma}$ demand signals from the stick were not recorded but consisted of demands corresponding to 10° or 20° doublets in all cases. The most common pilot comment for demands on $\dot{\gamma}$ was that at 60 or 80 knots the response was good but perhaps a little sluggish. At speeds of 100 knots or greater the response became somewhat unpredictable. At all three speeds the coupling of $\dot{\gamma}$ demands into VT was satisfactory, as pilots would generally not comment on VT variations of less than 2 knots. At 80 knots the pilots had no difficulty in acquiring desired flight path angles but the responses were

Pilot	VT Demand	γ Demand	CHR	Pilot Comments
1	80 knots	$\pm 10°$	3	- slightly sluggish response
2			3	- response is fine although a little sluggish
1	60 knots	$\pm 10°$	3	- response more sluggish but better than at 100 knots
2				- not a great deal of difference
1	100 knots	$\pm 10°$	4	- response less predictable
2			4	- γ carries on after demand is over, leading to an overshoot tendency

Table 5.1: Piloted simulation trial results for $\dot{\gamma}$ demands.

slightly sluggish and this led to Cooper-Harper Ratings (CHRs) of 3 for this task. Throughout the manoeuvre α was held to variations of approximately $\pm 1°$. Figure 5.37 shows the four internal engine variables that were controlled explicitly. The limit for each variable is shown in each figure as a dash-dot $(-\cdot)$ line. It can be seen that low-pressure spool speed $(NLPC)$ and the high- and low-pressure surge margins $(HPSM$ and $LPSM)$ are all maintained within their respective limits. However, $T10$ transiently exceeds its maximum allowable value during the γ manoeuvre at 100 knots. At this flight condition the CHR rating from each pilot shifted from level 1 (pilot ratings 1 to 3) to level 2 (pilot ratings 4 to 6), demonstrating that the achieved performance required some improvement. The different robustness characteristics of the controller for speeds below and above the design point arises from the change in control effectiveness of the actuators as the aircraft accelerates from jet-borne towards wing-borne flight. These results indicate the need for controller scheduling at higher speeds in order to retain desired handling qualities over the full flight envelope.

Pilot comments and CHRs for demands on VT are shown in Table 5.2.

Pilot	VT Demand	γ Demand	CHR	Pilot Comments
1	80 ± 10knots	$0°$	1	- snappy VT response
2			2	- a little overshoot
1	80 ± 20knots	$0°$	1	- slight cross coupling with θ
				- nothing to do but set the throttle
				- ratings didn't really consider the γ interactions
2			3	- coupling into γ

Table 5.2: Piloted simulation trial results for VT demands.

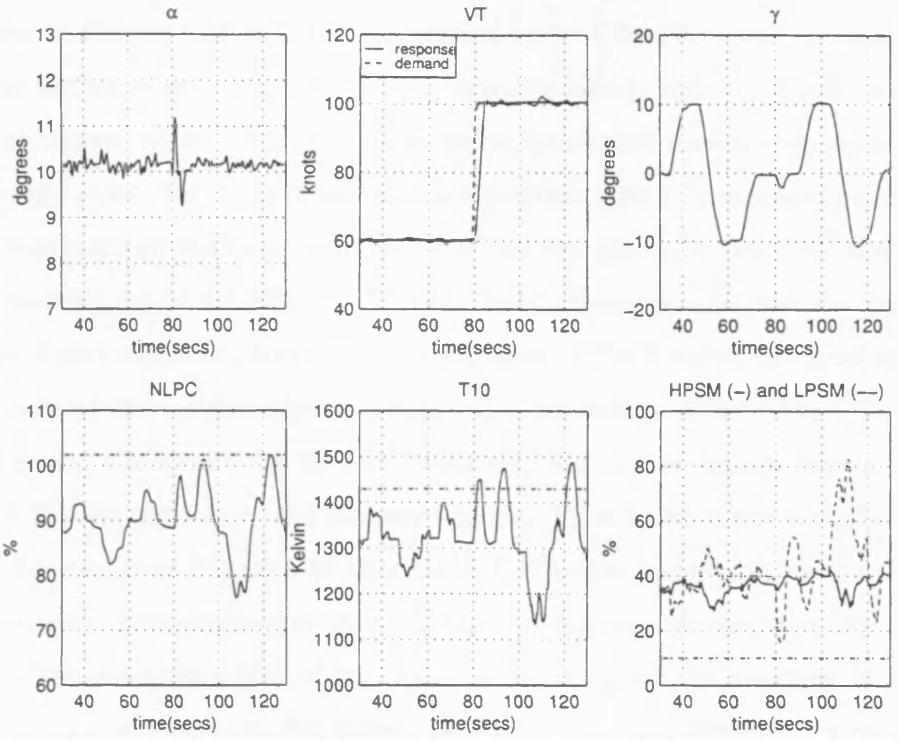

Figure 5.37: Airframe and engine responses to a $\dot{\gamma}$ demands at 60 and 100 knots for the piloted simulation.

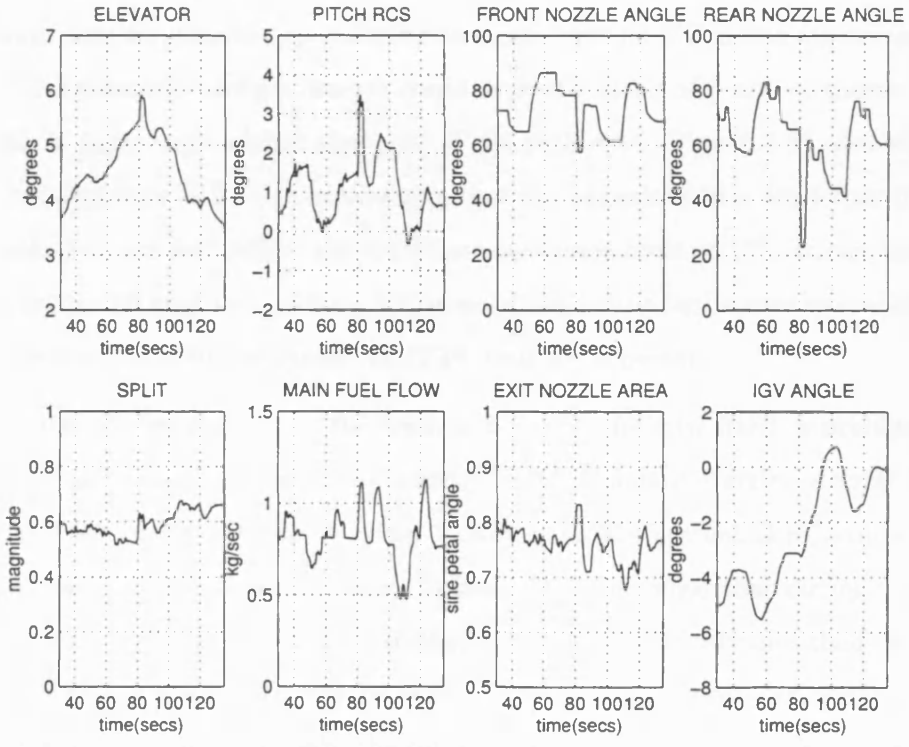

Figure 5.38: Actuator responses to a $\dot{\gamma}$ demands at 60 and 100 knots for the piloted simulation.

Time histories of aircraft states, engine variables and actuators for 10 and 20 knot doublets on VT are shown in Figures 5.39 to 5.40. The general view of the pilots was that the response to VT demands was excellent, but the resulting coupling into γ, and so altitude, was too large. This coupling became worse for larger VT steps as the aircraft moved away from the 80 knot controller design point. For both 10 and 20 knot doublets pilot 1 commented on the sharpness of the VT response and tracking. Comments of 'no compensation required' and 'nothing to do but set throttle' led to a CHR of 1 for both tasks. However, the pilot did comment that these ratings didn't really consider the γ interactions. Pilot 2 noted the good speed control but also considered the γ interactions when giving their ratings. This led to CHRs of 2 for 10 knot demands and CHRs of 3 for 20 knot demands. The poorer ratings were a result of the coupling of VT demands into γ and therefore height. Pilot 2 also noted a slight overshoot in VT. It can be seen from Figure 5.39 that the VT response appears to have a time delay of around 0.8 seconds. Examination of the actuator time histories showed that this delay mostly resulted from the dead-band (6°) of the nozzle actuators generally resulting in a 0.5 second delay in nozzle reaction to controller demands. Figure 5.39 also shows the γ response to the VT changes. It can be seen that there is significant coupling into γ with a maximum drop of greater than 1° for 10 knot demands and greater than 2° for 20 knot demands. Unexpectedly, this coupling into γ is negative for both positive and negative speed changes, leading to height reductions of approximately 3m and 8m respectively (Figure 5.41). Improvements in the level of γ interaction may be possible by reducing the rate limit on VT below its current value of 10 knots s^{-1}. Alternatively, future designs could control a blend of $\dot{\gamma}$ and γ, thereby explicitly controlling flight path angle rather than just flight path rate. Figure 5.39 also shows the α response to the demands in VT. It can be seen that the variations in α are larger than for the $\dot{\gamma}$ demand tasks but are still below the maximum incidence limit of 12°. It can be seen from Figure 5.39 that for 10 knot demands on VT none of the engine limits were exceeded, but that the 20 knot doublet manoeuvres caused the $T10$ limit to be broken.

During the course of the work-up to the trial, and during the trial itself, a problem with the front and rear nozzle positions became evident. In the standard Harrier aircraft both front and rear nozzles are fixed to the same value. However, for the current IFPC study the nozzle angles were allowed to vary independently. In general it was found that during a manoeuvre the front and rear nozzle angles would transiently vary independently and then return to the same value as another trim state was reached. However, after large demands in VT or $\dot{\gamma}$, especially at higher than design speeds, the controller sometimes trimmed the aircraft with the nozzles at different angles. This can be seen in Figure 5.38 where the front and rear nozzles come to rest approximately 5° apart. Once the nozzles had split in a steady flight regime they

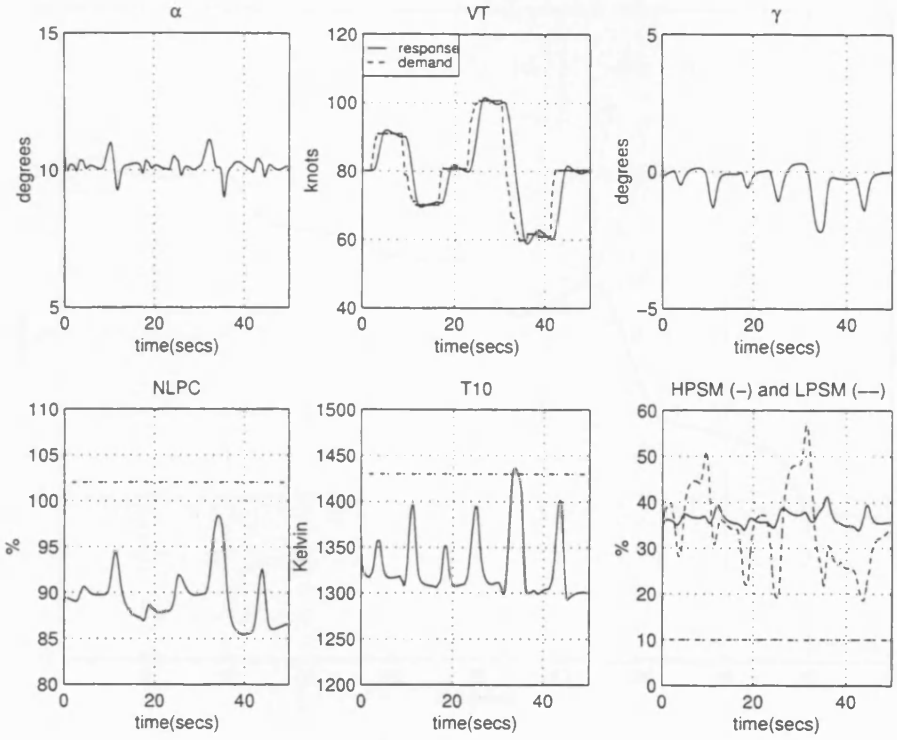

Figure 5.39: Airframe and engine responses to 10 knot and 20 knot demands on VT for the piloted simulation.

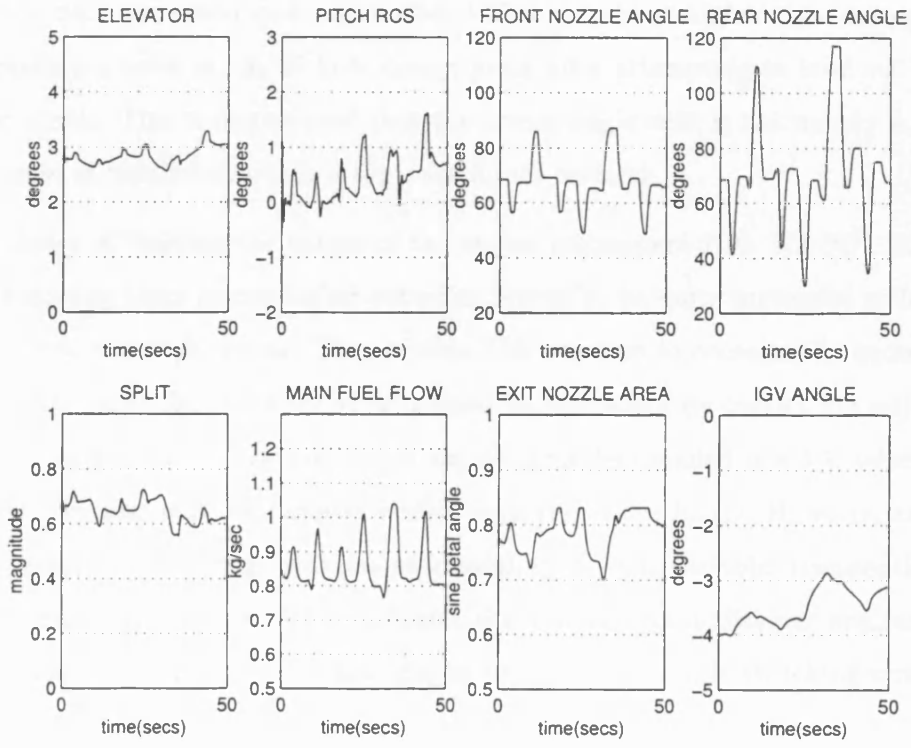

Figure 5.40: Actuator responses to 10 knot and 20 knot demands on VT for the piloted simulation.

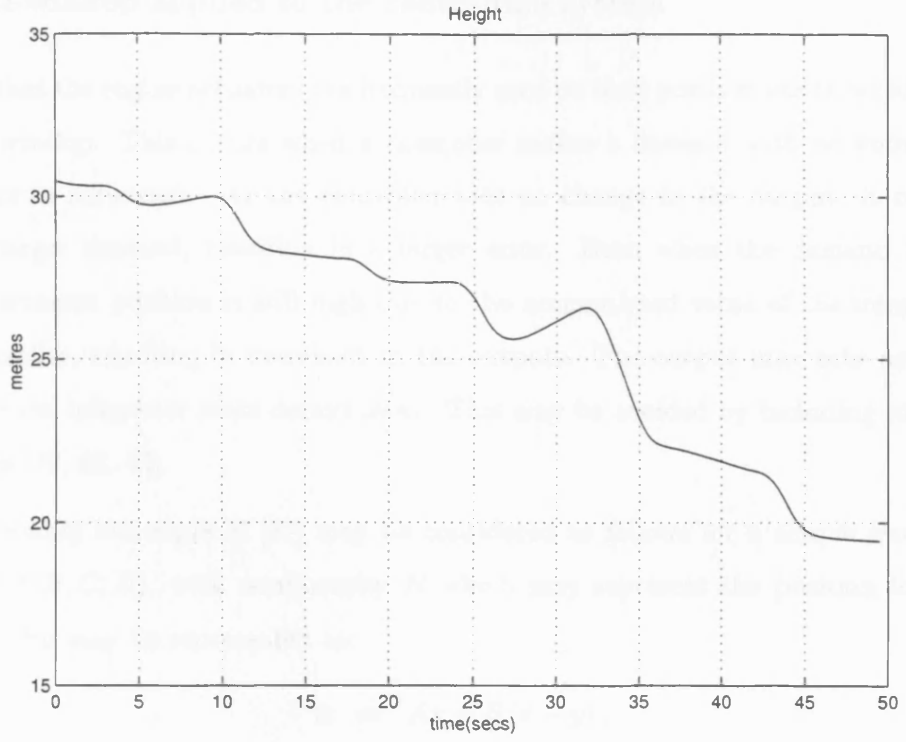

Figure 5.41: Height response to 10 knot and 20 knot demands on VT for the piloted simulation.

tended to drift further apart over time with the result that they began to oppose each other. Although this effect was noted more often after VT or $\dot{\gamma}$ tasks at higher than design speeds, it also occasionally occurred at the 80 knot design point after attempting to level out or descend after a steep climb. This demonstrated that the nozzle angle drift is not merely a robustness issue. The cause of this behaviour is a topic for future research.

The design choice of limiting the values of the engine parameters $T10$, $NLPC$, $HPSM$, and $LPSM$ by including them as controlled variables proved to be quite successful at both the 80 knot design speed and at 60 knots. The variable $T10$ was seen to occasionally exceed its limit at 80 knots, but this could probably be addressed by tightening its control via adjustment of the $\mathcal{H}_\infty$ weighting functions. The additional engine variables detailed in § 5.1, which were not controlled directly, also generally stayed within their respective limits. However, at speeds of 100 knots or greater the strategy was less successful, with some variables transiently breaking their limits during manoeuvres. This indicates the need for an additional nonlinear control scheme to guarantee hard limits, for example by using a multi-mode switching strategy as in [59].

5.6 Anti-windup applied to the centralised system

It is noted that the engine actuators are frequently used on their position limits, which instigates integrator windup. This occurs when a controller makes a demand with no knowledge that the actuator is saturated. As the controller sees no change in the output, it continues to request a larger demand, resulting in a larger error. Even when the demand is met, the requested actuator position is still high due to the accumulated value of the integrator state of the controller, resulting in overshoot in the outputs. The output may take some time to settle while the integrator state decays away. This may be avoided by including some form of anti-windup [70, 63, 44].

The conditioning technique of [37] may be considered as follows for a simple error feedback controller (A, B, C, D), with nonlinearity N which may represent the position limits of the actuators. This may be represented as:

$$\begin{aligned} \dot{x} &= Ax + B(r - y), \\ u &= Cx + D(r - y), \\ u_m &= N(u). \end{aligned} \tag{5.1}$$

where r and y represent the reference input and the plant output, and u_m represents the actual plant input. It is possible to implement a 'realisable' reference r_r to the controller such that the output of the controller is u_m, i.e., r_r is the reference signal that would make $u = u_m$ if applied to the controller state and output equations in place of r. Thus, we have:

$$\dot{x} = Ax + B(r_r - y), \tag{5.2}$$

$$u_m = Cx + D(r_r - y). \tag{5.3}$$

Using the state x as given by (5.2) and the actual reference input r to build the control u, we can write:

$$u = Cx + D(r - y). \tag{5.4}$$

From (5.3) and (5.4) we obtain the expression for r_r as:

$$r_r = r + D^{-1}(u_m - u), \tag{5.5}$$

where we assume that D is invertible. Combining equations (5.1), (5.2), (5.4), and (5.5), we get:

$$\dot{x} = \left(A - BD^{-1}C\right)x + BD^{-1}u_m, \tag{5.6}$$

$$u = Cx + D(r - y), \tag{5.7}$$

$$u_m = N(u). \tag{5.8}$$

This is the self-conditioned form of the controller. However, from the implementation structure, it can be seen that all of the integral action is contained in the weighting matrix W_1, given by:

$$W_1 := \left[\begin{array}{cc} A_{W_1} & B_{W_1} \\ C_{W_1} & D_{W_1} \end{array}\right] \tag{5.9}$$

This can instead be implemented in its self-conditioned form [37] as:

$$u = \left[\begin{array}{c|cc} A_{W_1} - B_{W_1} D_{W_1}^{-1} C_{W_1} & 0 & B_{W_1} D_{W_1}^{-1} \\ \hline C_{W_1} & D_{W_1} & 0 \end{array}\right] \left[\begin{array}{c} u_s \\ u_r \end{array}\right] \tag{5.10}$$

as shown in Figure 5.42.

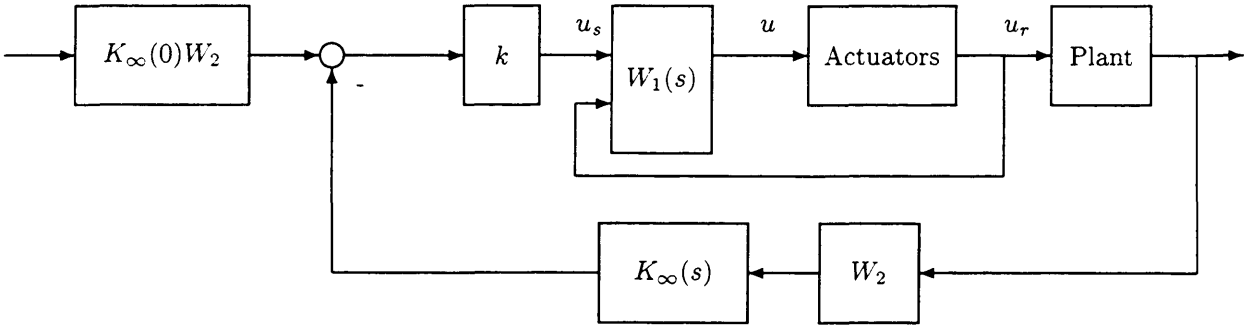

Figure 5.42: Anti-windup Implementation.

This means windup is prevented by keeping the states of W_1 consistent with the actual plant input at all times. When there is no saturation $u_r = u$, the dynamics of W_1 remain unaffected and (5.10) simplifies to (5.9). But when $u_r \neq u$ the dynamics are inverted and driven by u_r so that the states remain consistent with the actual plant input u_r. This implementation requires W_1 to be invertible and minimum phase.

The results for a demand on $\dot{\gamma}$ are shown in Figures 5.43 and 5.44. It can be seen that the $\dot{\gamma}$ tracking is improved, with $\dot{\gamma}$ decreasing more rapidly when the anti-windup is applied. This also has the benefit of reducing the engine variables sooner. It is clear these benefits result from the main fuel flow and the exit nozzle area coming out of saturation sooner. The results for a demand on VT are shown in Figures 5.45 and 5.46. It can be seen here that the $\dot{\gamma}$ coupling is slightly reduced.

5.7 IFPC using pure thrust vectoring

A logical continuation of this work is to implement some form of controller scheduling, as it was seen during the piloted simulation that performance, in particular at velocities above the nominal design point of 80 knots, was degraded [33]. One consideration was the possibility of adjusting the weights on the actuators to reflect the effectiveness of each actuator across the

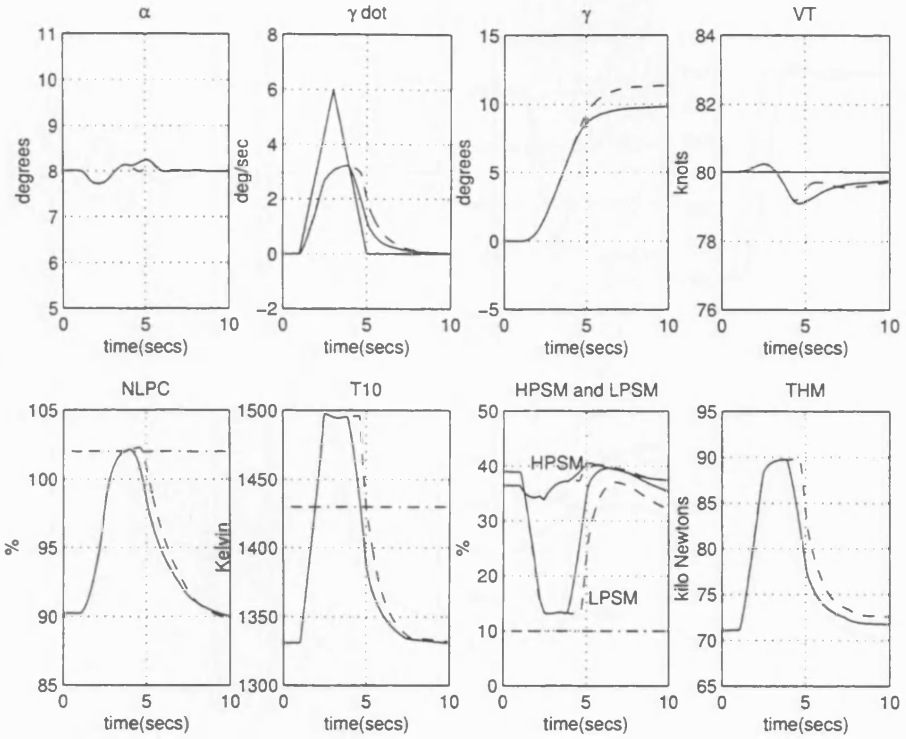

Figure 5.43: Airframe and engine responses to a demand on $\dot{\gamma}$ for the centralised system with anti-windup (—) and for the centralised system (- -).

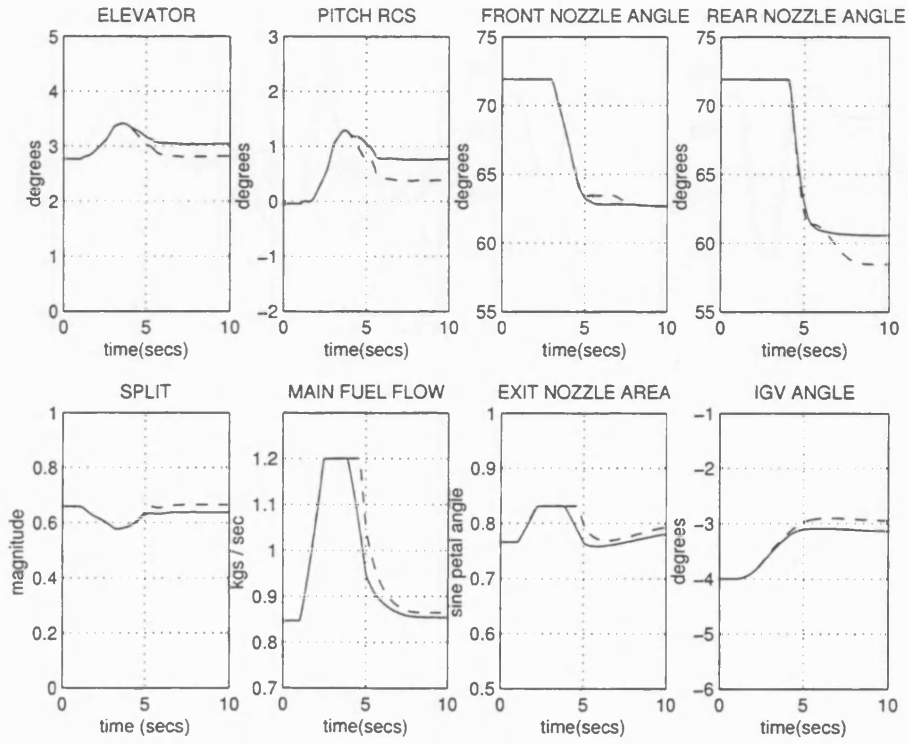

Figure 5.44: Actuator responses to a demand on $\dot{\gamma}$ for the centralised system with anti-windup (—) and for the centralised system (- -).

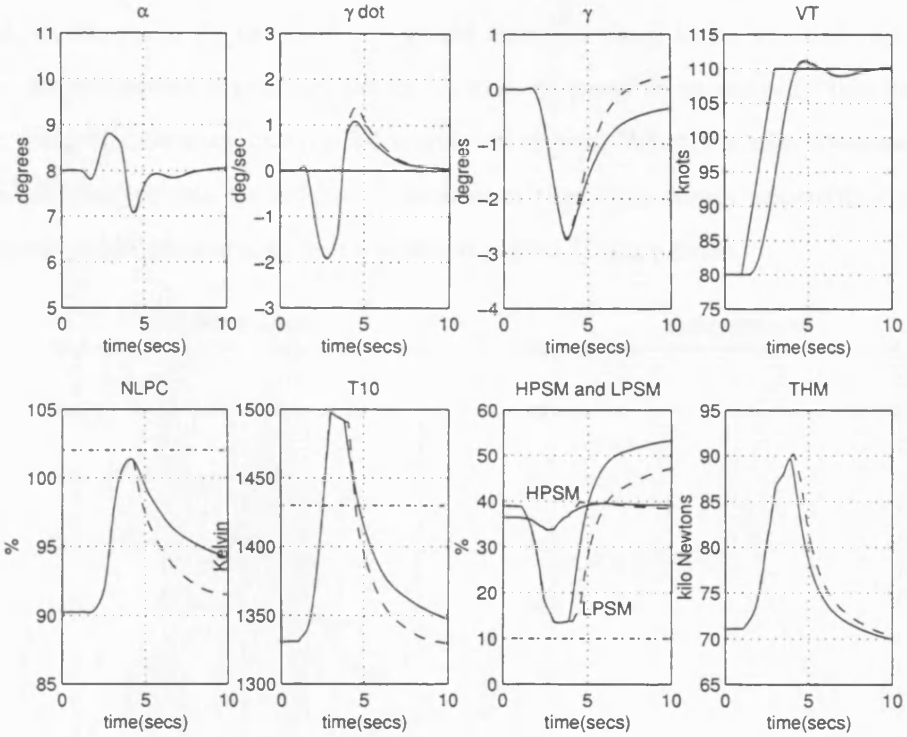

Figure 5.45: Airframe and engine responses to a demand on VT for the centralised system with anti-windup (—) and for the centralised system (- -).

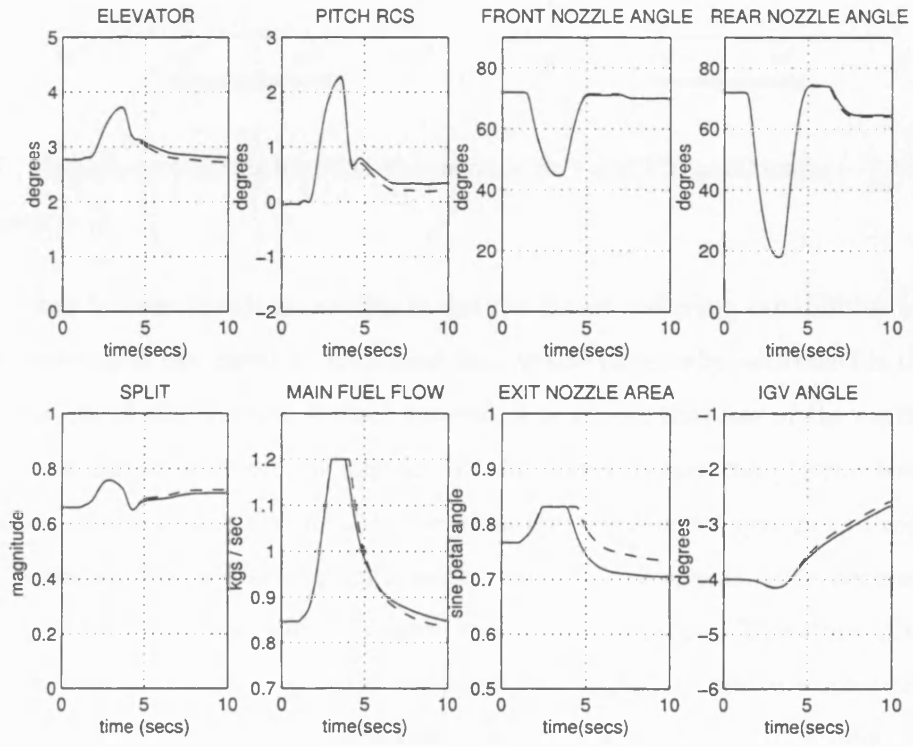

Figure 5.46: Actuator responses to a demand on VT for the centralised system with anti-windup (—) and for the centralised system (- -).

flight envelope. It can be seen that the elevator becomes increasingly ineffective as hover is approached, as shown in Figure 5.47. It would therefore need to be blended out accordingly throughout the transition region. However, it may be possible to simplify this scheduling by simply not using the elevator during the transition region. When the effectiveness of the front and rear nozzle angles was considered, it was seen that they retain approximately the same level of control power throughout the transition region (Figure 5.48).

