TUNING HYBRID INCREMENTAL DYNAMIC INVERSION CONTROL LAWS USING H_∞ LOOP-SHAPING

LEONARDO ENCARNAÇÃO | Master's Thesis in Aerospace Engineering

TU Delft

Tuning Hybrid Incremental Dynamic Inversion Control Laws using H_∞ Loop-Shaping

Thesis Report

by

Leonardo Encarnação

to obtain the degree of Master of Science
at the Delft University of Technology
to be defended publicly on May 8th, 2025 at 09:00

Thesis committee:	
Chair:	Dr. ir. Erik-Jan van Kampen
Supervisors:	Dr. Eng. Spilios Theodoulis
	Dr. ir. Tijmen Pollack
External examiner:	Dr. ir. Xuerui Wang
Additional examiner:	Dr. ir. Gertjan Looye
Place:	Faculty of Aerospace Engineering, Delft
Project Duration:	March, 2024 - May, 2025
Student number:	5845157

Cover image: illustration of the Grumman X-29 experimental aircraft by Vasco Viegas, 2025.

An electronic version of this thesis is available at https://repository.tudelft.nl/.

Faculty of Aerospace Engineering · Delft University of Technology

Abstract

Nonlinear Dynamic Inversion (NDI) control techniques provide a conceptually simple and modular control framework, making it an attractive technique for designing flight control laws with shorter design cycles. However, its lack of inherent robustness guarantees shifts the burden of the design from the synthesis to the analysis part. Conversely, $\mathcal{H}_\infty$ Loop-Shaping provides controllers with robust stability guarantees. This work proposes a novel framework leveraging the $\mathcal{H}_\infty$ Loop-Shaping Design Procedure to optimise a structured linear variant of Incremental Nonlinear Dynamic Inversion (INDI) control, a Hybrid IDI controller. The Hybrid IDI controller consists of a blend between classical model-based DI and sensor-based IDI. The proposed methodology is validated through the design of a pitch-rate controller for NASA's X-29 experimental aircraft. Results demonstrate that the approach achieves robustness guarantees comparable to standard full-order $\mathcal{H}_\infty$ controllers while maintaining the simplicity and modular architecture of NDI-like structures, thereby combining the advantages of both techniques.

Keywords: Robust control, H-infinity Loop-Shaping, Generalised stability margin, Nonlinear Dynamic Inversion, Incremental Nonlinear Dynamic Inversion, Hybrid Incremental Nonlinear Dynamic Inversion, Flight control.

Acknowledgements

This thesis marks the end of my student journey, and what a ride it has been. From my initial move from Algarve to Lisbon for my bachelor's studies at Instituto Superior Técnico, to pursuing my master's at TU Delft in the Netherlands, and later undertaking an internship in a different country in Munich, Germany, each step has contributed immensely to my growth.

For the past year, I have poured my heart and soul into this thesis project. However, I could not have done it without the guidance and support of my supervisors, Dr. Eng. Spilios Theodoulis and Dr. ir. Tijmen Pollack, to whom I am profoundly grateful. Thank you for being so supportive, approachable and committed to my work. I will always remember fondly the many after-hours discussions about the control 'mafias'.

I am also grateful for all the incredible friendships I made in the past five and a half years, many of whom, despite being hundreds of kilometres away, still feel like they are next to me. This would have been nothing without all of you.

Moreover, I am deeply grateful to my family. To my parents, Luís and Rosa, and to my sister Daniela, thank you for inspiring me, being my role models and for continuously pushing me to do better. I could not have done it without you.

Lastly, I am beyond lucky to have met such an incredible girl, with whom I had the privilege of sharing these roller-coaster experiences. Thank you, Bea, for all your support and for making every moment brighter. Thank you for your constant smile and all your love.

Contents

Nomenclature

List of Abbreviations

AoA	Angle-of-attack
ARE	Algebraic Riccati Equations
CAP	Control Anticipation Parameter
DI	Dynamic Inversion
DOF	Degree of freedom
FBW	Fly-by-wire
FCSD	Flight Control System Design
HIRM	High Incidence Research Model
INDI	Incremental Nonlinear Dynamic Inversion
LFT	Linear Fractional Transformation
LHP	Left-Half Plane
LMIs	Linear Matrix Inequalities
LOES	Low Order Equivalent System
LPF	Low-Pass Filters
LPV	Linear Parameter-Varying
LQ	Linear Quadratic
LQG	Linear Quadratic Gaussian
LQR	Linear Quadratic Regulator
LSDP	Loop-Shaping Design Procedure
LTI	Linear Time-Invariant
MIMO	Multi-Input Multi-Output
NCF	Normalized Coprime Factorization
NDI	Nonlinear Dynamic Inversion
OBAC	On-Board AirCraft
OBM	On-board Model
PID	Proportional-Integral-Derivative
RHP	Right-Half Plane
RSS	Relaxed Static Stability
SB	Sensor-based
SISO	Single-Input Single-Output
SSV	Structured Singular Value
SV	Singular Values
SVD	Singular Values Decomposition
UAV	Unmanned Aerial Vehicle
VTOL	Vertical Take-off and Landing

1

Introduction

Flight Controls are a fundamental aspect of modern aircraft, which rely on full-authority digital fly-by-wire (FBW) flight control systems for safe and efficient operation. It contributes to enhancing the aircraft's natural stability and control characteristics, decreasing pilot workload, and improving the overall safety of flight transportation.

The General Dynamics F-16 was the first fighter aircraft intentionally designed to be slightly aerodynamically unstable, also known as Relaxed Static Stability (RSS), which made it rely entirely on active stabilisation by the flight control system for safe operation (Droste & Walker, 2010). Ever since, FBW flight control systems have become indispensable for ensuring the operability of fighter jets, designed for agility and supermaneuverability, at the cost of stability (Gal-Or, 1990). Consequently, the development of robust flight control laws for fighter jets (and other aircraft) presents a highly compelling and challenging research domain.

For decades, linear design and analysis techniques have formed the foundation of the divide-and-conquer approach in control design (Balas, 2003). This includes classical methods such as pole placement and root locus analysis, as well as modern multivariable control techniques like LQR/LQG and $\mathcal{H}_\infty$-synthesis, which address the multi-input-multi-output (MIMO) nature of flight control (Balas, 2003; Blight et al., 1994; Honeywell Technology Center, Lockheed Martin Skunk Works and Lockheed Martin Tactical Aircraft Systems, 1996). These methods enable the design of control laws that effectively meet requirements within a linear time-invariant (LTI) framework. However, due to the inherently nonlinear nature of system dynamics, addressing robust performance often requires multiple local designs and employing interpolation or scheduling strategies (Rugh & Shamma, 2000). While this approach has long been standard in industry and has led to numerous successes, such as the F-22 Raptor and the Eurofighter (Kim et al., 2023), it remains an intensive and time-consuming process.

Nonlinear Dynamic Inversion (NDI) and its incremental counterpart (INDI) emerged as alternatives to the traditional divide-and-conquer paradigm (Enns et al., 1994). By leveraging the concept of feedback linearization, NDI-based controllers inherently address the nonlinear nature of system dynamics, resulting in self-scheduling control laws (Khalil, 2002). This approach provides a transparent and modular framework for flight control design, enabling the decoupling of airframe-dependent and flying-qualities-dependent components. Its conceptual simplicity and modularity make it particularly attractive for designing flight control laws with shorter development cycles. However, the lack of inherent robustness guarantees requires rigorous a posteriori analysis to ensure robustness criteria are met, effectively shifting the design

burden from synthesis to analysis (Hyde & Papageorgiou, 2001). Despite this challenge, the increasing adoption of NDI has led to its implementation in state-of-the-art aircraft such as the Lockheed Martin F-35 Lightning II strike fighter jet (Canin, 2019; Harris, 2018).

Conversely, within the realm of $\mathcal{H}_\infty$ Robust Control, $\mathcal{H}_\infty$ Loop-Shaping provides controllers with robust stability guarantees without requiring explicit uncertainty models (McFarlane & Glover, 1992). The designer shapes the broken-loop response magnitude based on classical insights using weighting filters and then makes use of $\mathcal{H}_\infty$-synthesis to maximise robust stability against normalised coprime factor (NCF) uncertainty (McFarlane & Glover, 1992). Therefore, it can be said that this method reconciles classical control methods such as loop-shaping with $\mathcal{H}_\infty$ optimisation. These tools emerged around the same time as NDI techniques, rendering the approach comparably modern, and also enable, to some extent, the decoupling of airframe-dependent and flying-qualities-dependent components (Limebeer et al., 1993). The method has been successfully demonstrated in a flight test campaign with the VAAC Harrier (Hyde, 1995) and with the Bell 205 helicopter (Postlethwaite et al., 2005). This LTI control technique can be extended to LPV control as done in a later flight test campaign with the VAAC Harrier in G. Papageorgiou and Glover, 1999; G. Papageorgiou et al., 2000.

The recent advent of non-convex, non-smooth formal synthesis techniques allowed solving $\mathcal{H}_\infty$-synthesis problems, unstructured in nature, under structural constraints (Apkarian & Noll, 2006a). These advances allow control law structures to be specified a priori, ensuring greater flexibility in controller design. As a result, modern robust control synthesis methods facilitate the systematic robust design of structured controllers, such as PID and NDI-based control laws, expanding the applicability of $\mathcal{H}_\infty$-synthesis and unlocking new possibilities for robust control in practical applications (Apkarian & Noll, 2017).

The present study aims to combine the $\mathcal{H}_\infty$ Loop-Shaping Design Procedure (LSDP) to optimise linear variants of NDI/INDI control, referred to as Dynamic Inversion (DI) and Incremental Dynamic Inversion (IDI), using non-smooth non-convex formal synthesis techniques. The goal is to take advantage of the conceptually simple and modular control framework of NDI/INDI controllers while addressing its inherent absence of robustness guarantees by employing a robust $\mathcal{H}_\infty$ control synthesis procedure. The proposed methodology is validated by designing a digital pitch-rate controller for NASA's X-29 experimental aircraft in a severely unstable flight condition.

1.1. Report Structure

The thesis is organised into eight additional chapters, each building on the concepts and findings of the previous ones.

Chapter 2 entails a literature study of $\mathcal{H}_\infty$ control, namely $\mathcal{H}_\infty$ Loop-Shaping, and NDI/INDI control. From the identified research gap in the literature review, a research proposal is outlined to answer the research objective and the research questions.

Chapter 3 lays the theoretical foundation by providing a concise description of the underlying principles of Dynamic Inversion and $\mathcal{H}_\infty$ Loop-Shaping control. Building on this foundation, Chapter 4 offers an in-depth analysis of G. Papageorgiou and Polansky, 2009, identified as one of the most relevant articles in Chapter 2 on the topic of combining Dynamic Inversion with $\mathcal{H}_\infty$ Loop-Shaping.

Chapter 5 delves into the robustness properties of various Dynamic Inversion methods, including Model-Based, Sensor-Based, and Hybrid-Based approaches, providing insights into the advantage of using Hybrid-based architectures. Following this, a novel approach for tuning a Hybrid Incremental Dynamic Inversion (IDI) controller using the $\mathcal{H}_\infty$ Loop-Shaping Design Procedure is introduced.

Chapter 6 demonstrates the application of the proposed methodology by tuning a Hybrid IDI pitch-rate digital controller for NASA's X-29 experimental aircraft. Extending this work, Chapter 7 explores the potential for applying the procedure to Multi-Input Multi-Output (MIMO) systems and investigates the implications of using non-diagonal weighting filters in the tuning process.

The thesis concludes with Chapter 8, which addresses the research questions posed at the outset, summarises the study's main findings, and presents recommendations for future research.

2

Research Proposal

2.1. Literature Review

This chapter aims to provide a literature review of Nonlinear Dynamic Inversion (NDI) Control and $\mathcal{H}_\infty$ Loop-Shaping Control. After this chapter, it should come naturally that these two methods represent sensible approaches to the design of flight control systems and that there is a lot of potential in combining these two methods. The research gap, objective, questions, and project plan for this thesis should follow logically.

2.1.1. Practical Considerations for the Application of Linear and Nonlinear Control Techniques

The adoption of digital fly-by-wire flight control systems by fighter aircraft and civil aircraft has become a standard in the industry. The methods applied to the design of flight control laws for fighter jets have experienced an evolution, from SISO classical control methods used for the F-16 (Figure 2.14), to more advanced multivariable control law design techniques, such as Eigen-structure Assignment for the F-22 Raptor (Figure 2.1), LQ control for the JAS-39 Gripen (Figure 2.2), differential PI gains for the Eurofighter (Figure 2.3), simplified Dynamic Inversion control for the T-50 (Figure 2.4), and, more recently, Nonlinear Dynamic Inversion techniques for the three versions of the F-35 Lightning II (Figure 2.8) (Kim et al., 2023). With the exception of the latest, all the other techniques are linear control techniques that loosely fall within a linearization-based gain-scheduling approach. Thus, linear control techniques have demonstrated a long and successful history of practical application in fighter jet flight control systems.

Due to the need to design control laws for three different versions, each with different mission requirements and aerodynamic characteristics, the Lockheed Martin control law team chose a methodology with a control structure that could support all versions. The three versions are the F-35A conventional takeoff and landing (CTOL) variant, the F-35B short takeoff and vertical landing (STOVL) variant, and the F-35C carrier-based variant (CV). In light of these requirements, the NDI approach emerged as the chosen approach since it can directly accommodate system nonlinearities, avoid the need for time-consuming gain scheduling, and its modular approach allows it to be easily transferable across different versions, making the design process more streamlined and cost-effective. Despite these advantages, the control team had to deal with challenges in the computing demand of the algorithms, the invertibility of the system, numerical stability and sensitivity to unmodeled dynamics and practical limitations on the size of the model that can be embedded in the software (the F-35 onboard models consist of approximately 3 million data

Figure 2.1: Lockheed Martin F-22 Raptor [Credit US Air Force (USAF)].

Figure 2.2: Saab JAS 39 Gripen [Credit Tuomo Salonen / SIM Finnish Aviation Museum].

points) (Harris, 2018). In order to simplify and accelerate the control law software development process, MATLAB® Autocode was used to auto-generate code. In this process of designing NDI control laws, the role of the control law engineer changes from focusing on developing gain schedules based on linear models to understanding the aerodynamics of the aircraft, developing the onboard model, and also working as an embedded software engineer, responsible for software development in accordance with established industry standards for safety-critical systems (Canin, 2019; Harris, 2018).

Another good reference on the industry perspective and practical considerations of different control methods is given in the work of Magni et al., 1997. It presents an extremely detailed overview of the main conclusions of the Research Civil Aircraft Model (RCAM) initiative, a collaborative effort of European Industry and Academia, organised by the Group for Aeronautical Research and Technology in Europe (GARTEUR). The goal of this project was to advance flight control technology by providing a standardised model for researchers to develop and test robust flight control designs and analyse them in relation to industrial experience and practices. Naturally, as this project dates back to the 90s, some of the comments are no longer applicable, but they offer a lot of valuable insights into the challenges of a wider adoption of $\mathcal{H}_\infty$ robust control tools at the time and the slight preference for NDI tools from a practical viewpoint. Among the tested techniques, applications of μ-synthesis and $\mathcal{H}_\infty$ Loop-Shaping are found, and other techniques, like the Multi-Objective Parameter Synthesis (MOPS) by the German Aerospace Centre (DLR), stand out.

Despite the acknowledgement of the results of $\mathcal{H}_\infty$ tools entries as among the best, the author summarizes a few points that might limit its adoption in industry: requirement of advanced mathematical knowledge, lack of visibility and high complexity of the control structure, implementation issues due to the need of discretizing high-order controllers and re-design challenges in case the aircraft dynamics' models are modified (Magni et al., 1997). On the other hand, the MOPS approach does not face these issues. This technique can include various types of design objectives and make use of multi-objective synthesis tuning by min-max parameter optimisation. As it is applicable for every structure that can be described mathematically, and every requirement can be used as long as it can be expressed mathematically by a cost function, it was deemed one of the best approaches (Joos, 1999).

Additionally, Magni et al., 1997 also reviews the main outlooks of the High Incidence Research Model (HIRM) design challenge. The HIRM was also a European initiative focused on enhancing the understanding of fighter aircraft behaviour at high angles of attack. The design entries included approaches such as LQ control, a combination of Nonlinear Dynamic Inversion and LQG, Robust Inverse Dynamics Estimation

(RIDE), μ-synthesis and $\mathcal{H}_\infty$ Loop-Shaping.

From these approaches, the main conclusions from the evaluation team (flight control designers from British Aerospace Military Aircraft, SAAB Military Aircraft and British Aerospace Dynamics) on the NDI with LQG approach and $\mathcal{H}_\infty$ Loop-Shaping are highlighted. The NDI approach is deemed as handling the time response requirements well, but as robustness requirements are not addressed directly, it possibly leads to multiple design iterations. For the $\mathcal{H}_\infty$ Loop-Shaping, similar conclusions to the ones described in the RCAM initiative were found: issues associated with the mathematical theory challenges and lack of state-space block's visibility, but, on a positive note, the degree of commonality with classical control (loop shaping principles) and sound theoretical linear robustness guarantees were highlighted.

The authors of Magni et al., 1997 argue that different control approaches can be organised into three categories according to the design philosophy employed. The first one is the control theory approach, which is centred around the use of a theoretical control synthesis method, such as $\mathcal{H}_\infty$-synthesis. The second one is the flight physical approach, which is centred around the flight dynamics, mission requirements and physical reasoning. NDI and INDI control would fall within this category. Finally, there is the process-oriented approach, which focuses on the design process and the design engineer by making use of automation software tools. The MOPS approach is an example of the latter approach.

The primary motivation behind the aeronautical industry's adoption of advanced flight control design methods is driven by economic considerations. As engineers face increasing demands to develop control laws more quickly for increasingly complex systems, they naturally seek more efficient and automated design approaches. Such improvements aim to streamline the design process, making it faster and more cost-effective, while continuing to ensure the highest standards of safety. Therefore, for a new flight control method to be accepted by any industrial flight control systems design community, it is essential that it integrates well with their existing design process. Additionally, practical considerations, such as implementation in the aircraft's flight control computer, must be thoroughly addressed. Motivated by this, controllers based on μ-synthesis and $\mathcal{H}_\infty$ LS were deemed too 'academic' at the time, as the lack of visibility and high complexity of the control structure was considered a big issue for practical applications. In this context, the adoption of NDI control laws for the F-35 Joint Strike Fighter Program (JSF) might be easily understood.

The practical challenges that limited the use of $\mathcal{H}_\infty$ control in industry in the 1990s have since been overcome. Research efforts towards $\mathcal{H}_\infty$ static output-feedback control (Gadewadikar et al., 2007; Gadewadikar et al., 2006) and the continuous research efforts towards non-smooth $\mathcal{H}_\infty$-synthesis (Apkarian & Noll, 2006a) have resulted in the development of tools for the design of structured controllers under the $\mathcal{H}_\infty$ framework, now publicly available for control engineers via MATLAB® `hinfstruct` and `systune`. The advent of these tools, namely `hinfstruct`, was essential to the success of the Rosetta mission of the European Space Agency (Apkarian et al., n.d.). In 2011, following the failure of a thruster, new attitude control systems were urgently needed. The mission's original software proved inadequate for designing these replacements; it was the employment of `hinfstruct` that ultimately enabled the successful development of new, reliable attitude controllers. Thus, the majority of barriers to the practical implementation and use in the industry of $\mathcal{H}_\infty$-synthesis methods described in Magni et al., 1997 are no longer an issue. Similar to how the MOPS technique was deemed as a process-oriented approach, $\mathcal{H}_\infty$-synthesis is, today, also a process-oriented approach, though anchored in a sound theoretical mathematical background. The main challenge left to the designer is how to translate multiple system requirements, such as ideal time-domain responses, certification requirements (for instance, minimum gain and phase margins), and others, into frequency-domain requirements in the form of weighting filters needed in the $\mathcal{H}_\infty$-synthesis process.

Figure 2.3: Eurofighter Typhoon [Credit Chris Lofting].

Figure 2.4: KAI T-50 Golden Eagle [Credit Philippine Air Force (PAF)].

2.1.2. Motivation for Linear Flight Control Techniques and LPV Extensions

Linear control techniques have been widely studied and accepted in flight control applications. As practical applications of flight control have to cover a wide flight envelope, the most standard approach is to use gain scheduling based on local linear designs of LTI systems obtained from the Jacobian linearization of the nonlinear dynamics around chosen operating points.

In this case, the gain-scheduled controller, referred to as Linearization-Based Gain-Scheduling (LBGS), can be obtained based on ad-hoc interpolation methods. These methods vary in complexity and include controller switching, controller blending, ZPK interpolation, transfer function coefficient interpolation, state-space matrix interpolation, observer/state feedback interpolation, and other interpolation schemes (Theodoulis, 2008).

While the increased linear control theory maturity, especially with the development of robust $\mathcal{H}_\infty$ tools, allows the design of linear controllers with robustness guarantees in the linear domain, translation to robustness guarantees in the nonlinear domain of the interpolation-based gain-scheduled controller is not assured. In this context, the interpolation procedure for the construction of the global gain-scheduled is critical since robust local linear designs can even result in instability of the closed-loop system. Thus, without an appropriate theoretical analysis, these gain-scheduled controllers come with no guarantees on the performance, robustness, and even nominal stability (Shamma & Athans, 1991).

In the 80s and 90s, there was an increased interest in studying the stability and robustness guarantees of LBGS. The work of Shamma and Athans, 1991 establishes conditions for which robust stability and performance of parameter-varying linear systems in the context of gain-scheduling are guaranteed. Namely, it demonstrates that the gain-scheduled controller's desirable feedback properties for all parameters' frozen values are maintained if the parameter time variations are sufficiently slow. Another very relevant work is the benchmark paper of Nichols et al., 1993 for the gain-scheduled autopilot of an air-to-air missile, which proposes conditions for the gain-scheduled controller realisation to avoid hidden coupling terms. The existence of these hidden coupling terms is detrimental since they can result in the instability of the closed-loop system. The full theoretical results are later given by the work of Lawrence and Rugh, 1995. Moreover, the work of Stilwell and Rugh, 1999 develops a method of interpolating LTI controllers with observer state feedback structures that guarantees closed-loop stability of the Linear Parameter-Varying system within a certain range of variation rates of the scheduling vector. Posteriorly, in Stilwell and Rugh, 2000, the same authors generalise these results to two other cases: state space matrix interpolation and interpolation of stable coprime factors of transfer functions.

In many cases, the absence of robustness guarantees of the nonlinear gain-scheduled controller is ad-

dressed by a thorough a-posteriori robustness analysis of the controller. In order to verify that the controller meets performance and robustness requirements across the flight envelope, Monte Carlo simulations of hundreds or thousands of operating points with perturbations on the aerodynamics coefficients, addition of wind gust disturbances, and consideration for other effects are often used. Other robustness tests, for instance, include the linearization of the total open loop of the system plus the gain-scheduled controller in order to inspect its stability margins for frozen values of time. Moreover, another important aspect to consider is that linear analysis tools are quite well-established, especially compared to nonlinear ones. In fact, aerospace industry safety certification requirements are often based on linear methods, such as gain and phase margins, which can readily be met with $\mathcal{H}_\infty$ linear control tools.

While the previous approach to gain-schedule control falls within the divide-and-conquer design philosophy, alternative approaches were developed in the late 80s and 90s, referred to as Linear Parameter-Varying (LPV) gain-scheduling. These methods make use of LPV plant descriptions of the nonlinear dynamics, allowing the direct synthesis of a gain-scheduled controller rather than its construction from a family of local linear controllers designed based on LTI plants. The great advantage of LPV control is that it offers the potential of both stability and performance guarantees, both within the neighbourhood of selected operating conditions and throughout transient phases. However, this approach can introduce a certain degree of conservatism into the control solution. These methods widely employ the $\mathcal{L}_2$ norm as a performance measure since this enables a degree of continuity with linear $\mathcal{H}_\infty$ control. Consequently, LPV control can be seen as an extension of linear $\mathcal{H}_\infty$ control to LPV systems (Leith & Leithead, 2000; Rugh & Shamma, 2000).

There are multiple approaches to LPV control design, such as small-gain LFT (Apkarian & Gahinet, 1995) and Lyapunov-based LPV approaches, namely, Quadratic Lyapunov function approaches (Apkarian et al., 1995) and Parameter-dependent Lyapunov function approaches (Wu et al., 1996). These techniques have found application in multiple high-performance aircraft, such as the F-14 (Balas et al., 1997), F-16 (Lu et al., 2006), F/A-18 (Balas et al., 1999) and the VAAC Harrier (Figure 2.7) (G. Papageorgiou & Glover, 1999; G. Papageorgiou et al., 2000). More recently, LPV control was also flight-tested in the Cessna Citation II (Figure 2.12) aircraft (Weiser et al., 2020).

As a final remark, it is important to acknowledge that the popularity of nonlinear control techniques, such as Nonlinear Dynamic Inversion and Incremental Nonlinear Dynamic Inversion, motivated some scepticism towards the merits and successes of linear control and gain-scheduled approaches (Kim et al., 2023). Nevertheless, the discussion above supports the idea that LBGS is a practical and effective method for nonlinear control design and that LPV control design extensions, despite the increased complexity, present systematic tools for the design of flight control laws with stability and performance guarantees. For a thorough treatment of the 80s and 90s developments in LBGS and LPV control design, the reader is invited to read the two survey papers of Leith and Leithead, 2000 and Rugh and Shamma, 2000.

In this context, it is important to stress some articles that acknowledge the usefulness of robust control tools to analyse and synthesise NDI/INDI control laws, a design philosophy that will be followed in the present work. The work of Hyde and Papageorgiou, 2001 recognises the conceptual simplicity of NDI control laws but also the fact that the technique gives no robustness guarantees, proposing the use of frozen point and LPV methods to analyse the stability properties of NDI control. More specifically, it makes use of a normalised coprime factor stability test for the frozen point analysis, a concept which directly ties to $\mathcal{H}_\infty$ Loop-Shaping control, and an $\mathcal{L}_2$ extension of this test for LPV analysis. Later, C. Papageorgiou and Glover, 2005 extends the robustness analysis tests of NDI control laws by making use of quasi-LPV models and considering the presence of time-varying and parametric uncertainty. These tests are formulated as

linear minimisation problems subject to LMI constraints and solved via a gridding procedure, providing sufficient conditions for the analysis of the nonlinear system's performance in the presence of bounded disturbances and uncertainty. More recently, the work of Pollack, 2024 challenges the recent claims about the increased robustness of INDI control laws over NDI laws by employing μ-analysis to linear variants of those techniques (DI and IDI control). The research demonstrates that the two techniques have complementary robustness properties, and, as such, a hybrid schematic which combines both can be leveraged to augment the robustness properties of the inversion inner loop. Moreover, a μ-synthesis approach is developed to systematically tune the Hybrid IDI controller in light of non-smooth, multi-objective $\mathcal{H}_\infty$-synthesis.

Several conceptual bridges can be established between these articles. The articles point out that the advantages of NDI control laws lie in their architectural value and modular simplicity. However, recognising the lack of robustness guarantees, $\mathcal{H}_\infty$ linear synthesis and analysis tools and LPV extensions are proposed to address this issue. In this context, the idea of leveraging the $\mathcal{H}_\infty$ Loop-Shaping Design Procedure to tune a Hybrid IDI control law arises. Going back into the three design philosophy categories of control methods mentioned in the previous section, this concept fulfils the three of them: it is a flight physics approach in terms of its controller structure, it is a control theory approach in terms of how the gains of the controller are optimized using the $\mathcal{H}_\infty$ Loop-Shaping, and it is a process oriented approach in the sense that many aspects are automated using `systune` and the designer interacts more with the process by translating the design requirements to frequency domain ones and weighting filters.

Figure 2.5: NASA F/A-18 High Alpha Research Vehicle (HARV) [Credit NASA].

Figure 2.6: McDonnell Douglas F/A-18 Hornet [Credit USAF].

2.1.3. Robust Flight Control Framework

Historical Background and Developments in $\mathcal{H}_\infty$ Control

Control theory had its origin in the 40s, a time referred to as the classical control era, where methods like Bode plots, root locus, frequency domain analysis, and SISO loops were used to design controllers. Classical control comes to its limits when controllers for multivariable systems with high internal coupling are to be designed. Consequently, time-domain modern control methods were developed during an "elegant" modern control era, which had its origins in the 60s. It has its foundation in state space methods, which are used to design optimal controllers such as Linear Quadratic Regulators (LQR) and Linear Quadratic Gaussian (LQG), which consist of multivariable control synthesis techniques. However, with great optimism and overconfidence came a rude surprise: the fundamental problem of system uncertainty was largely ignored by modern control theorists, which led to supposedly 'optimal' controllers with poor performance,

or even instability, in real-world applications due to high sensitivity to modelling errors. In fact, in the 70s, John Doyle, a PhD student at the time, demonstrated that the traditional LQR guaranteed margins vanish when an estimator is introduced, as it occurs in the case of the full LQG, with a Kalman filter in the loop (Doyle, 1978)[1].

Robust Control Theory originated in the 70s and was born out of the realisation of the inherent flaws of time-domain modern control methods. Robust control, or in other words, $\mathcal{H}_\infty$ control design and analysis methods are direct and natural developments of the classical frequency domain control techniques still in use by much of the aerospace industry today (Balas, 2003). It can be seen as a return to the frequency domain philosophy of classical control, with the primary development being to allow for truly multivariable design in the presence of system and signal uncertainty. The theory is based on minimising the $\mathcal{H}_\infty$ norm.

The $\mathcal{H}_\infty$ problem was formally presented by Zames, 1981, but the origins of the $\mathcal{H}_\infty$ can be traced back to the 1960s when the same author discovered the small gain theorem (Zames, 1966). In the last decades, the $\mathcal{H}_\infty$ problem was solved multiple times, in the sense that the problem was progressively solved from a least complex one, with a full-order unstructured controller, to a more complex one, with a controller under structural constraints (Apkarian & Noll, 2006a). The first (extremely unpractical) solution to the problem presented by Doyle, 1984, and later refined by Glover and Doyle, 1988 and by the celebrated DGKF paper Doyle et al., 1989[2], shows that the $\mathcal{H}_\infty$ problem requires the solution of two Algebraic Riccati Equations (AREs). In 1994, Gahinet and Apkarian, 1994a gave a solution for the $\mathcal{H}_\infty$ problem by reducing it to Linear Matrix Inequalities (LMIs). In the same year, Iwasaki and Skelton, 1994 also provided a similar solution to the $\mathcal{H}_\infty$ problem in terms of LMIs. Both AREs and LMIs solutions result in full-order controllers, but the latter suggested the possibility of using LMIs to synthesise reduced-order controllers. However, synthesising a reduced-order controller leads to Bilinear Matrix Inequalities (BMIs) rather than LMIs, resulting in a non-convex minimisation problem, which, at the time, did not guarantee global convergence (Apkarian & Noll, 2006a). Moreover, BMI approaches run into numerical difficulties even for problems of moderate size (Apkarian & Noll, 2006a). In Apkarian and Noll, 2006a, the $\mathcal{H}_\infty$ problem was finally solved for the structured reduced-order case with non-smooth optimisation techniques. This solution became available to control engineers in 2011 via the `MATLAB® hinfstruct` tools and later in 2015 via `MATLAB® systune`. The significance of this discovery is profound and was considered to be the 'holy grail' of robust control synthesis. Before, synthesised full-order controllers had to go through a process of model-order reduction to ensure implementability in the design. Today, a control engineer has at his disposal the tools to directly synthesise structured and reduced controllers. For instance, the very much cherished in industry PID controllers are structured controllers that can now be optimised under the $\mathcal{H}_\infty$ framework instead of tuned using heuristic techniques such as Ziegler Nichols (Ziegler & Nichols, 1942) or others.

$\mathcal{H}_\infty$ Control Methods

$\mathcal{H}_\infty$ control has been widely used to design flight controllers, especially for fighter jets. Amidst the different $\mathcal{H}_\infty$ methods, three deserve to be highlighted: $\mathcal{H}_\infty$ mixed sensitivity, μ-synthesis, and $\mathcal{H}_\infty$ Loop-Shaping.

Mixed sensitivity $\mathcal{H}_\infty$ consists of shaping the singular values of various closed-loop transfer functions through the use of weighting filters. Control system requirements, like attenuation of low-frequency disturbances at the plant output, minimisation of high-frequency actuator usage, and robust stability to additive uncertainty at high frequencies, among other requirements, can be translated to desired closed-

[1] This bold and disruptive work, published at the IEEE as the paper with the shortest abstract ever, earned him quite the 'reputation' at the time.

[2] In the inControl Podcast, Doyle jokingly says this is a really bad paper and advises his students not to read it.

loop transfer functions. The filters are then carefully designed to reflect these objectives, and then a controller K is synthesised based on a $\mathcal{H}_\infty$ minimisation problem to meet those criteria.

In Voulgaris, 1988 and later in Voulgaris and Valavani, 1991, Mixed sensitivity $\mathcal{H}_\infty$ control was used to design a longitudinal MIMO controller for an F-18 High Alpha Research Vehicle (HARV) at an operating point within the high "alpha" regime. In Ikeda et al., 1990, mixed sensitivity control was applied to the longitudinal dynamics of a fighter aircraft, and the results were compared to the LQG design. In Garg and Ouzts, 1991 and Garg, 1993, mixed sensitivity was applied on a supersonic short take-off and vertical landing (STOVL) to operate in transition flight. In Chiang et al., 1993, a mixed sensitivity $\mathcal{H}_\infty$ controller was designed for a supermaneuverable fighter flying the Herbst manoeuvre.

The structured singular value, μ, introduced by Doyle, 1982, is a powerful tool for analysing robustness performance with a certain controller. The above result defines a test for stability (robustness measure) of a closed-loop system subject to structured uncertainty in terms of the maximum structured singular value of a particular part of the closed-loop system (Bates & Postlethwaite, 2002). This structured singular value can also be used for controller synthesis in a process called μ-synthesis, which aims at directly satisfying the specifications imposed by μ in terms of robustness performance. It can allow for the synthesis of controllers that are less conservative than those produced with either $\mathcal{H}_\infty$ Loop-Shaping or Mixed Sensitivity, but care must be taken. For example, the $D-K$ iteration, a sub-optimal method that combines $\mathcal{H}_\infty$ and μ-analysis, does not guarantee joint convexity (for both K-step and D-step) and, consequently, the iterations may converge to a local optimum (Skogestad & Postlethwaite, 2005). Furthermore, the computational burden of the algorithms that compute the exact value of μ is an exponential function of the size of the problem, which can be problematic for large systems (Bates & Postlethwaite, 2002).