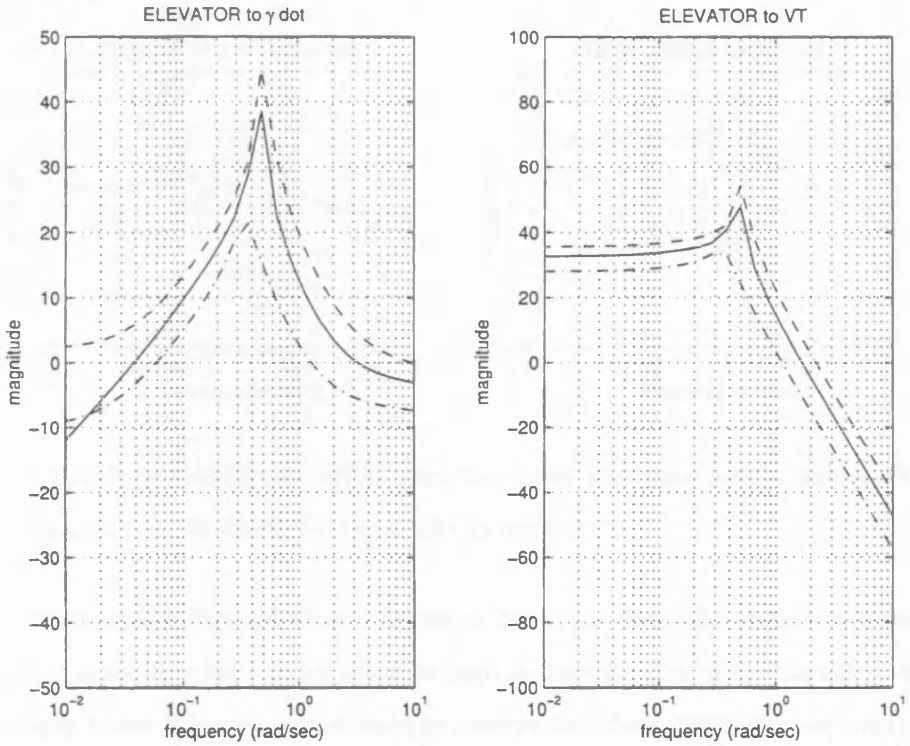

Figure 5.47: Open-loop transfer function the elevator to $\dot{\gamma}$ and VT at 50 knots (-·), 80 knots (—) and 120 knots (- -).

From this, it can be seen that it is possible to use the thrust vectoring capabilities to effectively replace the function of the elevator. One issue that would have to be addressed is the resulting increase in the use of the reaction control system. It is known that use of the reaction control system creates a disturbance on the engine, via the high pressure compressor bleed flow air. It has also been seen through testing that the pitch reaction control system can easily become saturated, degrading the achieved performance. At very low speeds, it is necessary to hold power in reserve for use of the reaction control system manoeuvring. Therefore, if the reaction control system were not in use, this additional power that is held in reserve will become available to the thrust vectoring nozzles. This additional power held in reserve will also affect the payload and range performance of the aircraft. Use of reaction control systems which bleed air from the compressor for manoeuvring control also carries significant penalties in terms of both engine and airframe performance. In hover, the entire control capability for the aircraft is derived from the

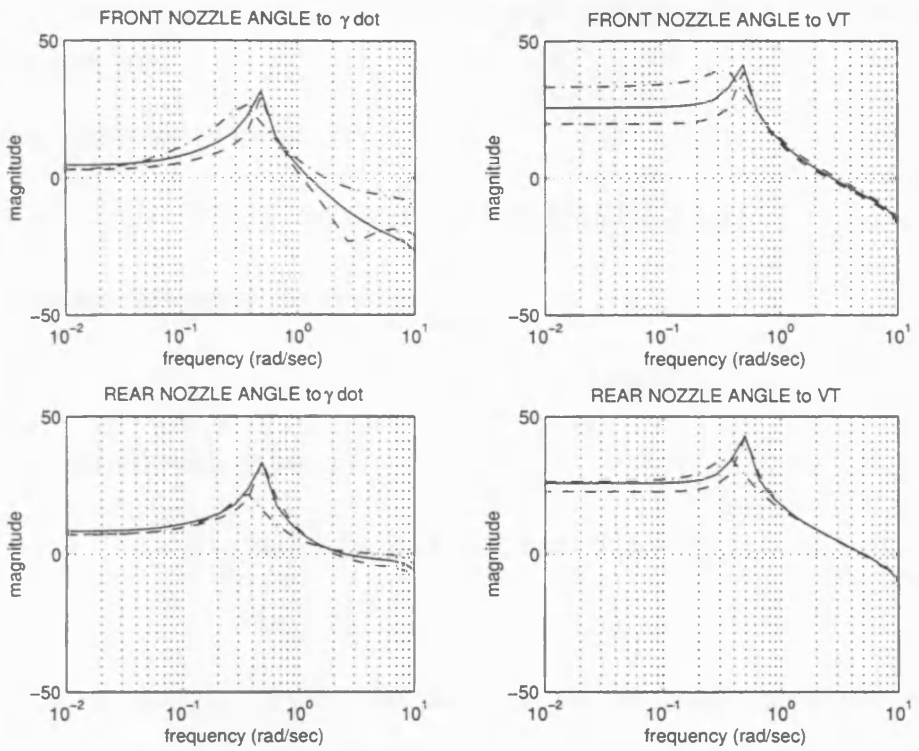

Figure 5.48: Open-loop transfer function from the front and rear nozzle pair deflections to $\dot{\gamma}$ and VT at 50 knots (-·), 80 knots (—) and 120 knots (- -).

propulsion system, either through direct thrust or bleed air from the engine compressor. With requirements for manoeuvring control often as high as one-quarter of the thrust-to-weight ratio needed to simply hover [51] any thrust held in reserve for these tasks can seriously reduce the aircraft's payload and range performance. The requirements on the propulsion system can thus be significantly alleviated if the aircraft can be controlled via thrust vectoring and thrust split (between front and rear nozzles), without recourse to bleed-air actuation systems. The final motivation for this study is to investigate the actual degree of actuator redundancy present in the proposed aircraft configuration, in order to understand the possibilities for controller reconfiguration during different failure modes. Indeed the proposed control scheme can be considered as an emergency control system which can be 'switched in' in the event of damage to the elevator or propulsion system. In this section, an integrated flight and propulsion control scheme is presented which provides full-authority flight-path and velocity control using only differential thrust vectoring. In view of this, it was decided to design a purely thrust vectoring controller without the use of either the elevator or the pitch reaction control system. Another use for a pure thrust vectoring control scheme is in the event of actuator failure. If an airframe control surface were to fail, it is then possible to use the thrust vectoring nozzles and still achieve the required performance. This is a clear advantage of adopting a IFPC approach, as there are many combinations of actuator positions that will produce the same overall effect on

the airframe. Therefore, if one actuator fails, the remaining actuators will ensure the desired performance is still met.

The remaining inputs to the system are:

$$u = [FNOZ, RNOZ, SPLIT, MFF, ENOZA, IGV]$$

which are to be used to control the following:

$$z = [\alpha, VT, \dot{\gamma}, NLPC, T10, HPSM]$$

The weighting matrices were chosen as:

$$k = diag\,(0.25,\ 0.23,\ 0.3,\ 0.25,\ 0.25,\ 0.2)\,,$$

$$W_1 = \frac{s+1}{s} \times I_{6\times 6}, \quad W_2 = I_{6\times 6}$$

which achieved a value of $\gamma_{opt} = 3.94$. It was noted that the thrust vectoring controller responses also resulted in the engine actuators frequently being used on their limits. The anti-windup protection scheme as described in § 5.6 is therefore also implemented. In this section we seek to replicate the performance results achieved in § 5.5 without using either the elevator or reaction control system as control inputs.

Performance results for the IFPC system in nonlinear simulation are presented in Figures 5.49 to 5.52. Figures 5.49 and 5.50 show the responses of controlled variables and actuators respectively to a 15 knot demand in VT. The coupling into γ is just outside the 0.3° specification for both the centralised and thrust vectoring controller, but is still quite small for such a large velocity demand. The specified limits on α and the internal engine variables are seen to be respected at all times. One effect of not using the bleed air for the reaction control system is lower temperatures within the engine. The VT response is only slightly slower for the thrust vectoring controller.

Figures 5.51 to 5.52 give the same information for a demand on $\dot{\gamma}$. The γ change is slightly slower for the thrust vectoring controller. Coupling into VT and α is seen to be very small, but the large thrust change required does cause both $T10$ and $NLPC$ to transiently exceed their specified limits. It can be seen that the absence of the elevator and pitch reaction control actuators in the new design has been compensated for through modified use of the vectored nozzles and thrust split capability. This is not, however, simply a case of moving the nozzles more - it is a question of how the thrust is distributed in terms of both magnitude and direction. Overall, the two sets of results are remarkably similar, with just marginally increased coupling into α and internal engine variables for the differential thrust design.

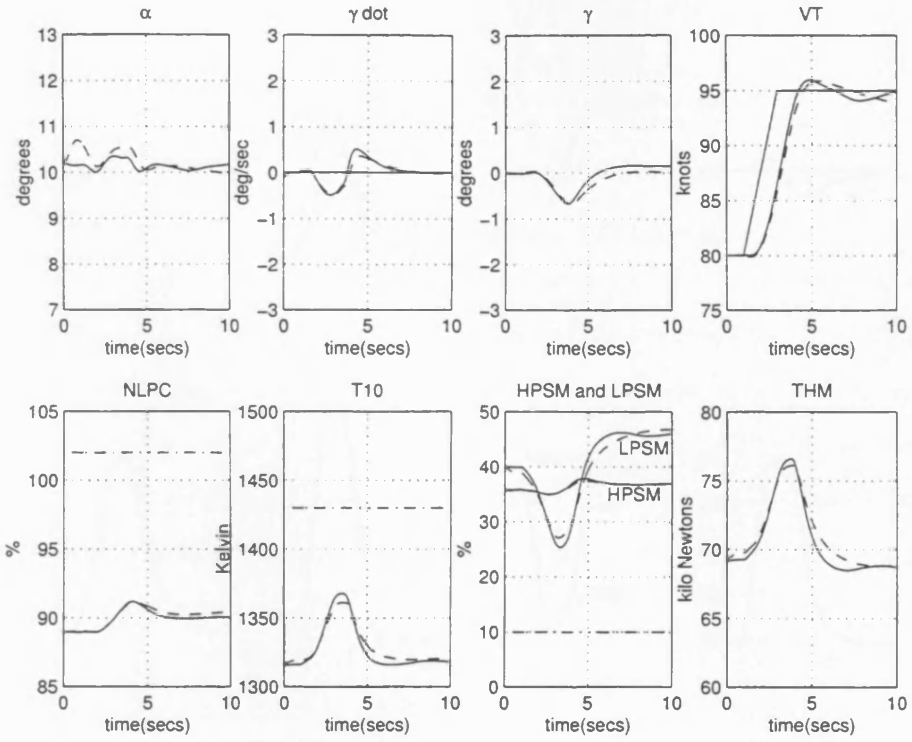

Figure 5.49: Airframe and engine responses to a VT demand for the centralised controller (—) and the thrust vectoring controller (- -) at 80 knots.

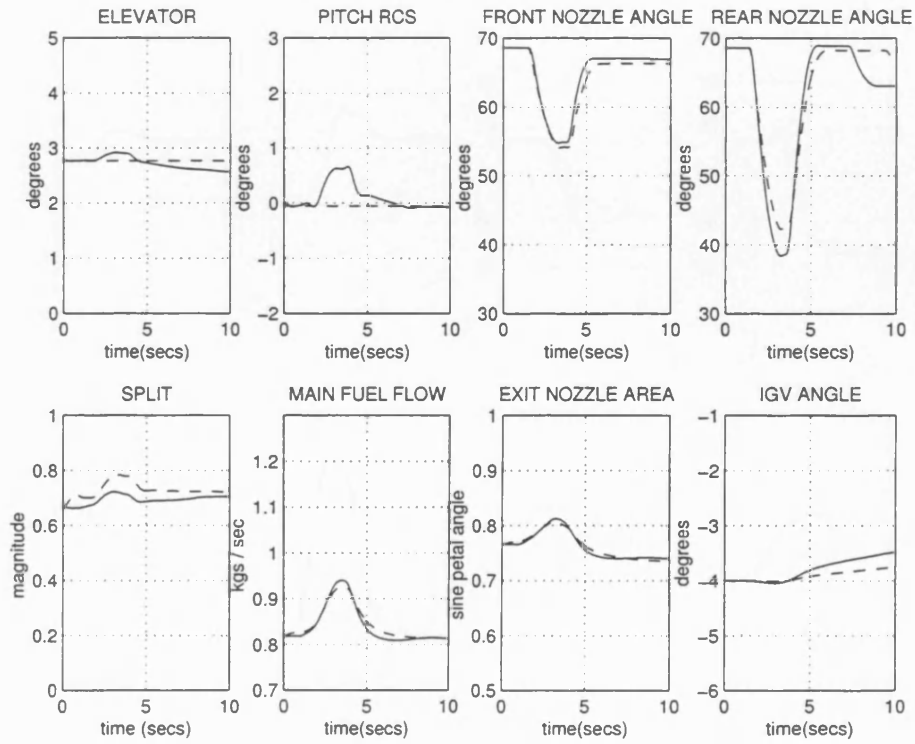

Figure 5.50: Actuator responses to a VT demand for the centralised controller (—) and the thrust vectoring controller (- -) at 80 knots.

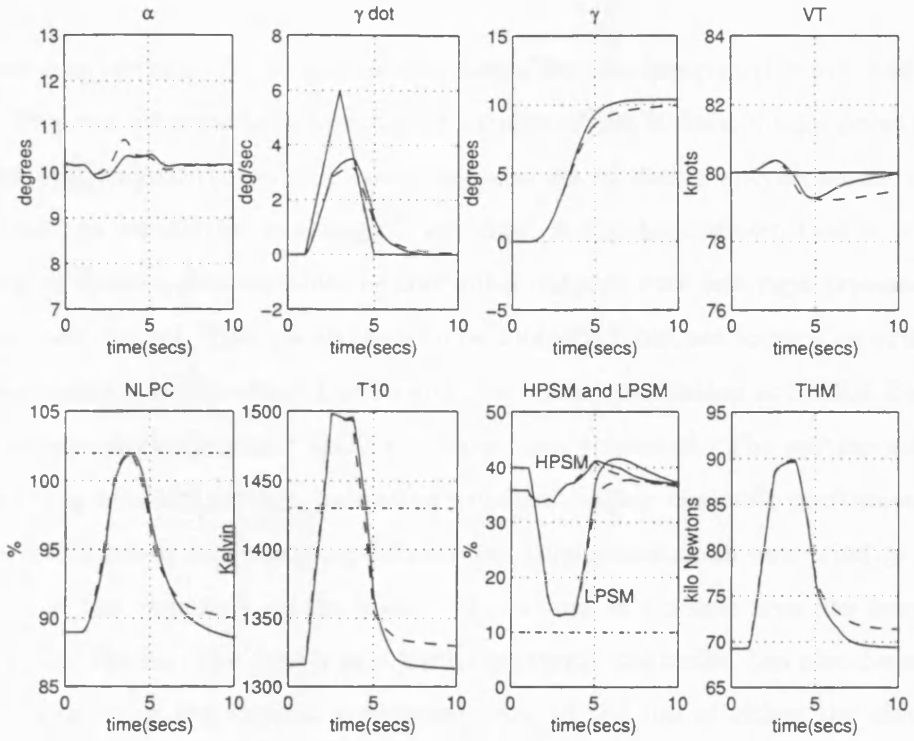

Figure 5.51: Airframe and engine responses to a $\dot{\gamma}$ demand for the centralised controller (—) and the thrust vectoring controller (- -) at 80 knots.

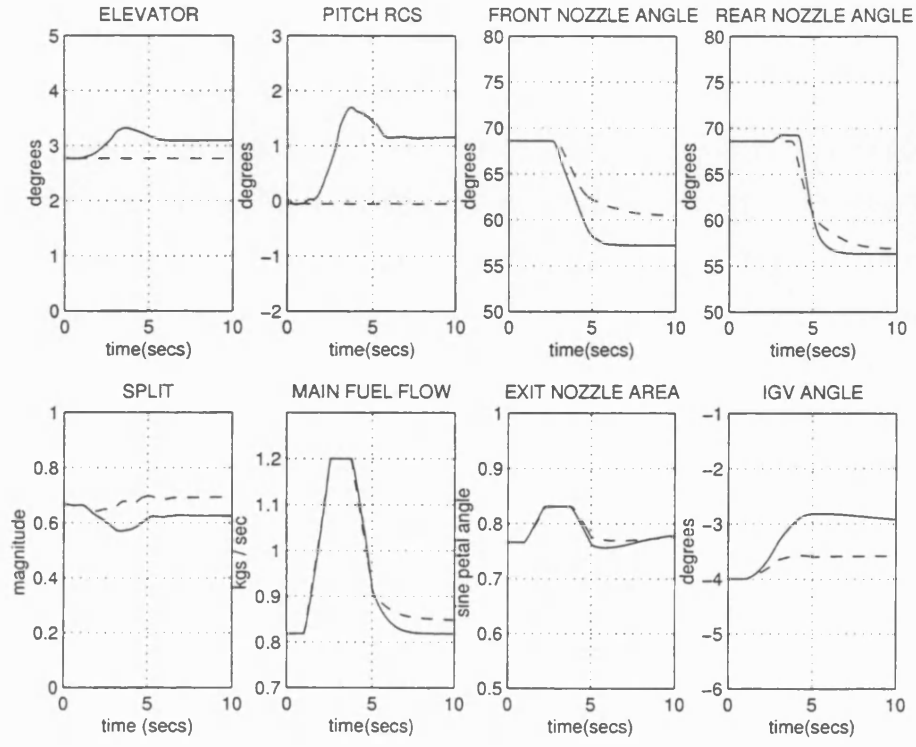

Figure 5.52: Actuator responses to a $\dot{\gamma}$ demand for the centralised controller (—) and the thrust vectoring controller (- -) at 80 knots.

5.8 Summary

This chapter has introduced the centralised design for the integrated flight and propulsion controller. This was achieved by designing a controller at the 80 knot design point in the flight envelope for longitudinal control. A comprehensive set of design specifications were stated, including limits on certain internal engine variables. It has been shown that it is possible to include some of these engine variables as controlled outputs with less rigid transient demands placed upon their control. This allows them to be controlled, but not to such an extent that the airframe performance is degraded. The results of a piloted simulation at DERA Bedford using the real-time all-vehicle simulator (RTAVS) have been presented. The system achieved level 1 and 2 handling qualities ratings, indicating excellent, highly desirable performance levels to performance with minor but annoying deficiencies. Implementing an anti-windup scheme was seen to improve the responses, as the main fuel flow and exit nozzle area are frequently used on their position limits. The design of a thrust vectoring controller has also been presented, whereby the control of the aircraft is achieved without the use of either the elevator or the pitch reaction control system. Nonlinear simulation results have been presented that show good performance across the flight envelope.

Chapter 6

Partitioning of the centralised IFPC system

Partitioning the centralised controller is a crucial part of this approach to IFPC. A centralised controller has been designed as this provides a benchmark for performance that can be achieved, in which all of the interactions between the subsystems are considered. However, it would not be practical to implement a controller with such a structure, due to impediments that include the high order of the resulting centralised controller, and the use of feedback paths which are unrealistic in the context of available sensor and databus technology. Complex centralised controllers are also difficult to certify for flight-clearance, as they inherently suffer from a lack of transparency with regard to subsystem functionality. A further complication arises from the fact that the aircraft's airframe and engine are generally designed by different manufacturers. Since the engine manufacturers must be able to guarantee satisfactory operation of the engine propulsion system, a separate engine subcontroller is required for integrity and performance testing. For example, engine safety considerations require that some engine variables never exceed certain critical limits. This requirement generally necessitates the addition of some nonlinear switching logic to the engine control system - a task which would be very difficult in the context of one centralised IFPC system. Furthermore, the control requirements of the airframe and engine are inherently different, and the control system designers have significant experience in their respective fields. If separate airframe and engine subcontrollers are used, such experts may re-design the subcontrollers using their expert knowledge in order to ensure that the subcontrollers produce the best performance.

Chapter 2 discussed the different approaches to IFPC design - the design of a centralised controller, a decentralised controller or a direct decentralised controller. Both the centralised and decentralised methods had advantages and disadvantages, so a combination of the two was seen to provide the most favourable solution. First a centralised controller is designed that accounts for all of the interactions within the subsystems. This centralised controller may then be partitioned into subcontrollers that would be easier to implement, validate and

re-design. It is this partitioning procedure that is described in this chapter. § 6.1 introduces some of the previously used approaches to partitioning the centralised controller. § 6.2 presents the partitioning approach used, along with details of how the individual subcontrollers may be modified while retaining the desired performance. § 6.3 compares the nonlinear simulation results of the centralised system with the partitioned system. Finally, § 6.4 provides a summary of this chapter.

6.1 Initial partitioning methods

Several alternative methods were used prior to the final version presented here. This work has been published in the IFAC Journal of Control Engineering Practice [32]. The first involved assigning all controlled outputs and inputs as either airframe or engine based on control effectiveness rather than physical location. More details of how this type of selection is made will be given in § 6.2. The most significant of these assignments was that of $\dot{\gamma}$ as an engine controlled output. The motivation behind this is the fact that, although the flight path angle is physically an airframe control variable, it is primarily affected by the engine at the nominal 80 knot design point. The resulting partitioned system was also tested in the piloted simulation described in Chapter 5. The pilots were unable to differentiate between the centralised and the partitioned systems, giving both level 1 or 2 Cooper-Harper handling quality ratings. While this procedure allows the centralised IFPC system to be partitioned in a particularly simple fashion, the allocation of airframe controlled variables such as $\dot{\gamma}$ to the engine subsystem reduces the level of transparency of the overall partitioned system. The approach is also inherently limited to the particular regions of the flight envelope where the assumptions about airframe variables being primarily influenced by engine effectors hold true. An alternative partitioning method also reported in [32] involved assigning the actuators as either airframe, engine or interface actuators. The controlled interface variable was selected to be thrust magnitude, THM and the interface actuator was MFF. This approach had the obvious limitation of attempting to control the thrust magnitude using only the main fuel flow, and the remaining engine variables using only the exit nozzle area and the inlet guide vane. This again could be seen as a serious limitation of the method as, for example, $T10$ also depends on MFF. What is required is a method that is sufficiently general that can be used regardless of the point in the flight envelope chosen for the nominal design, while retaining adequate control of all controlled outputs.

6.2 Partitioning the centralised controller

The particular implementation structure for the centralised $\mathcal{H}_\infty$ loop shaping controller as shown in Figure 4.3 is quite distinct from the standard one degree-of-freedom mixed sensitivity $\mathcal{H}_\infty$ controller structure, which normally can only be implemented as a single block $K(s)$ in the forward loop [35]. While $\mathcal{H}_\infty$ loop shaping controllers can also be implemented as a single block in the forward loop by setting $K(s) = kW_1(s)K_\infty(s)W_2$, there exists some extra flexibility since the $K_\infty(s)$ block and the loop weighting functions can be kept separate. It has been found in practice [44] that this implementation structure gives better time domain performance, in particular resulting in reduced overshoot in response to reference demands. This is because with this structure the references do not directly excite the dynamics of K_∞, which can often result in large overshoot (classical derivative kick), since K_∞ has been designed in the frequency domain for robustness objectives rather than time-domain performance. This structure, together with the fact that $\mathcal{H}_\infty$ loop shaping controllers can be implemented as exact plant observers with state feedback [66], is also very useful in making more transparent the functionality of the overall controller, i.e. it is not simply one 'black-box' state space block. When it comes to the process of partitioning the centralised IFPC system into separate engine and airframe subcontrollers, however, the multi-block structure of the $\mathcal{H}_\infty$ loop shaping design creates some additional complications which are not addressed explicitly in the controller partitioning procedure of [27]. In particular, notice that the $\mathcal{H}_\infty$ loop shaping implementation structure of Figure 4.3 corresponds to a static two degree-of-freedom design, and that a full dynamic two degree-of-freedom $\mathcal{H}_\infty$ loop shaping controller would have essentially the same structure - the constant prefilter $K_\infty(0)W_2$ being replaced by a dynamic compensator. Thus, the controller partitioning procedure described in this section can be regarded as an extension of the method of [27] to two degree-of-freedom control systems. The particular structure required for the partitioned IFPC system depends on the nature of the interactions between the airframe and engine subsystems, as well as on the various requirements arising from industrial and commercial constraints, e.g. the engine manufacturers need to run closed-loop tests prior to integration of the overall system. In practice, a hierarchical structure has been found to be most appropriate, with the airframe subcontroller generating commands both for the airframe control surfaces, and for those propulsion system variables which directly effect the airframe. A step-by-step procedure for partitioning centralised IFPC systems with a $\mathcal{H}_\infty$ loop shaping implementation structure can be summarised as follows:

Step 1.

Assign the plant's inputs and controlled outputs to be either engine or airframe variables. The

two subcontrollers cannot both have authority over the same control actuator, to allow for independent control implementation and subsystem validation. For the control inputs, this decision is based on control effectiveness (in terms of capability to directly control the controlled outputs) rather than the physical location of the actuator. For example, although the front and rear nozzle pairs are physically a part of the propulsion system, the angle of these nozzle pairs directly affects airframe quantities such as α and VT, and so it is natural to consider them as airframe actuators. Although other engine control inputs may also appear to directly affect airframe variables, these changes are in fact achieved via other variables which are denoted as interface variables. For example, a change in main fuel flow may be seen to produce a change in velocity. However, this velocity change is more directly due to the resulting change in thrust magnitude, which can thus be denoted as an interface variable, leaving main fuel flow to be considered as part of the engine subsystem. The controlled outputs are selected to be mutually exclusive, and their assignment is based on traditional definitions of airframe and engine variables. The interface variables represent propulsion system quantities that affect the airframe. These are generally selected to be propulsion system generated forces and moments. Proceeding in this manner, the plant's inputs and outputs may be grouped in the following way, where the subscripts have the obvious meanings:

$$U_a = [ETAD, ETASTK, FNOZ, RNOZ, SPLIT]$$

$$U_e = [MFF, ENOZA, IGV]$$

$$Z_a = [\alpha, VT, \dot{\gamma}]$$

$$Z_{ea} = [THM]$$

$$Z_e = [HPSM, LPSM, NLPC, T10]$$

Note that only an interface variable from the engine to the airframe is considered, i.e. it is assumed that airframe/engine interactions are in one direction only. Although this is not strictly the case for this particular system (e.g. use of the reaction control system $ETASTK$ for airframe control effects the engine operating point), the approach is valid due to the qualitatively different nature of the interactions - interactions from the airframe to the engine are primarily considered as unwanted disturbances, which we expect the IFPC system to be robust enough to reject, whereas interactions from the engine to the airframe represent primary control power which is to be exploited as fully as possible by the IFPC system.

Step 2.

Based on the above assumptions, an internal engine subcontroller, K_{ee} is extracted from the

centralised IFPC system. This is done by considering each part of the centralised controller ($k, W_1(s), K_\infty(s), W_2$, and $K_\infty(0)$) and partitioning each part as shown for k:

$$k = \begin{pmatrix} kaa & kae \\ kea & kee \end{pmatrix}$$

The engine subcontroller is then made up of the $*_{ee}$ section of each. The order of the internal engine subcontroller was then be decreased by reducing the $K_{\infty ee}$ block from 43^{rd} order to 15^{th} order, resulting in an 18^{th} order internal engine subcontroller.

Step 3.

Analyse the closed-loop frequency responses from the airframe commands Z_{a_c} to the interface variable Z_{ea} with the centralised controller in order to identify the control requirements for the interface subcontroller for tracking the interface variable commands generated by the partitioned airframe subcontroller. One way to determine these requirements is to study the closed-loop frequency response from all the airframe commands to each individual element of the interface variables with the centralised controller. We would like the demand for response in interface variables required to track airframe commands to roll-off prior to loss in the capability of the engine subcontroller to track the corresponding command. Other limits on the minimum required tracking bandwidth for the interface variables imposed by subsystem specific performance requirements include disturbance rejection and performance robustness to low frequency model variations. Furthermore, control actuation limits and requirements of robustness to high frequency modeling errors will impose limits on the maximum achievable tracking bandwidth for the engine subsystem. It is important to consider all of the requirements in generating specifications for interface tracking.

Step 4.

Design an interface subcontroller K_{ea} to meet these control requirements. Any design method may be used to design the interface subcontroller - in this study the $\mathcal{H}_\infty$ loop shaping method was used to design a 13 state controller, with weighting functions:

$$k = diag\,(0.049,\ 0.049,\ 0.049)\,, \quad W_1 = \frac{s+1}{s} \times I_{3\times 3}$$

W_1 was chosen to provide zero steady state tracking error via integral action, and also to decrease the roll-off rate around cross-over, while k was chosen to meet actuator limitation requirements. These result in a value of $\gamma_{opt} = 1.53$, indicating good robustness properties. Following model order reduction, the interface subcontroller is 13^{th} order. Note that the interface subcontroller design problem is non-square, i.e. three control inputs are available to control

the single interface variable THM. This issue of controls' redundancy will be returned to in § 7.5.1. The full engine subcontroller KE is thus made up of the internal engine subcontroller K_{ee} and the interface subcontroller K_{ea} connected in parallel.

Step 5.

With the engine subsystem closed using the centralised controller, as shown in Figure 6.1, obtain a state-space representation of the airframe subcontroller block, K_a as a reduced order (25 states) approximation of the E to $[U_a\ Z_{ea}]$ transfer function matrix. The airframe subcontroller KA is then made up of the K_a block along with the blocks $K_{\infty a}W_{2a}$ and $K_{\infty a}(0)W_{2a}$ which are simply extracted from the original centralised system. Note the K_a block of the airframe subcontroller generates commands for the airframe actuators while also generating reference demands for the interface subcontroller.

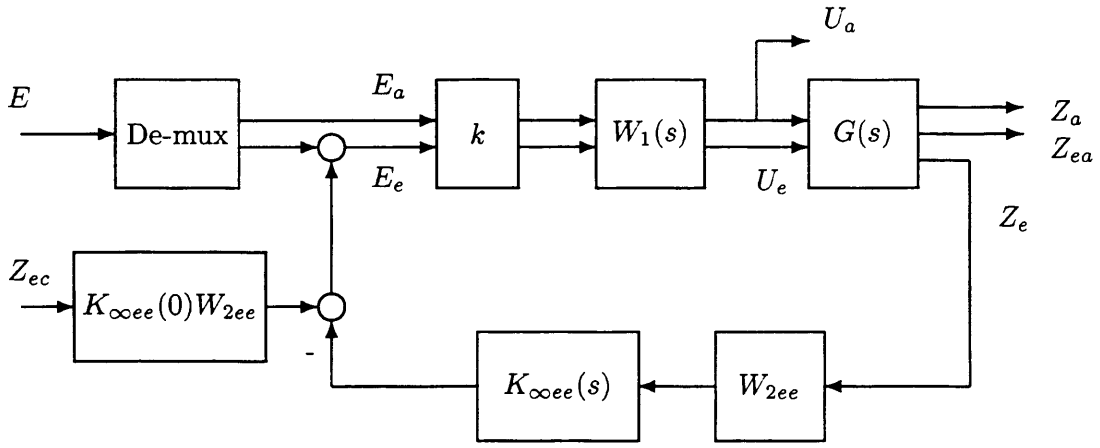

Figure 6.1: Configuration for computing K_a block of airframe subcontroller.

Step 6

Design a lead filter to compensate for the limited z_{ea_c} tracking bandwidth of the engine subsystem. K_a as described above generates the desired response in the interface variables to airframe controlled variable commands such that the integrated system achieves the specified tracking and decoupling response. If these were used directly as commands for the interface variables then the actual z_{ea} response with the partitioned subcontrollers would lag the desired response due to the limited tracking bandwidth of the engine subsystem, thus resulting in deterioration in integrated system performance. In general its choice is based on practical considerations between the amount of lead compensation in K_{lead}, and the z_{ea_c} tracking bandwidth of the engine subsystem. High lead compensation is undesirable as it can result in saturation of the engine actuators due to command magnification, whereas low lead compensation will require large z_{ea_c} tracking bandwidth. Since the interface subcontroller portion of the engine controller

provides decoupled tracking of z_{ea_c}, K_{lead} can simply be of the form

$$K_{lead}(s) = diag\left[\frac{s+a_i}{a_i}\frac{b_i}{s+b_i}\right],\ a_i < b_i$$

with a_i and b_i chosen based on the amount of lead desired. These were selected as $a_i = 1$, $b_i = 15$.

The final partitioned system is shown in Figure 6.2. The original centralised IFPC system has been partitioned into separate engine and airframe subcontrollers connected via a hierarchical structure. As may be expected, the overall order of the partitioned IFPC system is greater than for the centralised design. Each individual subcontroller is, however, of low order, leading to easier implementation and greater transparency. Finally, note that the partitioning procedure described above is not dependent on the static nature of the $\mathcal{H}_\infty$ loop shaping controller pre-filter block, i.e. the procedure can be applied without modification to control systems with a general two degree-of-freedom structure.

6.3 Nonlinear simulation results for the partitioned system

Inceptor demands corresponding to manoeuvres flown in piloted simulation trials of the centralised IFPC system, [7], were used in all cases. Again, as with the centralised system, it was noted that the partitioned system frequently operates with the engine actuators on their position limits. In particular, this saturation is due to the demands made by the interface subcontroller. It was therefore decided to implement the weighting matrix W_{1ea} in its self-conditioned form, as discussed in § 5.6. Figures 6.3 and 6.4 show the aircraft responses to a demand on $\dot{\gamma}$, for both the centralised and partitioned IFPC systems, each implemented with an anti-windup scheme. It can be seen that the γ steady-state response for the partitioned system is slightly greater than that for the centralised system. This is because the partitioned system tracked the positive $\dot{\gamma}$ demand slightly more quickly that the centralised system, and it tracked the negative $\dot{\gamma}$ demand slightly more slowly. It can also be seen that there is no increase in the α or VT coupling. The main difference between the two systems is the $LPSM$ steady-state value.

Figures 6.5 and 6.6 show similar information for a demand on VT. The results for the partitioned system resulted in reduced $\dot{\gamma}$, and so γ, and α coupling. Again, the most significant difference can be seen in the steady-state values of the engine variables.

These show that the responses of the airframe variables for the partitioned IFPC system in nonlinear simulation closely match those achieved with the centralised design.

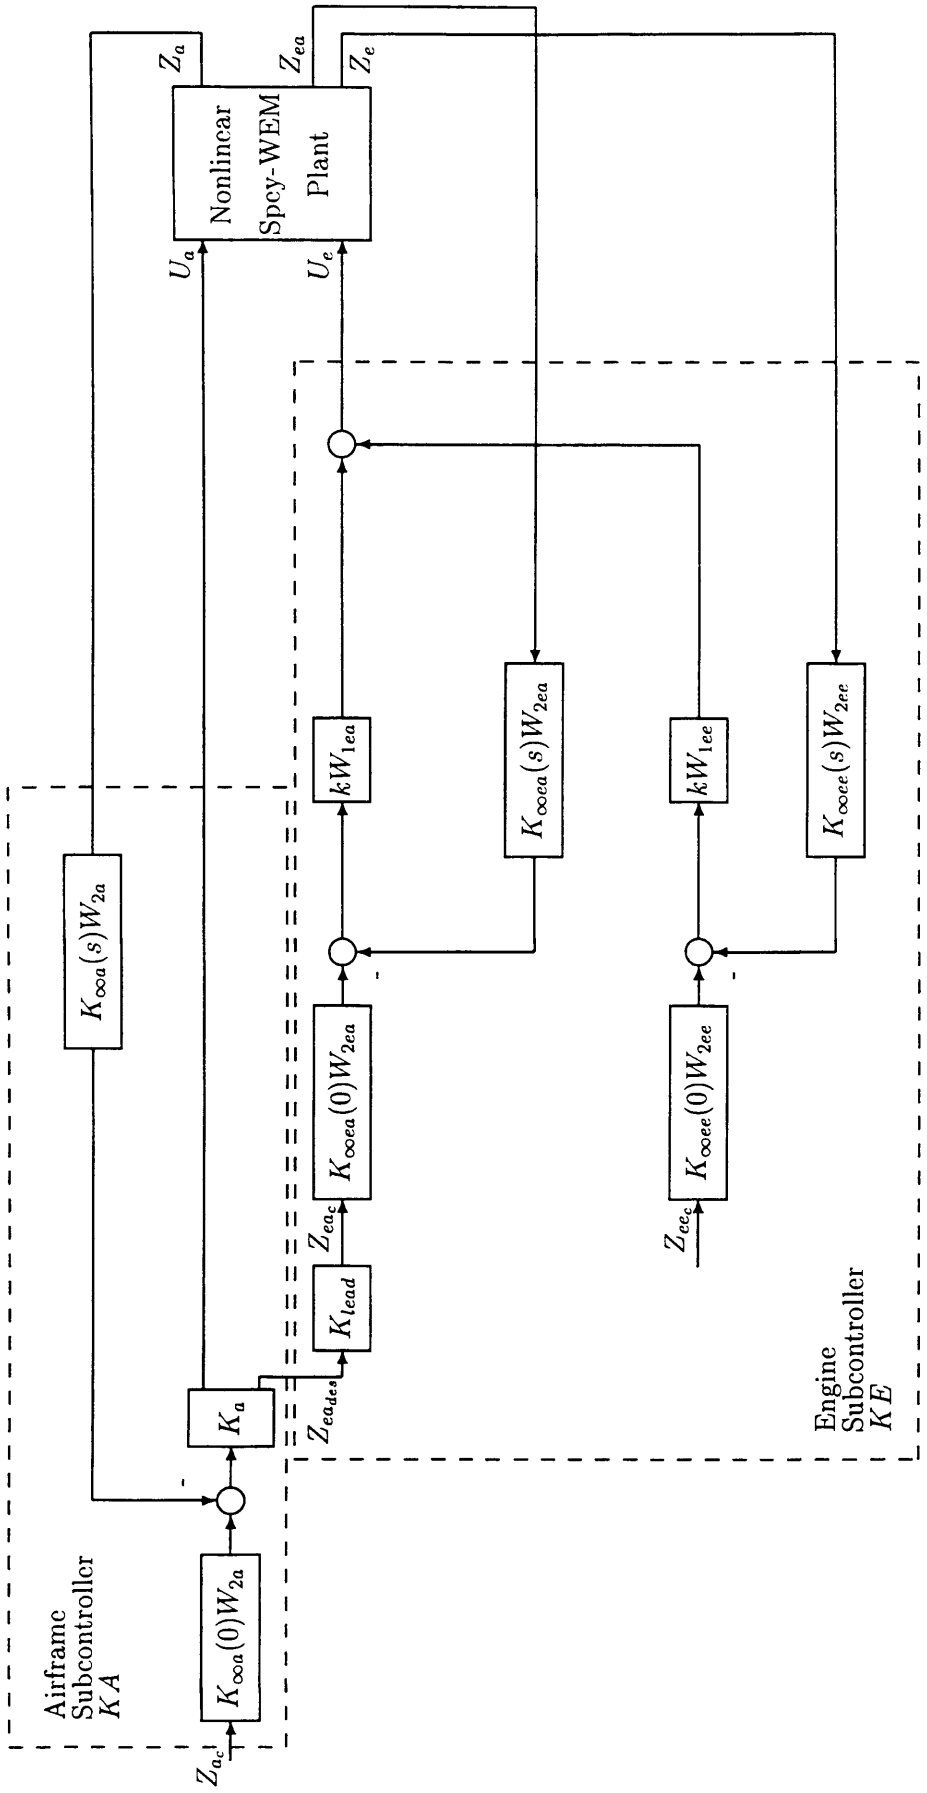

Figure 6.2: Partitioned IFPC system.

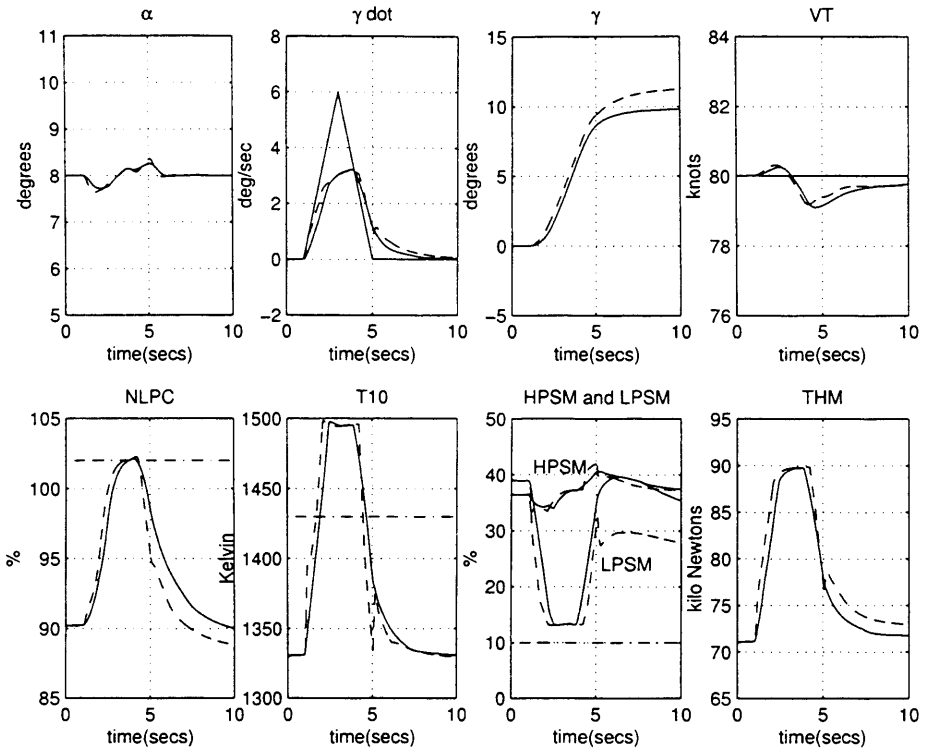

Figure 6.3: Airframe and engine responses to a $\dot{\gamma}$ demand for the centralised system (—) and the partitioned system (- -) at 80 knots.

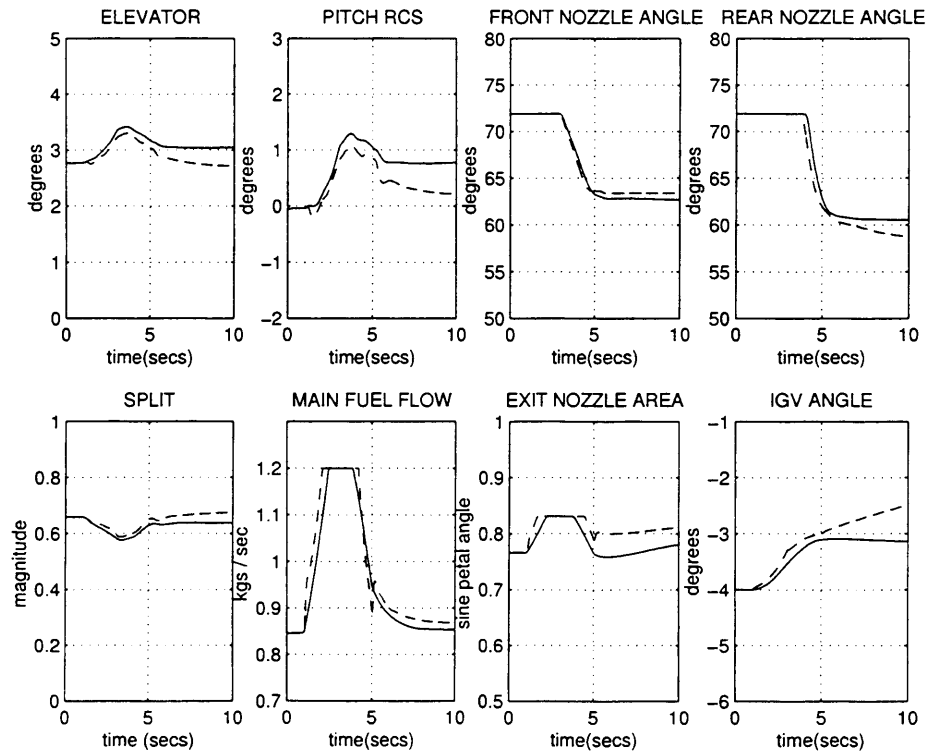

Figure 6.4: Actuator responses to a $\dot{\gamma}$ demand for the centralised system (—) and the partitioned system (- -) at 80 knots.

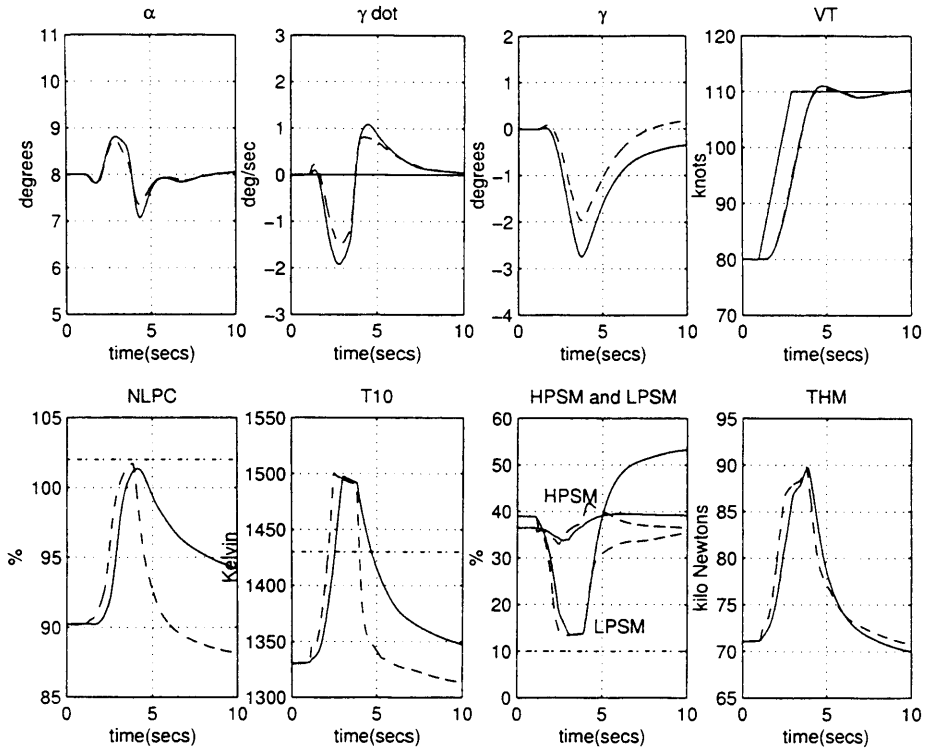

Figure 6.5: Airframe and engine responses to a VT demand for the centralised system (—) and the partitioned system (- -) at 80 knots.

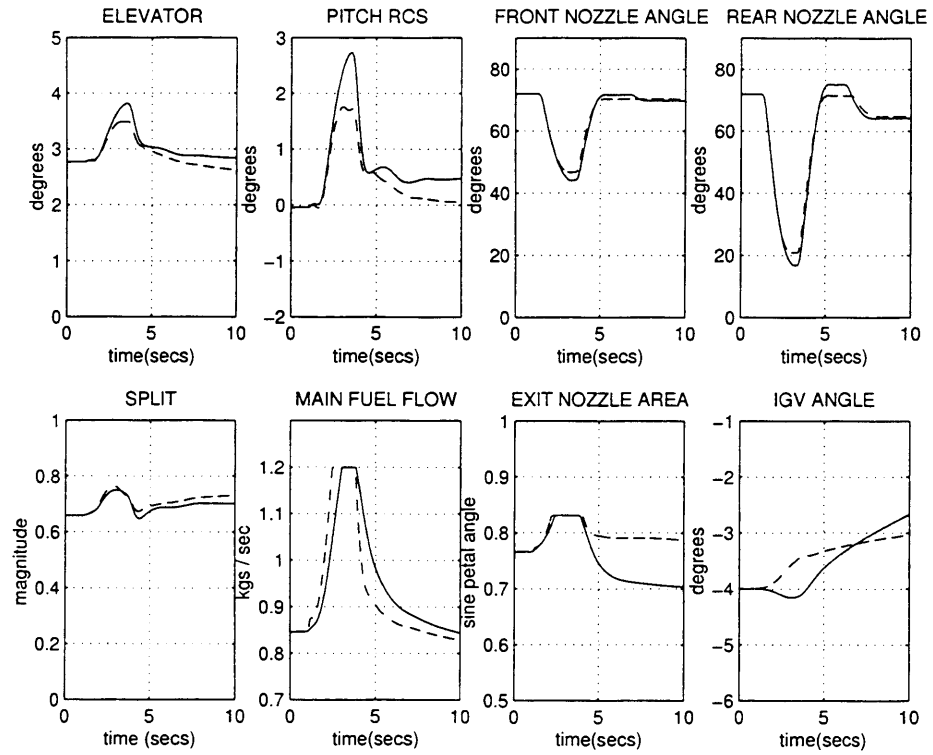

Figure 6.6: Actuator responses to a VT demand for the centralised system (—) and the partitioned system (- -) at 80 knots.

6.4 Summary

Although the design of a centralised controller that accounted for the interactions between subsystems was seen to be desirable, it would not be practical to implement. The method used to partition the centralised controller into airframe and engine subcontrollers has been introduced in this chapter. This can be seen as an extension of that of [27], as the implementation structure of the $\mathcal{H}_\infty$ loop shaping controller may represent that of a two-degree of freedom controller. Partitioning the centralised controller will enable the engine manufacturer to carry out the necessary tests to ensure safe operation of the engine. The partitioned system nonlinear simulation responses compare favourably to the original centralised system in nonlinear simulations

Chapter 7

Robustness analysis

The application of linear controller design techniques to highly nonlinear systems such as vertical/short take-off and landing (V/STOL) aircraft generally results in controllers which produce acceptable flying qualities only at, and in a limited region around, a particular point in the flight envelope. The recently developed linear parameter varying (LPV) control techniques [67, 46, 24] offer the potential to directly design controllers with guaranteed performance and robustness properties over the full flight envelope. In order to use these techniques, however, a quasi-LPV model of the plant dynamics must first be developed, and computationally intensive linear matrix inequality (LMI) based optimisation problems must be solved to calculate the controller. For many aerospace systems, and in particular for the V/STOL aircraft configuration considered in this work, both of these requirements are problematic - the nature of the aircraft model (interacting engine and airframe sub-models, developed independently and subsequently integrated) makes accurate LPV models very difficult to construct, and the high order of the overall system makes LMI-based optimisation computationally expensive.

A more practical solution to the problem of constructing full-envelope multivariable IFPC systems for such an aircraft model could be to use nonlinear control logic to switch (or blend) between a set of linear controllers that have been designed at particular points in the operating envelope, as in [44] for example. An essential prerequisite for constructing efficient (i.e. minimal) schemes of this kind, however, is to be able to analyse over which portions of the flight envelope each linear controller retains adequate stability and performance properties.

While μ-analysis has become well established as a powerful technique for analysing the robustness of aerospace systems [22], there is still a surprising lack of examples in the literature of its application to complex industrial control problems [40, 41, 30]. In particular, analysis of multivariable flight control laws of the type considered here, where the effect of variations in aircraft parameters and dynamics over the flight envelope is modelled as real state-space uncertainty, causes problems for the efficient computation of tight bounds on μ. This compu-

tational difficulty is addressed here by comparing two different methods for generating tight bounds on real μ by introducing small levels of complex uncertainty in the uncertainty description. Alternative approaches to computing tight bounds on real μ are described in [23, 42]. The robustness analysis problem considered here also poses a challenge to the more recent ν-gap techniques, which have heretofore largely been applied to relatively low-order systems [79]. Despite the modelling and computational difficulties involved in their application to such a complex problem, both techniques are seen to provide much useful information about the robustness properties of the integrated flight and propulsion control system. Furthermore, the results of the analysis are seen to correspond closely with the evaluation of the controller from pilot comments reported in [7].

The lay-out of this chapter is as follows. § 7.1 will introduce the use of μ to analyse the stability of the centralised system due to parameter variations. § 7.2 will consider the stability of the centralised system due to uncertain dynamics, due to flying at points in the flight envelope other than the nominal design point. § 7.3 presents the results of the μ-analysis of the centralised system, and § 7.4 presents a ν-gap analysis of the centralised system. § 7.5 gives the robustness analysis results of the partitioned system. Finally, § 7.6 presents a summary of this chapter.

7.1 Robust stability analysis of centralised system due to parameter variations

In this section, 'uncertainties' in the nominal plant model arising from two distinct sources are considered: variations in the values of the aircraft mass and inertia, and variations in the airframe/engine dynamics over different regions of the V/STOL flight envelope. In order to use μ to analyse the robustness of the IFPC system, it is necessary to generate linear fractional transformation (LFT) based parametric uncertainty descriptions for the model and its associated uncertainty. The approach used for generating such descriptions will necessarily depend on the nature (and size) of the plant model under consideration. For a general nonlinear plant of the form

$$\begin{aligned} \dot{x}(t) &= F(x(t), u(t), p) \\ y(t) &= G(x(t), u(t), p) \end{aligned} \tag{7.1}$$

where x, u, y are the state-, input- and output-vectors respectively, and p is a vector of uncertain parameters, symbolic linearisation methods [75] can be used to generate linear, rationally parameter-dependent representations about some equilibrium point of the form

$$\begin{aligned} \delta\dot{x} &= A(p)\delta x + B(p)\delta u \\ \delta y &= C(p)\delta x + D(p)\delta u. \end{aligned} \tag{7.2}$$

The model (7.2) is non-conservative in that the entries of the state-space matrices (which are rational functions of the uncertain parameters p) preserve an exact description of joint parametric dependencies in the original nonlinear model. Transformation of the model (7.2) into an LFT form for μ-analysis is then relatively straightforward. For plants of the type considered here, however, several problems arise with the above approach. Firstly, due to the high order of the plant, the computational burden associated with the required symbolic manipulations (even for relatively few parameters in p) can easily become prohibitive. Secondly, precise models relating the effect of the parameters p on the nonlinear system, as in (7.1), are not readily available. Thirdly, models of the form (7.2) are only valid in a neighbourhood around a particular linearisation point, and so are not generally appropriate for investigating robustness over large regions of the operating envelope. An alternative approach to uncertainty modelling which avoids these problems, at the expense of a certain amount of conservatism, is described next.

Given a nonlinear software model (in any form) of the plant, repeatedly perform numerical linearisations over several points in the operating envelope, and/or over all combinations of the extreme points of the uncertain parameters. For an uncertain parameter vector p of size n, with each parameter lying between some minimum and maximum value, a set of 2^n linear models are thus generated. These models form a so-called multi-model state description

$$\begin{aligned} \delta\dot{x}(t) &= A_i\delta x(t) + B_i\delta u(t) \\ \delta y(t) &= C_i\delta x(t) + D_i\delta u(t). \end{aligned} \tag{7.3}$$

For each varying element of each state-space matrix it is now possible calculate its minimum (e.g. a_{ij}^{min}), maximum (e.g. a_{ij}^{max}), and nominal (e.g. $(a_{ij}^{max}+a_{ij}^{min})/2$) values. It is thus possible to replace the multi-model system (7.3) by an affine parameter dependent representation of the form

$$\begin{aligned} \delta\dot{x}(t) &= A_{p_A}\delta x(t) + B_{p_B}\delta u(t) \\ \delta y(t) &= C_{p_C}\delta x(t) + D_{p_D}\delta u(t). \end{aligned} \tag{7.4}$$

where the state matrices are in the form

$$\begin{bmatrix} A_{p_A} & B_{p_B} \\ C_{p_C} & D_{p_D} \end{bmatrix} = \begin{bmatrix} A_{p0} & B_{p0} \\ C_{p0} & D_{p0} \end{bmatrix} + \sum_{i=1}^{n_A} \Delta_i \begin{bmatrix} A_{pi} & 0 \\ 0 & 0 \end{bmatrix}$$

$$+ \sum_{i=n_A+1}^{n_B} \Delta_i \begin{bmatrix} 0 & B_{pi} \\ 0 & 0 \end{bmatrix} + \sum_{i=n_B+1}^{n_C} \Delta_i \begin{bmatrix} 0 & 0 \\ C_{pi} & 0 \end{bmatrix} + \sum_{i=n_C+1}^{n_D} \Delta_i \begin{bmatrix} 0 & 0 \\ 0 & D_{pi} \end{bmatrix}$$

Thus, for example, for each varying entry in the matrices A_i, there results a Δ_i and an A_{pi} in

the above expression, where Δ_i is an uncertain real scalar parameter which varies between 1 and -1, and A_{pi} is equal to $((a_{ij}^{max} - a_{ij}^{min})/2)$. Each of the matrices associated with each Δ_i has rank one and can be factored using the singular value decomposition into row and column vectors:

$$\begin{bmatrix} A_{pi} & 0 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} E_i \\ F_i \end{bmatrix} \begin{bmatrix} G_i & H_i \end{bmatrix}$$

Now proceeding according to the method of [52], if the linear system P is defined with extra inputs and outputs via the equations

$$\begin{bmatrix} \dot{x} \\ y \\ z_1 \\ \vdots \\ \vdots \\ z_n \end{bmatrix} = \begin{bmatrix} A_0 & B_0 & E_1 & \cdots & E_n \\ C_0 & D_0 & F_1 & \cdots & F_n \\ G_1 & H_1 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ G_n & H_n & 0 & \cdots & 0 \end{bmatrix} \begin{bmatrix} x \\ u \\ w_1 \\ \vdots \\ \vdots \\ w_n \end{bmatrix}$$

where $n = n_D$, the closed-loop interconnection structure shown in Figure 7.1 can be formed. It is then trivial to convert this system into the standard form for μ-analysis shown in Figure 7.2.

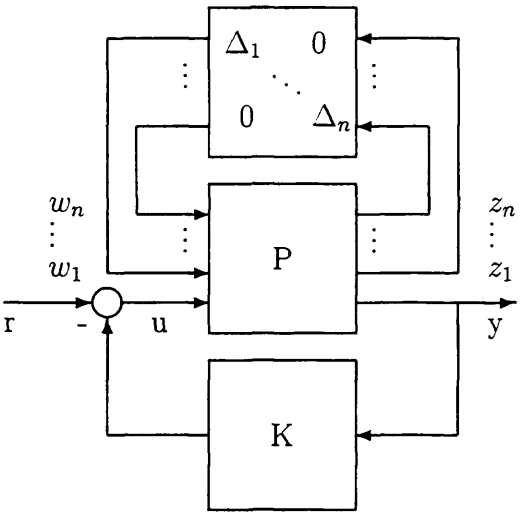

Figure 7.1: Interconnection structure of uncertain closed-loop system.

The main points to be made about this approach to uncertainty modelling are the following. Assuming that efficient trimming and linearisation routines are available for the nonlinear model, generation of the multi-model uncertainty description (7.3) is fast, easy and can be automated. Subsequent transformation of this description to the standard form for μ-analysis is then also more or less automatic. The main problem with this approach is that the Δ parameters in the final model are artificially introduced, and thereby reflect the uncertainty due to the actual uncertain parameters p only implicitly. Thus the size of the final Δ matrix in Figure 7.2 depends on the size of the system state-space matrices, as well as on the number

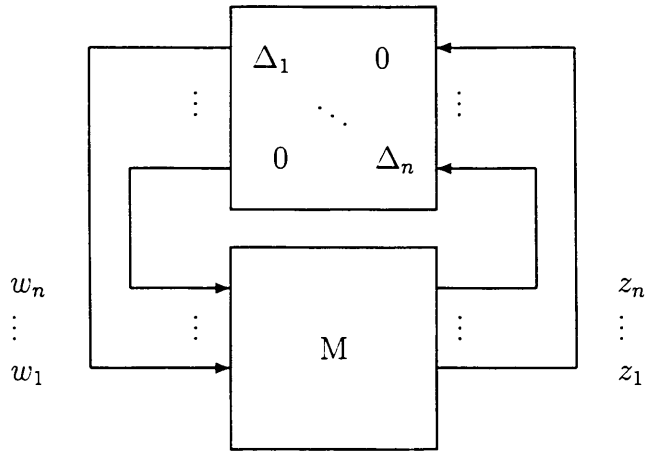

Figure 7.2: Standard block diagram for μ-analysis.

of parameters in the vector p. Although the parametric linear model (7.4) certainly covers all possible linearisations arising from the uncertain nonlinear model, it will be conservative, since it ignores possible joint parameter dependencies in the model. Thus only positive analysis results can be proposed, i.e. it is possible to prove robust stability for a given level of uncertainty, but if the results indicate the existence of a destabilising uncertainty it is not possible to say anything definite. This is because it cannot be sure if the destabilising uncertainty is actually present in the system or has been introduced by the approach used to generate the LFT-based uncertainty model. While this is certainly a drawback of the proposed approach, the results to be presented indicate that the amount of conservatism is in practice quite small, and good agreement with the results of the ν-gap analysis and piloted simulation trials is noted. Finally, one advantage of the proposed approach is that the resulting uncertainty models do not contain any repeated parameters. This may be in contrast to exact uncertainty models where repeated parameters may result in extra conservatism in the generation of upper bounds on μ.

Another complication arising from the proposed approach to uncertainty modelling is that the resulting Δ matrix is purely real, thus making it difficult to calculate lower bounds on μ. This issue is treated in detail in § 7.3. In this study, the above method is used to generate an uncertainty description for the Spey-WEM model, in order to assess the robustness of its IFPC system to variations in the aircraft's mass, m ($\pm 20\%$), and inertia, I_{yy} ($\pm 5\%$), i.e. $p = [m, I_{yy}]$. These were selected based on a consideration of the nonlinear system, whereby the effects of varying several parameters were noted.