In the late 90s, Paul et al., 1997 and Tu et al., 1999 used μ-synthesis to design a robust lateral-directional flight controller for a tailless fighter aircraft and the F/A-18 fighter (Figures 2.5 and 2.6), respectively. More recently, μ-synthesis has found many application cases in the design of flight controllers for UAVs (Michailidis et al., 2018; Panza et al., 2020).

$\mathcal{H}_\infty$ Loop-Shaping Control

$\mathcal{H}_\infty$ Loop-Shaping has its foundation in the concept of Normalised Coprime Factorisation (NCF) of linear systems introduced by Vidyasagar et al., 1982 and Vidyasagar and Kimura, 1986. Posteriorly, Glover and McFarlane, 1988, McFarlane et al., 1988, and McFarlane and Glover, 1990 expanded the concept of NCF Robust Stabilisation, which consists of finding the controller K that robustly stabilises a plant to NCF uncertainty. $\mathcal{H}_\infty$ Loop-Shaping is formally defined in McFarlane and Glover, 1992 (and hence the also common designation of McFarlane-Glover Loop-Shaping). This method uses shaping filters W_1 and W_2 on the plant G to obtain a desirable loop (gain) shape in terms of its Singular Values, and then applies NCF Robust Stabilisation. Therefore, it can be said that this method reconciles classical control methods such as Loop-Shaping with $\mathcal{H}_\infty$ optimisation. The described Loop-Shaping Design Procedure (LSDP) was then extended from a 1DOF to a 2DOF by Limebeer et al., 1993. The 2DOF LSDP, with the extra DOF given by the independent design of the pre-filter gain, allows the trade-off of robust stability and performance properties of the closed-loop system.

$\mathcal{H}_\infty$ Loop-Shaping design method has several properties that make it particularly interesting to address control problems for aerospace applications: no gamma iteration, robustness to unstable perturbations and uncertainty in the location of lightly damped resonant poles, no pole-zero cancellations of the stable plant poles, clear management of conflict specifications around crossover, balanced robustness and

Figure 2.7: VAAC Harrier XW175 [Credit Bedford Aeronautical Heritage Group].

Figure 2.8: F-35A Lightning II Joint Strike Fighter [Credit USAF].

performance properties at the plant input and output, and can be implemented as an exact plant observer plus state feedback, which can lead to simplified gain-scheduling (Bates & Postlethwaite, 2002).

One of the most challenging aspects of the procedure is the design of the weighting filters W_1 and W_2. The work of Hyde, 1995 provides a method to design diagonal filters W_1 and W_2, which works particularly well for plants with dominant diagonal terms. However, since diagonal weights do not work particularly well for plants that have strong cross-coupling between channels, the design of the weights for such plants can be difficult and time-consuming. The work of G. Papageorgiou, 1998 addresses these issues by proposing a systematic procedure to design non-diagonal filters W_1 and W_2, which allows the designer to specify the desired loop shape more accurately as he exploits more degrees of freedom.

The $\mathcal{H}_\infty$ Loop-Shaping Design Procedure has been used to design flight control laws for fighter jets. In Hyde and Glover, 1993, a flight control law for the longitudinal motion of a Harrier jump jet has been developed using $\mathcal{H}_\infty$ Loop-Shaping. Following this, G. Papageorgiou and Glover, 1999 and G. Papageorgiou et al., 2000 extend the procedure to LPV gain-scheduling and flight-tested it in the Harrier jump jet. In Bates and Postlethwaite, 2002, a multivariable wing-leveller flight controller for a linearised model of the lateral dynamics of an F-16 aircraft is designed using both a static $\mathcal{H}_\infty$ Loop-Shaping controller and LQ optimisation, and the results are compared. Again, in Bates and Postlethwaite, 2002, an $\mathcal{H}_\infty$ Loop-Shaping controller is used to design a centralised, integrated flight and propulsion control system for the VAAC Harrier aircraft (WEM model), considering the longitudinal and lateral control problems simultaneously. Recent approaches to $\mathcal{H}_\infty$ Loop-shaping control of fighter aircraft are scarce. However, $\mathcal{H}_\infty$ Loop-shaping has found a lot of recent applications to the design of robust agile vehicle control, which shares some similarities with fighter aircraft in terms of longitudinal control. In Friang et al., 1998, the control of a bank to turn a missile with very lightly damped bending modes is synthesised using $\mathcal{H}_\infty$ Loop-Shaping, and robustness against large modelling uncertainties is investigated using the ν-tool. A systematic methodology for the synthesis of global gain-scheduled controllers for nonlinear time-varying systems, using reduced-order, static $\mathcal{H}_\infty$ Loop-Shaping controllers, whose set of operating points is chosen based on the gap metric, is presented in Theodoulis and Duc, 2009. In Kanade and Mathew, 2013, the performance of an $\mathcal{H}_\infty$ Loop-Shaping and an $\mathcal{H}_\infty$ Mixed-Sensitivity controller is analysed when applied to a pitch model of a rocket. Finally, a gain-scheduled structured $\mathcal{H}_\infty$ Loop-Shaping autopilot is designed for spin-stabilized canard-guided projectiles in Sève et al., 2017.

2.1.4. Nonlinear Dynamic Inversion Control Framework

Historical Background and Developments in NDI Control

Nonlinear Dynamic Inversion (NDI), also known as Feedback Linearisation among the control community, was developed as an alternative to the divide-and-conquer paradigm, which results in gain-scheduled controllers. The general idea of NDI control law is to transform selected input-output channels into a chain of integrators by using state feedback information (hence the designation Feedback Linearisation). This inner dynamic inversion loop linearises the nonlinear dynamics, transforming the nature of the system from nonlinear to linear. As a result, an outer virtual control loop may employ linear control design methods to control this linearised system, resulting in a self-scheduled controller.

Inspired by Meyer and Cicolani, 1975, the work on state-space linearisation of scalar-input systems by Brockett, 1978, solutions for general multi-input systems established by Su, 1982, and the work on zero dynamics by Byrnes and Isidori, 1984 laid the theoretical foundations for what would be later be known as Nonlinear Dynamic Inversion control.

However, it was around the early 1990s that NDI flight control started to take off. The Honeywell Technology Centre in Minneapolis, MA, USA, took a fundamental role in that aspect, with the works of Elgersma, 1988, Bugajski et al., 1990, Bugajski and Enns, 1992, and Snell et al., 1992. Later, in 1996, a combined research effort by Honeywell Technology Centre, Lockheed Martin Tactical Aircraft Systems, and Lockheed Martin SkunkWorks led to the release of the design guidelines for various multivariable flight control methods, including NDI (Honeywell Technology Center, Lockheed Martin Skunk Works and Lockheed Martin Tactical Aircraft Systems, 1996). Furthermore, Honeywell developed a proprietary tool, MACH (Multi-Application Control), for the systematic design of NDI control laws for aerospace applications (Balas, 2003). The success of this methodology led to its adoption in the STOVL X-35 version (Figure 2.9) as part of the Joint Strike Fighter (JSF) program, and later to the Lockheed Martin F-35 Lightning II production aircraft (Canin, 2019; Harris, 2018; Walker & Allen, 2002).

Classical NDI Control

The aforementioned development of Nonlinear Dynamic Inversion control falls within the classical NDI control. It is a model-based approach that requires an On-board model (OBM) of the system dynamics to do the inversion.

One of the main advantages of NDI control is that it allows for a direct treatment of nonlinear systems. The dynamic inversion inner loop allows the system to be treated from an outer loop perspective as a chain of integrators (whose order depends on the relative degree of the plant). This means that a single virtual control law may be employed to control the system throughout the entire flight envelope. The fact that this controller results in a self-scheduled controller can lead to shorter design cycles. Moreover, NDI enables the decoupling of the control design's airframe-dependent and flying qualities-dependent parts. Therefore, it can be seen as a transparent and modular approach to flight control design, allowing re-use and easier adaptation of the control design across different flight control problems. Besides the already mentioned ones, the multiple case applications for the X-38 Space Station Crew Return Vehicle (Figure 2.10) (Valasek et al., 2001), large flexible aircraft (Gregory, 1998), and helicopters (Enns & Keviczky, 2006) verify the success of this methodology.

Nevertheless, several aspects require consideration. First of all, the model-dependency nature of this dynamic inversion approach implies precision and accuracy of the OBM. As a result, an unreliable OBM can lead to poor dynamic inversion and, consequently, poor flight control system performance.

Secondly, the NDI control framework does not provide a priori robustness guarantees (G. Papageorgiou & Polansky, 2009). This does not mean that it can not be made robust (G. Papageorgiou & Polansky, 2009); however, it might require further robust analysis and redesign to meet closed-loop requirements. Thirdly, one fundamental assumption made when designing a DI control is the need for the zero dynamics of the system to be stable (Byrnes & Isidori, 1984). This is particularly important for fighter jet aircraft as they often exhibit a non-minimum phase zero behaviour at particular flight conditions (Hauser et al., 1992).

Figure 2.9: Lockheed Martin X-35 concept demonstrator [Credit USAF].

Figure 2.10: NASA X-38 Space Station Crew Return Vehicle [Credit NASA].

Incremental Nonlinear Dynamic Inversion (INDI) Control

Incremental Nonlinear Dynamic Inversion (INDI) is an incremental approach to the classical model-based NDI. This approach attempts to solve NDI detailed and accurate model dependency by using sensor measurements of the state derivatives (hence the alternative designation of sensor-based (SB) INDI). As a result, the only onboard model information required is the plant control effectiveness, making the control law design less expensive and time-consuming.

Its origin traces back to the work of Smith, 1998 at the UK Defence Evaluation and Research Agency (DERA) in 1998. Dedicated flight test trials with the VAAC Harrier by the DERA were run to assess the potential of this new method, and the results turned out promising. This method was further studied in terms of its potential for fault-tolerant control in Bacon et al., 2000, Bacon and Ostroff, 2000, and Bacon et al., 2001. The work of Sieberling et al., 2010 from the Delft University of Technology marks the revival of the sensor-based INDI approach, which had been relatively dormant in the past years. It showed that robustness to aerodynamic model uncertainty comes at the expense of increased sensitivity to sensor delay, which can be tackled with a predictive filter. This study was followed by other researchers at TU Delft working on further developing the theory (Koschorke et al., 2013) and exploring other applications (Acquatella et al., 2012; Simplício et al., 2013). Given its advantages and the described progress, INDI control has found its application to many flight control problems, from VTOL UAVs (Lombaerts et al., 2020; Raab et al., 2018; Smeur et al., 2020) to helicopter applications (Pavel et al., 2020), and flight test campaigns on the PH-LAB Cessna Citation II (Grondman et al., 2018).

As mentioned before, SB INDI aims to retain the advantages of NDI control while decreasing the model dependency of NDI. However, a drawback of this method is that it has relatively small stability robustness margins against singular dynamic perturbations, like time delays, compared to model-based NDI (Pollack & Van Kampen, 2023; Van't Veld et al., 2018). For instance, the work of Cordeiro et al., 2022 describes how actuator or filter dynamics affect the time delay margin of SB INDI, and the work of Huang et al., 2022

details how control effectiveness mismatch influences it. Smeur et al., 2016 tackled this issue, introducing the concept of signal synchronisation. Despite the recognised success of the approach, the question of how to improve the robustness properties of SB INDI remains, to the authors' knowledge, an open question.

Figure 2.11: SIMONA research simulator [Credit TU Delft].

Figure 2.12: Cessna Citation II PH-LAB [Credit TU Delft].

Hybrid Incremental Nonlinear Dynamic Inversion

The limitations of SB INDI led to the idea of reintroducing model information of the bare airframe dynamics in the form of a complementary augmentation element, a method designated by Hybrid INDI (Pollack & Van Kampen, 2023). The concept was originally proposed by Jiali and Jihong, 2016, used in Yang et al., 2020, and formulated in Kim et al., 2021 and Kumtepe et al., 2022. The work of Pollack, 2024 marks a significant step in the development of this method by quantitatively assessing the robustness properties of the different inversion structures, like NDI, SB INDI, and Hybrid INDI. It also proposes a structured way to design the synchronisation filter of the INDI controllers under a robust control framework using μ-synthesis. Finally, the properties, limitations, and advantages of each of the three inversion-based methods are demonstrated with a case study design of a C^* controller for a simulation model of the Boeing 747-100/200 using a multi-objective structured robust $\mathcal{H}_\infty$-synthesis against mixed uncertainty to optimise the design parameters of each control architecture.

2.1.5. Combination of Dynamic Inversion Control with Robust Control

Despite the surprisingly lack of research towards combining $\mathcal{H}_\infty$ synthesis methods with DI control, these ideas are not exactly new. This section describes multiple approaches to the topic found in the literature.

The first attempt at combining both frameworks can be traced back to Adams and Banda, 1993, which proposes the use of dynamic inversion in combination with μ-synthesis and μ-analysis to design a lateral directional flight controller for an F-16 Variable Stability In-Flight Simulator Test Aircraft (VISTA) model. The idea consists of designing a desired equalised inner loop (based on the regulated dynamics of a central plant with an LQR) across the flight envelope through dynamic inversion on the first step. In the second step, an outer loop compensator is designed to achieve the required flying qualities and robustness using μ-synthesis.

Following up on this work, Reigelsperger and Banda, 1998 extends it to the longitudinal case. Dynamic uncertainty weighting filters are introduced to the actuators and sensors in the μ-synthesis framework, and a command pre-filter is added. Nevertheless, the general guidelines of designing first a desired equalised inner loop with DI and an outer loop with μ-synthesis are maintained.

Fer and Enns, 1997 apply feedback linearisation (dynamic inversion) in combination with Loop-Shaping LQR design to achieve a nonlinear regulation law with locally optimal LQ full-state feedback properties. The regulated output matrix C (referred to in the study as H) has a free parameter, which can be tuned to obtain a desired loop shape at the plant input. This approach is applied to a nonlinear short-period model of the F-18 HARV. Under this framework, the point is made that by equating the Loop-Shaping LQR gains to the feedback linearisation gains, the nonlinear feedback linearisation control laws result in a closed-loop system with both local guaranteed stability and robustness properties in terms of gain and phase margins of LQR.

In Smit and Craig, 1998, the authors propose a methodology that uses the ideas of plant-desired equalised inner loop developed by Adams and Banda, 1993. However, in the second step, instead of μ-synthesis, $\mathcal{H}_\infty$ Loop-Shaping is used to design the outer loop controller. The case study consisted of a lateral/directional axis flight control system for an unmanned air vehicle. To the author's knowledge, this last method is the first formal attempt at combining $\mathcal{H}_\infty$ Loop-Shaping control with Nonlinear Dynamic Inversion control.

Furthermore, Tierno and Glavaski, 1999 designs an $\mathcal{H}_\infty$ Loop-Shaping compensator for a fixed architecture control law based on a dynamic inversion of the short-period model. This procedure has its application in high-performance twin jet models. The $\mathcal{H}_\infty$ controller designed takes the place of a standard phase lead compensator and aims at improving the robustness of the DI controller due to possible additional unconsidered dynamics.

An extensive work on this subject was developed by Hyde and Papageorgiou, 2001. The authors analyse the stability of NDI-based flight controllers with frozen point analysis (i.e. Linear Time-Invariant (LTI) analysis) and LPV analysis. In this setup, normalised coprime factor uncertainty is used as the basis for both methods, and the case study is carried out in DERA Harrier WEM (Wide Envelope Model). The analysed NDI methods encompass standard NDI, Honeywell NDI (particular implementation of NDI by Honeywell Technology Centre and Lockheed Martin Skunk Works), P+I NDI, Robust Inverse Dynamics Estimation (RIDE) resembling an SB INDI control approach, μ-synthesis NDI like the one described in Adams and Banda, 1993, and a novel approach of cascading NDI with an $\mathcal{H}_\infty$ compensator in the inner loop.

The work of Valasek et al., 2001 suggests a new way to address robustness issues in DI control by using an outer-loop Linear Quadratic Gaussian controller to add robustness to the Dynamic Inversion inner loop. This approach is used to design a lateral-directional controller for the X-38 vehicle, and the performance and robustness of the synthesised controller are evaluated using a time domain performance criterion proposed by Ghaoui et al., 1992. In this method, the desired dynamics of the DI controller are specified by the LQG through a Loop-Shaping process.

An alternative way to feedback linearisation is described in Guillard and Bourles, 2000. Instead of classical feedback linearisation, which is used by NDI and INDI methods, the authors propose to transform the original nonlinear system by feedback into its linear approximation around a given operating point. It is argued that this method allows for maintaining the good robustness properties obtained by a linear control around this operating point and can be combined with $\mathcal{H}_\infty$ type control laws. Thus, it is referred to as robust feedback linearisation. Posteriorly, in Franco et al., 2006, the same authors, together with other researchers, extend the previous work by combining the McFarlane-Glover $\mathcal{H}_\infty$ controller with robust feedback linearisation. This method and the classical feedback linearisation NDI are applied to a magnetic bearing system control case, and the results are compared.

G. Papageorgiou and Polansky, 2009 proposes a new method to tune dynamic inversion controllers using McFarlane-Glover ($\mathcal{H}_\infty$) Loop-Shaping. The method is used to design a pitch axis control for an LTI Boeing 747 model, obtained by trimming and linearising a high-fidelity model of the Boeing 747 (Figure 2.13) for a certain dynamic pressure. In its essence, this work aims at finding an equivalence between DI control and McFarlane-Glover Loop-Shaping so that a DI controller can be tuned in the $\mathcal{H}_\infty$ Loop-Shaping framework, which comes with robustness guarantees. The selection of the OBAC (On-board-Aircraft dynamics) is based on the highest stability margin obtained with the McFarlane-Glover Loop-Shaping controller from the different operating points of the flight envelope. Posteriorly, the desired dynamics parameters of all design points are designated using the desired dynamics parameters of the OBAC. A key aspect of this work is that, similarly to the work of Fer and Enns, 1997, the regulated output matrix C is a free parameter tuned in the synthesis process for robustness.

Figure 2.13: British Airways Boeing 747-400 [Credit Adrian Pingstone, Public Domain].

More recently, in Lesprier et al., 2014, researchers at ONERA make use of the key insights provided by the robust feedback linearisation method in Guillard and Bourles, 2000 to present a new methodology combining partially linearising inner-loops with structured and robust outer-loops, designed using a non-smooth multi-model $\mathcal{H}_\infty$ optimisation approach. At first, the method is used to synthesise a longitudinal controller for a nonlinear longitudinal civil aircraft model. Then, the same method is used to synthesise a lateral controller based on a nonlinear lateral aircraft model. Finally, the longitudinal and lateral controllers are tested on a realistic landing application.

In the same year, Lanzon et al., 2014 used the same method of robust feedback linearisation combined with $\mathcal{H}_\infty$ Loop-Shaping to design a flight controller for a quadrotor subject to rotor failure.

More recently, a paper was presented by Demirkiran et al., 2024 that extends the work of G. Papageorgiou and Polansky, 2009 to the design of a pitch rate controller for a fighter aircraft with a feed-forward modification to address handling qualities. Internal stability, robustness, and handling qualities of the flight control law augmented system are analysed in addition to Monte Carlo-type uncertainty analyses to assess robustness.

Discussion

As was described above, there were some efforts towards combining the $\mathcal{H}_\infty$ framework from Robust Control Theory and Dynamic Inversion control from Nonlinear Control Theory after both frameworks became recognised as sensible approaches to the design of flight controllers. However, perhaps motivated by a division in the design philosophy, there seems to be a distinction between the robust control ‘mafia’ and the nonlinear control ‘mafia’, and the research on the combination of these methods has not exactly been developed as one would expect.

The most promising approaches entail the combination of $\mathcal{H}_\infty$ Loop-Shaping and Dynamic Inversion. In all cases, Dynamic Inversion is not approached from a global perspective of linearising the nonlinear dynamics to behave like a chain of integrators (like NDI and INDI approaches), but rather from a linear DI standpoint.

One of the approaches (Guillard and Bourles, 2000) employs a robust feedback linearisation to transform the nonlinear system into a linear system equal to the linear approximation of the original nonlinear system around a nominal operating point. The authors argue that classical feedback linearisation, which brings the system to Brunovsky canonical form, is vulnerable to uncertainties and has infinite equilibrium points and no physical meaning (two different nonlinear systems with the same dimensions will have the same Brunovsky form), making it difficult to obtain a robust controller. On the other hand, the authors prove that the robustness properties of the controller obtained by the method of McFarlane–Glover Loop-Shaping for the linearised system are kept when this controller is applied, together with the robust feedback linearisation, to the nonlinear system by making use of the concept of "local W-stability". The method is applied to a magnetic bearing control system case, and the robust feedback linearisation is compared with the classical feedback linearisation (both using a linear $\mathcal{H}_\infty$ Loop-Shaping controller and weighting filters such that the frequency response of the broken-loop gain L_i and L_o for both cases is similar). The results of the case study highlight that under parametric variations, robust feedback linearisation can maintain a response close to the nominal one, while classical feedback linearisation does not, and it even results in an unstable response under some combinations of parametric variation.

The other approach from G. Papageorgiou and Polansky, 2009 focuses more on finding an equivalence between DI control and McFarlane-Glover Loop-Shaping control. By finding this equivalence, a DI controller can be synthesised under the $\mathcal{H}_\infty$ Loop-Shaping framework, which comes with robustness guarantees. An interesting concept of this approach is that the inverted model ($\hat{C}$), also referred to as the On-Board Aircraft (OBAC), can be different from the output matrix C, as it is a non-standard tuning parameter resulting from the McFarlane-Glover Loop-Shaping synthesis. In this work, 12 LTI Boeing 747 TRue AirCraft (TRAC) models have been generated by trimming and linearising a high-fidelity nonlinear model at approximately the same dynamic pressure but for different combinations of Mach number, altitude, mass, and centre-of-gravity location. Then, the tuned DI with McFarlane-Glover Loop-Shaping is used to design 12 controllers (one for each model). Then, each controller is used on the other TRAC models, and the controller that performs better (in terms of the generalised stability margin) in the worst-case scenario is chosen as the OBAC. All the other TRACs use the inverted dynamics (C hat) of the OBAC. With this procedure, all the TRACs achieve the desired performance. As a result, even if this approach does not result in a self-scheduled controller like NDI and INDI do, it allows for a systematic design approach which, given the linear $\mathcal{H}_\infty$ control design framework, comes with robustness guarantees. Nevertheless, in this work, the feedforward and integral gain of the DI are constrained to $f_c + \frac{f_i}{f_C} = 1$, such that the transfer function from y_c to y corresponds to a first-order model. When more demanding handling qualities are required, like in the case of fighter jets, this may not suffice. For this reason, Demirkiran et al., 2024 addresses this limitation by modifying the feedforward to address the handling qualities of a fighter jet.

On the other hand, the work of Pollack, 2024 also constitutes a central piece approach to the design of nonlinear dynamic inversion with robust control considerations. Unlike the previously described approach, Hybrid INDI falls within a global perspective of linearisation of the nonlinear dynamics. However, the author focuses on $\mathcal{H}_\infty$ robustness assessments to perform quantitative comparisons of the robust stability and performance characteristics of model-based (classical NDI), sensor-based INDI, and Hybrid INDI flight control laws in a linear context. This is done by studying the characteristics of these different inversion

strategies in the frequency domain and analysing the broken-loop response at the plant input resulting from the different inversion strategies. Furthermore, the different design parameters of each control law are designed using multi-model, multi-objective $\mathcal{H}_\infty$-synthesis. The analysis of the fundamental robustness properties of the inversion strategy (inner loop) of nonlinear dynamic inversion control is quite a relevant topic. As Hyde and Papageorgiou, 2001 states in their work when analysing an inner loop NDI with an outer loop $\mathcal{H}_\infty$ controller designed independently: **"It is not always possible to correct for a poorly designed inner loop with an outer loop"**. A parallel can be established between the additional filter augmentation techniques required for IDI control and the process of Loop-Shaping in the $\mathcal{H}_\infty$ Loop-Shaping control, which raises the question about the possibility of combining Hybrid IDI with the $\mathcal{H}_\infty$ Loop-Shaping Design Procedure. In contrast to the μ-synthesis approach used in Pollack, 2024, which relies on explicit uncertainty modelling and iterative $D-K$ optimization, the $\mathcal{H}_\infty$ LSDP has direct connections with classical control concepts and provides robustness against a general description of uncertainty (NCF), making it easier to learn and use. In this context, the goal would be to use the $\mathcal{H}_\infty$ LSDP to optimise the different Hybrid IDI gains.

2.1.6. Handling qualities

The nature of the operational mission is directly linked to the formulation of the flight control law design goals. In this context, Anonymous, 1990 specifies aircraft class designation (I-IV) and identifies different flight phases (A-C). Under these criteria, the F-16 Fighting Falcon (Figure 2.14) and the NASA's X-29 experimental aircraft (Figure 2.15), the two aircraft models used in the present study, fall within the Class IV and Flight Phase Category A. Flying qualities are defined as "those characteristics of the complete air vehicle/system which allow the pilot/operator to perform to his/her satisfaction the flying tasks required to safely accomplish the mission, with an acceptable workload, while operating in the real world environment for which it is intended to operate" (Anonymous, 2002). The flying quality levels can be classified into three different levels (I-III), and it is desired to accomplish Level I satisfactory flying quality levels across the operational flight envelope.

Figure 2.14: General Dynamics F-16 Fighting Falcon [Credit USAF].

Figure 2.15: Grumman X-29 Experimental Aircraft [Credit NASA].

Based on this, Anonymous, 1990 describes desired modal response characteristics. These modal parameters are given as Low Order Equivalent System (LOES) models of the desired response of certain controlled variables to pilot demands. The Control Anticipation Parameter (CAP) is one of the most important LOES-based pitch criteria, and it will be one of the criteria used to define the performance criteria the longitudinal controller should be able to attain. In his dissertation, Gibson, 1999, the author of the Gibson Criteria, presents core discoveries of his new methodology for the handling qualities design of

control laws. The insights provided by his new methods and his analysis of the CAP criteria will be used to design the requirements on the short-period frequency w_{sp}, damping ζ_{sp}, and numerator time constant T_{θ_2}. The 2DoF LSDP explicitly attempts to meet the performance requirements imposed by these ideal response models.

2.2. Research Proposal

2.2.1. Research Gap

In the previous section, a thorough literature review of robust control theory was conducted, with a special focus on $\mathcal{H}_\infty$ Loop-Shaping and Nonlinear Dynamic Inversion theory with application to the design of flight control laws for aircraft. Both methods are popular and have proven successful in the design of flight control laws. However, they are the result of different philosophies: DI control is the product of time-domain thought, and $\mathcal{H}_\infty$ Loop-Shaping is the product of frequency-domain thought (G. Papageorgiou & Polansky, 2009). Both methods have established strengths, which naturally raises the question of how to combine both frameworks. In an attempt to answer this question, several authors delved into this problem around the time $\mathcal{H}_\infty$ control and Nonlinear Dynamic Inversion control became popular approaches to flight control. However, the recent advent of non-smooth, non-convex $\mathcal{H}_\infty$-synthesis tools, which enable the design of optimal controllers under structural constraints, opened the door for numerous research opportunities.

The work from G. Papageorgiou and Polansky, 2009 constitutes a centrepiece on how to combine both approaches. However, the robustness implications of the equivalence established between DI control and McFarlane-Glover Loop-Shaping still need further consideration to assess if the robustness guarantees of McFarlane-Glover Loop-Shaping are preserved with this approach. This constitutes the first identified research gap.

The work from Guillard and Bourles, 2000 proposes an alternative to the classical feedback linearisation of NDI and INDI methods, which the author designates robust feedback linearization. According to the author, this method preserves, to some extent, the controller's robustness properties obtained by the $\mathcal{H}_\infty$ Loop-Shaping process on the linearised system when this controller is applied, together with the robust feedback linearisation, to the nonlinear system. However, this method has found little application in the design of flight controllers. This constitutes the second identified research gap.

Finally, the work from Pollack, 2024 sheds new light on the robustness assessment of the properties of the different inversion strategies (model-based NDI, sensor-based INDI, and Hybrid INDI) and proposes a multi-model, multi-objective $\mathcal{H}_\infty$-synthesis to optimise the different design parameters of the controllers. Given the parallel established between the design of the filter augmentation techniques of IDI control, taking into consideration the nominal loop shape and the $\mathcal{H}_\infty$ LSDP, the question is raised whether there is a way to synthesise the Hybrid IDI gains under this framework. This constitutes the third identified research gap. As this approach heavily relies on the use of more modern, non-smooth, non-convex $\mathcal{H}_\infty$-synthesis tools, multiple research extensions arise, such as 2DoF multi-objective structured $\mathcal{H}_\infty$ Loop-Shaping.

2.2.2. Research Objectives

The general goal of this thesis is to investigate how to combine robust control and nonlinear dynamic inversion control with application to flight controls. As a result, the first research objective is defined as:

Research Objective

To investigate how $\mathcal{H}_\infty$ Loop-Shaping and Nonlinear Dynamic Inversion control can be combined with application to flight control laws.

However, given the broadness of possible approaches to the problem, the conducted literature review served as a way to identify suitable approaches to the problem. From the identified research gap, the following additional research objectives are defined:

Research Objective

To investigate the robustness implications of the proposed method by G. Papageorgiou and Polansky, 2009 to tune a Dynamic Inversion controller using McFarlane-Glover Loop-Shaping.

Research Objective

To investigate the potential of synthesising Hybrid IDI control laws under the $\mathcal{H}_\infty$ Loop-Shaping Design Procedure framework using non-smooth non-convex $\mathcal{H}_\infty$-synthesis.

Research Objective

To investigate the potential of using Robust Feedback Linearisation with $\mathcal{H}_\infty$ Loop-Shaping applied to a flight control problem and how it compares with classical NDI/INDI approaches.

2.2.3. Research Questions

Taking into consideration the research objectives, the research questions can be defined as:

Research Question 1

To what extent does the proposed procedure to combine $\mathcal{H}_\infty$ Loop-Shaping and a structured Dynamic Inversion controller preserve the robustness guarantees of $\mathcal{H}_\infty$ Loop-Shaping control?

1.1. How can the inverted model $\hat{C}$ be used to configure extra degrees of freedom to DI-based architectures to address robustness?
1.2. How can the procedure be extended to a 2DoF setup and address handling qualities requirements?

Research Question 2

How can a Hybrid IDI structured controller be optimised using the $\mathcal{H}_\infty$ Loop-Shaping Design Procedure?

2.1. What are the robustness implications of the hybrid inversion scheme?
2.2. What are the robustness implications of designing a controller based on the virtual law location?
2.3. How can the $\mathcal{H}_\infty$ Loop-Shaping problem be framed as a non-smooth, non-convex $\mathcal{H}_\infty$ problem under the structural constraints imposed by Hybrid IDI architectures?
2.4. How can the design procedure be extended to include model-following requirements in accordance with flying and handling qualities?
2.5. To what extent can the procedure be used to develop a digital pitch-rate (SISO) controller for a highly unstable aircraft?
2.6. What parallelism can be established between the proposed procedure and the stability assessment of Dynamic Inversion-based control laws using coprime factors in Hyde and Papageorgiou, 2001?
2.7. How can the procedure be extended to a MIMO case?

Research Question 3

How can Robust Feedback linearisation be combined with $\mathcal{H}_\infty$ Loop-Shaping for flight control applications?

3.1. What are the robustness implications of this inversion approach compared to standard NDI/INDI approaches?
3.2. How can the increase of robustness over linear $\mathcal{H}_\infty$ Loop-Shaping implementations be quantified?

2.2.4. Research Plan

The research plan is divided into three different phases. The first focuses on recreating the results of G. Papageorgiou and Polansky, 2009 and assessing the robustness implications of the proposed method to tune a Dynamic Inversion controller using McFarlane-Glover Loop-Shaping. The second dives into the potential of synthesising Hybrid IDI control laws under the $\mathcal{H}_\infty$ Loop-Shaping Design Procedure framework using non-smooth, non-convex $\mathcal{H}_\infty$-synthesis. The midterm deliverable should encompass the results from analysing G. Papageorgiou and Polansky, 2009 and have a preliminary approach to the optimisation of Hybrid IDI gains using the $\mathcal{H}_\infty$ LSDP. In the third phase, the practical implementation of the latter procedure to design longitudinal flight control laws for the X-29 aircraft should be considered. Depending on the implications of the results, a choice will be made between expanding on this research or addressing Research Question 3. The research culminates in the Dissemination, in which the report is submitted and the defence is prepared. The detailed Gantt Chart of the project is displayed in Figure 2.16.