7.2 Robust stability analysis of centralised system at different velocities using μ

The approach to evaluating the robustness of the IFPC system to variations in the aircraft's dynamics over the flight envelope requires a slight modification to the above procedure. In this analysis, we seek to represent differences between the linearised model of the aircraft at the 80 knot design point and linearised aircraft models at various other points in the envelope, as structured uncertainty about the 80 knot model. In standard μ-analysis the uncertainty matrix Δ is required to be normalised so that each Δ_i can assume any value in the interval [-1,+1]. For the particular problem under consideration this is not appropriate, however, since it is sought to represent changes in the plant state-space matrices from 80 knots to 120 knots (for example), and not from 80 Knots $\pm$ 40 knots. In effect it is required that a 'one-sided' uncertainty is considered. Hence the following modifications are made to the standard configuration in order to prevent μ considering an uncertainty set that is too large:

1. Construct the standard diagonal perturbation formulation (DPF) with an additive perturbation matrix where $\Delta_i = 0.5 \ \forall i$.

2. Scale the outputs from Δ to M by 0.5. This has the effect of constraining $\Delta_i \in [-0.5, 0.5]$.

M is now in a suitable form so that μ can be determined with a 'correct' uncertainty set. An alternative procedure would be to use repeated real parameters to ensure that $\Delta_i \in [0, 1]$, as shown in Figure 7.3. This is achieved by ensuring that the output from one Δ_i forms the input to a similar (repeated) Δ_i. The overall perturbation gain will then correspond to a Δ where every $\Delta_i \in [0, 1]$. One drawback to the use of repeated real uncertainty is that the corresponding μ-analysis necessarily involves a larger DPF matrix $\tilde{M}(s)$ and therefore the generation of extra constraints by the μ software. For example, if the first (non-repeated, scaled only real parameter) solution outlined above contains q uncertain parameters, then $2q$ constraints must be generated by the μ-analysis optimisation software. However, an equivalent repeated real parameter μ-analysis will require $4q$ constraints, and thus results in bounds on μ which are more conservative than those achieved with the former method.

The former method of determining the uncertainty associated with the aircraft linearisations at different points in the flight envelope may be verified by comparing the closed-loop systems of the 80 knot plant 'hit' with the appropriate Δ with, for example, the actual 50 knot plant. This is shown in Figures 7.4 for the 80 knot plant with a zero uncertainty matrix and 7.5 for each linearised point in the flight envelope.

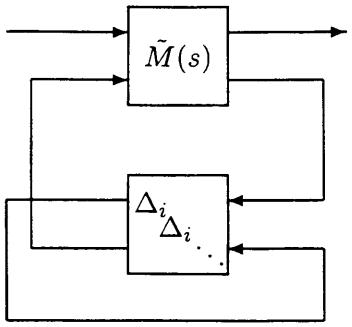

Figure 7.3: Repeated real approach for a one-sided analysis problem.

It can be seen that there is no difference in the singular values for the two systems, showing this is an effective method of representing the uncertainties.

7.3 μ-analysis results

The results presented from this point are those found using a slightly redesigned controller. The nominal 80 knot linear plant was re-trimmed in order to achieve consistent values of α. The performance characteristics of the original controller were maintained. Using the method described in § 7.1, the final uncertainty description for variations in p gave a real Δ matrix of size 76, and so calculating bounds on μ is settled for, using standard software routines [4]. An upper bound for μ, calculated using the mixed μ algorithm of [81] is shown in Figure 7.6. From the graph it is clear that the IFPC system is guaranteed to be stable for all uncertainties associated with the parameter p.

To understand exactly how much bigger these uncertainties can be before instability could occur, however, a lower bound on μ must be computed. Now, it is known that in the case of purely real uncertainties, μ is not necessarily a continuous function [56], i.e. it is possible for the robustness margin for real parameter uncertainty to change abruptly for infinitesimal changes in the problem data. These discontinuities can cause problems in the convergence of the lower bound mixed-μ algorithm of [81]. In fact, for this particular problem, this algorithm does not converge. Furthermore, the high order of the system and the resulting size of the Δ matrix also rules out the use of exponential time algorithms for calculating lower bounds on real μ, such as the one described in [19].

One way of getting around these difficulties is to introduce small levels of complex uncertainty into the uncertainty model, in order to improve the convergence properties of the lower bound mixed-μ algorithm of [81]. Two alternative ways of introducing such uncertainty are considered here. The first approach, suggested in [4], is to add small amounts of complex uncertainty onto each uncertain real parameter. The idea is to introduce just enough complex uncertainty

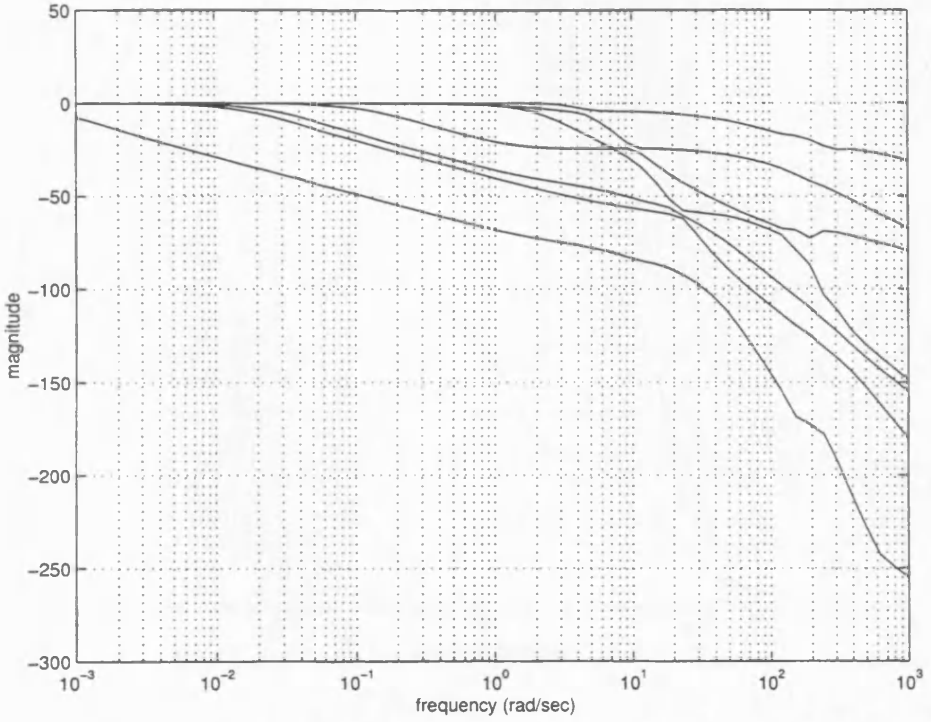

Figure 7.4: 80 knot plant (—) and 80 knot plant 'hit' with zero Δ (- -).

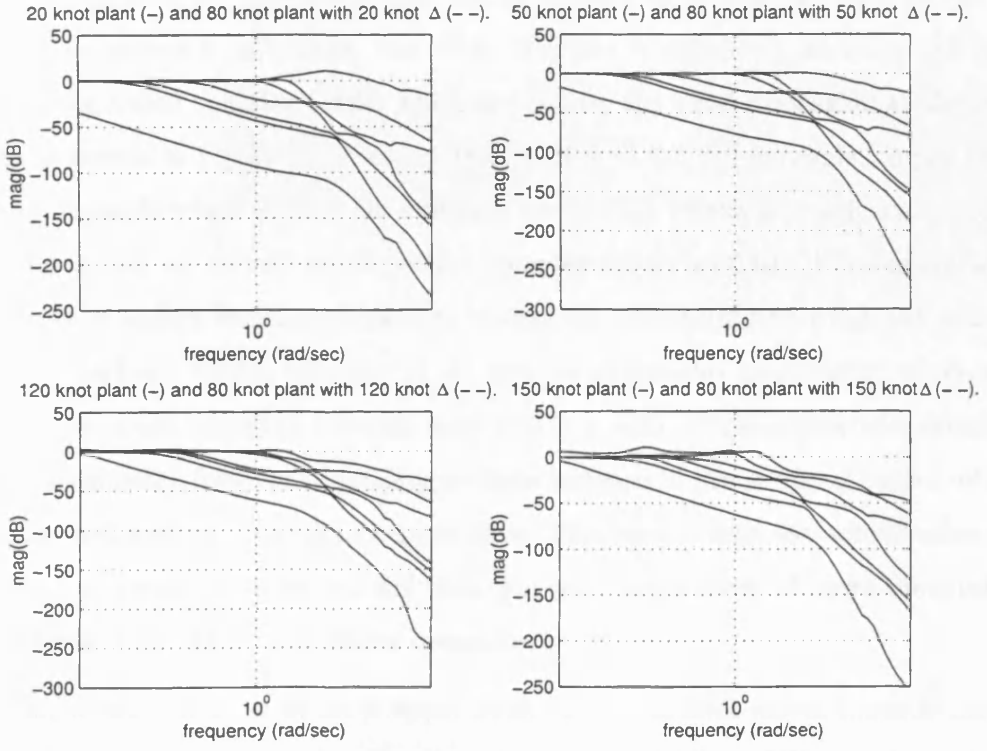

Figure 7.5: Nominal plants (—), and 80 knot plant with appropriate Δ (- -).

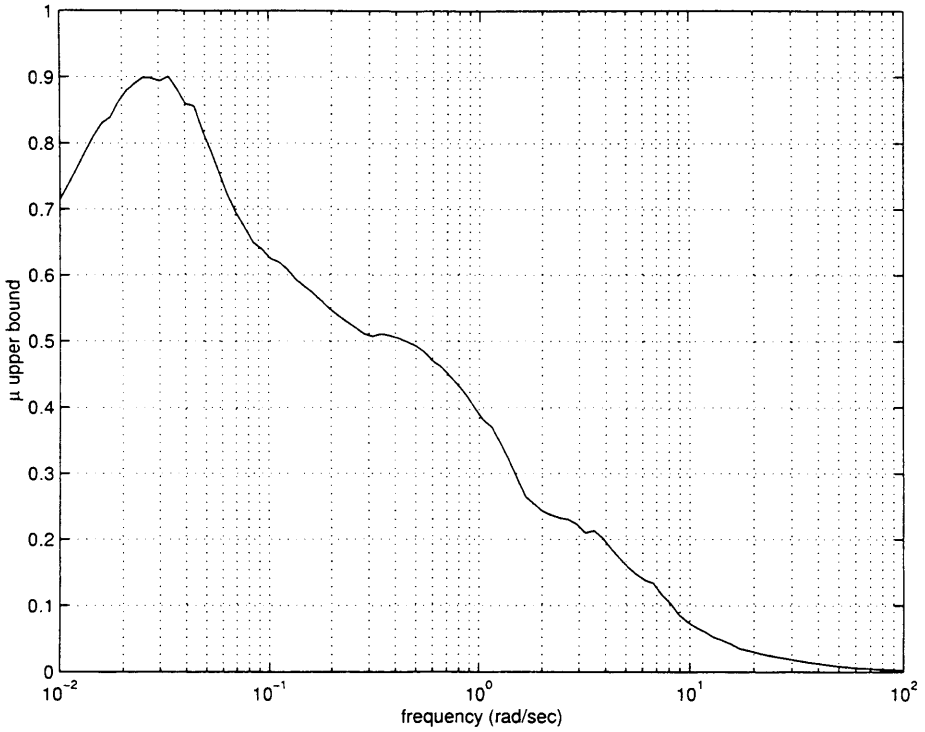

Figure 7.6: μ upper bound for the IFPC system at 80 knots for 2 uncertain parameters (m and I_{yy}).

to get the mixed μ lower bound software to converge, without significantly increasing the fundamental uncertainty set considered (this can be checked by noting the effect on the upper bounds). The procedure is as follows: the block structure is expanded (to twice the original size) by including scaled complex blocks which are exactly the same dimension as the original real blocks, as shown in Figure 7.7. Scaling factors of $\chi = 0.1$, for example, imply that the input/output channels which provide the complex uncertainty effects are each scaled down by a factor of 10, giving an overall scaling of the complex blocks of 0.01. Figures 7.8 and 7.9 show the effects of adding increasing amounts of complex uncertainty on the lower and upper bounds for μ. The lower bound improves as the amount of complex uncertainty increases, but it also causes the upper bound to increase until it is > 1 with 16% complex uncertainty. For 9% complex uncertainty however, there is very little increase in the μ upper bound while the lower bound is well behaved, and in fact quite tight. This reveals that the actual value of μ is quite close to 1 at certain frequencies, and thus that only small levels of extra uncertainty in p could be tolerated by the system before instability could occur.

Figure 7.10 shows the variation in the μ upper bound as the aircraft moves towards the hover region of the flight envelope. From the plots it can be seen that while stability can be guaranteed down to 50 knots, at 20 knots it cannot. Figure 7.11 shows the variation in the upper bound as

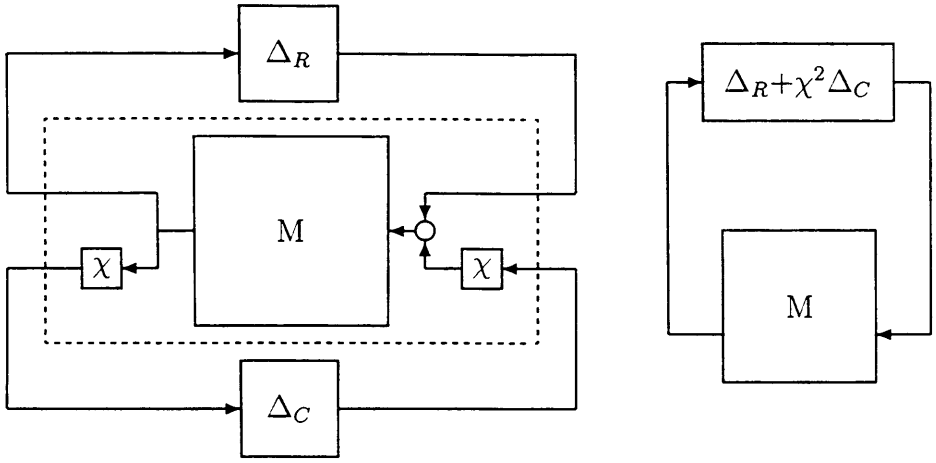

Figure 7.7: Replacing real uncertainty with real + complex uncertainty.

the aircraft moves towards the fully wingborne region of the flight envelope for the 120 knots and 150 knots points in the flight envelope. In this case it can be seen that at neither point can robust stability be guaranteed. To prove *instability* at these points in the flight envelope lower bounds for μ must be calculated. To do this we proceed again by adding small amounts of complex uncertainty. Figures 7.12 to 7.13, and 7.14 to 7.15 show the resulting μ lower and upper bounds for uncertainty corresponding to variations in the aircraft dynamics between 80 knots and 20 knots, and between 80 knots and 120 knots. From the figures it can be seen that μ lower bounds of greater than 1 are achieved for each case, for even 1% added complex uncertainty, and thus it can be concluded that stability is not preserved at these points in the flight envelope. It is recalled at this point, however, that the 'black-box' procedure for modelling uncertainty neglects inter-dependencies in the elements of the state-space matrices, and as a result the robustness results are likely to be quite conservative.

An alternative approach to introducing complex uncertainty into the system is to include uncertainty models for the dynamics of the aircraft actuators. Although the dynamics of aircraft actuators can often be modelled to a high accuracy via wind tunnel testing, some 'uncertainty' will always be present in linear models of such systems, due to (a) the neglect of nonlinear dynamics such as rate limits and deadzone/hysteresis effects, and (b) the neglect of high-frequency dynamics. For the aircraft model used in this study the actuators with the most complex behaviour are the thrust vectoring nozzles, which exhibit significant nonlinear dynamics of the type mentioned above. Complex uncertainty models were therefore placed on these actuators, corresponding to 20% uncertainty in the position of the nozzles at low frequencies, increasing to 100% at approximately 100 rads^{-1}. Similar models corresponding to 1% uncertainty at low frequency, rising to 100% at 200 rads^{-1} were placed on all the other aircraft actuators. Figure 7.16 compares the μ bounds for the uncertainty in mass and inertia,

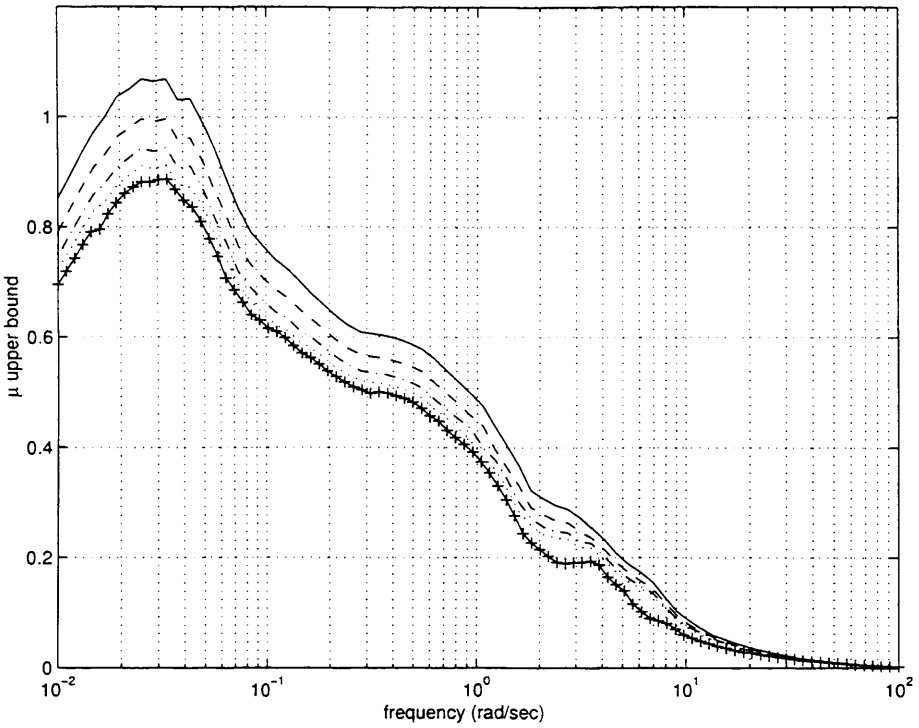

Figure 7.8: μ upper bounds for the centralised system at 80 knots for 2 uncertain parameters with 0% (+), 1% (··), 4% (-·), 9% (- -) and 16% (—) added complex uncertainty.

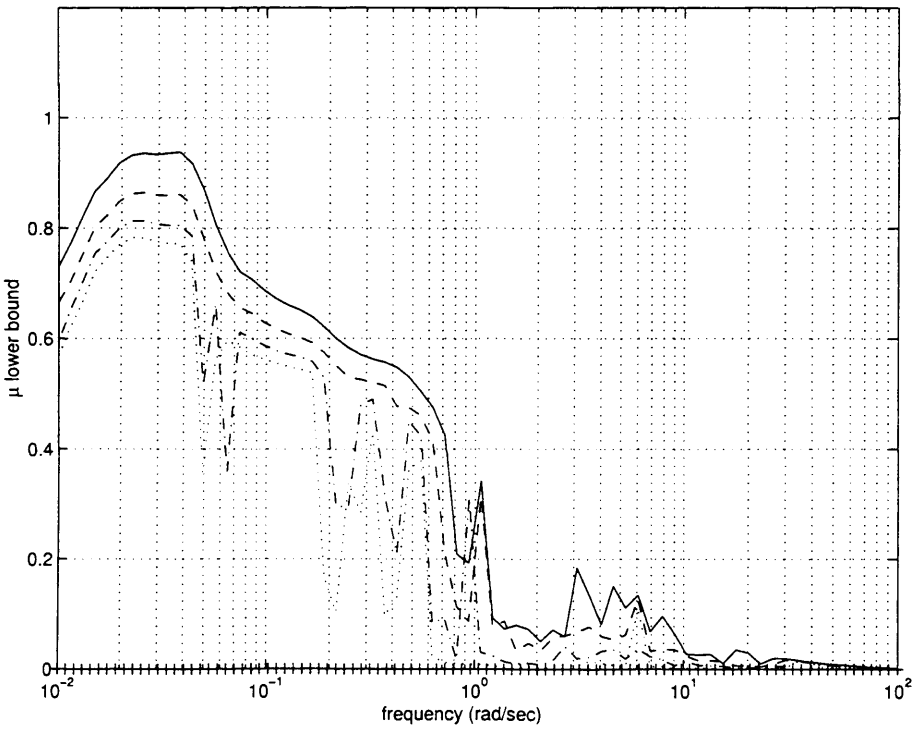

Figure 7.9: μ lower bounds for the centralised system at 80 knots for 2 uncertain parameters with 0% (+), 1% (··), 4% (-·), 9% (- -) and 16% (—) added complex uncertainty.

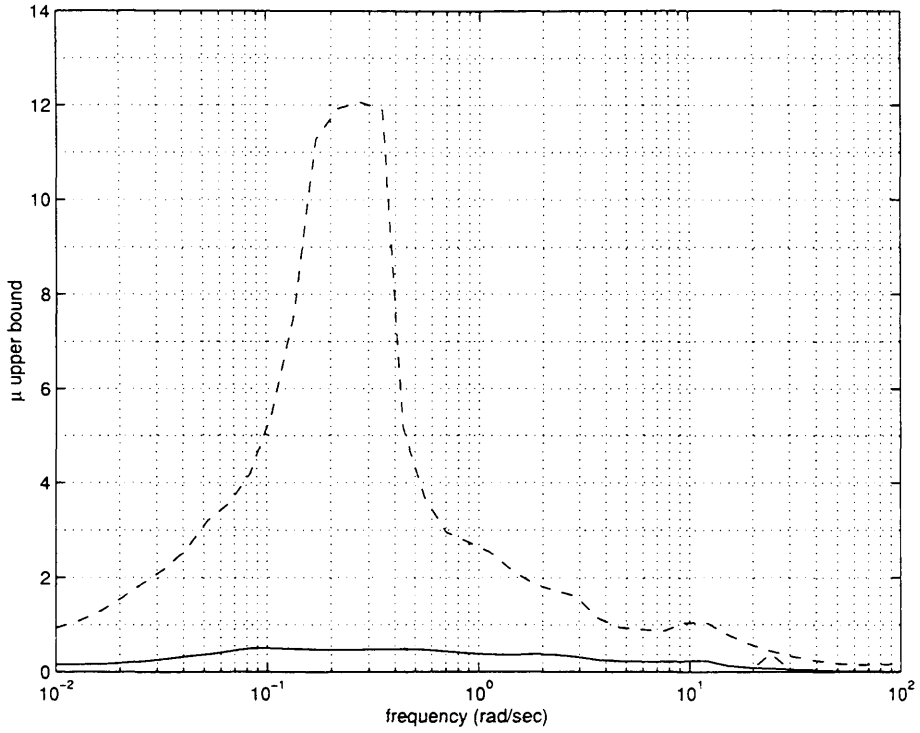

Figure 7.10: μ upper bounds for 80-50 knots (—) and 80-20 knots (- -) dynamic variations.

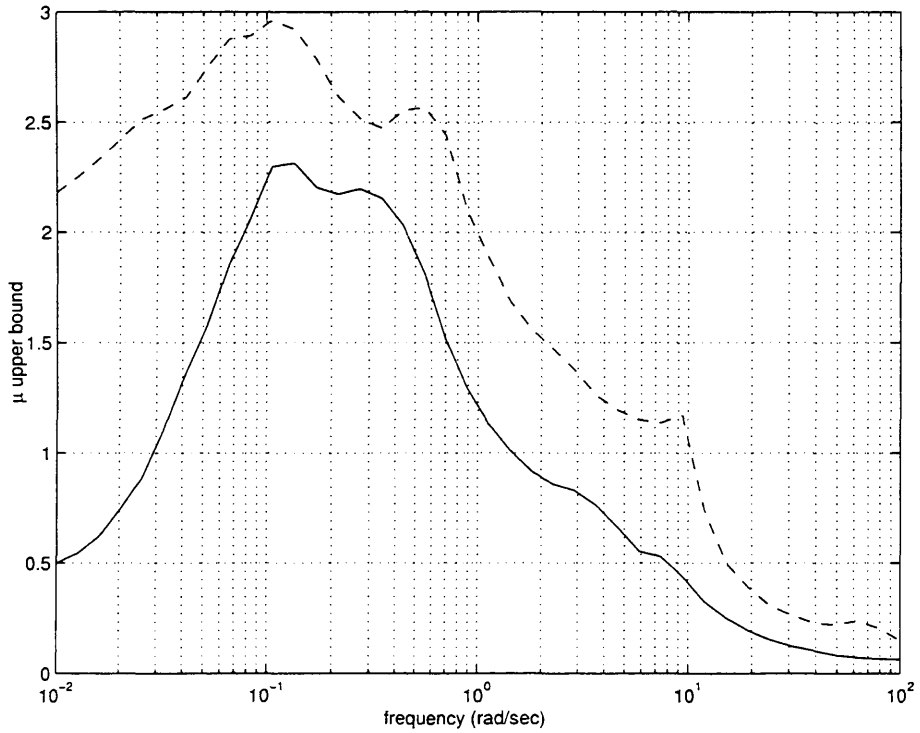

Figure 7.11: μ upper bounds for 80-120 knots (—) and 80-150 knots (- -) dynamic variations.

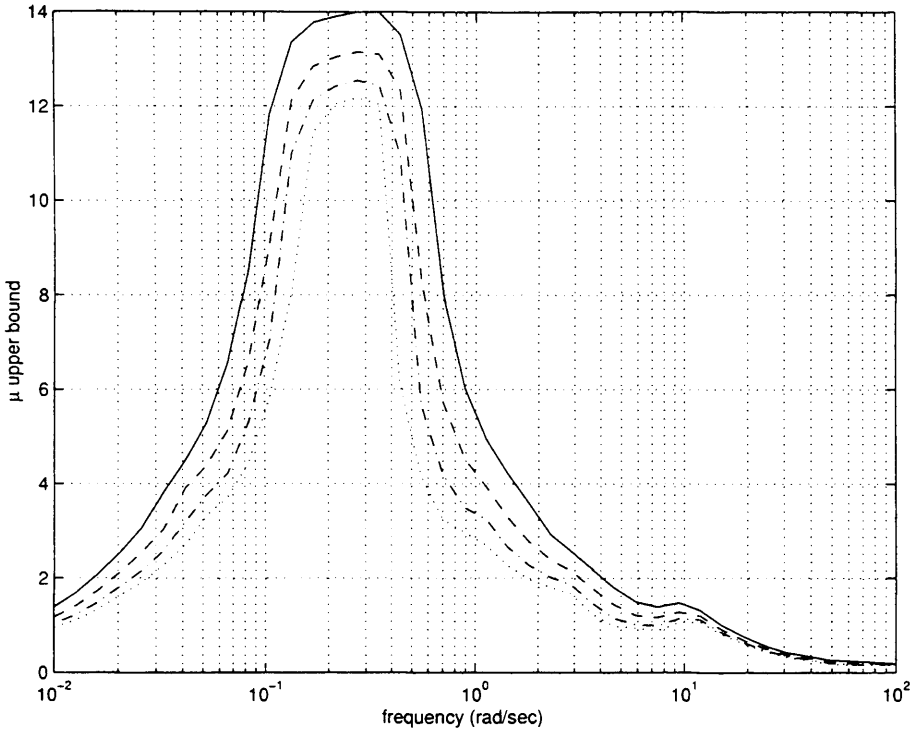

Figure 7.12: μ upper bounds for the centralised system for 80-20 knots dynamic variations with 1% (··), 4% (-·), 9% (- -) and 16% (—) added complex uncertainty.

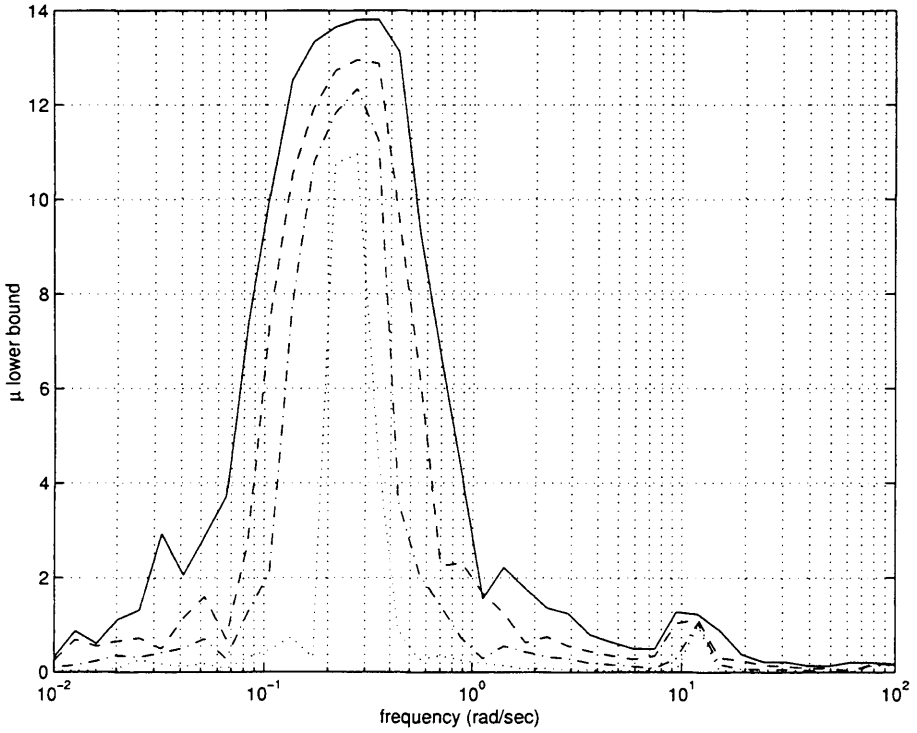

Figure 7.13: μ lower bounds for the centralised system for 80-20 knots dynamic variations with 1% (··), 4% (-·), 9% (- -) and 16% (—) added complex uncertainty.

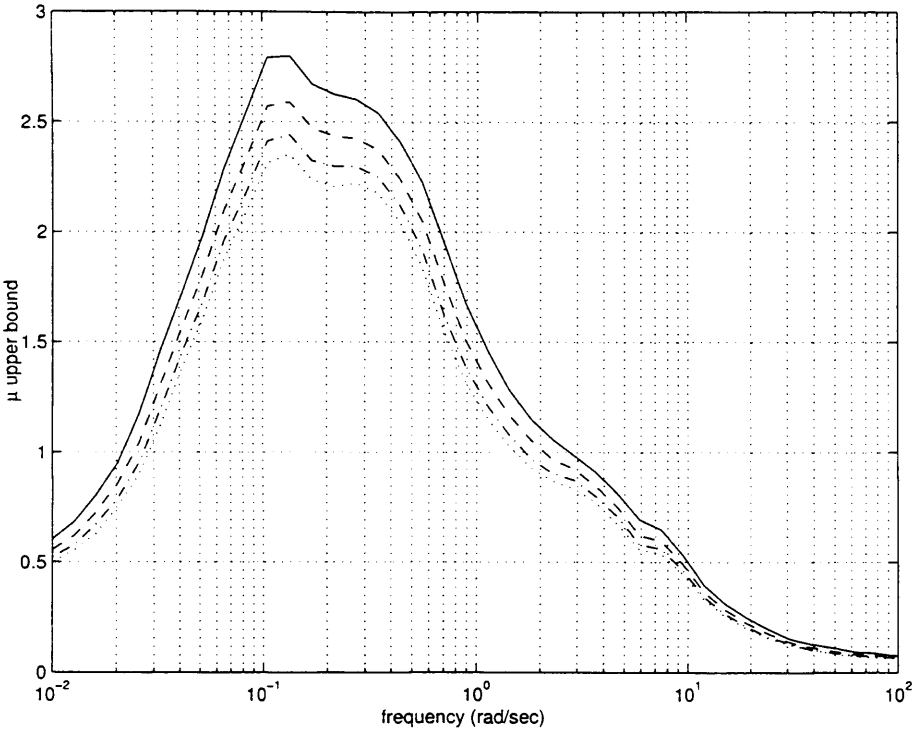

Figure 7.14: μ upper bounds for the centralised system for 80-120 knots dynamic variations with 1% (··), 4% (-·), 9% (- -) and 16% (—) added complex uncertainty.