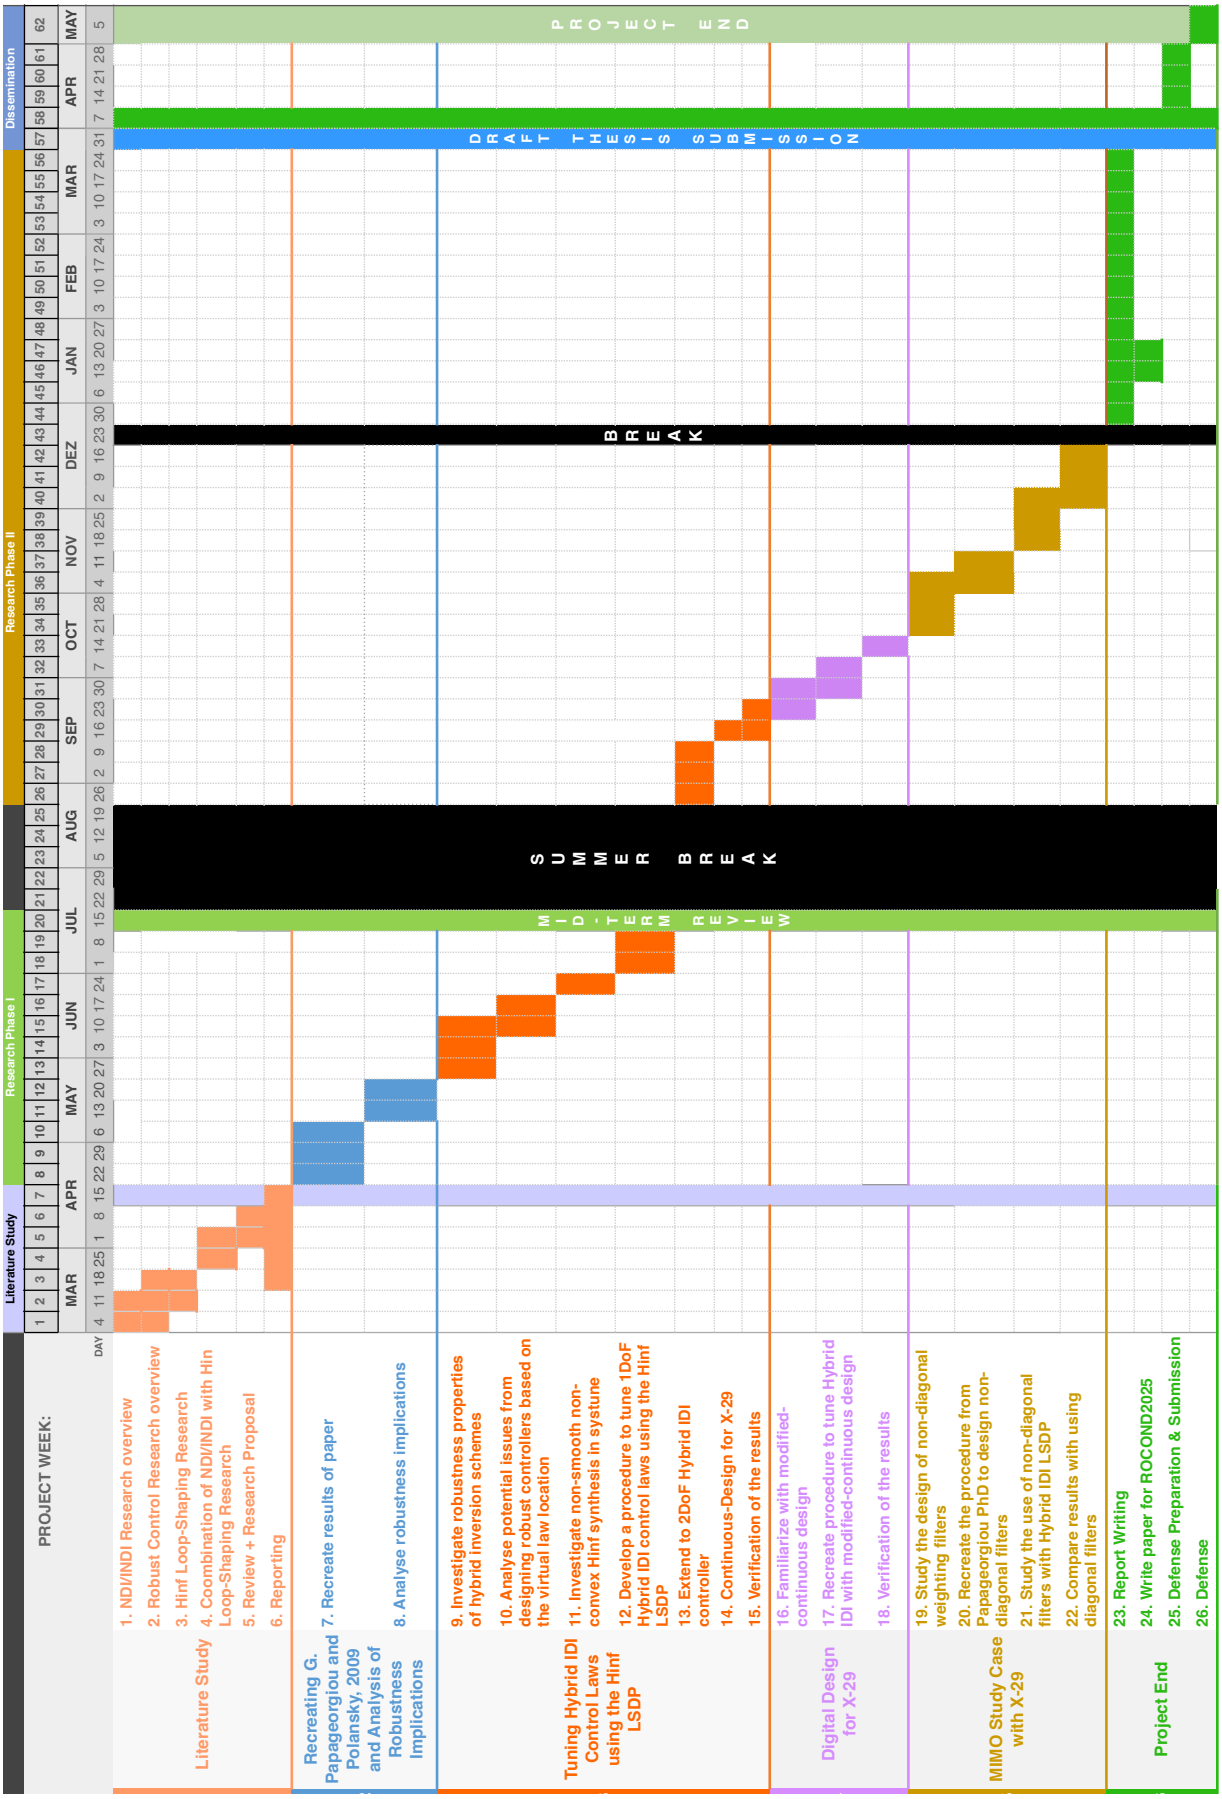

Figure 2.16: Project Plan Gantt Chart.

3

Theoretical Background

The present chapter aims to present the reader with the necessary theoretical background to grasp all the concepts described throughout the rest of the thesis. The use of the nomenclature G and P to refer to the plant's dynamics is used interchangeably, as often found in formal robust control synthesis and practical flight control applications.

3.1. Robust Control Framework

The robust control framework can be summarised by means of the generalised interconnection structure shown in Figure 3.1. The goal is to design a control law K that suppresses the effects of uncertainty, captured by the Δ-block, on control system performance. These considerations can be formulated in terms of minimisation of system norms ($\mathcal{H}_2$ or $\mathcal{H}_\infty$ norms), forming the basis for multivariable control techniques. In this context, the task of the control engineer is more about the interaction with the set of design objectives and requirements, whereas the controller tuning is the result of a formal synthesis algorithm that minimises the specified system norm.

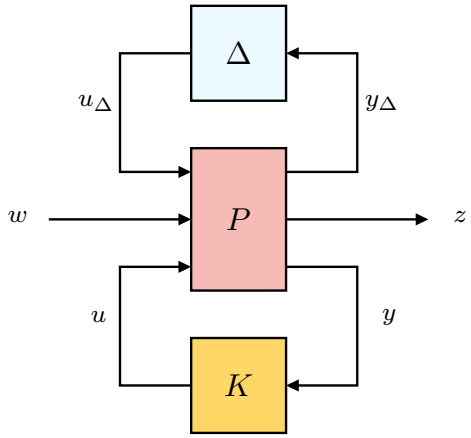

Figure 3.1: Generalised control diagram.

3.1.1. Singular Value Decomposition

Any $l \times m$ constant complex matrix G can be decomposed into its singular value decomposition (Bates & Postlethwaite, 2002):

$$G = U\Sigma V^H \tag{3.1}$$

where Σ is an $l \times m$ matrix with $k = min(l, m)$ non-negative singular values σ_i, arranged in descending order along its main diagonal (the other entries are zero). The singular values are the positive square roots of the eigenvalues of $G^H G$, where G^H is the complex conjugate transpose, i.e.

$$\sigma_i(G) = \sqrt{\lambda_i(G^H G)} \tag{3.2}$$

U is an $l \times l$ unitary ($U^H = U^{-1}$) matrix whose columns are the output singular vectors of u_i, with the vectors u_i orthonormal. V is an $m \times m$ unitary matrix whose columns are the input singular vectors v_i, with the vectors v_i orthonormal.

To make the concept of singular values clearer, consider the following geometric interpretation:

$$G = \begin{bmatrix} 0.96 & 1.72 \\ 2.28 & 0.96 \end{bmatrix} = U\Sigma V^T = \begin{matrix} \\ \\ \end{matrix} \begin{matrix} u_1 & u_2 \end{matrix} \begin{bmatrix} 0.6 & -0.8 \\ 0.8 & 0.6 \end{bmatrix} \begin{bmatrix} 3 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0.8 & 0.6 \\ 0.6 & -0.8 \end{bmatrix} \begin{matrix} v_1 \\ v_2 \end{matrix}$$

Any real matrix G of dimension 2×2 maps a unit circle into an ellipsoid, where the singular values σ of G correspond to the lengths of semi-axes of the ellipsoid. The output singular vectors u_i, where $U = [u_1\ u_2]$, correspond to the mutually orthogonal directions of the semi-axes of the ellipsoid, while the input singular vectors v_i, where $V = [v_1\ v_2]$, correspond to the directions in the unit circle which are then mapped to u_i vectors of the elipsoid with gain σ_i, such that $Gv_i = \sigma_i u_i$ (Bates & Postlethwaite, 2002).

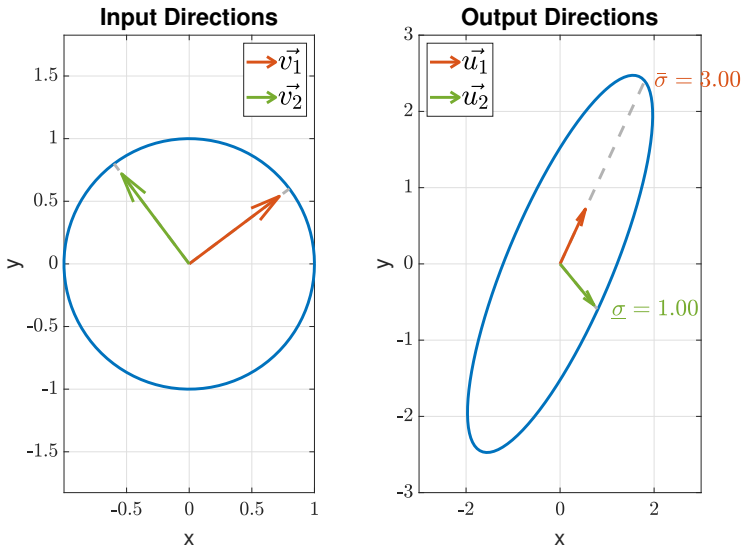

Figure 3.2: Geometric interpretation of singular vectors.

Naturally, these concepts also extend to higher-dimensional matrices and vectors: instead of a unit disc, there is a sphere or hypersphere, and the ellipsoid becomes a hyperellipsoid.

Consider now the frequency response of a system $G(jw)$. At each frequency w_i, it is possible to perform a singular value decomposition of $G(jw_i)$. Then, if the values of the singular values σ are plotted across frequency, it is possible to directly visualise how σ varies across frequency, but also the maximum $\overline{\sigma}$ and minimum $\underline{\sigma}$ of the entire system. This is important to understand the concept of $\mathcal{H}_\infty$ norm.

3.1.2. System's Norms

System norms play a central role in multivariable control techniques, particularly the $\mathcal{H}_\infty$ and $\mathcal{H}_2$ norms, which refer to two different Hardy spaces. $\mathcal{H}_\infty$ refers to the Hardy space with functions that are analytic and bounded in the open right-half plane. The real rational subspace of $\mathcal{H}_\infty$ is denoted by $\mathcal{RH}_\infty$, which consists of all proper and real rational stable transfer matrices. Similarly, $\mathcal{H}_2$ refers to the Hardy space with functions that are analytic in the open right-half plane. The real rational subspace of $\mathcal{H}_2$ is denoted by $\mathcal{RH}_2$, which consists of all strictly proper and real rational stable transfer matrices.

The $\mathcal{H}_2$ norm of a stable, strictly proper system $G(s)$ can be defined as:

$$\|G(s)\|_2 \triangleq \sqrt{\frac{1}{2\pi}\int_{-\infty}^{+\infty} \|G(j\omega)\|_F^2 \, d\omega} = \sqrt{\frac{1}{2\pi}\int_{-\infty}^{+\infty} \mathsf{Trace}(G(j\omega)^H G(j\omega)) \, d\omega} \tag{3.3}$$

where the Frobenius norm $\| * \|_F$ of a general complex matrix A is defined as:

$$\|A\|_F \triangleq \sum_{i,j} |a_{i,j}| = \sqrt{\mathsf{Trace}(A^H A)} \tag{3.4}$$

The lack of an interpretation of the $\mathcal{H}_2$ system norm as an induced norm, however, **is a fundamental disadvantage in the context of robust stability using the Small Gain Theorem** (Bates & Postlethwaite, 2002).

The $\mathcal{H}_\infty$ norm of a stable multivariable LTI system $G(s)$ is defined as the peak value over frequency of the largest singular value of the frequency response $G(jw)$:

$$\|G(s)\|_\infty = \max_\omega \sigma(G(j\omega)) \tag{3.5}$$

This norm is induced by the $\mathcal{L}_2$ signal norm, which means for an LTI system with input w, output z and transfer function matrix $G(s)$:

$$\|G(s)\|_\infty = \sup_{w \neq 0} \frac{\|z(t)\|_2}{\|w(t)\|_2} \tag{3.6}$$

Thus, minimisation of the $\mathcal{H}_\infty$ norm of a system corresponds to minimising the energy of the worst-case output signal vector, i.e. the output energy is minimised over all non-zero finite energy input signals.

3.1.3. Formal Statement of the $\mathcal{H}_\infty$ Control Problem

Given a real rational transfer matrix of the plant $P(s)$ and a space $\mathcal{K}$ of real rational transfer matrices $K(s)$, named the controller space, find the optimal solution $K^* \in \mathcal{K}$ to the following problem:

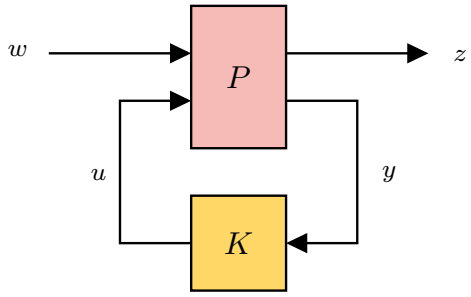

Figure 3.3: Standard LFT model of the $\mathcal{H}_\infty$ control problem.

$$\begin{aligned} \text{minimize} \quad & \|T_{w\to z}(P,K)\|_\infty \\ \text{subject to} \quad & K \text{ stabilizes } P \text{ internally} \\ & K \in \mathcal{K} \end{aligned} \tag{3.7}$$

where the plant $P(s)$ has a state-space representation of the form:

$$P: \begin{cases} \dot{x} = Ax + B_1 w + B_2 u \\ z = C_1 x + D_{11} w + D_{12} u \\ y = C_2 x + D_{21} w + D_{22} u \end{cases} \qquad P(s): \left[\begin{array}{c|cc} A & B_1 & B_2 \\ \hline C_1 & D_{11} & D_{12} \\ C_2 & D_{21} & D_{22} \end{array}\right] \tag{3.8}$$

where $x \in \mathbb{R}^{n_x}$ is the state, $u \in \mathbb{R}^{n_u}$ the control, $y \in \mathbb{R}^{n_y}$ the measured output, $w \in \mathbb{R}^{n_w}$ the exogenous input, and $z \in \mathbb{R}^{n_z}$ the regulated output.

Solutions of the $\mathcal{H}_\infty$-control problem in the full-order case, i.e. within the space $\mathcal{K}_{full}$ of full-order controllers, where $\mathcal{K}_{full} \cong \mathbb{R}^N$ and $N \triangleq n_x^2 + n_x(n_y + n_u) + n_y n_u$, are convex optimisation problems. The solution to these optimisation problems can be computed via Algebraic Riccati Equations (AREs) (Zhou et al., 1996) or via Linear Matrix Inequalities (LMIs) (Gahinet & Apkarian, 1994b), but the resulting controller is unstructured (Apkarian & Noll, 2017). For the sake of succinctness, the general $\mathcal{H}_\infty$ solution will not be provided, and rather, the focus will be merely on the $\mathcal{H}_\infty$ Loop-Shaping solution. Nevertheless, the reader is invited to read Zhou et al., 1996 for more details on the general $\mathcal{H}_\infty$ solution.

3.2. Robust Stability and Uncertainty Models

The generalised control diagram displayed in Figure 3.1 can be rearranged into the $N - \Delta$ description in Figure 3.4, where N represents the known part of the system (plant and controller) and Δ represents the uncertainty present in the system.

The uncertain closed-loop transfer function from w to z ($z = F(\Delta)w$) is related to N and Δ by an upper Linear Fractional Transformation (LFT) (Skogestad & Postlethwaite, 2005):

$$F = F_u(N, \Delta) = N_{22} + N_{21}\Delta(I - N_{11}\Delta)^{-1}N_{12} \tag{3.9}$$

If only robust stability is required for analysis, it is possible to rearrange the system in Figure 3.4 as the one in Figure 3.5, where $M = N_{11}$ is the transfer function from the output to the input of the perturbations.

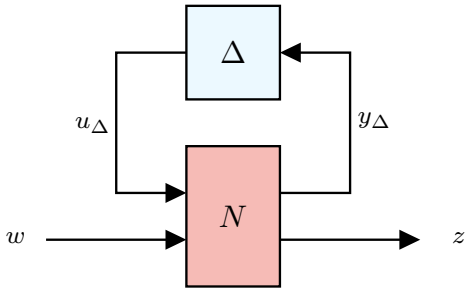

Figure 3.4: $N\Delta$ structure for robust performance analysis.

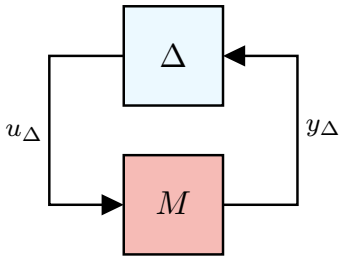

Figure 3.5: $M\Delta$ structure for robust stability analysis.

Assuming that the nominal closed-loop in Figure 3.4 is asymptotically stable and that Δ is a complex unstructured uncertainty block, the Small Gain Theorem dictates that (Bates & Postlethwaite, 2002):

The closed-loop system in Figure 3.4 is stable if and only if

$$\overline{\sigma}(\Delta(jw)) < \frac{1}{\overline{\sigma}(N_{11}(j\omega)} = \frac{1}{\overline{\sigma}(M(j\omega)} \quad \forall\, \omega \tag{3.10}$$

The above result defines a test for stability (and thus a robustness measure) for a closed-loop system subject to unstructured uncertainty in terms of the maximum singular value of the matrix N_{11} (Bates & Postlethwaite, 2002).

3.2.1. Unstructured Uncertainty

Unstructured uncertainty refers to a type of uncertainty which assumes unknown yet bounded uncertainty models. For an unstructured uncertain block $\Delta(s)$ that satisfies $\|\Delta(s)\|_\infty \leq 1$, the uncertain model $G_p(s)$ can be modelled in various ways (Skogestad & Postlethwaite, 2005):

- Additive uncertainty: $G_p = G_0 + w_A\Delta_A$
- Multiplicative input uncertainty: $G_p = G_0(I + w_I\Delta_I)$
- Multiplicative output uncertainty: $G_p = (I + w_O\Delta_O)G_0$
- Inverse additive uncertainty: $G_p = G_0(I - w_{iA}\Delta_{iA})^{-1}$
- Inverse multiplicative input uncertainty: $G_p = G_0(I - w_{iI}\Delta_{iI})^{-1}$
- Inverse multiplicative output uncertainty: $G_p = (I - w_{iO}\Delta_{iO})^{-1}G_0$

Over the set of multiple unstructured uncertainty descriptions, there is one additional one whose properties stand out: coprime factorisation.

Coprime Factorisation

Coprime factorisation is another form of representing unstructured system uncertainty, which has many advantages over additive and multiplicative uncertainty and leads directly to the $\mathcal{H}_\infty$ Loop-Shaping Design Procedure, which is at the heart of the present work.

Given a system G, the left coprime factorisation is given by:

$$G(s) = \tilde{M}^{-1}(s)\tilde{N}(s) \tag{3.11}$$

where $\tilde{M}(s)$ and $\tilde{N}(s)$ are stable coprime functions. The stability condition dictates that:

- $\tilde{N}(s)$ must encompass all the right-half plane (RHP) zeros of $G(s)$
- $\tilde{M}(s)$ should include all the RHP poles of $G(s)$

The coprimeness condition ensures that $\tilde{N}(s)$ and $\tilde{M}(s)$ have no common RHP zeros, thereby preventing pole-zero cancellations when forming $\tilde{M}^{-1}(s)\tilde{N}(s)$. In other words, coprimeness requires the existence of stable transfer functions or matrices, denoted as $U(s)$ and $V(s)$, such that the following equation holds:

$$\tilde{M}(s)V(s) + \tilde{N}(s)U(s) = I \tag{3.12}$$

where I represents the identity matrix.

Introducing the operator M^* defined as $M^*(s) = M^T(-s)$, $G(s) = \tilde{M}^{-1}(s)\tilde{N}(s)$ is called a normalized left coprime factorization iff:

$$\tilde{M}\tilde{M}^* + \tilde{N}\tilde{N}^* = I \tag{3.13}$$

From this description of a plant G, an uncertain plant model G_p can be written as (McFarlane & Glover, 1990):

$$G_p = \{(\tilde{M} + \Delta_M)^{-1}(\tilde{N} + \Delta_N) : ||[\Delta_N \ \Delta_M]||_\infty < \varepsilon\} \tag{3.14}$$

where Δ_M and Δ_N are stable unknown transfer matrices which represent the uncertainty in the nominal plant model G (see Figure 3.6).

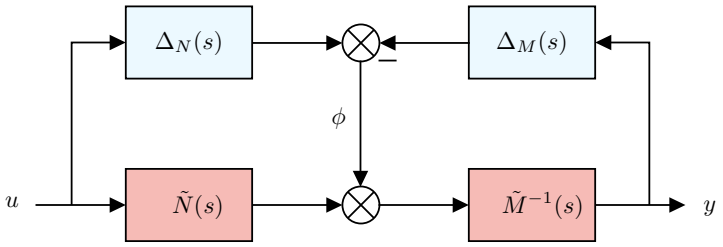

Figure 3.6: Coprime factor uncertainty representation of an uncertain plant model G_p.

Despite the fact that this model appears to have a less physically intuitive description of uncertainty than additive and multiplicative uncertainty, it has a number of important advantages, such as (Bates & Postlethwaite, 2002):

- Since it is not constrained to be a stable transfer matrix, it can be used to represent a situation where parameter variations result in a stable system becoming (open-loop) unstable.
- It allows for a sensible description of uncertainty in the locations of lightly damped resonant poles.

3.2.2. Structured Uncertainty

Structured uncertainty descriptions are more precise than unstructured descriptions and may model uncertainty directly on the parameters of the system, such as mass, inertia, and aerodynamic coefficients. This is referred to as parametric uncertainty.

Parametric uncertainty is quantified by assuming that each uncertain parameter is bounded within some region $[\alpha_{min}, \alpha_{max}]$, which results in the following uncertainty description (Skogestad & Postlethwaite, 2005):

$$\alpha_p = \overline{\alpha}(1 + r_\alpha \Delta) \tag{3.15}$$

where $\overline{\alpha}$ is the mean parameter value, $r_\alpha = (\alpha_{max} - \alpha_{min})/(\alpha_{max} + \alpha_{min})$ is the relative uncertainty in the parameter, and Δ is any real scalar satisfying $|\Delta| \leq 1$.

3.2.3. Structured Singular Value Robustness Measure μ

In cases where it is possible to generate more structured models of uncertainty, i.e. Δ has structure, it is possible to reduce the level of conservatism in the robustness analysis. This motivated the introduction of the Structured Singular Value (SSV) μ, defined as:

$$\mu(M) \triangleq \frac{1}{min\{k_m | det(I - k_m M\Delta) = 0 \text{ for structured } \Delta, \overline{\sigma}(\Delta) \leq 1\}} \tag{3.16}$$

where M corresponds to the block in Figure 3.5, and $\Delta = diag\{\Delta_i\}$ denotes a set of complex matrices with $\overline{\sigma}(\Delta) \leq 1$ and with a given block-diagonal structure. A value of $\mu = 1$ means there exists a perturbation with $\overline{\sigma}(\Delta) = 1$, which is just large enough to make $I - M\Delta$ singular. A larger value for μ is undesirable as it means that a smaller perturbation makes $I - M\Delta$ singular, while a smaller value of μ is positive. The SSV value μ forms the basis for the $\mathcal{H}_\infty$ μ-synthesis control technique (Skogestad & Postlethwaite, 2005).

3.3. $\mathcal{H}_\infty$ Loop-Shaping

3.3.1. NCF Robust Stabilization

$\mathcal{H}_\infty$ Loop-Shaping is anchored in the concept of NCF robust stabilisation, that is, to design a controller which not only nominally stabilises the nominal plant G but also robustly stabilises the entire set of NCF perturbed plants G_p.

The NCF robust stabilisation theorem dictates that the controller $K(s)$ which robustly stabilises the family of NCF perturbed plants $G_p(s)$ for all NCF perturbations $\Delta(s)$ with $\|\Delta(s)\|_\infty = \|[\Delta_N \ \Delta_M]\|_\infty < \epsilon$ satisfies the following two conditions (McFarlane & Glover, 1990):

- The controller $K(s)$ internally stabilizes $G(s)$
- The following condition holds: $\left\| \begin{bmatrix} -K(I_p + GK)^{-1}\tilde{M}^{-1} \\ (I_p + GK)^{-1}\tilde{M}^{-1} \end{bmatrix} \right\|_\infty \leq \epsilon^{-1}$

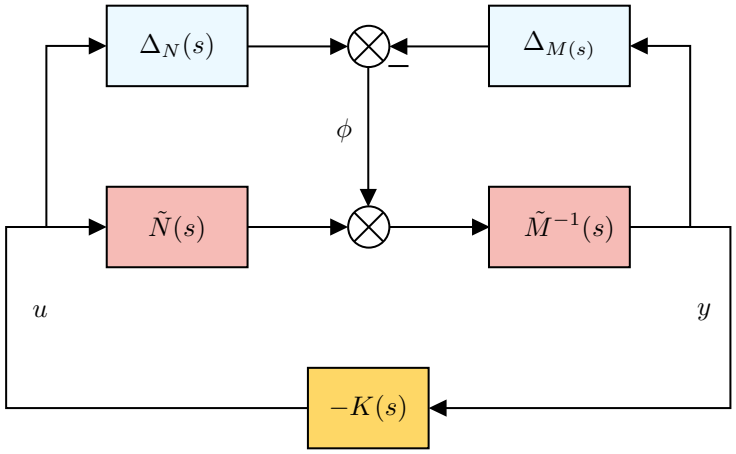

Figure 3.7: NCF perturbed plant G_p with controller.

It is possible to demonstrate with some manipulations that the NCF robust stabilisation can be rewritten as a 4-block problem, as follows (Zhou & Doyle, 1998):

$$\left\| \begin{bmatrix} -K(I_p + GK)^{-1}\tilde{M}^{-1} \\ (I_p + GK)^{-1}\tilde{M}^{-1} \end{bmatrix} \right\|_\infty \leq \epsilon^{-1} \Leftrightarrow \left\| \begin{bmatrix} S_o & S_oG \\ -KS_o & -T_i \end{bmatrix} \right\|_\infty \leq \epsilon^{-1} \tag{3.17}$$

where the obtained closed-loop 2×2 system is basically the transfer function from the output and input disturbance $[d_o\ d_i]^T$ to the system and controller outputs $[y\ u]^T$, respectively:

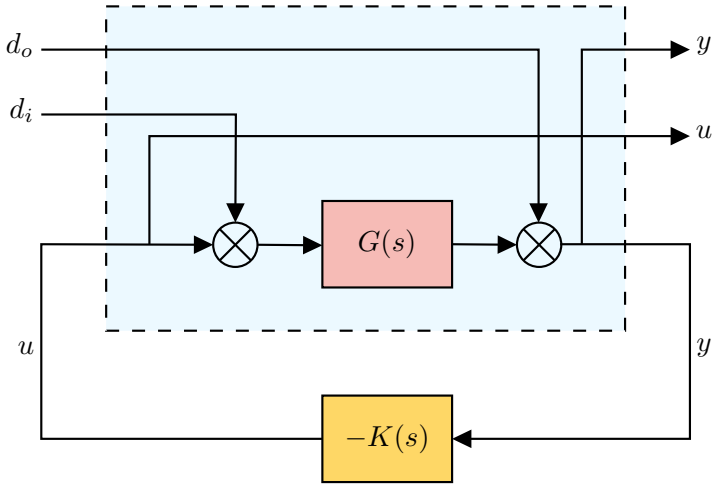

Figure 3.8: NCF robust stabilisation written as a 4-block problem.

What makes NCF robust stabilisation particularly appealing is that the maximum stability margin ϵ_{max} or the minimum performance level γ_{min} can be computed directly from:

$$\gamma_{min} = \epsilon_{max}^{-1} = \left\{ 1 - \left\| \begin{bmatrix} N & M \end{bmatrix} \right\|_H^2 \right\}^{-\frac{1}{2}} = (1 + \rho(XZ))^{\frac{1}{2}}, \tag{3.18}$$

where $\|\cdot\|_H$ denotes the Hankel norm, ρ denotes the spectral radius, and for a minimal state-space realisation (A,B,C,D) of G, Z and X are the unique positive solution to the following algebraic Riccati equations:

$$(A - BS^{-1}D^TC)Z + Z(A - BS^{-1}D^TC)^T - ZC^TR^{-1}CZ + BS^{-1}B^T = 0 \tag{3.19}$$

$$(A - BS^{-1}D^TC)^TX + X(A - BS^{-1}D^TC) - XBS^{-1}B^TX + C^TR^{-1}C = 0 \tag{3.20}$$

where R and S are defined as:

$$R = I + DD^T, \quad S = I + D^TD \tag{3.21}$$

The formulae simplify considerably for a strictly proper plant, i.e. when $D = 0$. The $\mathcal{H}_\infty$ controller can be parameterised in infinitely different ways, according to the choice of the Youla Parameterisation Q (Zhou et al., 1996). If the Youla Parameterisation Q is set equal to 0, the solution to the central controller is obtained for a specified suboptimal $\gamma > \gamma_{\text{min}}$:

$$K = \begin{bmatrix} A + BF + \gamma^2(L^T)^{-1}ZC^T(C + DF) & \gamma^2(L^T)^{-1}ZC^T \\ B^TX & -D^T \end{bmatrix} \tag{3.22}$$

$$F = -S^{-1}(D^TC + B^TX) \tag{3.23}$$

$$L = (1 - \gamma^2)I + XZ. \tag{3.24}$$

However, robust stabilisation against NCF uncertainty alone is of limited practical value because it does not allow the designer to specify any performance requirements (McFarlane & Glover, 1992). The multiple control requirements are reflected by the broken-loop response of the controlled system. The $\mathcal{H}_\infty$ Loop-Shaping Design Procedure (LSDP) leverages exactly this property to present a systematic way of designing robust controllers.

In a first step, the plant is shaped with weighting filters W_1 and W_2 such that the shaped plant $G_s = W_2GW_1$ accommodates the desired Loop-Shaping response. In the second step, a robust controller K_∞ is computed, which robustly stabilises the shaped plant G_s with respect to NCF uncertainty. It is possible to demonstrate that the degradation of the initial loop shape caused by the $\mathcal{H}_\infty$ controller K_∞ is limited and, therefore, the weighting filters W_1 and W_2 translate directly to the fulfilment of design requirements.

3.3.2. Design Requirements

An appropriately designed control system must satisfy a set of design requirements. Apart from closed-loop stability, certain closed-loop relationships should meet some frequency domain criteria. Consider the following one-degree-of-freedom feedback configuration:

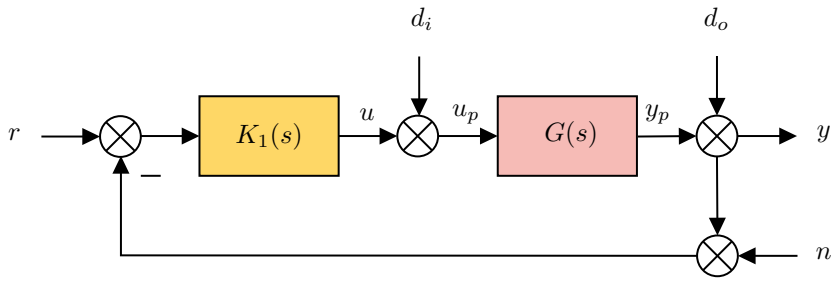

Figure 3.9: Generalised feedback configuration.

which has the following closed-loop transfer functions:

$$\begin{bmatrix} y \\ u_p \\ u \\ e \end{bmatrix} = \begin{bmatrix} S_oG & S_o & -T_o & T_o \\ S_i & -KS_o & -KS_o & KS_o \\ -T_i & -KS_o & -KS_o & KS_o \\ -S_oG & -S_o & T_o & S_o \end{bmatrix} \begin{bmatrix} d_i \\ d_o \\ n \\ r \end{bmatrix} \tag{3.25}$$

The general design requirements impose the following:

- Disturbance rejection and reference tracking
  - Reduce the effect of disturbances d_i, d_o to y (S_oG,S_o) and to u_p (S_i,KS_o)
  - Reduce the 'effect' of r to the tracking error e (S_o).
- Noise attenuation
  - Reduce the effect of noise n to u, y (KS_o, T_o)
- Control signal reduction
  - Reduce the effect of d_i and d_o, r to u (T_i, KS_o)

Therefore, meeting the design requirements boils down to adequately shaping the 'gang of six' CL transfer functions S_i, S_o, T_i, T_o, S_oG and KS_o. Some design requirements can have a conflicting nature, since it is not possible to simultaneously make $\overline{\sigma}[S_o] \ll 1$ for disturbance rejection and $\overline{\sigma}[T_o] \ll 1$ for noise attenuation, given that $S_o + T_o = I$. Nevertheless, different design objectives are often important over different frequency ranges, meaning that feedback controller design is a problem of managing trade-offs between conflicting design objectives. For instance, disturbance rejection and reference tracking are typically important at low frequencies, while noise attenuation is often relevant only at high frequencies.