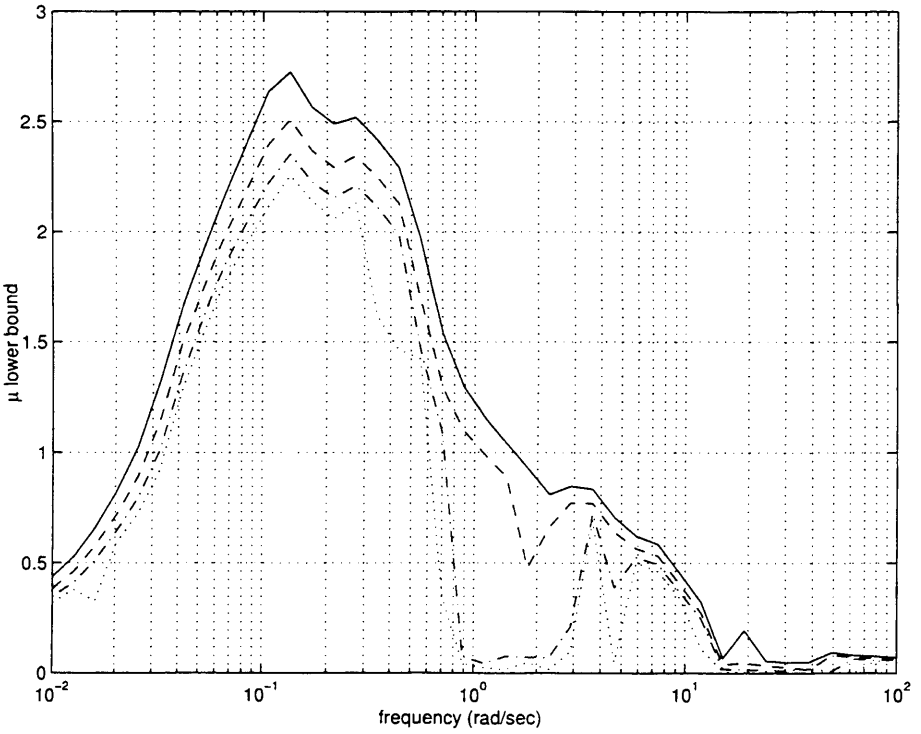

Figure 7.15: μ lower bounds for the centralised system for 80-120 knots dynamic variations with 1% (··), 4% (-·), 9% (- -) and 16% (—) added complex uncertainty.

for complex uncertainty introduced via the actuators, and via the addition of 1% complex uncertainty to each real uncertainty element. Figures 7.17 to 7.20 provide a similar comparison for variations in aircraft dynamics over the flight envelope. In general, the use of the actuator uncertainty approach is seen to produce tighter bounds, but at the cost of greater increases in the upper bounds, especially at higher frequencies.

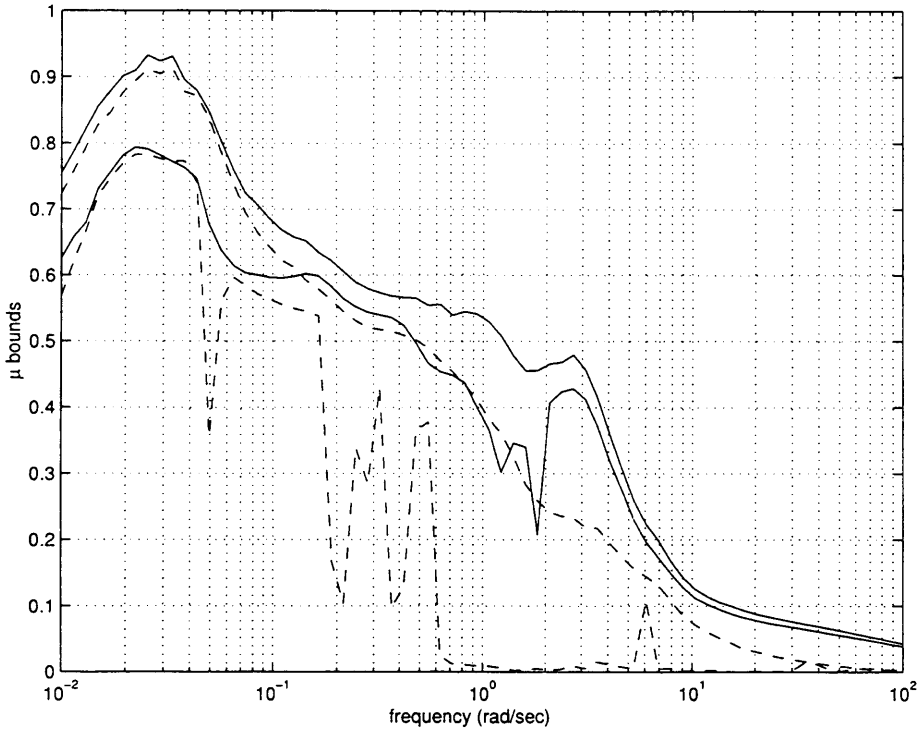

Figure 7.16: μ bounds for 2 uncertain parameters with actuator uncertainties (—) and with real + 1% complex uncertainties (- -) for the centralised system.

7.4 Stability analysis of centralised system at different velocities using the ν-gap metric

§ 4.3 introduced the ν-gap metric and described how it may be used to test if the nominal controller can stabilise plants other than that which it was designed for. It was noted that this was of particular use if using the $\mathcal{H}_\infty$ loop shaping design procedure. The 80 knot controller now achieves an optimal value of $\gamma = 4.02$. The suboptimal controller was constructed, giving a stability margin $B_{P_{80knots},K}$ of 0.226. Looking at Table 7.1, it may be concluded that stability cannot be guaranteed for any of the selected points away from the nominal, as the value of $\delta_\nu(P_{80knots}, P_\star)$ is greater that the value of $B_{P_{80knots},K}$ for all plants where $\star$ = 20 knots, 50 knots, 120 knots and 150 knots.

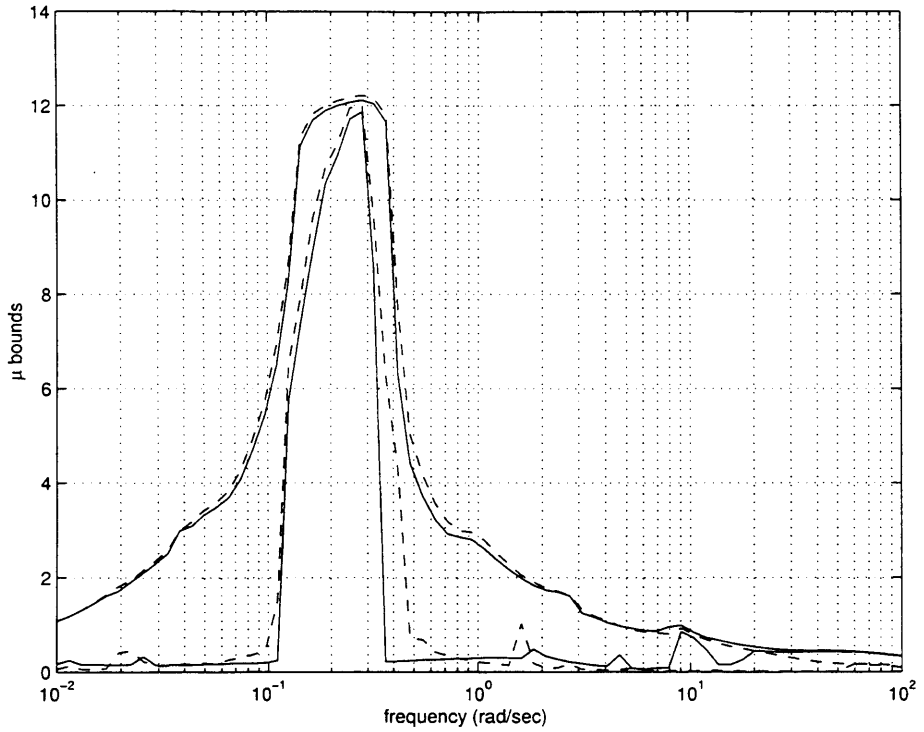

Figure 7.17: μ bounds for 80-20 knot dynamics variations with actuator uncertainties (—) and with real + 1% complex uncertainties (- -) for the centralised system.

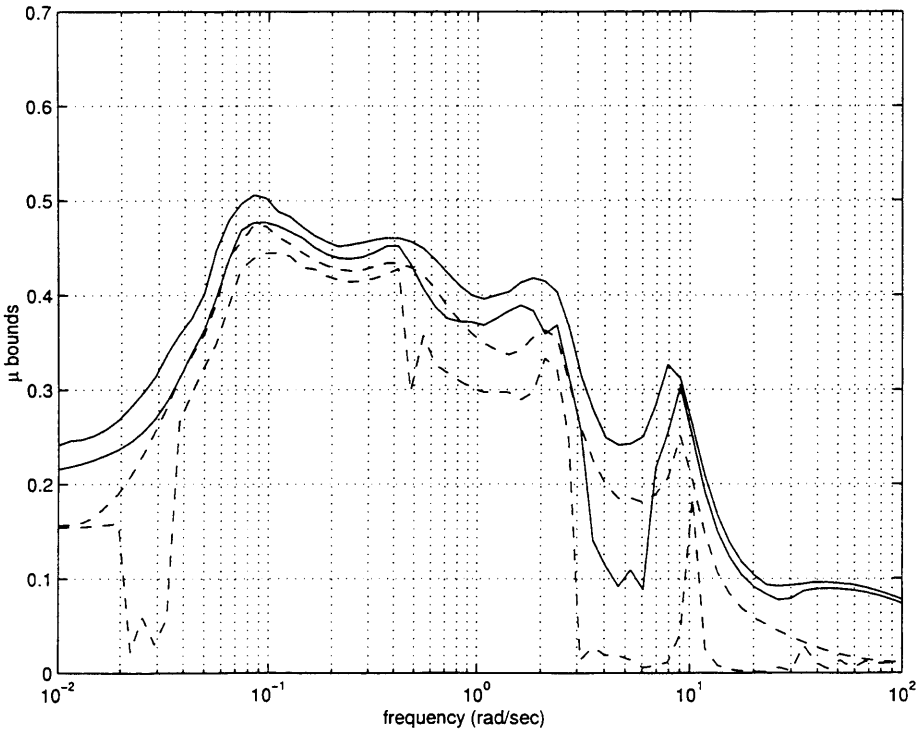

Figure 7.18: μ bounds for 80-50 knot dynamics variations with actuator uncertainties (—) and with real + 1% complex uncertainties (- -) for the centralised system.

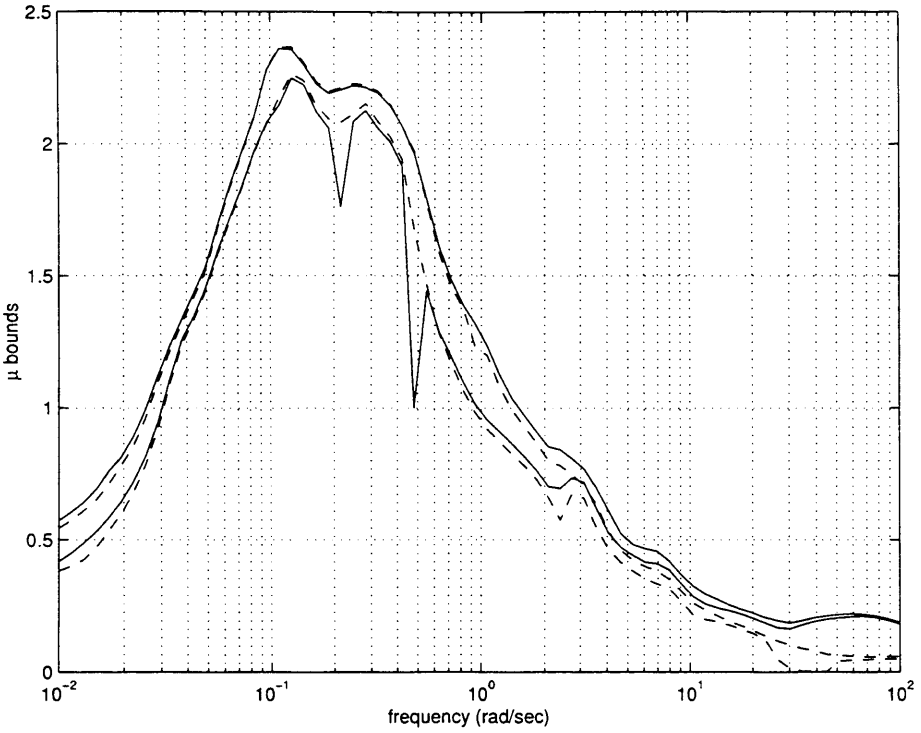

Figure 7.19: μ bounds for 80-120 knot dynamics variations with actuator uncertainties (—) and with real + 1% complex uncertainties (- -) for the centralised system.

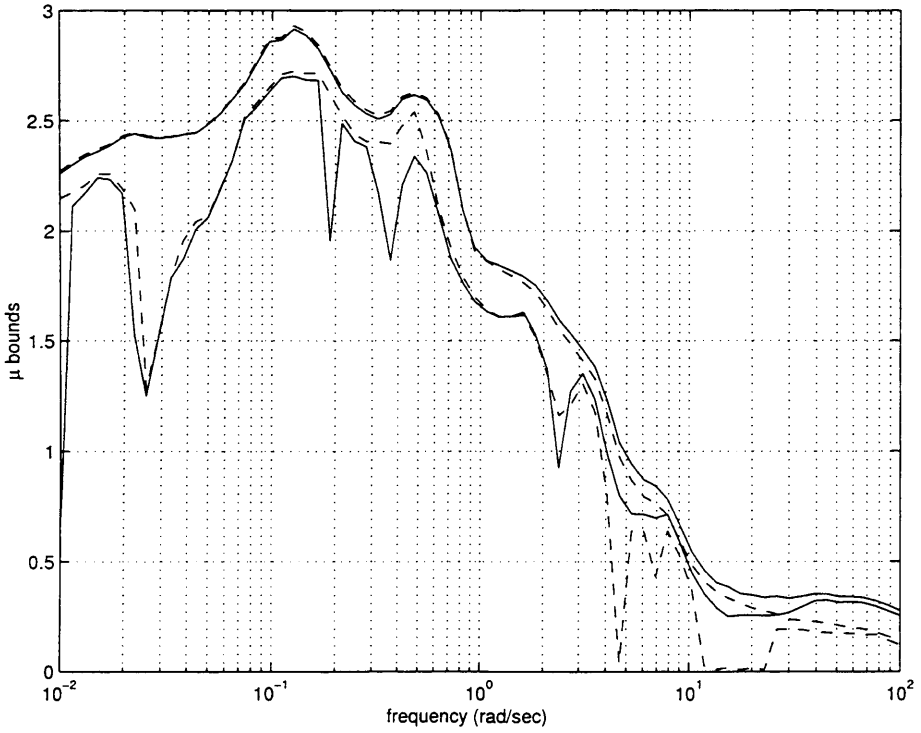

Figure 7.20: μ bounds for 80-150 knot dynamics variations with actuator uncertainties (—) and with real + 1% complex uncertainties (- -) for the centralised system.

Plant ($P_\star$)	max μ U.B.	$\delta_\nu(P_{80knots}, P_\star)$
20 knots	12.10	0.67
50 knots	0.45	0.31
120 knots	2.34	0.41
150 knots	2.88	0.78

Table 7.1: δ_ν and μ across the flight envelope.

However, a significant advantage of using the ν-gap metric is the fact it can be plotted across frequency, as shown in Figure 7.21. This shows $\psi(P_{80knots}(j\omega), P_\star(j\omega))$ as well as the robustness margin $b_{P_{80knots},K}(\omega)$ for the IFPC system at the nominal 80 knots design point.

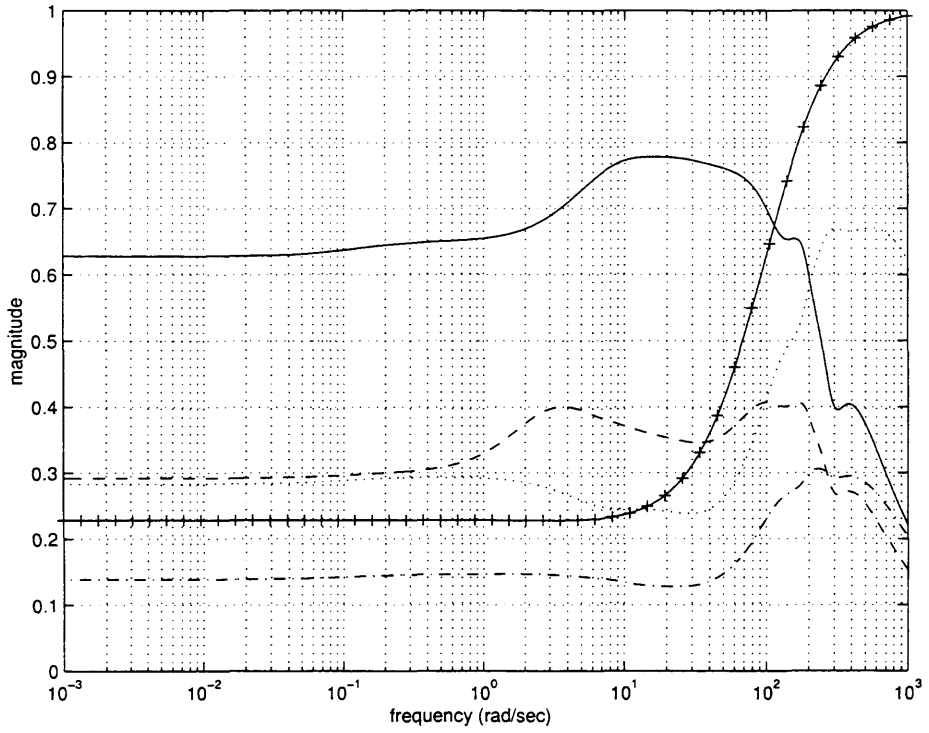

Figure 7.21: $\psi(P_{80knots}(j\omega), P_\star(j\omega))$, where $\star$ = 20 knots (··), 50 knots (-·), 120 knots (- -) and 150 knots (—), and $b_{P_{80knots},K}(\omega)$ (+).

From this figure it can be seen, as found in the previous μ-analysis, that the centralised IFPC system loses stability and performance robustness below 50 knots and at and above 120 knots. Note the usefulness of the frequency domain interpretation of the ν-gap metric for this example

$$B_{P_{80knots},K} < \delta_\nu(P_{80knots}, P_{50knots})$$

but

$$b_{P_{80knots},K}(\omega) > \psi(P_{80knots}(j\omega), P_{50knots}(j\omega)) \ \forall\omega$$

This frequency domain interpretation of the ν-gap metric can also provide much useful infor-

mation for controller design. For example, plots of $\psi(P_{80knots}(j\omega), P_\star(j\omega))$ for a particular selection of weighting functions can reveal the frequency regions over which robust performance is most degraded, and the weighting functions can then be modified accordingly over those critical frequencies.

Finally, the consistency of the μ- and ν-gap analyses is noted, both with each other, in Table 7.1, and with the results of piloted simulation trials carried out on DERA Bedford's RTAVS simulation facility and reported in [7]. As an example, Figure 7.22 shows the response of the IFPC system in nonlinear simulation to a demand on $\dot{\gamma}$ at 20 knots, 50 knots, 120 knots and 150 knots. Figure 7.23 shows the levels of coupling into γ caused by a demand on VT at the same points in the flight envelope. The level of performance degradation for each manoeuvre at each point of the flight envelope correlates reasonably well to the corresponding value of μ and the ν-gap in the frequency domain analyses described above. For example, the significant degradation in the performance of the IFPC system at the 150 knot point in the envelope is reflected in the substantial increase in the values of both μ and the ν-gap in the corresponding robustness tests.

7.5 Robustness analysis of the partitioned system

Nominal and robust performance analysis was carried out in § 6.3 by comparing the flying qualities of the partitioned IFPC system in nonlinear simulation with those achieved by the initial centralised design. In addition to preserving the time domain performance properties of the centralised design, the partitioned IFPC system must also exhibit the same levels of robustness to variations in aircraft/engine dynamics and disturbances over the V/STOL flight envelope. In order to investigate this issue, a detailed analysis of the stability robustness characteristics of the two systems using the structured singular value μ was conducted [4]. In this analysis, as for the centralised system, uncertainties in the nominal plant model arising from two distinct sources are again considered: variations in the values of aircraft parameters (aircraft mass, m and inertia, I_{yy}), and variations in the airframe/engine dynamics as the aircraft moves away from the 80 knot design point of the V/STOL flight envelope [35, 5]. Upper bounds on μ for variations in airframe/engine dynamics between the 80 knot design point and 20, 50, 120 and 150 knots are shown in Figures 7.24 to 7.25, for both the centralised and partitioned IFPC systems. μ upper bounds for uncertainty due to variations in aircraft mass and inertia are shown in Figure 7.26.

The figures reveal a dramatic degradation in the robustness properties of the partitioned IFPC system, as compared with the initial centralised design. Such a degradation in robustness is

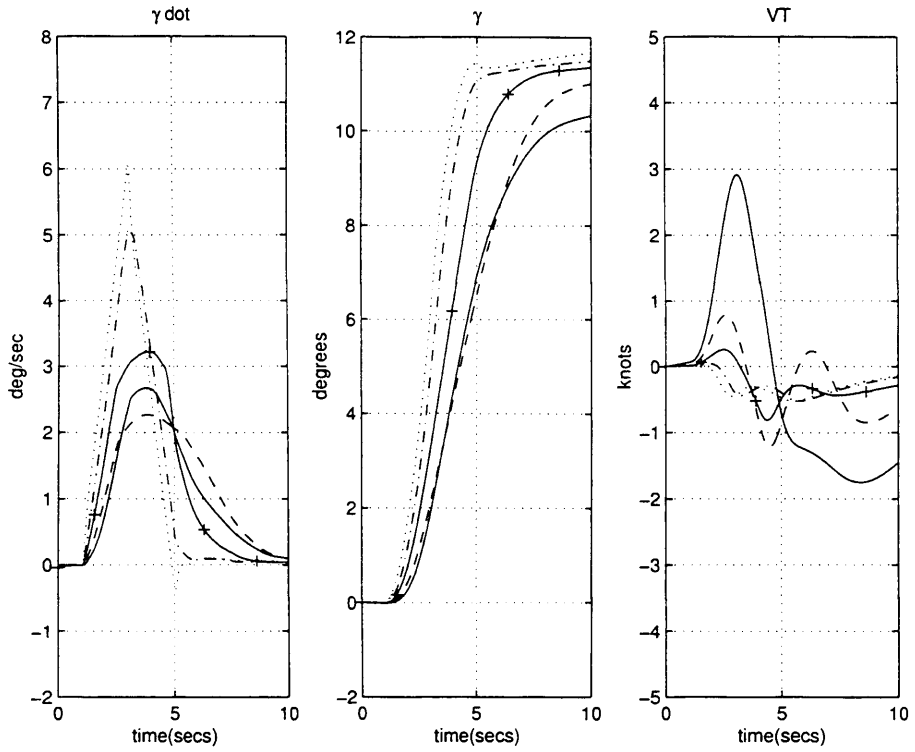

Figure 7.22: $\dot{\gamma}$, γ and VT responses to a demand on $\dot{\gamma}$ at 20 knots (··), 50 knots (-·), 80 knots (+), 120 knots (- -) and 150 knots (—) for the centralised system.

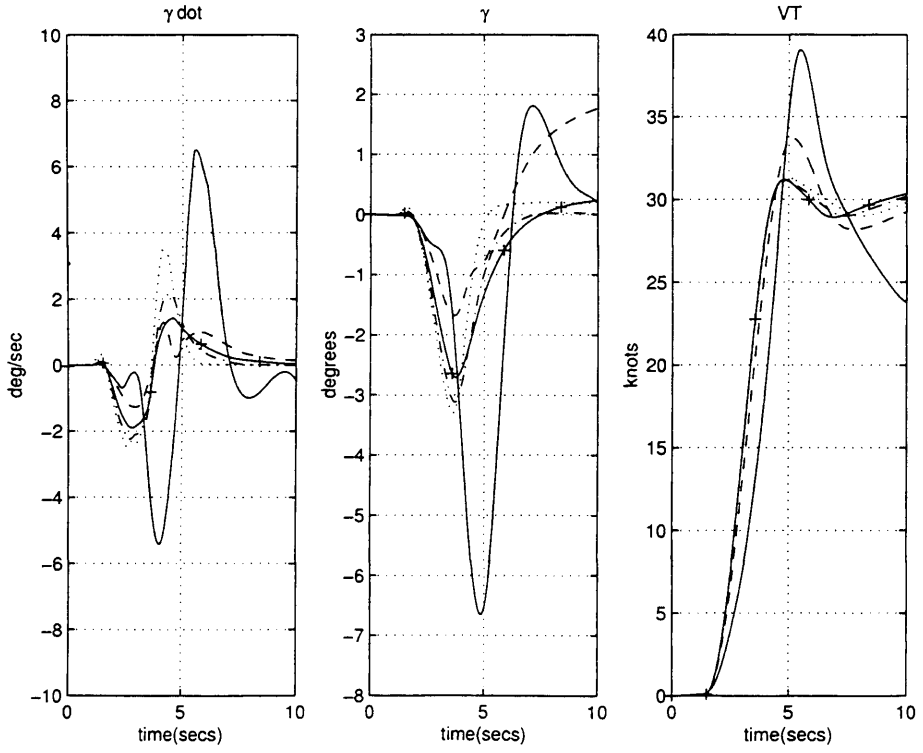

Figure 7.23: $\dot{\gamma}$, γ and VT responses to a 30 knots demand on VT at 20 knots (··), 50 knots (-·), 80 knots (+), 120 knots (- -) and 150 knots (—) for the centralised system.

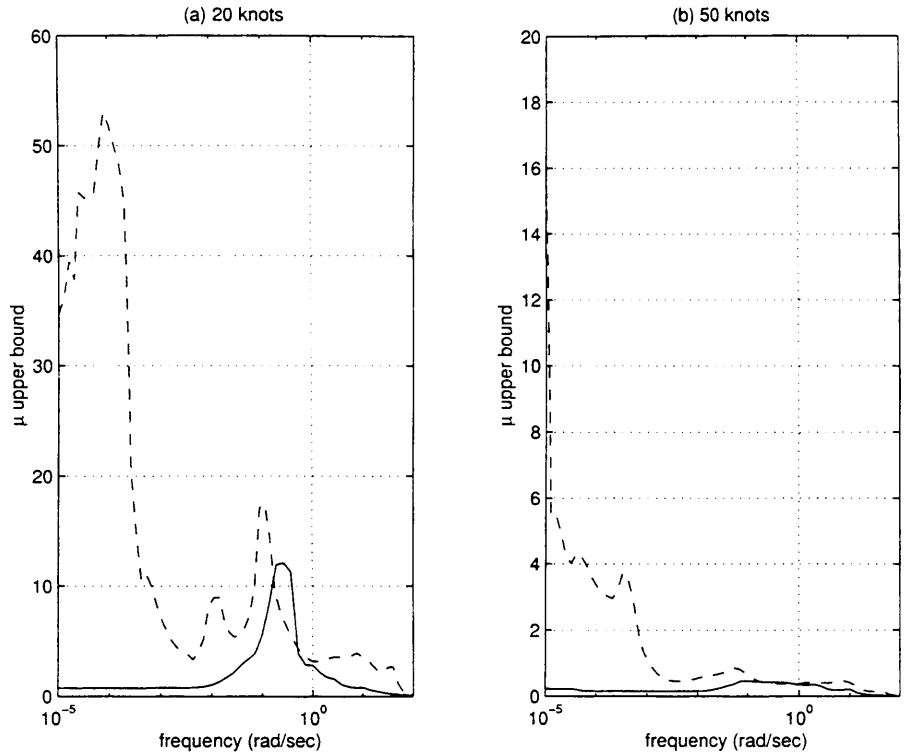

Figure 7.24: μ upper bounds for variations in airframe/engine dynamics between 80 knots and (a) 20 knots, (b) 50 knots, centralised (—) and partitioned (- -) IFPC systems.

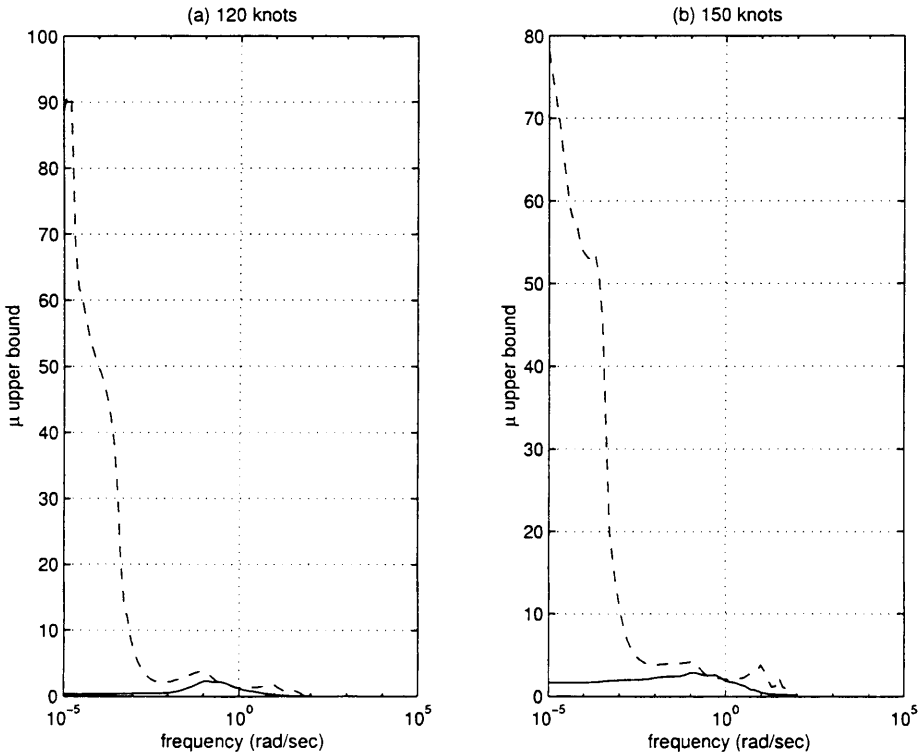

Figure 7.25: μ upper bounds for variations in airframe/engine dynamics between 80 knots and (a) 120 knots, (b) 150 knots, centralised (—) and partitioned (- -) IFPC systems.

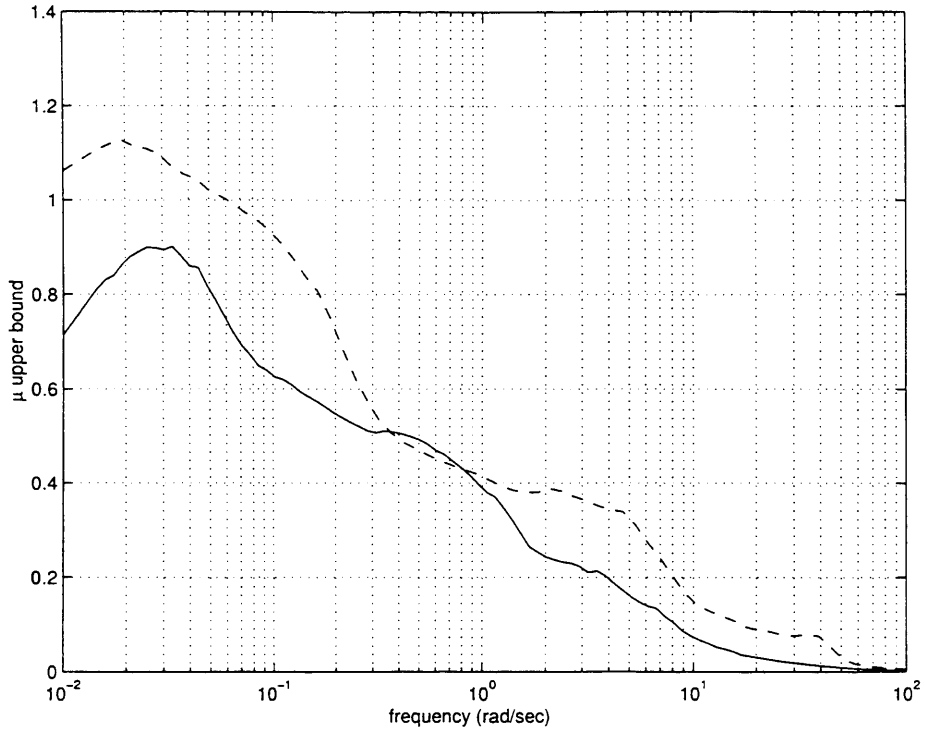

Figure 7.26: μ upper bound for uncertainty due to parameter variations (m and I_{yy}), centralised (—) and partitioned (- -) IFPC systems.

unacceptable, and points to the need for further refinement of the partitioned IFPC system. An advantage of the approach to controller partitioning described in the previous section is that it readily allows the subsequent re-design of the various subcontrollers. This design flexibility is exploited in the next section, in order to improve the performance and robustness properties of the partitioned IFPC system.

7.5.1 Re-design of the engine subcontroller

A key requirement of any partitioning procedure for integrated flight and propulsion control systems is to provide independent integrity of the individual subcontrollers. This is necessary for commercial reasons, since different subsystems may be designed and built by independent suppliers who will need to retain responsibility for implementation and testing up to final system integration. It is also necessary, however, in order to allow the re-design of the individual subcontroller blocks subsequent to partitioning. This re-design may be required for a variety of reasons, e.g. a change in a subsystem or in a subsystem's performance specifications. It may also be required to address any degradation in performance or robustness caused as a result of the partitioning process itself, as is the case here. One approach to the problem of re-designing the partitioned IFPC system consists of using numerical optimisation methods to adjust the

parameters of the subcontrollers so that some cost function which describes the 'error' between the performance and/or robustness properties of the centralised and partitioned systems is minimised [65]. An alternative strategy is to directly re-design the engine subcontroller, in order to make more efficient use of the available engine subsystem control effectors. The motivation for this approach comes from the observation that the major difference in nonlinear simulations between the partitioned and centralised systems seems to be in the use of the engine actuators, see Figures 6.4 and 6.6. Recall that in the interface subcontroller design step of the partitioning procedure it was noted that the control problem was non-square, i.e. three engine control inputs are available to control one interface variable THM. This control power can be used more efficiently if some of the other internal engine variables are also included in the interface subcontroller design, since this subcontroller should ideally provide *decoupled* tracking of demands on Z_{ea}. The engine variables $T10$ and $NLPC$ are therefore re-assigned to Z_{ea}, leaving only the two surge margins $HPSM$ and $LPSM$ in Z_{ee}. The interface subcontroller design problem is now square, and a new K_{ea} (10 state) controller was computed, again using the method of $\mathcal{H}_\infty$ loop shaping, in order to meet the performance specifications (taken from the centralised design) on all three interface variables. The weighting functions were selected as

$$k = diag\,(0.05,\ 0.05,\ 0.05)\,, \quad W_1 = \frac{s+1}{s} \times I_{3\times 3},$$

resulting in a value of $\gamma_{opt} = 1.54$. Nonlinear simulation results for the re-designed partitioned IFPC system, Figures 7.27 to 7.30, show the partitioned system achieves good performance and, for the demand on VT, improved decoupling between the controlled variables.

Robustness analysis of the re-designed system, however, reveals a dramatic improvement over the initial partitioned system, resulting in an almost complete recovery of the robustness properties achieved by the centralised design, Figures 7.31 to 7.33.

A final robustness test can be performed by using μ-analysis to calculate the smallest perturbation in the closed-loop performance of the interface subsystem which will cause the overall partitioned system to become unstable. The closed-loop singular values for the resulting 'worst-case' interface subsystem can then be compared with the 'desired' singular values (i.e. the specifications generated from the centralised design) and the singular values actually achieved by the interface subcontroller. The resulting plot shows how 'close' the partitioned system is to both the best-case (perfect matching of centralised performance) and worst-case (overall system instability) scenarios. As shown in Figure 7.34, the closed-loop singular values for the interface subsystem of the re-designed partitioned system show a good match to the desired singular values, and are significantly different from those corresponding to overall system instability. Tests such as the above can provide much useful information for iterative re-design

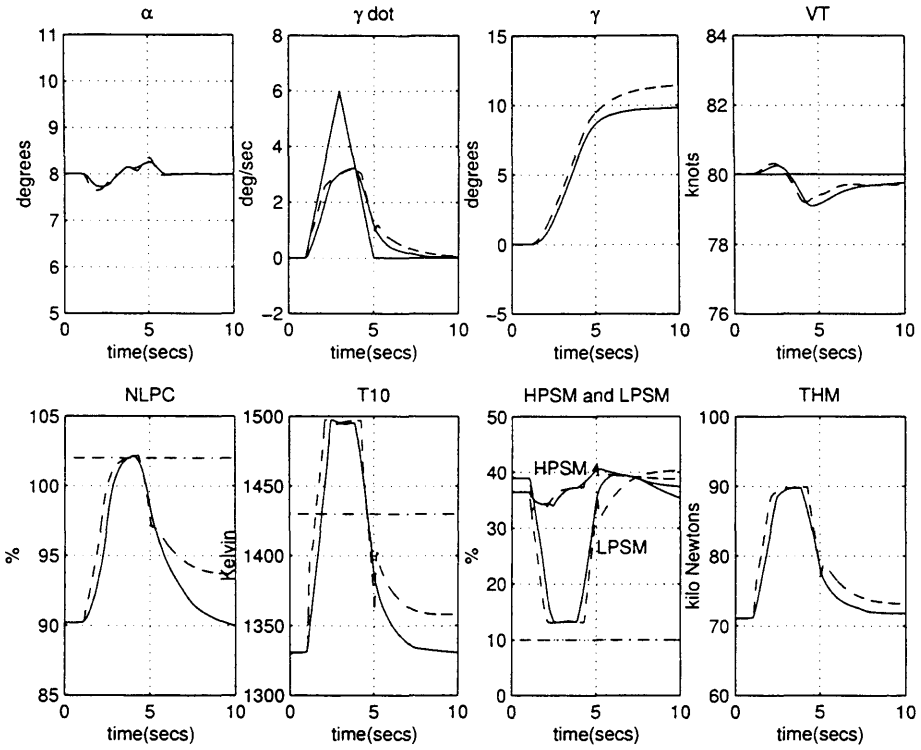

Figure 7.27: Airframe and engine responses to a $\dot{\gamma}$ demand for the centralised system (—) and the re-designed partitioned system (- -) at 80 knots.

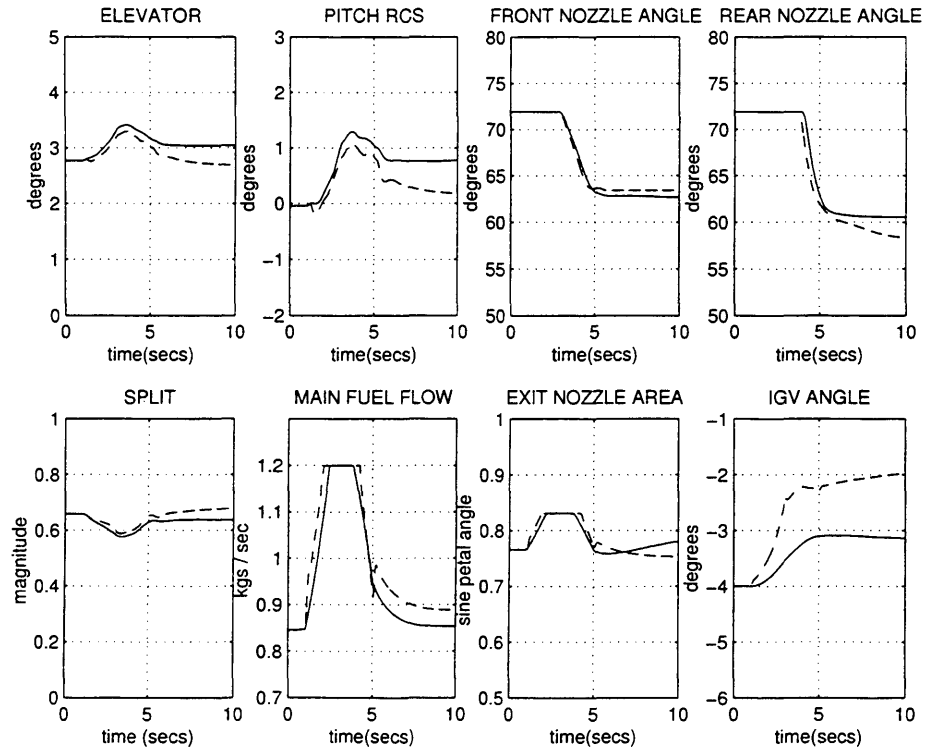

Figure 7.28: Actuator responses to a $\dot{\gamma}$ demand for the centralised system (—) and the re-designed partitioned system (- -) at 80 knots.

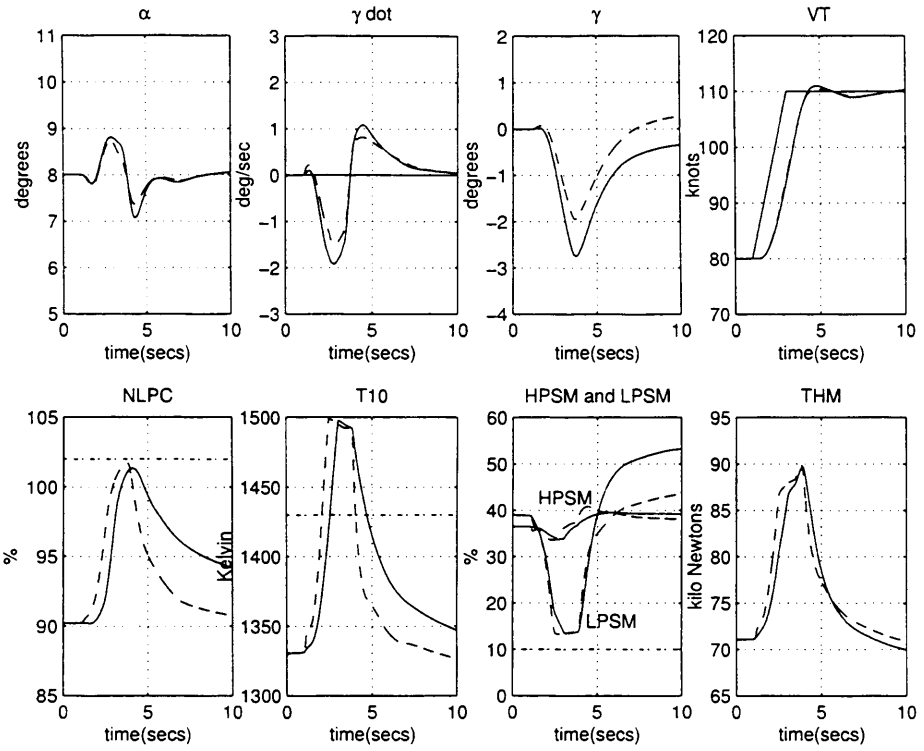

Figure 7.29: Airframe and engine responses to a VT demand for the centralised system (—) and the re-designed partitioned system (- -) at 80 knots.

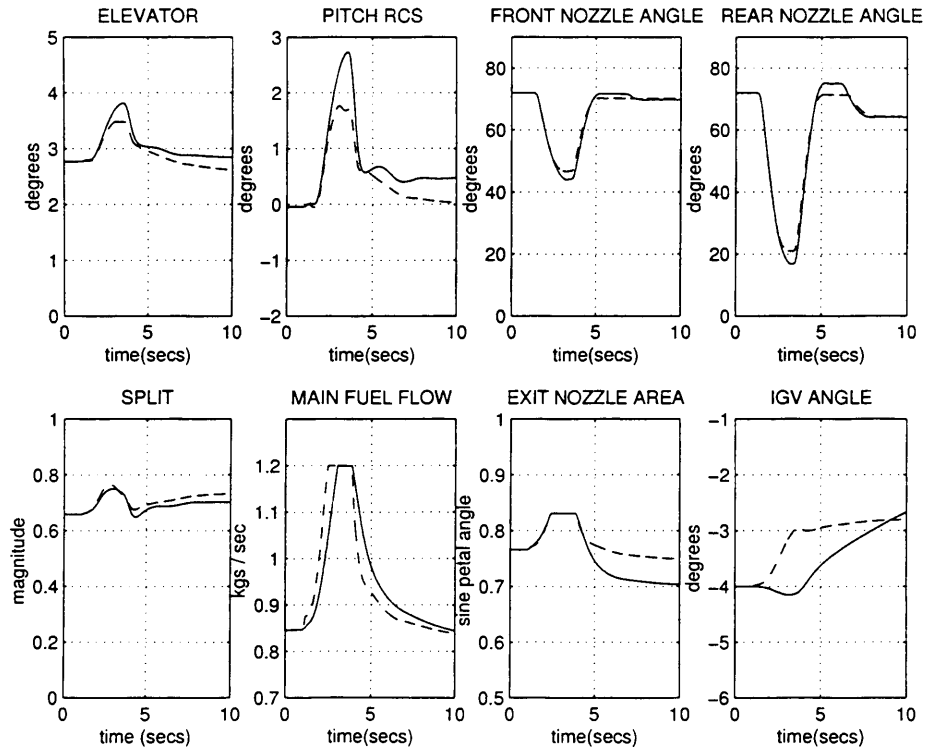

Figure 7.30: Actuator responses to a VT demand for the centralised system (—) and the re-designed partitioned system (- -) at 80 knots.

ERROR: timeout
OFFENDING COMMAND: timeout

STACK:

-dictionary-
-dictionary-
false
[1 0 0 1 0 0]

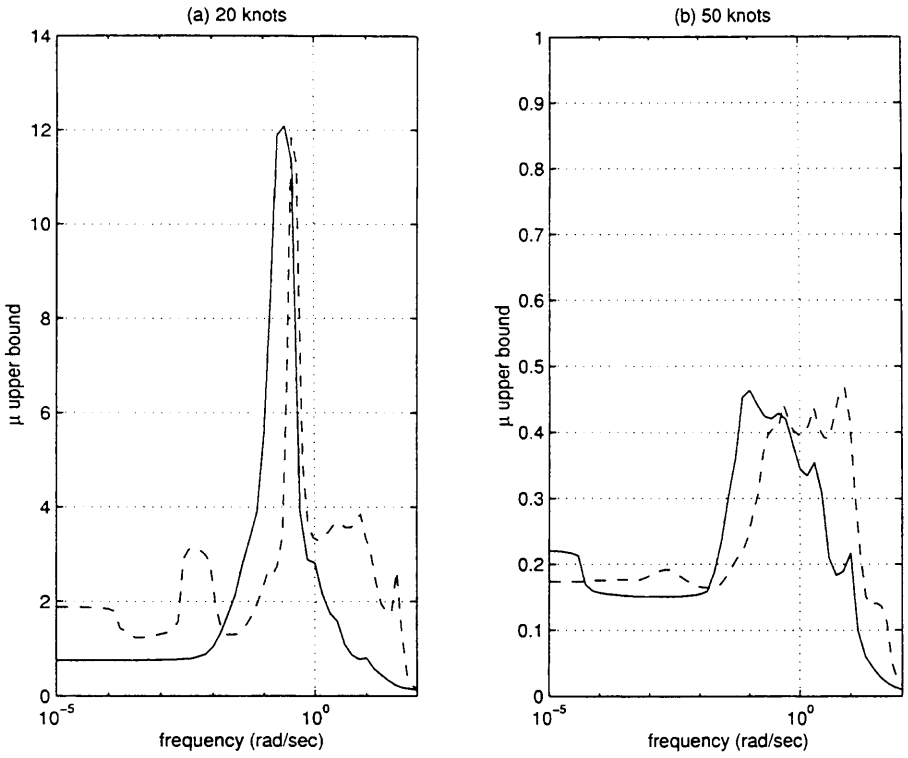

Figure 7.31: μ upper bounds for variations in airframe/engine dynamics between 80 knots and (a) 20 knots, (b) 50 knots, centralised (—) and re-designed partitioned (- -) IFPC systems.

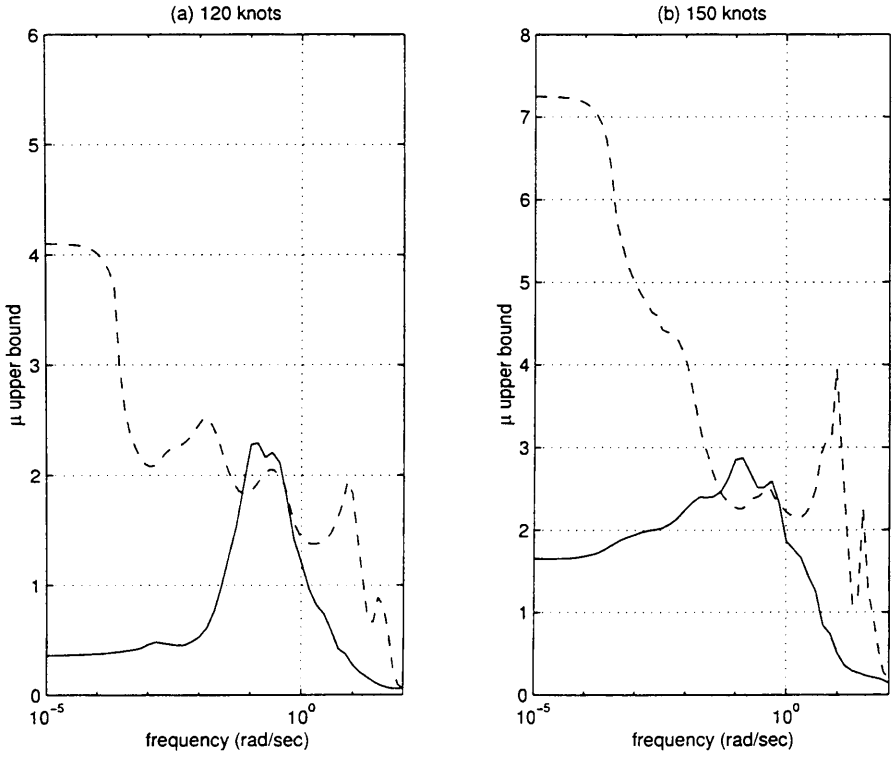

Figure 7.32: μ upper bounds for variations in airframe/engine dynamics between 80 knots and (a) 120 knots, (b) 150 knots, centralised (—) and re-designed partitioned (- -) IFPC systems.

of subcontrollers. If, for example, the desired, achieved and worst-case singular values all correspond closely, the subcontroller could be re-designed to increase robustness at the expense of sacrificing some matching between the centralised and partitioned systems.

7.6 Summary

Chapters 5 and 6 presented good nonlinear simulation results for the centralised and partitioned systems. However, it is not expected that a single nominal controller will be able to guarantee stability across the flight envelope, and that some form of controller scheduling/switching will be required. In order to implement this, it is important to know where in the flight envelope the nominal controller can no longer guarantee stability. A μ-analysis of the system has been presented that considers the stability of the nominal system to perturbations originating from parameter variations such as mass and inertia, or from differences due to the airframe/engine dynamics according to the point in the flight envelope. One drawback of using μ-analysis in this way is that the perturbations will always be real, leading to difficulties in the lower bound in converging. Two methods were considered that introduced some form of complex uncertainty into the system in order to allow a lower bound to be calculated. Further, the ν-gap metric has been used as an alternative method of calculating stability away from the nominal design point. It has been shown that although the ν-gap metric may be greater than the stability margin found for the $\mathcal{H}_\infty$ loop shaping controller in a frequency independent test, if its value is considered on a frequency by frequency basis it may be concluded that stability is in fact guaranteed. These were applied to both the centralised and partitioned systems. The partitioned system displayed significant robustness degradation. An important advantage of the partitioning method has been demonstrated, whereby the subcontrollers may be easily redesigned. It has been seen that it is not sufficient to ensure the performance of the partitioned system meets that of the centralised system - its robustness properties must also be preserved. Following a μ-analysis of the initial partitioned system, it was clear that all robustness had been sacrificed. However, a redesign of the engine subcontroller resulted in good performance as shown in nonlinear simulations, but with an almost complete recovery of the robustness of the centralised system.

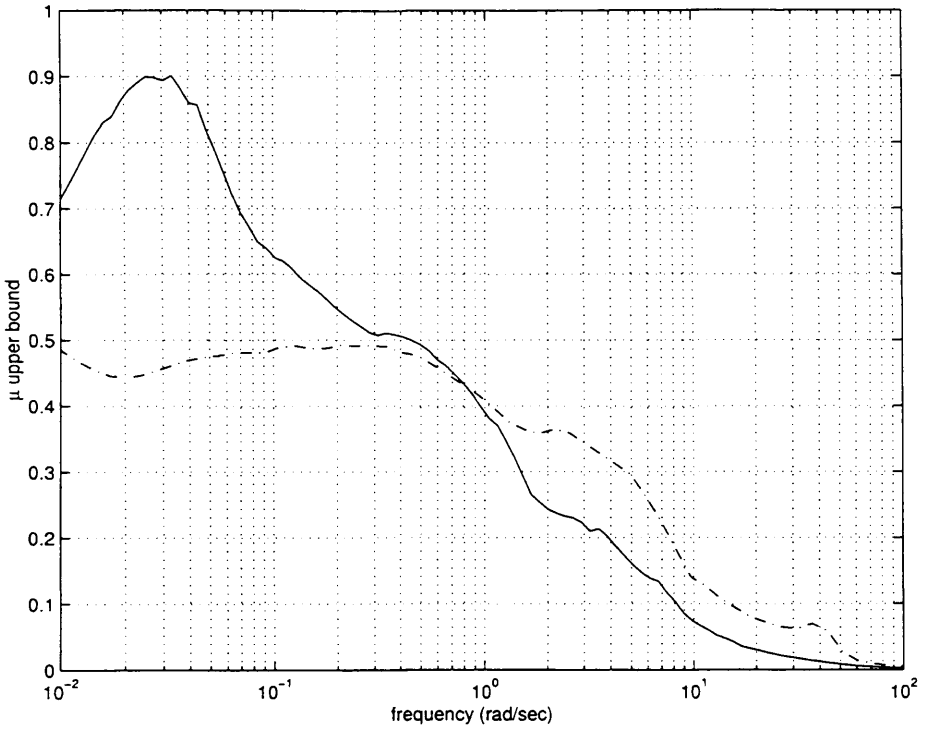

Figure 7.33: μ upper bound for uncertainty due to variations in m and I_{yy}, centralised (—) and re-designed partitioned (-·) IFPC systems.

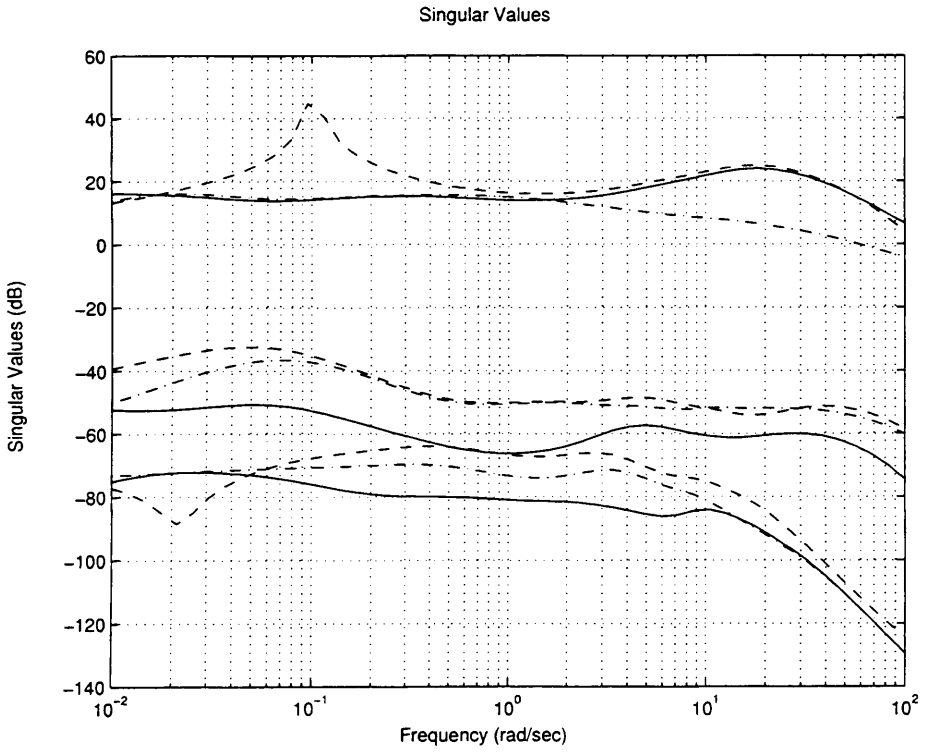

Figure 7.34: Desired (-·), achieved (—) and worst case (- -) closed-loop singular values for the interface subsystem.

Chapter 8

Limiting of safety-critical engine variables

This chapter will describe the approach used to address the engine variable limiting problem [32]. § 8.1 gives a brief description of why it is important to have some method of restricting the internal engine variables. § 8.2 introduces a new engine limiting method, which was reasonably successful. The approach presented here is based on the strategy of imposing maximum position limits on engine actuators, via a switching function, in the case of a particular engine output variable reaching its specified safety limit. The maximum limits to be imposed on the engine actuators are calculated on-line, and correspond to the values of the actuators just prior to the engine safety limit being violated. The actuators are subsequently 'released' to resume normal operation as soon as the pilot's demands have returned to a level which will not cause any engine safety limits to be exceeded. Finally, § 8.3 provides a summary of this chapter.

8.1 Need for engine limiting strategies

§ 1.1.7 and § 1.1.8 introduced some details of what the engine control problem involves, and some advanced engine control strategies. Current engine control generally involves working on limits - temperatures must be limited in order to maintain material strength and to minimise creep, the surge margins must be kept sufficiently high in order to avoid surge, and spool speeds must be limited in order to keep the rotor stress within safe margins. Restrictions such as these are important in order to maximise the life of the engine while still achieving maximum possible performance and maintaining safety. These limits are particularly important due to the nature of the use of the engine. High fluctuating loads, high temperatures and temperature gradients, frequent starts and stops and stress concentrations resulting from complex geometrical shapes and from surface discontinuities produced by service - these conditions all contribute to making the components fatigue-prone. Overall, it is desired that the engine will provide performance subject to these engine limits and restrictions. The engine will have some safety margin which may be used occasionally for short periods if required. If this short time is exceeded, however,

the engine will need to be dismantled and the affected parts replaced. Exceeding temperatures may significantly shorten the life of an engine, but exceeding rotational speeds may result in failure of parts, or in excessive vibrations being transmitted throughout the airframe where additional damage may be caused. Excessive temperatures, pressures and rotational speeds within the engine can have an adverse effect on its life span, and in the extreme case can cause catastrophic failure or engine blowup. A complicating factor is due to conflicting requirements. For example, in order to increase the efficiency of the engine, the turbine inlet temperature should be high. However, this will exceed the maximum temperature allowed by the materials the turbine is constructed of. Therefore, either some cooling strategy is required or the inlet temperature must be limited. For every 20°K increase in temperature near its rated operating point the blade life is typically halved [63].

8.2 An integrated engine-limiting strategy

Specifications of maximum limits for certain engine variables were provided by DERA in order to protect the engine from dangerous over stress and over temperature [12]. It was noted at the beginning of this study that in some areas of the flight envelope the flying qualities requirements specified may lie outside the aircraft's capability. In such cases, it was stated that the control law should govern aircraft motion to give the maximum performance obtainable while staying within all specified safety limits. In order to enforce this priority, additional nonlinear control logic must be 'wrapped around' the engine subcontroller. This problem was addressed in [28] by incorporating nonlinear control logic to impose an accel/deccel schedule on engine actuators. This approach also necessitates the addition of integrator wind-up protection logic [13, 20] in order to accommodate the effects of the enforced actuator limits.

Through running nonlinear simulations, it was noted that of all the engine variables that are to remain within specific limits, $T10$ was generally the first to violate its limit [34]. This temperature is closely related to the use of main fuel flow and exit nozzle area. There will be many combinations of engine actuator positions that will cause each of the internal engine variables to exceed limits. Moreover, manufacturing tolerances will mean that each new engine will differ slightly from the others, and worn engines will differ further. In order to account for all of these sources of error, it was decided to apply a limiting scheme that calculated and imposed actuator limits on line. This is achieved by monitoring $T10$ constantly. If it approaches within 10K of its maximum limit of 1430K, the positions of each of the engine actuators are noted. If $T10$ continues to increase, these noted actuator positions are used as temporary artificial position limits, as it is known that this combination of actuator positions is sufficient to keep $T10$ 10K below its maximum limit. Once the demands to the actuators fall

below these 'false' limits, they are released and are free to move as required.

It has been noted from previous nonlinear simulation results that the system frequently operates on engine actuator position limits. Implementation of the engine limiting scheme will result in the actuators also operating on additional 'artificial' limits. Therefore, implementing an anti-windup scheme is necessary to maintain the good performance properties of the system, both due to actuator saturation and due to artificial position limits imposed on the system from the use of the engine limiting strategy. As noted in 6.3, it is the interface part of the engine subcontroller that makes the large demands causing integrator windup, therefore it is necessary to implement W_{1ea} in its self-conditioned form. The final system is shown in Figure 8.1.

The approach has been seen to work well in nonlinear simulation, as shown in Figures 8.2 to 8.4 for a demand on $\dot{\gamma}$. These results compare the nonlinear responses of the partitioned IFPCS system with the engine limiting strategy with those of the centralised system. It is clearly shown that $T10$ exceeds its maximum specified value for the centralised system. However, for the partitioned system with the additional engine control logic, $T10$ is held just below its specified limit until the reference demand has decreased sufficiently. Note also from Figure 8.2 that the satisfaction of the specifications on maximum engine limits has been achieved with little degradation in overall Mission level flying qualities. It is seen from Figures 8.2 that by directly limiting $T10$, the engine variable $NLPC$ has also been kept within its maximum limit - thus it seems likely that in practice it will suffice to limit only a subset of engine variables of interest. Figure 8.3 shows the responses of the additional internal engine variables that were included in the advisory section of the specifications. It can be seen that by limiting $T10$, these variables are also effectively kept within their specified limits. From Figure 8.4 it can also be seen that, as it is the engine actuators that are affected directly by the engine limiting strategy, the airframe actuator positions are similar for both the partitioned system with engine limits and for the centralised system.