At low frequencies, disturbance rejection and reference tracking dictate that:

- From $d_o \to y$ and $r \to e$: $\overline{\sigma}[S_o] \ll 1$
- From $d_i \to u_p$: $\overline{\sigma}[S_i] \ll 1$
- From $d_i \to y$: $\overline{\sigma}[S_oG] \ll 1$
- From $d_o \to u_p$: $\overline{\sigma}[KS_o] \ll 1$

At high frequencies, noise attenuation dictates:

- From $n \to y$: $\overline{\sigma}[T_o] \ll 1$
- From $n \to u$: $\overline{\sigma}[KS_o] \ll 1$

3.3.3. Fundamental Trade-offs in terms of L

The fundamental design trade-offs presented above can be written in terms of the input and output open-loop transfer matrices $L_i = KG$ and $L_o = GK$. The derivations are retrieved from Zhou et al., 1996 and make use of the following properties:

$$\overline{\sigma}(A^{-1}) = \frac{1}{\underline{\sigma}(A)} \tag{3.26}$$

$$\underline{\sigma}(A) - 1 \leq \frac{1}{\overline{\sigma}(I + A)^{-1}} \leq \underline{\sigma}(A) + 1 \tag{3.27}$$

Given that:

$$\overline{\sigma}[S_o] = \overline{\sigma}[(I_p + L_o)^{-1}] = \frac{1}{\underline{\sigma}[I_p + L_o]} \tag{3.28}$$

then, using Equation 3.27:

$$\frac{1}{\underline{\sigma}\,[L_o] + 1} \leq \overline{\sigma}\,[S_0] \leq \frac{1}{\underline{\sigma}\,[L_o] - 1} \text{ with } L_o = GK \tag{3.29}$$

Similarly, for S_i:

$$\frac{1}{\underline{\sigma}\,[L_i] + 1} \leq \overline{\sigma}\,[S_i] \leq \frac{1}{\underline{\sigma}\,[L_i] - 1} \text{ with } L_i = KG \tag{3.30}$$

At low frequencies, this results in:

$$\overline{\sigma}\,[S_o] \ll 1 \Leftrightarrow \underline{\sigma}\,[L_o] \gg 1$$
$$\overline{\sigma}\,[S_i] \ll 1 \Leftrightarrow \underline{\sigma}\,[L_i] \gg 1$$

Nevertheless, as shown in Section 3.3.2, disturbance rejection at low frequencies still requires that $\overline{\sigma}[S_oG] \ll 1$ and $\overline{\sigma}[KS_o] \ll 1$.

Assuming K and G are invertible:

$$\overline{\sigma}[S_oG] = \overline{\sigma}[(S_oGK)K^{-1}] = \overline{\sigma}[T_oK^{-1}] = \overline{\sigma}[(I_p - S_o)K^{-1}] \tag{3.31}$$

Using the fact that $\overline{\sigma}[S_o] \ll 1$ and $\overline{\sigma}[K^{-1}] = 1/\underline{\sigma}[K]$, for frequencies $0 \leq w \leq w_l$, this reduces to:

$$\overline{\sigma}[S_oG] \simeq \frac{1}{\underline{\sigma}[K]} \tag{3.32}$$

Hence, to obtain $\overline{\sigma}[S_oG] \ll 1$, large control gains $\underline{\sigma}[K] \gg 1$ is required (for instance, integral control).

For KS_o, using the following simplifications:

$$\overline{\sigma}[KS_o] = \overline{\sigma}[G^{-1}(GK_{S_o})] = \overline{\sigma}[G^{-1}T_o] = \overline{\sigma}[G^{-1}(I_p - S_o)] \tag{3.33}$$

and using the fact that $\overline{\sigma}[S_o] \ll 1$ and $\overline{\sigma}[G^{-1}] = 1/\underline{\sigma}[G]$ for frequencies $0 \leq w \leq w_l$, this reduces to:

$$\overline{\sigma}[KS_o] \simeq \frac{1}{\underline{\sigma}[G]} \tag{3.34}$$

Hence, to obtain $\overline{\sigma}[KS_o] \ll 1$, a large plant gain (sufficient control authority) is required, and this relationship is independent of the controller.

Therefore, for low frequencies to guarantee good disturbance rejection and tracking qualities, it is necessary to have:

- Sufficient control authority $\underline{\sigma}[G] \gg 1$
- High open-loop gain $\underline{\sigma}[L_o] \gg 1$, $\underline{\sigma}[L_i] \gg 1$
- High control gain $\underline{\sigma}[K] \gg 1$

Similarly, for the high-frequency $w_h < w < +\infty$ requirements, using the following SV property:

$$1 - \overline{\sigma}[L_o] \leq \underline{\sigma}[I_p + L_o] \leq 1 + \overline{\sigma}[L_o] \tag{3.35}$$

or, after inverting:

$$\frac{1}{1 + \overline{\sigma}[L_o]} \leq \overline{\sigma}[S_o] \leq \frac{1}{1 - \overline{\sigma}[L_o]} \quad \text{with } L_o = GK \tag{3.36}$$

it is possible to bound T_o as:

$$\overline{\sigma}[T_o] = \overline{\sigma}[L_o S_o] \leq \overline{\sigma}[L_o]\overline{\sigma}[S_o] \tag{3.37}$$

Then, from Equation 3.36:

$$\overline{\sigma}[T_o] \ll 1 \Leftrightarrow \overline{\sigma}[L_o] \ll 1 \tag{3.38}$$

which is a necessary condition for noise attenuation at high frequency.

In order to meet the condition $\overline{\sigma}[-KS_o] \ll 1$, notice that:

$$\overline{\sigma}[KS_o] \leq \overline{\sigma}[K]\overline{\sigma}[S_o] \tag{3.39}$$

Using Equation 3.36 and the fact that $\overline{\sigma}[L_o] \ll 1$ at high frequency, then $\overline{\sigma}[S_o] \simeq 1$ and thus:

$$\overline{\sigma}[KS_o] \leq \overline{\sigma}[K] \tag{3.40}$$

Therefore, for high frequencies to guarantee good noise attenuation and low control action, it is necessary to have:

- Low open-loop gain $\underline{\sigma}[L_o] \ll 1$, $\underline{\sigma}[L_i] \ll 1$
- Low control gain $\underline{\sigma}[K] \ll K_{max}$ assuming K_{max} is not too large.

As a result, in order to meet all the design requirements, it is necessary to have an adequate plant which provides enough control authority, an open-loop response as exhibited in Figure 3.10 and an appropriate controller structure to ensure high gain of K at low frequencies (integral action) and reduced control action at high frequencies (through low-pass filters). In general, assuming the control engineer is given a feasible task (the plant G has enough control authority), meeting all the design requirements boils down to having an adequate open-loop gain response and control structure. The ability to translate multiple closed-loop design requirements into just two open-loop responses offers a significant advantage to control designers. This simplification forms the foundation of classical Loop-Shaping methods and serves as the cornerstone of the more advanced $\mathcal{H}_\infty$ Loop-Shaping Design Procedure (LSDP).

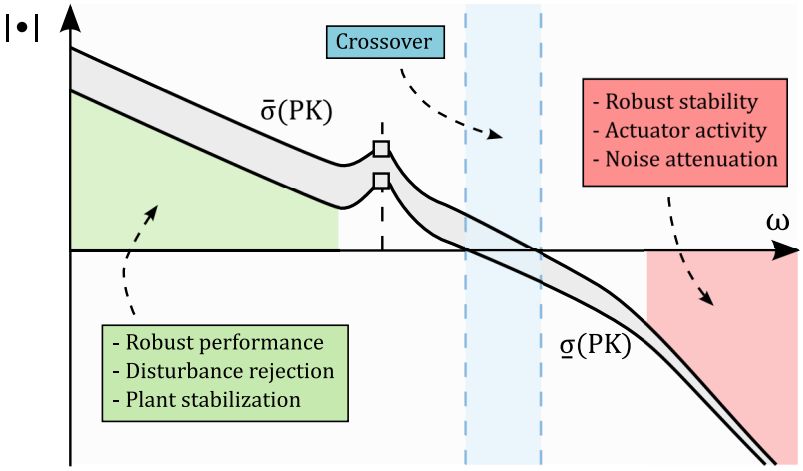

Figure 3.10: Singular value loop specifications for a successful design. Retrieved from Pollack, 2024.

3.3.4. $\mathcal{H}_\infty$ Loop-Shaping Design Procedure

The $\mathcal{H}_\infty$ Loop-Shaping Design Procedure, originally proposed by McFarlane and Glover, 1992, leverages the concepts of classical Loop-Shaping and the $\mathcal{H}_\infty$ machinery to design controllers which robustly stabilise a plant with respect to coprime factor uncertainty. The procedure is described as follows:

1. Shape the open-loop plant G, with frequency-dependent weights W_1 and W_2, according to closed-loop objectives (recall the discussion in Sections 3.3.2 and 3.3.3). The nominal plant G and the shaping functions W_1, W_2 are combined to form the shaped plant, G_s, where $G_s = W_2 G W_1$. It is assumed that W_1 and W_2 are such that G_s contains no hidden modes.

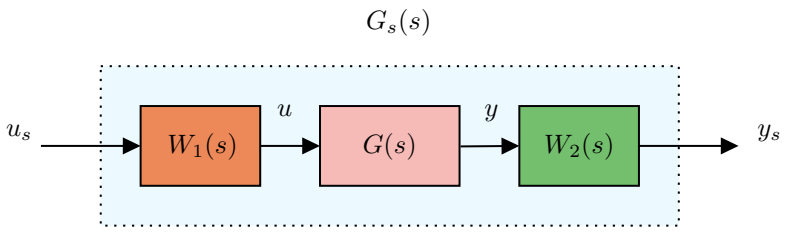

Figure 3.11: Step 1 - Shaping the open-loop plant G with W_1 and W_2.

2. Robust stabilisation:

   (a) Compute the ϵ_{max} associated with the shaped-plant G_s:

$$\epsilon_{max} \triangleq \gamma_{min}^{-1} = \left(\inf_{K \text{ stabilizing}} \left\| \begin{bmatrix} I \\ K \end{bmatrix} (I + G_s K)^{-1} \tilde{M}_s^{-1} \right\|_\infty \right)^{-1} = \sqrt{1 - \left\| \begin{bmatrix} \tilde{N}_s & \tilde{M}_s \end{bmatrix} \right\|_H^2} < 1 \tag{3.41}$$

   where $\tilde{M}_s, \tilde{N}_s$ define the normalized coprime factors of G_s such that $G_s = \tilde{M}_s^{-1}\tilde{N}_s$ and $\tilde{M}_s\tilde{M}_s^* + \tilde{N}_s\tilde{N}_s^* = I$. If $\epsilon_{max} << 1$ or $\gamma_{min} >> 1$ go back to step 1 and adjust W_1 and W_2. In fact, practical guidelines suggest that γ should be below 4 (Hyde, 1995). Therefore, if $\gamma > 4$, the designer should also go back to step 1 and adjust the weighting filters.

   (b) Select a sub-optimal $\epsilon < \epsilon_{max}$ or, in other words, $\gamma > \gamma_{min}$ and synthesize the robust stabilizing controller K_∞ that minimizes the following $\mathcal{H}_\infty$ norm (recall the solution in terms of two AREs in Section 3.3.1):

$$\left\| \begin{bmatrix} I \\ K_\infty \end{bmatrix} (I + G_s K_\infty)^{-1} \tilde{M}_s^{-1} \right\|_\infty \leq \epsilon^{-1} = \gamma \tag{3.42}$$

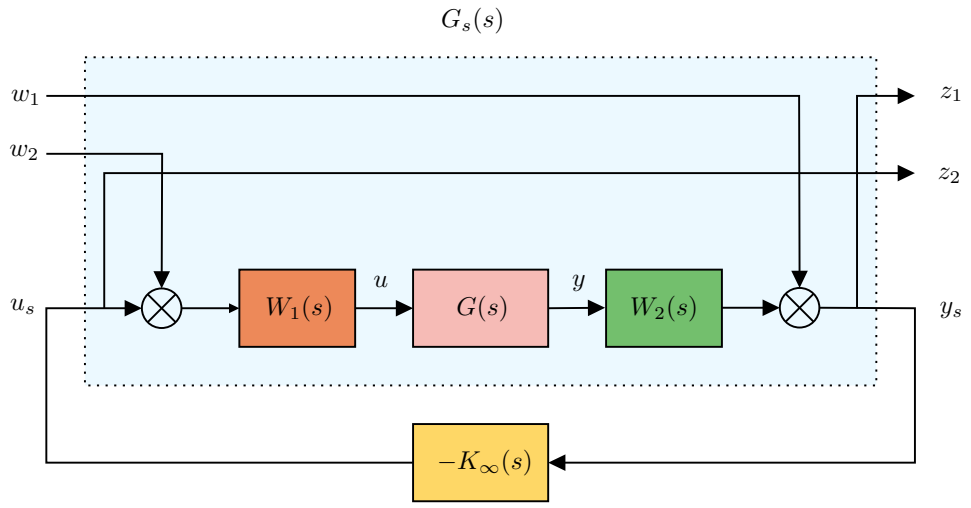

Figure 3.12: $\mathcal{H}_\infty$ Loop-Shaping as a 4-block problem.

3. Construct the final feedback controller K_s by combining the $\mathcal{H}_\infty$ controller K_∞ with the shaping functions W_1 and W_2, such that $K_s = W_1 K_\infty W_2$. When tracking of a reference signal r is required, there are three options for the choice of K_∞ in the loop: placed in the forward path, the feedback path or implemented in the observer form. For a more thorough discussion on this topic, refer to G. Papageorgiou, 1998.

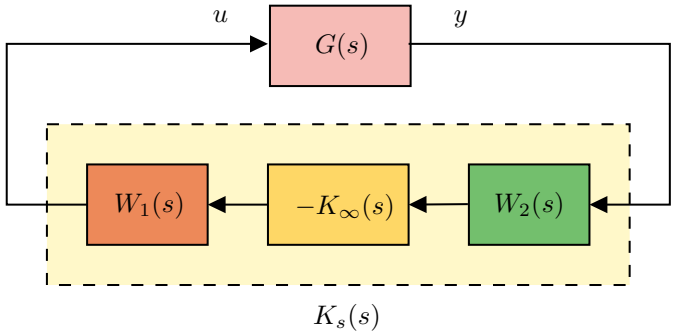

Figure 3.13: Final implemented robust controller consisting of $K_s = W_1 K_\infty W_2$.

The theoretical basis for $\mathcal{H}_\infty$ LS is that K_∞ does not modify the desired loop shape significantly at low and high frequencies if γ is not too large. Therefore, the shaping of G via W_1 and W_2 corresponds to shaping the loop gains GK_s and K_sG. Moreover, all closed-loop transfer functions can be bounded in terms of G, $W1$, $W2$ and γ (Zhou et al., 1996). The bounds on L_o and L_i are given by:

At Low-Frequencies (where $c(A) = \overline{\sigma}(A)/\underline{\sigma}(A)$):

$$\underline{\sigma}(L_o) = \underline{\sigma}(GK_s) \geq \underline{\sigma}(G_s)\underline{\sigma}(K_\infty)/c(W_2) \tag{3.43}$$

$$\underline{\sigma}(L_i) = \underline{\sigma}(K_sG) \geq \underline{\sigma}(G_s)\underline{\sigma}(K_\infty)/c(W_1) \tag{3.44}$$

At High-Frequencies:

$$\overline{\sigma}(L_o) = \overline{\sigma}(GK_s) \leq \overline{\sigma}(G_s)\overline{\sigma}(K_\infty)c(W_2) \tag{3.45}$$

$$\overline{\sigma}(L_i) = \overline{\sigma}(K_sG) \leq \overline{\sigma}(G_s)\overline{\sigma}(K_\infty)c(W_1) \tag{3.46}$$

This demonstrates that bounds on $\underline{\sigma}(K_\infty)$ and $\overline{\sigma}(K_\infty)$ are required to obtain a bound on the deterioration of the loop shape at low-frequency and high-frequency, respectively.

It is possible to bound $\underline{\sigma}(K_\infty)$ by (McFarlane & Glover, 1992):

$$\underline{\sigma}(K_\infty) \geq \frac{\underline{\sigma}(G_s) - \sqrt{\gamma^2 - 1}}{\underline{\sigma}(G_s)\sqrt{\gamma^2 - 1} + 1}, \quad \forall \omega \text{ such that } \underline{\sigma}(G_s) > \sqrt{\gamma^2 - 1} \tag{3.47}$$

Furthermore, if the loop shape G_s has very large gain at low-frequency $\underline{\sigma}(G_s) >> \sqrt{\gamma^2 - 1}$, then tighter bounds can be derived (McFarlane & Glover, 1992):

$$\underline{\sigma}(K_\infty) \gtrsim \frac{1}{\sqrt{\gamma^2 - 1}} \Rightarrow \underline{\sigma}(L_o) \geq \sigma(G_s)\left(\frac{1}{\sqrt{\gamma^2 - 1}} c(W_2)^{-1}\right) \tag{3.48}$$

Similarly, it is also possible to bound $\overline{\sigma}(K_\infty)$ by (McFarlane & Glover, 1992):

$$\overline{\sigma}(K_\infty) \leq \frac{\sqrt{\gamma^2-1}+\overline{\sigma}(G_s)}{1-\overline{\sigma}(G_s)\sqrt{\gamma^2-1}} \quad \forall\omega \text{ such that } \overline{\sigma}(G_s) < \frac{1}{\sqrt{\gamma^2-1}} \tag{3.49}$$

Moreover, if the loop shape G_s has very low gain at high-frequency $\overline{\sigma}(G_s) << \frac{1}{\sqrt{\gamma^2-1}}$, then (McFarlane & Glover, 1992):

$$\overline{\sigma}(K_\infty) \lesssim \sqrt{\gamma^2-1} \quad \text{as} \quad \underline{\sigma}(G_s) \to 0 \Rightarrow \overline{\sigma}(L_o) \leq \overline{\sigma}(G_s)\left(c(W_2)\sqrt{\gamma^2-1}\right) \tag{3.50}$$

Similar reasoning can be used to derive tighter bounds for $\underline{\sigma}(L_i)$ and $\overline{\sigma}(L_i)$ analogous to the ones in Equation 3.48 and 3.50.