Figures 8.5 to 8.7 compare the responses to a demand on VT of the partitioned system with engine limits with the centralised system. Again, it is shown that whereby for the centralised system, $T10$ exceeds its maximum specified value significantly, for the partitioned system with engine limits, all the engine variables are kept within safe limits while still achieving good performance properties.

Finally, the robust performance properties of the partitioned IFPC system are illustrated by checking its performance in nonlinear simulations away from the nominal 80 knot point of the flight envelope. Figures 8.8 and 8.9 show the performance of the partitioned system with engine limiting schemes and the centralised system to a $\dot{\gamma}$ demand at 50 knots. The coupling into all other outputs remains within specifications. Figures 8.10 and 8.11 compare the results for a

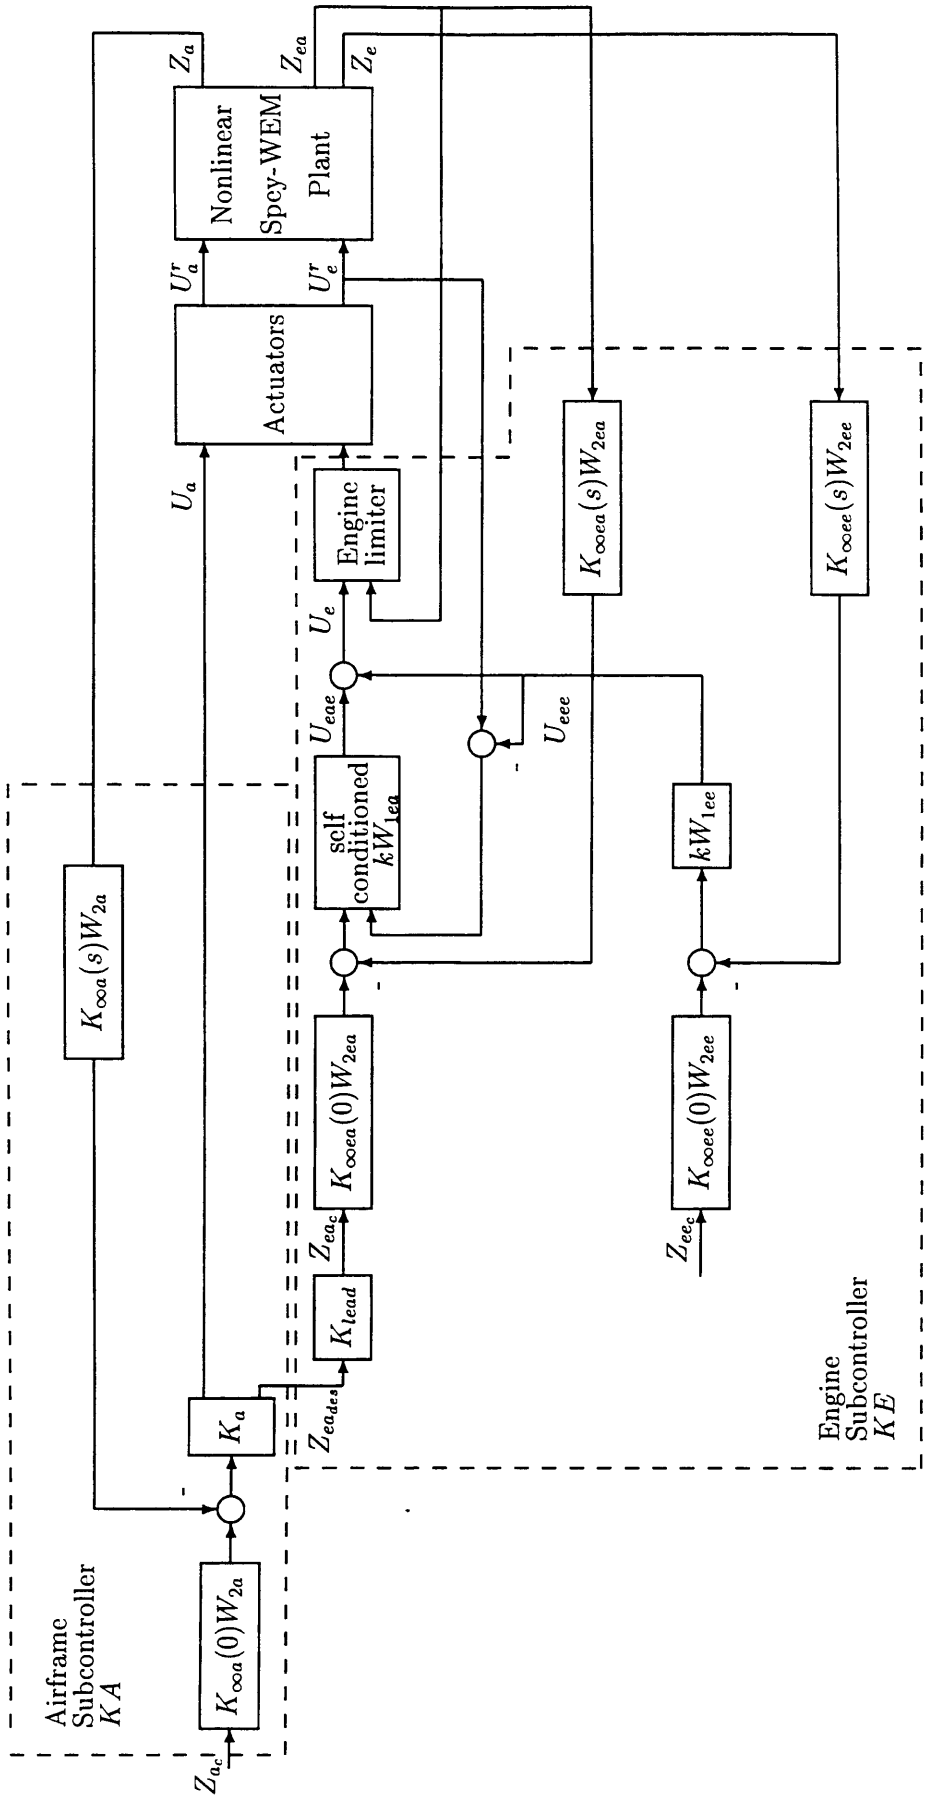

Figure 8.1: Partitioned IFPC system with engine limits and anti -windup.

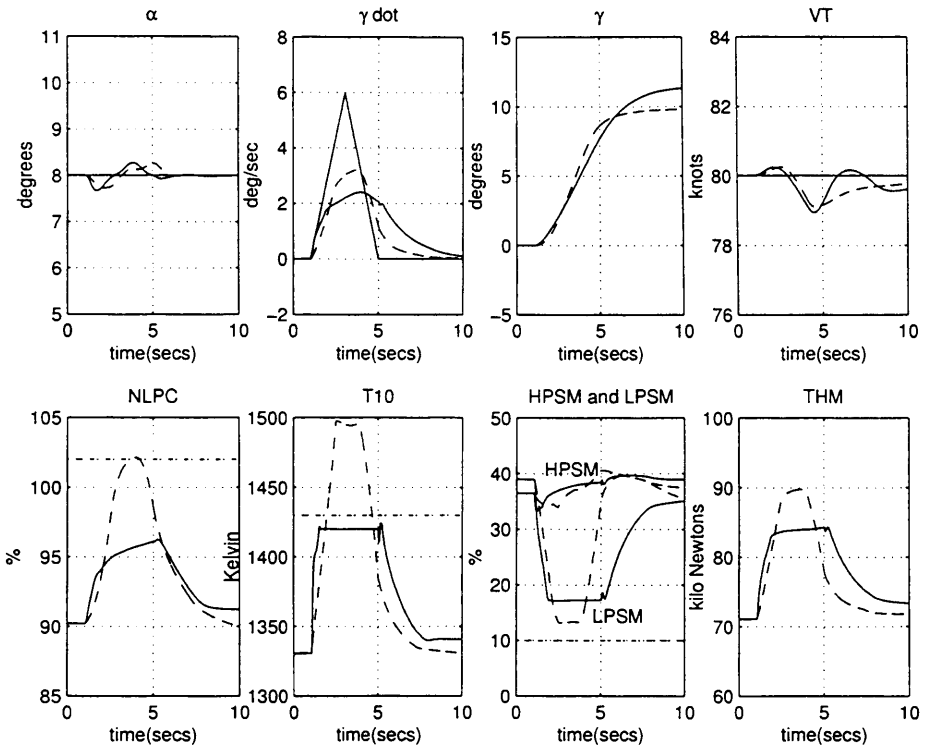

Figure 8.2: Airframe and engine responses to a demand on $\dot{\gamma}$ - general partitioned design with engine limits (—) and centralised design (- -) at 80 knots.

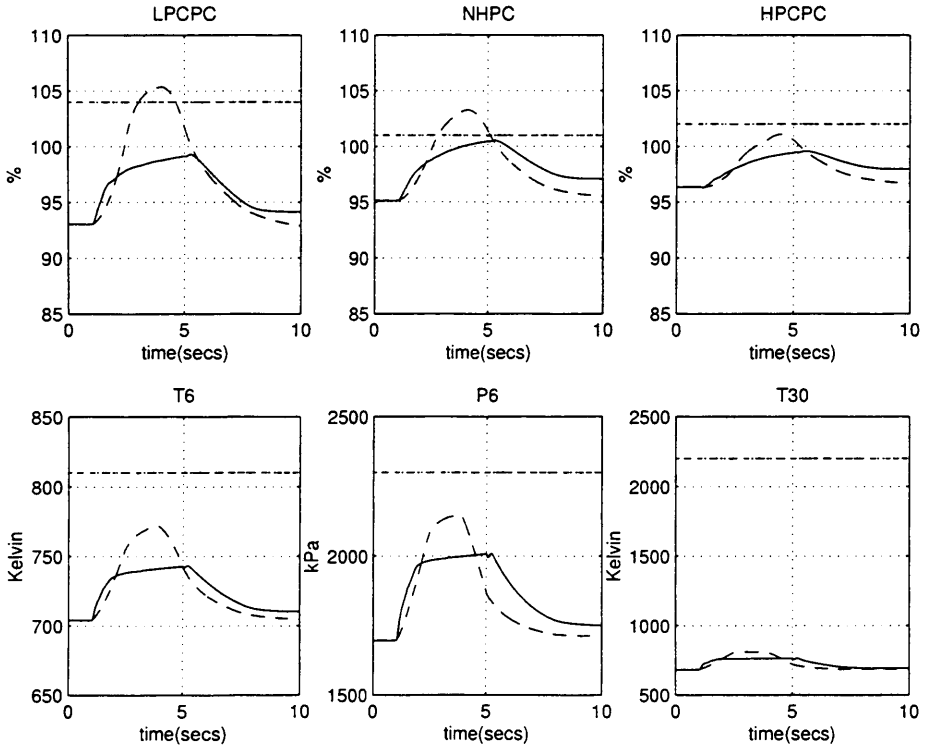

Figure 8.3: Internal engine variables responses to a demand on $\dot{\gamma}$ - general partitioned design with engine limits (—) and centralised design (- -) at 80 knots.

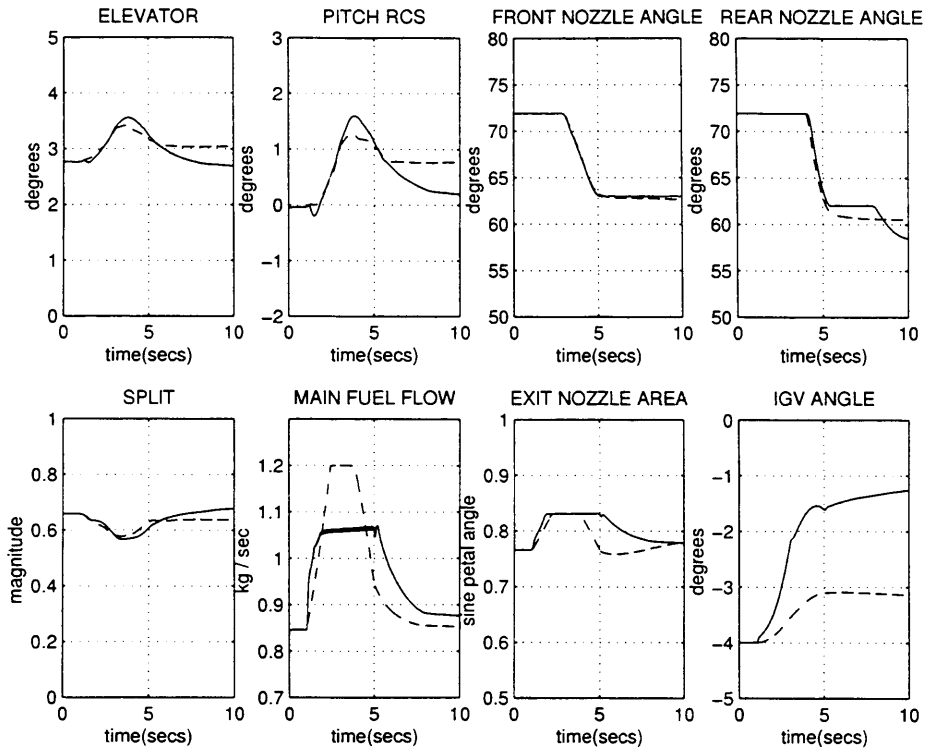

Figure 8.4: Actuator responses to a demand on $\dot{\gamma}$ - general partitioned design with engine limits (—) and centralised design (- -) at 80 knots.

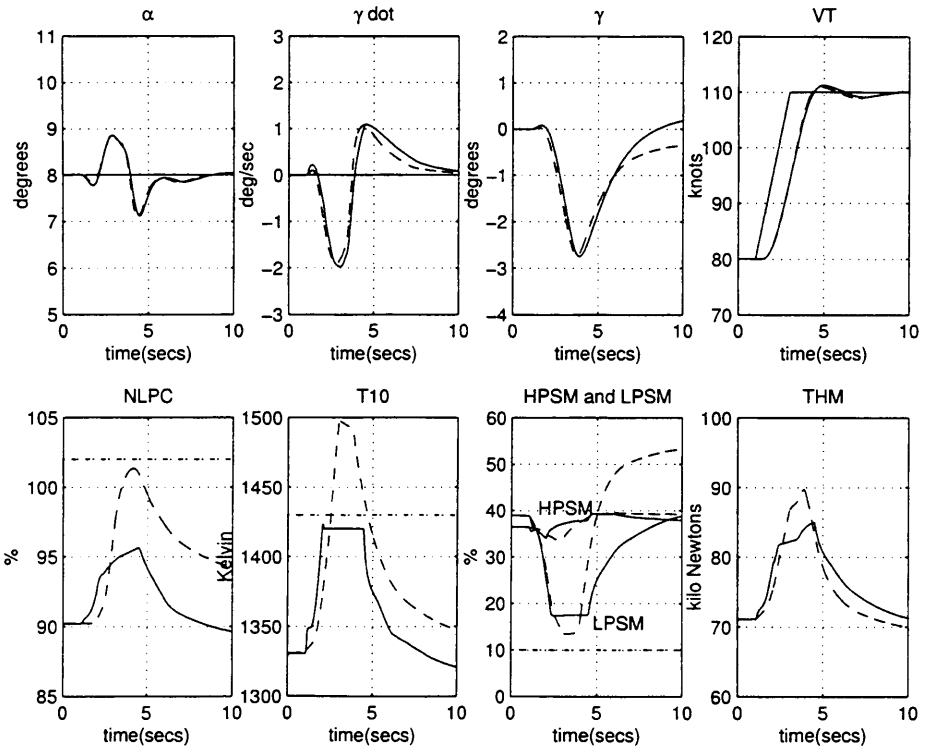

Figure 8.5: Airframe and engine responses to a demand on VT - general partitioned design with engine limits (—) and centralised design (- -) at 80 knots.

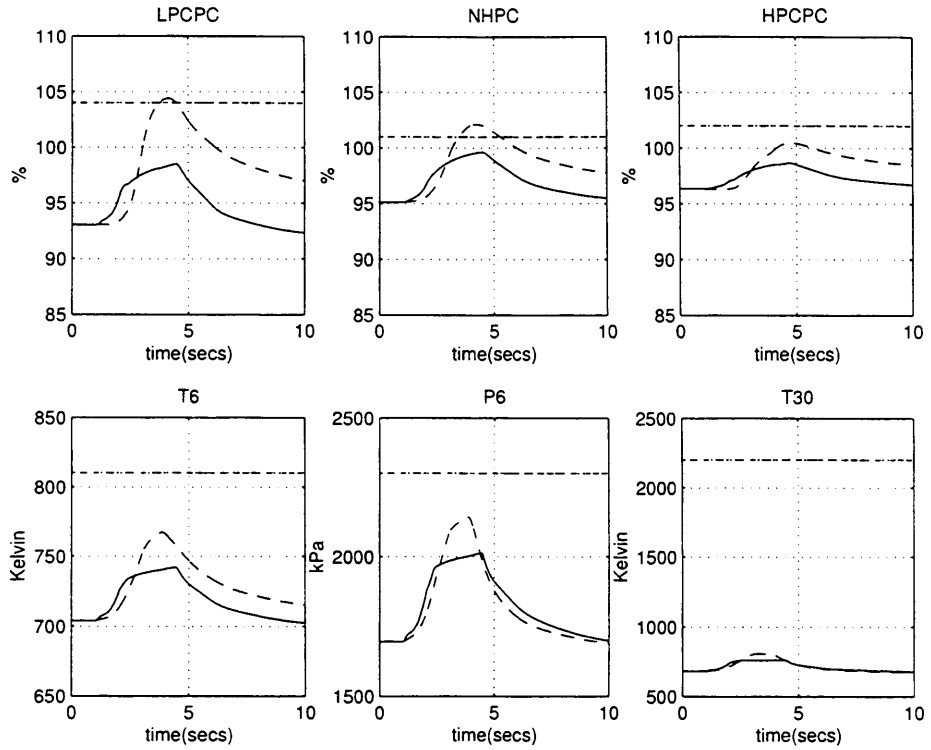

Figure 8.6: Internal engine variables responses to a demand on VT - general partitioned design with engine limits (—) and centralised design (- -) at 80 knots.

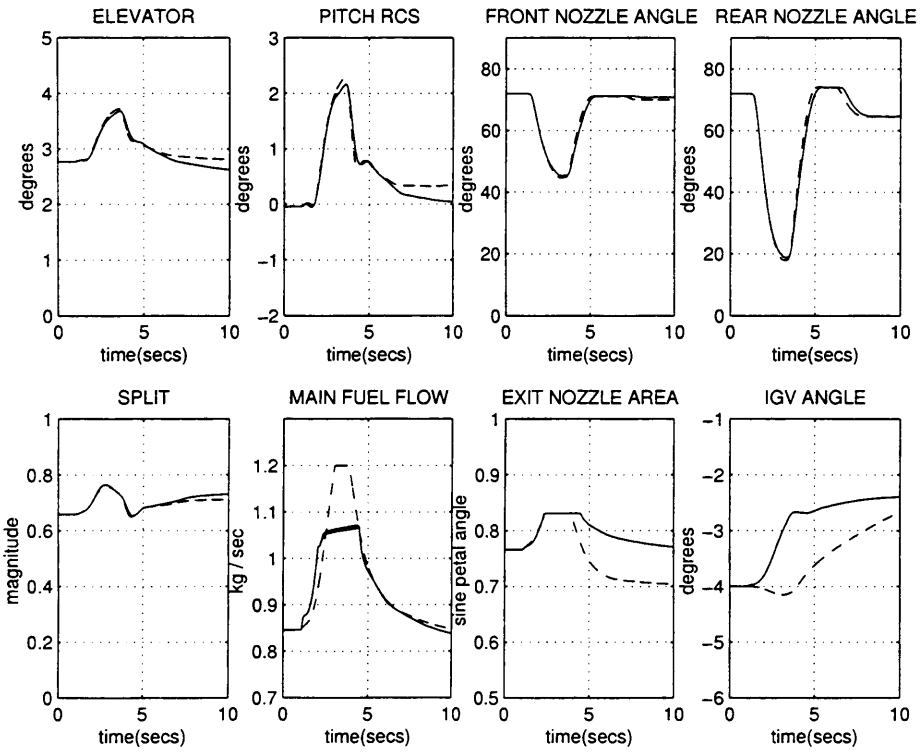

Figure 8.7: Actuator responses to a demand on VT - general partitioned design with engine limits (—) and centralised design (- -) at 80 knots.

demand on VT at 50 knots, where it can be seen that the VT response is very similar. Again, all the internal engine variables remain within their specified limits. Figures 8.12 to 8.15 show the same information for demands at 120 knots. While the γ response to a negative $\dot{\gamma}$ demand is slightly more sluggish for the partitioned limited system, all of the controlled outputs meet the specifications. The VT response for the partitioned limited system is slightly improved, with less overshoot, while the $\dot{\gamma}$ and α coupling is reduced.

8.3 Summary

An engine limiting scheme used to maintain internal engine variables within specified limits to ensure safe operation of the engine for the partitioned system has been introduced in this chapter. This is particularly flexible in that it accounts for differences in the engine due to, for example, manufacturing tolerances or internal wear. Some anti-windup is included, as the engine actuators often operate on limits, either physical or artificially imposed ones.

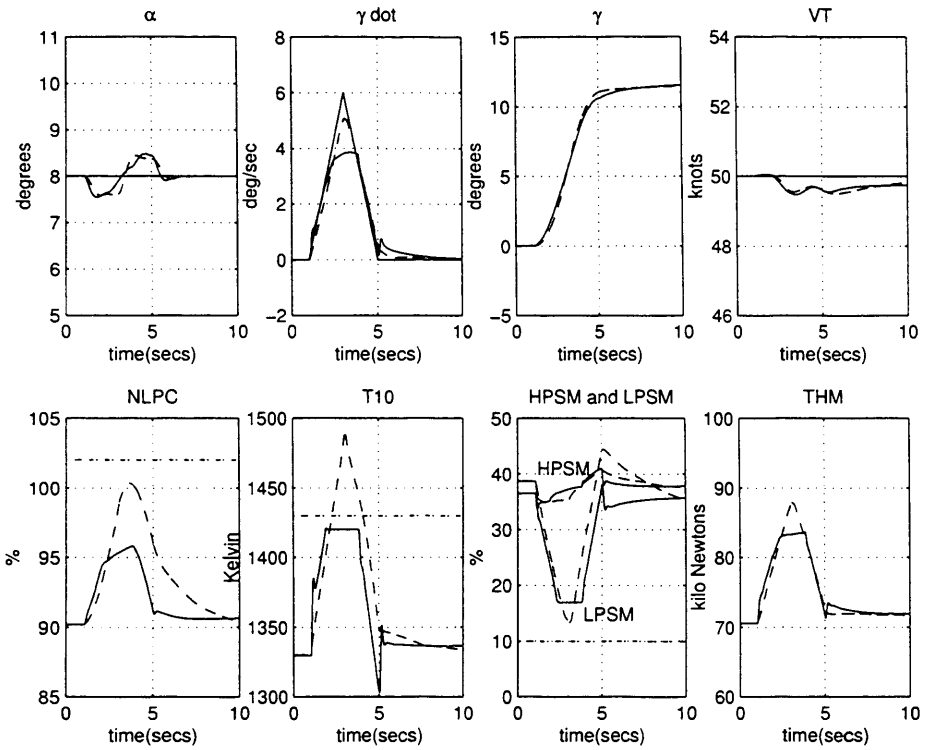

Figure 8.8: Airframe and engine responses to a demand on $\dot{\gamma}$ - general partitioned design with engine limits (—) and centralised design (- -) at 50 knots.

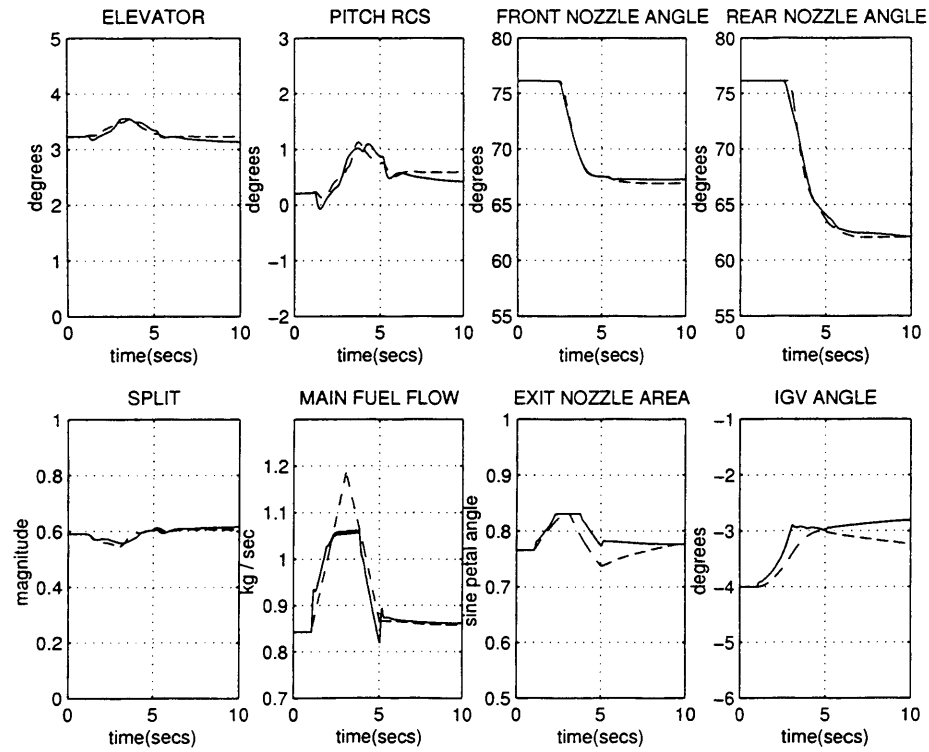

Figure 8.9: Actuator responses to a demand on $\dot{\gamma}$ - general partitioned design with engine limits (—) and centralised design (- -) at 50 knots.

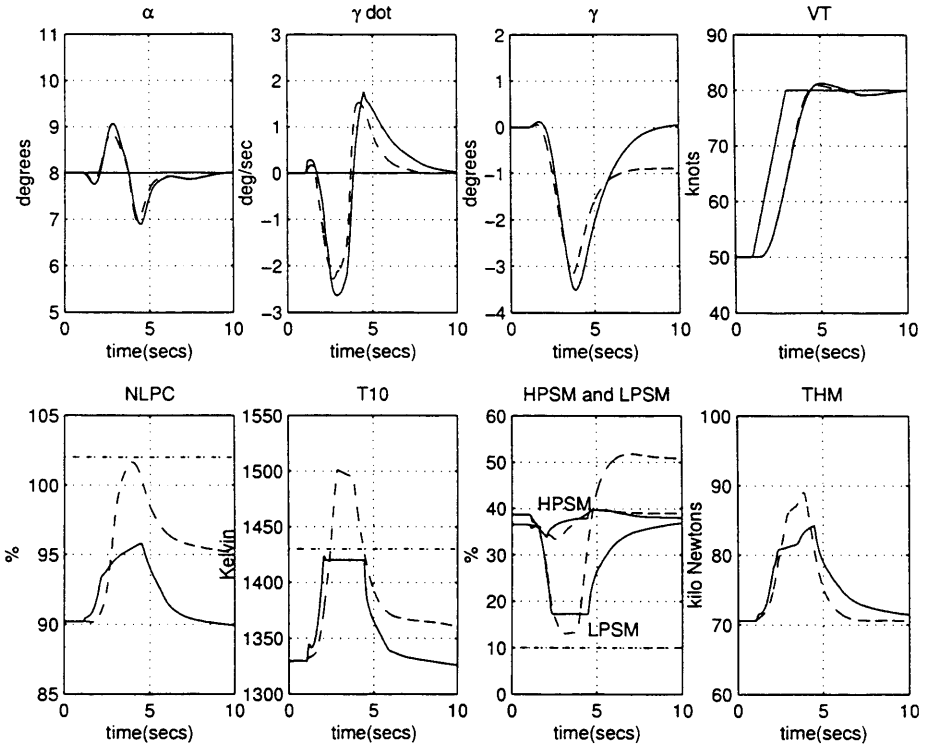

Figure 8.10: Airframe and engine responses to a demand on VT - general partitioned design with engine limits (—) and centralised design (- -) at 50 knots.

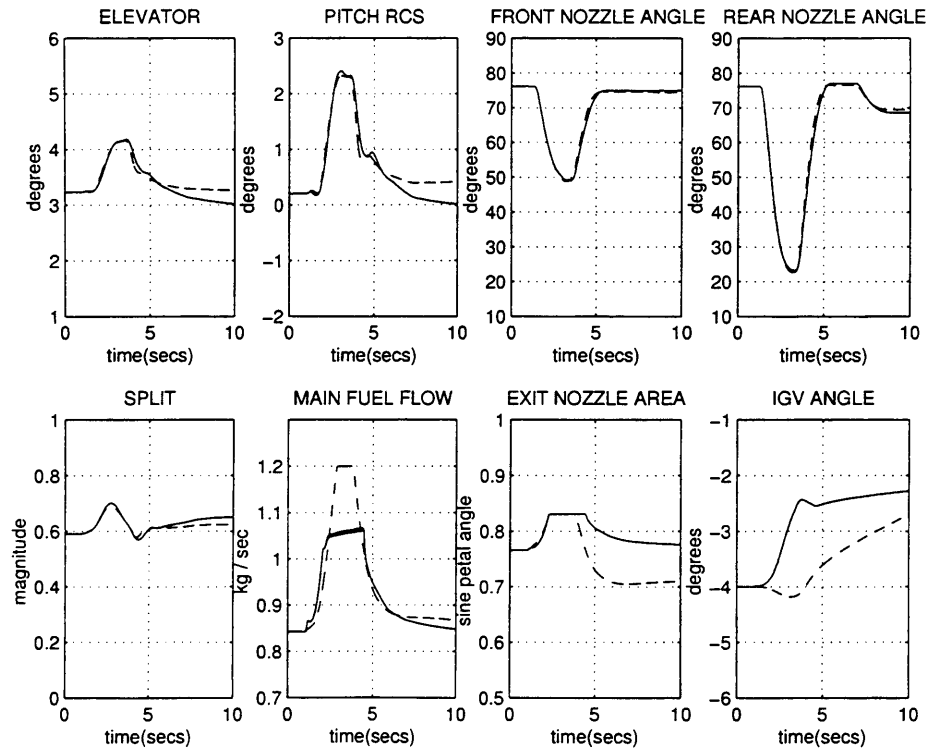

Figure 8.11: Actuator responses to a demand on VT - general partitioned design with engine limits (—) and centralised design (- -) at 50 knots.

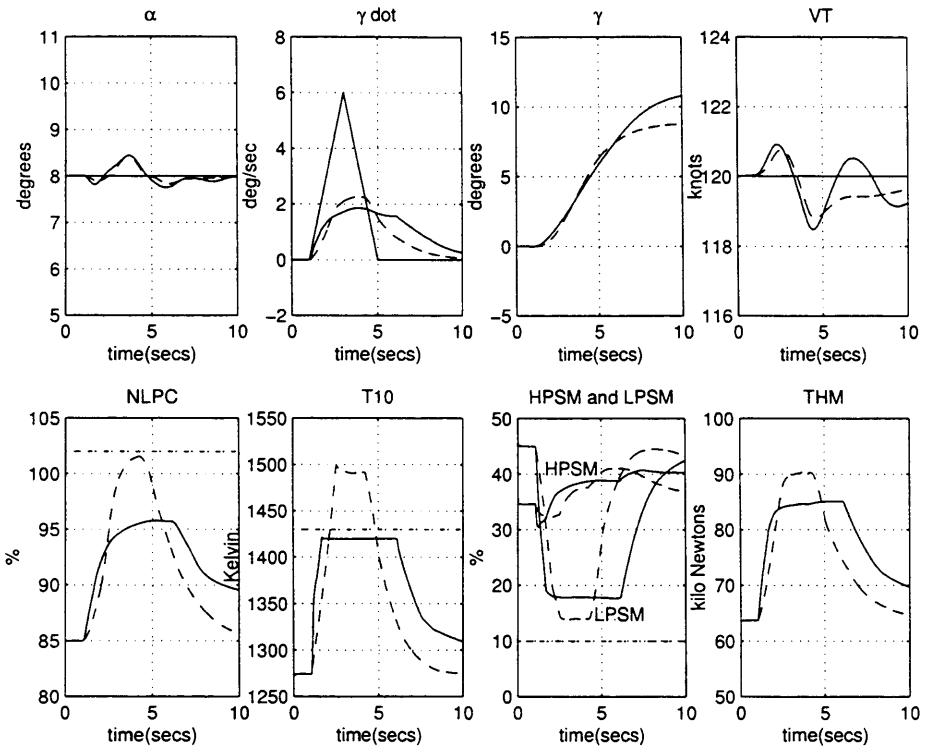

Figure 8.12: Airframe and engine responses to a demand on $\dot{\gamma}$ - general partitioned design with engine limits (—) and centralised design (- -) at 120 knots.

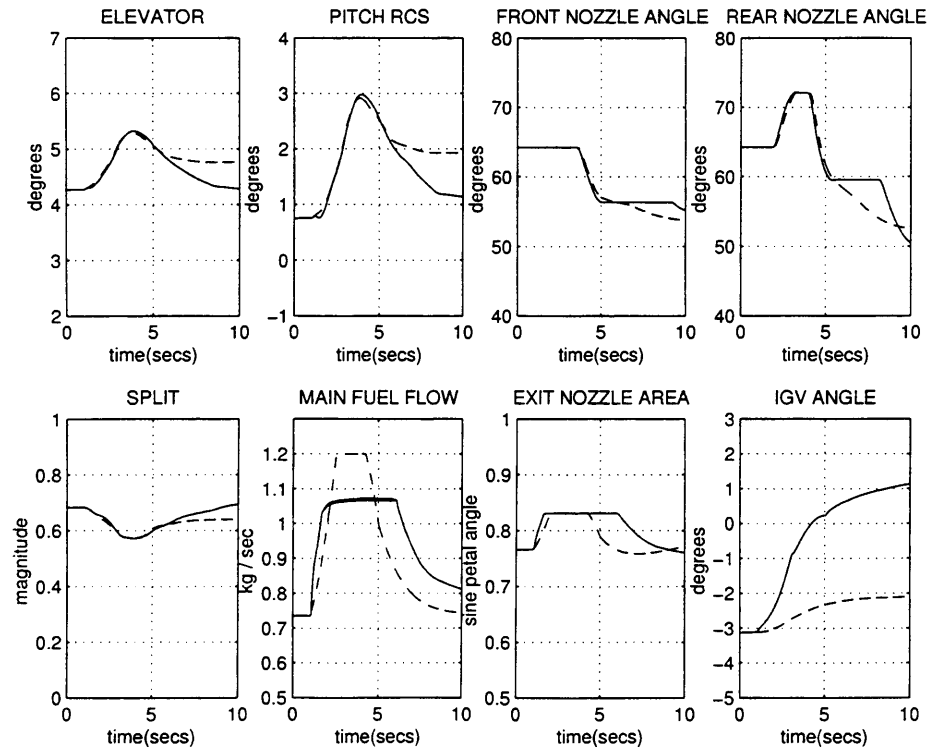

Figure 8.13: Actuator responses to a demand on $\dot{\gamma}$ - general partitioned design with engine limits (—) and centralised design (- -) at 120 knots.

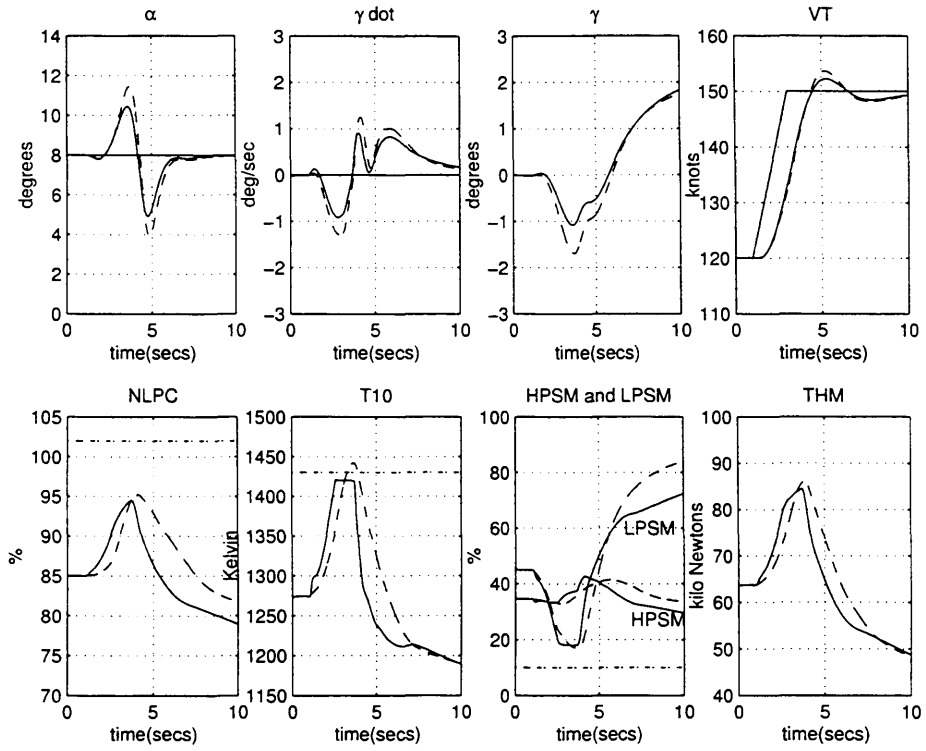

Figure 8.14: Airframe and engine responses to a demand on VT - general partitioned design with engine limits (—) and centralised design (- -) at 120 knots.

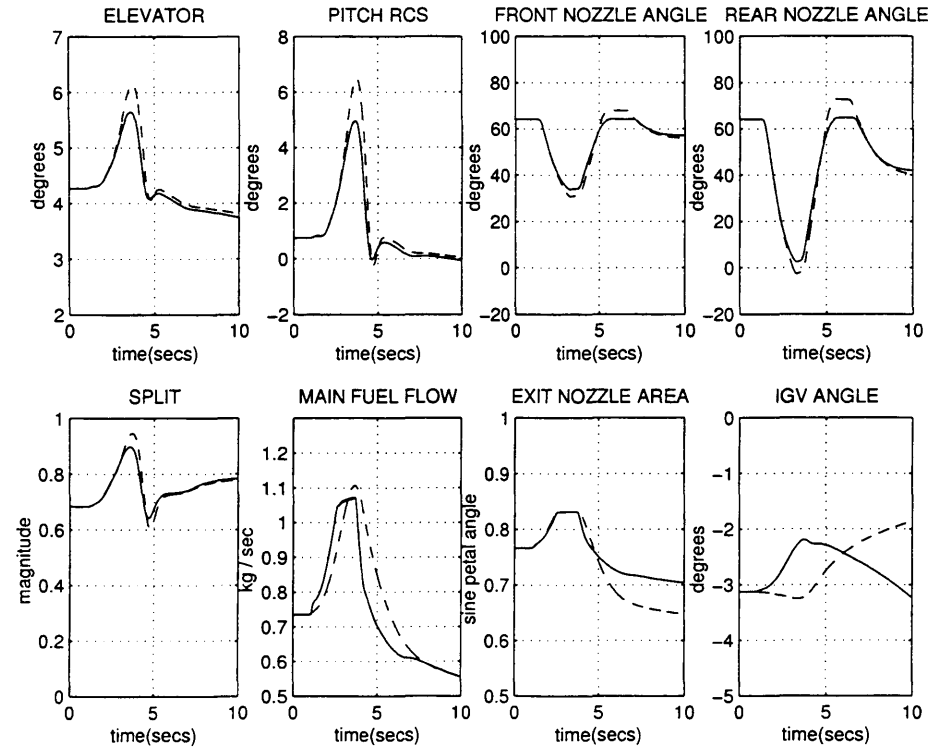

Figure 8.15: Actuator responses to a demand on VT - general partitioned design with engine limits (—) and centralised design (- -) at 120 knots.

Chapter 9

Conclusions and further work

This chapter will bring together the main conclusions from this thesis and will introduce further work required to expand the controlled flight envelope and to address some of the difficulties experienced.

9.1 Conclusions

This thesis has presented the design and analysis of an integrated flight and propulsion controller using the Spey-WEM aircraft model supplied by DERA Bedford. The nominal design point was chosen to be in the transition to hover region at 80 knots. At this point in the flight envelope, the propulsion system generated forces and moments have largely taken over those generated by the airframe and the aircraft is longitudinally unstable. The main conclusions are as follows:

- The designed centralised controller includes both airframe and engine variables as controlled outputs. This is in order to address the engine safety issues at the design stage. Output weighting is selected to ensure the internal engine controlled outputs are controlled, but not to the extent that the airframe responses are degraded. This alone was not expected to be sufficient to maintain the engine within specified limits at all times, and thus an additional nonlinear engine limiting scheme was included.

- The centralised controller was tested at DERA Bedford in piloted simulation using the real-time all-vehicle simulator. Overall, level 1 or level 2 handling qualities were achieved, indicating at best excellent, highly desirable characteristics, where pilot compensation is not a factor for desired performance to, at worst, minor but annoying deficiencies, where desired performance requires only moderate pilot compensation.

- A thrust vectoring controller was also designed which did not use the elevator or the pitch reaction control system. The longitudinal airframe changes are therefore dependent solely

on changes in thrust magnitude and thrust direction. This type of configuration may be useful in the event of actuator failure, indicating that safe operation of the aircraft would still be attainable.

- μ-analysis was shown to be a powerful tool for calculating at which points in the flight envelope stability may no longer be guaranteed. Perturbations to the nominal linear model were considered for both parameter variations, such as mass and inertia, and due to variations in the airframe and engine dynamics at different points in the flight envelope. When considering changes due to the position in the flight envelope, a 'one-sided' μ problem was considered, whereby, for example, the changes from 80 knots to 120 knots (80 knots + 40 knots) could be analysed rather than for 80 knots $\pm$40 knots.

- The ν-gap analysis was shown to provide much useful information about the robustness properties of the IFPC system. This is a natural companion to the $\mathcal{H}_\infty$ loop shaping design procedure. Using the ν-gap metric, it is possible to consider the closed-loop differences between plants across frequency, leading to less conservative results than would be found using, for example, the gap-metric. Thus, although the value of δ_ν between two plants may be greater than the value of $B_{P,K}$ calculated for the design plant and controller, this may occur at different frequencies. Therefore, if $\psi\left(P\left(j\omega\right), P2\left(j\omega\right)\right)$ is less than $b_{P,K}\left(\omega\right)$ when considered on a frequency by frequency basis, the nominal controller will in fact stabilise both plants.

- The partitioning method is based on that of [27], but has been adapted to consider the particular implementation structure of the $\mathcal{H}_\infty$ loop shaping design method. This method successfully retained the performance properties of the centralised controller. However, the robustness properties showed a dramatic degradation. This was addressed by redesigning the engine subcontroller to have a more intuitive structure, resulting in both the performance and robustness largely matching that of the centralised system.

- The additional nonlinear engine safety limiting scheme accounts for the multivariable nature of the engine, whereby many combinations of actuator inputs may result in equivalent internal conditions such as temperatures and speeds. Furthermore, the scheme allows for differences in the engine due to, for example, manufacturing tolerances or engine wear. The controlled engine variable $T10$ is monitored constantly. If it reaches within 10K of its maximum allowable temperature, the positions of the engine actuators are noted. This is because it is known that this particular combination of actuator positions will keep $T10$ below its limit. If $T10$ continues to increase, these noted actuator positions are used as temporary artificial position limits. When $T10$ drops below its trigger point, the

actuators are released.

- The partitioned scheme with engine limits will frequently require the actuators to be used on physical or artificial position limits. Operating on limits such as these is known to instigate integrator windup, resulting in degraded performance. This is avoided by including an anti-windup strategy on the interface part of the engine subcontroller.

9.2 Further work

Lateral/directional Control.

At present, only work in the longitudinal control has been investigated. The design of a lateral/directional controller will increase the controlled flight envelope.

Simultaneous $\dot{\gamma}$ and VT demand.

At the moment, if a simultaneous demand is made on both $\dot{\gamma}$ and VT, the VT demand, rather than the $\dot{\gamma}$ demand is met. Therefore, some form of priority setting is required.

Scheduling.

Possibly the most obvious extension of this work is to consider some form of controller scheduling in order to extend control across the flight envelope. One possibility that has been considered has been through varying the elements of the scaling weighting matrix k according to the effectiveness of each actuator at each point in the flight envelope. This would provide a simple, transparent approach which would simplify the implementation and validation of the overall system.

Engine limiting scheme.

The scheme applied in this thesis is not sufficient to guarantee all engine variables will remain within their specified limits at all times, although it does work adequately well in the nonlinear simulations tried so far.

Backlash in nozzles.

More consideration needs to be given to the problem of backlash. 6° of backlash is modelled in both the front and rear nozzles. This can be seen to cause limit cycling problems even when there is no demand made on the system. This is due to the fact that the nonlinear model will not be perfectly trimmed and so will tend to drift. However, due to the backlash in the front and rear nozzles, the controller is not able to compensate for this immediately, depending on where the actuator is initially within the backlash. The overall result is a limit cycle seen on

all of the actuators and so all of the outputs.

Efficient use of actuators.

It has been noted that the control scheme does not always select the most efficient combination of actuators. An example of this has been noted whereby the front and rear nozzles may be re-trimmed following a demand in such a way as to oppose each other. This may be a disadvantage of using the $\mathcal{H}_\infty$ loop shaping design procedure, as it is not possible to specify minimum actuator usage.

References

[1] *Simulink: Dynamic Simulation for MATLAB users manual.* The MathWorks Ltd., Natick, MA, 1997.

[2] Special session on Gas Turbine Testing and Modelling. *Proceedings of the UKACC International Conference on CONTROL '98*, 1998.

[3] R. D. Archer and M. Saarlas. *An introduction to aerospace propulsion.* Prentice Hall, New Jersey, 1996.

[4] G. J. Balas, J. C. Doyle, K. Glover, A. Packard, and R. Smith. *The μ-Analysis and Synthesis Toolbox.* MathWorks Inc., MA, USA, 1995.

[5] D. G. Bates, S. L. Gatley, and I. Postlethwaite. Partitioning of $\mathcal{H}_\infty$ loopshaping Integrated Flight and Propulsion Control Systems. *submitted to IEEE Transactions of Control Systems Technology.*

[6] D. G. Bates, S. L. Gatley, I. Postlethwaite, and A. J. Berry. Integrated Flight and Propulsion Control System design using $\mathcal{H}_\infty$ loop-shaping techniques. *Proceedings of the 38th IEEE Conference on Decision and Control (Phoenix, Arizona)*, pages 1523–1528, December 1999.

[7] D. G. Bates, S. L. Gatley, I. Postlethwaite, and A. J. Berry. Design and piloted simulation of a robust Integrated Flight and Propulsion Controller. *AIAA Journal of Guidance, Control and Dynamics*, 23(2):269–277, 2000.

[8] D. G. Bates, S. L. Gatley, I. Postlethwaite, L. Gillett, and D. Littleboy. Robust Integrated Flight and Propulsion Control of an experimental STOVL aircraft. *Proceedings of the 37th IEEE Conference on Decision and Control (Tampa, FA)*, pages 4220–4221, 1998.

[9] D. G. Bates and I. Postlethwaite. Robust Integrated Flight and Propulsion Control (IFPC) for the VAAC Harrier STOVL aircraft. *Proceedings of the 1998 International Conference on Control (Swansea, UK)*, (Part 2 (of 2)):1516–1521, September 1-4 1998.

[10] D. G. Bates and I. Postlethwaite. Robust Integrated Flight and Propulsion Control for the Harrier aircraft. *Journal of Guidance, Control and Dynamics*, 22(2):286–290, March-April 1999.

[11] D. Bennett, S. E. Burge, and A. Bradshaw. Design of a controller for a highly coupled V/STOL aircraft. *Transactions of the Institute of Measurement and Control*, (2/3):63–75, 1999.

[12] A. J. Berry. Leicester University IFPCS control law requirements document. DERA Technical Memorandum, 1998.

[13] P. J. Campo, M. Morari, and C. N. Nett. Multivariable anti-windup and bumpless transfer: A general theory. *Proceedings of the American Control Conference (Pittsburgh, USA)*, pages 1706–1711, 1989.

[14] M. Cantoni and G. Vinnicombe. Quantifying uncertainty and robust performance using the ν-gap metric. *AIAA*, (AIAA-99-4277):1624–1630, 1999.

[15] Z. Chen and P. G. Voulgaris. Decentralized design for Integrated Flight/Propulsion Control of aircraft. *Proceedings of the AIAA Conference on Guidance, Navigation and Control*, (AIAA-98-4504):1866–1872, 1998.

[16] M. V. Cook. On the design of command and stability augmentation systems for advanced technology aeroplanes. *Transactions of the Institute of Measurement and Control*, 21(2/3):85–98, 1999.

[17] G. E. Cooper and R. P. Harper Jr. The use of pilot rating in the evaluation of aircraft handing qualities. *NASA TN D-5152*, April 1969.

[18] G. J. Dadd, A. E. Sutton, and A. W. M. Greig. Multivariable control of military engines. *Proceedings of AGARD Conference No. 572 - Advanced Aero-Engine Concepts and Controls*, (28):1–12, 1996.

[19] R. Dailey. A new algorithm for the real structured singular value. *Proceedings of the American Control Conference (San Diego, USA)*, pages 3036–3040, 1990.

[20] C. Edwards and I. Postlethwaite. Anti-windup and bumpless transfer schemes. *Automatica*, 34(2):199–210, 1999.

[21] A. Emami-Naeini, R. P. Anex, and S. M. Rock. Integrated Control: A Decentralised Approach. *Proceedings of the* 24^{th} *Conference on Decision and Control (Ft. Lauderdale, U.S.A)*, pages 1836–1841, 1985.

[22] G. Ferreres. *A Practical approach to robustness analysis with aeronautical applications.* Kluwer Academic, New York, 1999.

[23] G. Ferreres and J. M. Biannic. A reliable computation of the robustness margin for a flexible transport aircraft. *Proceedings of the AIAA Conference on Guidance, Navigation and Control*, (4365), 2000.

[24] I. Fialho, G. Balas, A. Packard, J. Renfrow, and C. Mullaney. Linear fractional transformation control of the F-14 aircraft lateral-directional axis during powered approach landing. *Proceedings of the American Control Conference*, pages 128–132, 1997.

[25] James A. Franklin. Control of V/STOL aircraft. *Aeronautical Journal*, pages 157–173, May 1986.

[26] G. W. Gallops, C. F. Weiss, and R. A. Carlton. Integrated propulsion system requirements for control of STOVL aircraft. *Transactions of the ASME. Journal of Engineering for Gas Turbines and Power*, pages 60–67, January 1990.

[27] S. Garg. Partitioning of centralised Integrated Flight/Propulsion Control design for decentralised implementation. *IEEE Transactions on Control Systems Technology*, 1(2):93–100, 1993.

[28] S. Garg and D. Mattern. Application of an integrated methodology for propulsion and airframe control design to a STOVL aircraft. *Proceedings of the AIAA Conference on Guidance, Navigation and Control (Scottsdale, USA)*, pages 591–605, 1994.

[29] S. Garg, D. L. Mattern, M. Bright, and P. Ouzts. IMPAC - an Integrated Methodology for Propulsion and Aircraft Control. *Proceedings of the American Control Conference (Boston, USA)*, pages 747–754, June 1991.

[30] S. L. Gatley, D. G. Bates, M. J. Hayes, and I. Postlethwaite. Robustness analysis of an Integrated Flight and Propulsion Control System using μ and the ν-gap metric. *To appear the IFAC Journal of Control Engineering Practice.*

[31] S. L. Gatley, D. G. Bates, M. J. Hayes, and I. Postlethwaite. Robustness analysis of integrated flight and propulsion control systems for V/STOL aircraft. *To appear in the European Control Conference (Porto, Portugal)*, 2001.

[32] S. L. Gatley, D. G. Bates, and I. Postlethwaite. A partitioned Integrated Flight and Propulsion Control System with engine safety limiting. *IFAC Journal of Control Engineering Practice*, pages 845–859, 2000.

[33] S. L. Gatley, D. G. Bates, and I. Postlethwaite. Robust control of a V/STOL aircraft in transition flight using differential thrust vectoring. *Proceedings of UKACC International Conference on Control (Cambridge, UK)*, 2000.

[34] S. L. Gatley, D. G. Bates, and I. Postlethwaite. An engine limiting and anti-windup scheme for a partitioned IFPC system. *AIAA Conference on Guidance, Navigation and control (Montreal)*, (AIAA 2001-4385), August 2001.

[35] S. L. Gatley, D. G. Bates, and I. Postlethwaite. Partitioning and re-design of $\mathcal{H}_\infty$ loop-shaping integrated flight and propulsion control systems. *AIAA Conference on Guidance, Navigation and Control (Montreal)*, (AIAA 2001-4384), August 2001.

[36] M. J. Grimble, S. A. Carr, and M. R. Katebi. Integrated ship control using $\mathcal{H}_\infty$ robust design techniques. *IEEE Conference on Control Applications*, pages 1087–1091, 1994.

[37] R. Hanus, M. Kinnaert, and J.-L. Henrotte. Conditioning technique, a general anti-windup and bumpless transfer method. *Automatica*, 23(6):729–739, 1987.

[38] M. Harefors. Applications of $\mathcal{H}_\infty$ robust control to the RM12 jet engine. *Control Engineering Practice*, 5(9):1189–1201, 1997.

[39] R. T. C. Harman. *Gas Turbine Engineering - Applications, Cycles and Characteristics.* The Macmillan Press LTD, Hong Kong, 1981.

[40] M. J. Hayes, D. G. Bates, S. L. Gatley, and I. Postlethwaite. Stability robustness analysis of an Integrated Flight and Propulsion Controller using the structured singular value. *Proceedings of the 38th IEEE Conference on Decision and Control (Phoenix, USA)*, 5:4529–4530, 1999.

[41] M. J. Hayes, D. G. Bates, S. L. Gatley, and I. Postlethwaite. Performance robustness analysis of an Integrated Flight and Propulsion Controller using the structured singular value. *Proceedings of the UKACC International Conference on Control (Cambridge, UK)*, 2000.

[42] M. J. Hayes, D. G. Bates, and I. Postlethwaite. New tools for computing tight bounds on the real structured singular value. *to appear in AIAA Journal of Guidance, Navigation and Control.*

[43] R. H. Hoh and D. G. Mitchell. *Advances in Aircraft Flight Control*, chapter 1: Handling-qualities specification - a functional requirement for the flight control system, pages 3–33. Taylor and Francis, London, 1996.

[44] R. A. Hyde. *$\mathcal{H}_\infty$ Aerospace Control Design - A VSTOL Flight Application.* Springer-Verlag, London, 1995.

[45] A. C. Kermode. *Mechanics of Flight.* Addison Wesley Longman Limited, Essex, 10^{th} edition, 1996.

[46] L. H. Lee and M. Sillman. Control of slowly varying LPV systems: an application to flight control. *Proceedings of the AIAA Conference on Guidance, Navigation and Control (San Diego, USA)*, 1996.

[47] D. Mattern and S. Garg. Propulsion system performance resulting from an integrated flight/propulsion control design. *Proceedings of the AIAA Guidance, Navigation and Control conference*, (AIAA-92-4602-CP):1314–1323, August 1992.

[48] D. Mattern, S. Garg, and R. Bullard. Integrated Flight/Propulsion Control system design based on a decentralised, hierarchical approach. *Proceedings of the AIAA Guidance, Navigation and Control conference (Boston, MA)*, pages 1–12, August 1989.

[49] D. Mattern and P. Ouzts. A linear control design structure to maintain loop properties during limit operation in a multi-nozzle turbofan engine. *AIAA 27th Joint Propulsion Conference (Sacramento, CA)*, (AIAA-91-1997), June 1991.

[50] D. McFarlane and K. Glover. A loop shaping design procedure using $\mathcal{H}_\infty$ synthesis. *IEEE Transactions on Automatic Control*, 36:759–769, 1992.

[51] J. R. Mihaloew. Flight propulsion control integration for V/STOL aircraft. *AIAA*, (872330):303–315, 1987.

[52] B. Morton and R. McAfoos. μ-test for robustness analysis of a real-parameter variation problem. *Proceedings of the American Control Conference (Boston, USA)*, pages 135–138, 1985.

[53] L. P. Myers and K. R. Walsh. Performance improvements of an F-15 airplane with an integrated engine-flight control system. *Journal of Aircraft*, 28(12):812–817, 1991.

[54] J. S. Orme and T. R. Conners. Supersonic flight test results of a performance seeking control algorithm on a NASA F-15 aircraft. *30^{th} AIAA/ASME/SAE/ASEE Joint Propulsion Conference (Indianapolis, IN)*, (AIAA 94-3210), June 1994.

[55] A. Packard and J. Doyle. The complex structured singular value. *Automatica*, 29(1):71–109, 1993.

[56] A. Packard and P. Pandey. Continuity properties of the real/complex structured singular value. *IEEE Transactions on Automatic Control*, 38:415–428, 1993.

[57] B. N. Pamadi. *Performance, Stability, Dynamics, and Control of Airplanes.* American Institute of Aeronautics and Astronautics, Inc, Virginia, 1998.

[58] I. Postlethwaite and D. G. Bates. Integrated Flight and Propulsion Control System design. *IEE Colloquium on Integrated Systems in Aerospace (London, UK)*, (Digest No.1997/015):3/1–3, 1997.

[59] I. Postlethwaite, R. Samar, B. Choi, and D-W. Gu. A digital multi-mode $\mathcal{H}_\infty$ controller for the Spey turbofan engine. *Proceedings of the 3rd European Control Conference (Rome, Italy)*, 4:3881–3886, 1995.

[60] S. M. Rock. Integrated Flight/Propulsion Control : Requirements and issues. *IFAC Automatic Control in Aerospace (Tsukuba, Japan)*, pages 209–214, 1989.

[61] S.M. Rock, A. Enami-Naeini, and R.P. Anex. Propulsion control specifications in Integrated Flight/Propulsion Control systems. *AIAA*, pages 1–11, 1988.

[62] S.M. Rock, A. Enami-Naeini, and K. Neighbors. Integrated Flight/Propulsion Control: Subsystem specifications. *Journal of Guidance, Control and Dynamics*, 17(1):201–208, January-February 1994.

[63] R. Samar. *Robust multi-mode control of high performance aero-engines.* PhD thesis, University of Leicester, Department of Engineering, 1995.

[64] D. K. Schmidt. Integrated control of hypersonic vehicles - a necessity not just a possibility. *AIAA-93-3761-CP*, 1993.

[65] P. K. Schmidt, S. Garg, and B. Holowecky. A parameter optimisation approach to controller partitioning for Integrated Flight/Propulsion Control. *IEEE Transactions on Control Systems Technology*, 1(1):21–36, 1993.

[66] J. A. Sefton and K. Glover. Pole/zeros cancellations in the general $\mathcal{H}_\infty$ problem with reference to a two block design. *System and Control Letters*, 14:295–306, 1990.

[67] J. S. Shamma and J. R. Cloutier. Gain-scheduled missile autopilot design using linear parameter varying transformations. *AIAA Journal of Guidance, Dynamics and Control*, 16(2):256–263, 1993.

[68] G. T. Shanks, C. Fielding, S. J. Andrews, and R. A. Hyde. Flight control and handling research with the VAAC Harrier aircraft. *International Journal of Control*, 59(1):291–319, 1994.

[69] G. T. Shanks, S. L. Gale, C. Fielding, and D. V. Griffith. *Advances in Aircraft Flight Control*, chapter 6: Flight control and handling research with the VAAC Harrier aircraft, pages 159–186. Taylor and Francis, London, 1996.

[70] S. Skogestad and I. Postlethwaite. *Multivariable Feedback Control.* John Wiley, London, 1996.

[71] K. L. Smith, W B. Kerr, G. L. Hartmann, and C. Skira. Integrated flight/propulsion control - methodology, design and evaluation. *AIAA*, 1985.

[72] J. F. Stewart, F. W. Burcham Jr, and D. H. Gatlin. Flight-determined benefits of Integrated Flight-Propulsion Control systems. 18^{th} *ICAS Congress (Beijing, China)*, 2:1756–1777, September 1992.

[73] K. D. Tillman and T. J. Ikeler. Integrated Flight/Propulsion Control for flight critical applications: A propulsion system perspective. *Journal of Engineering for Gas Turbines and Power*, 114:755–762, October 1992.

[74] T. Trones. http://members.home.net/digitalviking/airpower.htm.

[75] A. Varga, G. Looye, D. Moormann, and G. Grubel. Automated generation of LFT-based parametric uncertainty descriptions from generic aircraft models. *GARTEUR TP-088-36*, 1997.

[76] G. Vinnicombe. Frequency domain uncertainty in the graph topology. *IEEE Transactions on Automatic Control*, 28(9):1371–1383, 1993.

[77] G. Vinnicombe. *Uncertainty and Feedback.* Imperial College Press, 2000.

[78] K. Zhou with J.C. Doyle. *Essentials of Robust Control.* Prentice-Hall, London, 1998.

[79] G. D. Wood and J. Woodhouse. An investigation into robust force control through a constrained flexible beam. *International Journal of Control*, 68(3):539–578, 1997.

[80] S. F. Yaros, M. G. Sexstone, L. D. Huebner, J. E. Lamar, R. E. McKinley Jr, A. O. Torres, C. L. Burley, R. C. Scott, and W. J. Small. Synergistic airframe-propulsion interactions and integrations. *Langley Research Center, Hampton, Virginia*, (NASA/TM-19998-207644).

[81] P. M. Young, M. P. Newlin, and J. C. Doyle. Computing the bounds for the mixed μ problem. *International Journal of Robust and Nonlinear Control*, 5:573–590, 1995.

[82] G. Zames and A. K. El-Sakkary. Unstable systems and feedback: The gap metric. *Proceedings of the Allerton Conference*, pages 380–385, 1980.

[83] K. Zhou, J. C. Doyle, and K. Glover. *Robust and Optimal Control.* Prentice-Hall, New Jersey, 1996.