Besides, assuming $\underline{\sigma}(G_s) >> 1$ at low-frequency and $\overline{\sigma}(G_s) << 1$ at high-frequency, the following upper bounds for the "gang of six" closed-loop transfer functions can be computed (McFarlane & Glover, 1990):

Table 3.1: Upper bounds on the different closed-loop transfer functions at Low-Frequency (LF) and High-Frequency (HF) (McFarlane & Glover, 1990).

Objective	LF Upper Bound	HF Upper Bound
$\overline{\sigma}(S_o) = \overline{\sigma}\left((I_p + GK_s)^{-1}\right)$	$\frac{\gamma}{\underline{\sigma}(G)\underline{\sigma}(W_1)\underline{\sigma}(W_2)}$	$\gamma c(W_2)$
$\overline{\sigma}(K_sS_o) = \overline{\sigma}\left(K_s(I_p + GK_s)^{-1}\right)$	$\frac{\gamma c(W_1)}{\underline{\sigma}(G)}$	$\gamma\overline{\sigma}(W_1)\overline{\sigma}(W_2)$
$\overline{\sigma}(T_i) = \overline{\sigma}\left(K_sG(I_m + K_sG)^{-1}\right)$	$\gamma c(W_1)$	$\gamma\overline{\sigma}(G)\overline{\sigma}(W_1)\overline{\sigma}(W_2)$
$\overline{\sigma}(S_oG) = \overline{\sigma}\left((I_p + GK_s)^{-1}G\right)$	$\frac{\gamma}{\underline{\sigma}(W_1)\underline{\sigma}(W_2)}$	$\gamma\overline{\sigma}(G)c(W_2)$
$\overline{\sigma}(S_i) = \overline{\sigma}\left((I_m + K_sG)^{-1}\right)$	$1 + \gamma c(W_1)$	$\gamma c(W_1)$
$\overline{\sigma}(T_o) = \overline{\sigma}\left(GK_s(I_p + GK_s)^{-1}\right)$	$\gamma c(W_2)$	$1 + \gamma c(W_2)$

The non-deterioration of the open-loop response after the robust controller K_∞ is added is a key aspect of the procedure. Since the weighting filters W_1 and W_2 embed multiple competing design goals, such as disturbance rejection, robustness, noise attenuation, and the crossover frequency, the non-deterioration properties guarantee that the controller does not undo or significantly compromise the design intent captured in the shaped plant.

As a final remark on the $\mathcal{H}_\infty$ LSDP, notice how the Loop-Shaping is done without explicit regard for the nominal phase information. This means that closed-loop stability requirements are disregarded at the initial stage. Overall, the **design procedure is both simple and systematic, assuming only elementary knowledge of Loop-Shaping principles by the designer** (Zhou & Doyle, 1998).

3.3.5. Guaranteed Robustness Properties of NCF Robust Stabilisation

Glover et al., 2000 has demonstrated that the Robust Stability to NCF uncertainty, evaluated through the coprime factor singular value robustness measure $\gamma = \epsilon^{-1}$ as described in Equation 3.42 can also be interpreted in terms of classical gain and phase margins for SISO and MIMO systems.

- SISO Case:

$$-20\log_{10}\left(\frac{1+\epsilon}{1-\epsilon}\right) \text{ dB} \leq GM_{ncf} \leq 20\log_{10}\left(\frac{1+\epsilon}{1-\epsilon}\right) \text{ dB} \tag{3.51}$$

$$-2\,\text{asin}\,\epsilon \text{ deg} \leq PM_{ncf} \leq 2\,\text{asin}\,\epsilon \text{ deg} \tag{3.52}$$

- MIMO Case, assuming diagonal weighting matrices (Glover et al., 2000):

$$-20\log_{10}\left(\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right) \text{ dB} \leq GM_{ncf} \leq 20\log_{10}\left(\sqrt{\frac{1+\epsilon}{1-\epsilon}}\right) \text{ dB} \tag{3.53}$$

$$-\,\text{asin}\,\epsilon \text{ deg} \leq PM_{ncf} \leq \text{asin}\,\epsilon \text{ deg} \tag{3.54}$$

Therefore, ϵ is a measure of how much simultaneous gain and phase variation the system can tolerate, and for the MIMO case, these margins are halved. This leads to guaranteed Nichols exclusion regions, which are extremely useful for flight control applications. Thus, the fact that $\mathcal{H}_\infty$ Loop-Shaping comes with guaranteed phase and gain margins makes it an appealing control technique for flight control laws design.

3.3.6. Connections to the gap metric and the ν-gap metric

The operator norm can be a poor measure of the distance between systems with respect to feedback control system design (Zhou & Doyle, 1998). This motivated the introduction of the gap metric into the control literature by El-Sakkary, 1985 and, subsequently, the ν-gap metric by Vinnicombe, 2000 as more appropriate tools for the study of uncertainty in feedback systems. These metrics quantify the distance between (potentially unstable) systems, even when they have different numbers of right-half-plane (RHP) poles.

The gap metric is defined as follows (El-Sakkary, 1985):

$$\delta_g(P_1, P_2) = max\left\{\tilde{\delta}(P_1, P_2), \tilde{\delta}(P_2, P_1)\right\} \tag{3.55}$$

where $\vec{\delta}_g$ represents the directed gap, which is computed as:

$$\vec{\delta}_g(P_1, P_2) = \inf_{Q\in\mathcal{H}_\infty} \left\| \begin{bmatrix} M_1 \\ N_1 \end{bmatrix} - \begin{bmatrix} M_2 \\ N_2 \end{bmatrix} Q \right\|_\infty \tag{3.56}$$

and where $P_1 = N_1 M_1^{-1}$ and $P_2 = N_2 M_2^{-1}$ are normalized right coprime factorizations.

An important connection between the uncertainties in the gap metric and the uncertainties characterised by the normalised coprime factors is that a ball of uncertainty in the directed gap is equivalent to a ball of uncertainty in the normalised coprime factors (Georgiou & Smith, 1989).

However, unlike the ν-gap metric, a shortcoming of the gap metric δ_g is that it is not easily related to the frequency response of the system. The ν-gap metric has a clear frequency domain interpretation and can, in general, be computed from the frequency response. The ν-gap metric between two plants P_1 and P_2 can be computed directly from the system transfer matrices without first finding the normalised coprime factorisations as (Vinnicombe, 2000):

$$\delta_\nu(P_1, P_2) = \max_\omega \left\| (I + P_2 P_2^*)^{-1/2}(P_1 - P_2)(I + P_1 P_1^*)^{-1/2} \right\|_\infty \tag{3.57}$$

provided that $det(I + P_2^* P_1)$ has the right winding number (Zhou & Doyle, 1998).

For both metrics, the following robust performance result holds (Zhou & Doyle, 1998):

$$\arcsin(b(P_2, K_2)) \geq \arcsin(b(P_1, K_1)) - \arcsin\big(\delta_{(.)}(P_1, P_2)\big) - \arcsin\big(\delta_{(.)}(K_1, K_2)\big) \tag{3.58}$$

where $b(P, K)$ refers to the generalised stability margin of the standard feedback interconnection between plant model P and controller K:

$$b(P,K) = \left\| \begin{bmatrix} I \\ K \end{bmatrix} (I + PK)^{-1} \begin{bmatrix} I & P \end{bmatrix} \right\|_\infty^{-1} = \left\| \begin{bmatrix} I \\ P \end{bmatrix} (I + KP)^{-1} \begin{bmatrix} I & K \end{bmatrix} \right\|_\infty^{-1} \tag{3.59}$$

and corresponds to the ϵ value, or γ^{-1}, associated with the NCF robust stabilization problem, i.e., the inverse of the gain from disturbances on the plant input and output to the input and output of the controller in Figure 3.12.

This result can be interpreted as follows: suppose that a nominal plant P_1 is stabilised by a controller K_1 with stability margin $b(P_1, K_1)$. Then, if P_1 is perturbed to P_2 and K_1 is perturbed to K_2, the stability margin is degraded by no more than the above formula. However, since the ν-gap metric is always less than or equal to the gap metric, its predictions using the above robustness results are tighter (Zhou & Doyle, 1998).

3.4. Non-Convex Optimization

The full-order solutions of the $\mathcal{H}_\infty$-control problem, computed via AREs or LMIs, are unstructured in nature (Apkarian & Noll, 2017). A downside is that it does not allow for control law structures to be specified a priori, which is often required in engineering practice. This motivated the development of non-convex, non-smooth formal synthesis techniques, which allow the solution of the $\mathcal{H}_\infty$ problem under structural constraints (Apkarian & Noll, 2006a). The framework of non-smooth $\mathcal{H}_\infty$-synthesis was proposed by Apkarian and Noll, 2006a and can be described as follows:

Assuming that $K(\boldsymbol{\kappa})$ is structured with parameter $\boldsymbol{\kappa} \in \mathbb{R}^n$, and that $D_{22} = 0$ in Equation 3.8, then the closed-loop transfer channel $w \to z$ in Figure 3.3 has the following state-space representation:

$$T_{w\to z}\left(P, K(\boldsymbol{\kappa})\right) := \begin{bmatrix} A(K(\boldsymbol{\kappa})) & B(K(\boldsymbol{\kappa})) \\ C(K(\boldsymbol{\kappa})) & D(K(\boldsymbol{\kappa})) \end{bmatrix} \tag{3.60}$$

Then, the $\mathcal{H}_\infty$-objective function in Equation 3.7 becomes:

$$\begin{aligned} \min_{\boldsymbol{\kappa}} f(\boldsymbol{\kappa}) &\triangleq \min_{\boldsymbol{\kappa}} \left\| T_{w\to z}\left(P, K(\boldsymbol{\kappa})\right) \right\|_\infty \\ &= \min_{\boldsymbol{\kappa}} \max_{\omega\in\mathbb{R}} \overline{\sigma}\left(C(K(\boldsymbol{\kappa}))\left(j\omega I - A(K(\boldsymbol{\kappa}))\right)^{-1} B(K(\boldsymbol{\kappa})) + D(K(\boldsymbol{\kappa})) \right) \end{aligned} \tag{3.61}$$

which is a non-smooth, non-convex function, which, in addition, is not defined everywhere. Its domain $D_f = \{\boldsymbol{\kappa} \in \mathbb{R}^n : f(\boldsymbol{\kappa}) < \infty\}$ contains the internally stabilizing set:

$$D_s = \{\boldsymbol{\kappa} \in \mathbb{R}^n : K(\boldsymbol{\kappa}) \text{ stabilizes } P \text{ internally}\} = \{\boldsymbol{\kappa} \in \mathbb{R}^n : A(K(\boldsymbol{\kappa})) \text{ stable}\} \tag{3.62}$$

The non-smooth, non-convex $\mathcal{H}_\infty$ problem can be solved by local function minimisation based on generalised gradients, and the tools to solve such a class of problems became available to the general public in 2011 through MATLAB® `hinfstruct` (Apkarian & Noll, 2017; Gahinet & Apkarian, 2011). The introduction of this tool marked a paradigm shift in the way controllers are synthesised, since it allows for structured controllers, such as PID controllers, to be formally optimised instead of tuned (Apkarian & Noll, 2017). Following this breakthrough, multiple extensions of the method were incorporated to MATLAB® `systune` (Apkarian et al., 2014), such as multi-objective H_2/H_∞ synthesis (Apkarian et al., 2008), multidisk design problems (Apkarian & Noll, 2006b; Apkarian et al., 2014), and limited frequency intervals (Apkarian & Noll, 2007). Other relevant developments include the concept of inner approximation to handle parametric uncertainty (Apkarian et al., 2015) and new applications to systematic gain scheduling (Gahinet & Apkarian, 2013). In this context, `systune`, which became available in 2015, can be seen as a successor of the `hinfstruct` tool. One of the key features that differentiates it from its predecessor is that it enables multi-model, multi-objective synthesis consisting of soft control objectives $f(\kappa)$ and hard constraints $g(\kappa)$ (Apkarian et al., 2014):

$$\min_{\kappa} \left(\max_{i \in [1,\ldots,n_f]} f_i(\kappa) \right) \text{ such that } \left(\max_{j \in [1,\ldots,n_g]} g_j(\kappa) \leq 1 \right) \tag{3.63}$$

which can be described in terms of the closed-loop map from signals w to z $T_{w \to z}$ as:

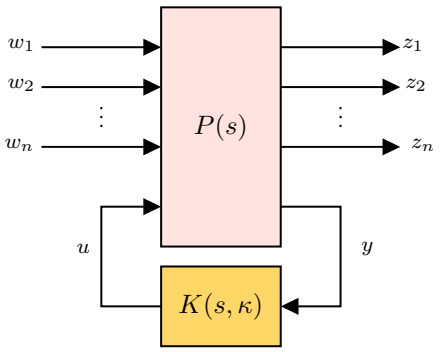

Figure 3.14: $\mathcal{H}_\infty$-synthesis against multiple requirements.

$$\begin{aligned} &\min_{\kappa} \max_{i} \left\| T_{w_i \to z_i}\left(K(s,\kappa)\right) \right\| \\ &\text{subject to } \max_{j} \left\| T_{w_j \to z_j}\left(K(s,\kappa)\right) \right\| \leq 1. \end{aligned} \tag{3.64}$$

where κ is the vector of tunable parameters and $\|\cdot\|$ denotes either the $\mathcal{H}_\infty$ or the $\mathcal{H}_2$ norm. This introduces greater flexibility to the formal robust synthesis of structured controllers, helping bridge the gap between formal robust control and control engineering practice.

3.5. Nonlinear Dynamic Inversion

3.5.1. Introduction

Nonlinear Dynamic Inversion (NDI), also known as Feedback linearisation among the Control community, was developed as an alternative to the divide-and-conquer paradigm, which results in self-scheduled controllers Enns et al., 1994. The general idea of NDI control law is to transform selected input-output channels into a chain of integrators by using state feedback information (hence the designation Feed-

back linearisation) (Khalil, 2002). This inner dynamic inversion loop linearises the nonlinear dynamics, transforming the nature of the system from nonlinear to linear. As a result, an outer virtual control loop may employ linear control design methods to control this linearised system, resulting in a self-scheduled controller. This approach results in a transparent and modular framework for flight control design, enabling the decoupling of airframe-dependent and flying-qualities-dependent components as demonstrated in Figure 3.15.

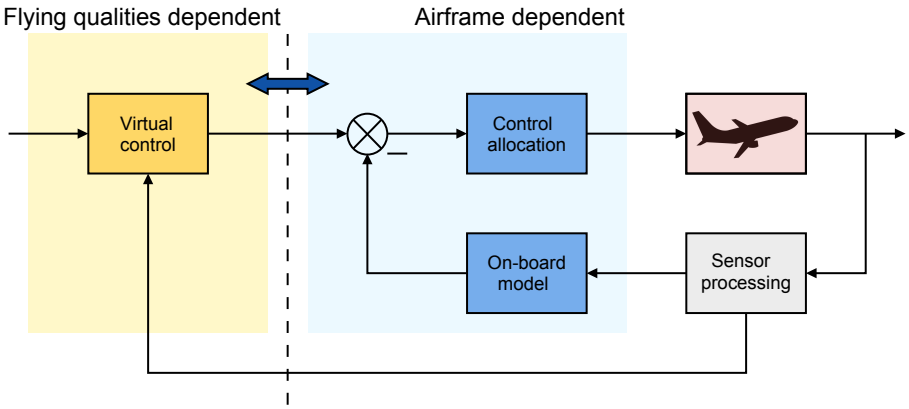

Figure 3.15: Basic NDI control structure, evidencing the modular nature of the control architecture. Figure recreated from Pollack, 2024.

3.5.2. Model-Based NDI

Classical NDI, which will be referred to as Model-based (MB) NDI, relies on On-Board (OBM) models of the plant dynamics to transform selected input-output channels into a chain of integrators of relative degree ρ according to the principles of feedback linearization.

Consider a multi-input, multi-output, input-affine, nonlinear system Σ with n states and m inputs described by:

$$\Sigma : \begin{cases} \dot{x} = f(x) + G(x)u \\ y = h(x) \end{cases} \tag{3.65}$$

where $x \in \mathbb{R}^n$ represents the state vector, $u \in \mathbb{R}^m$ the input vector, $y \in \mathbb{R}^m$ the observation vector and f, G, and h are smooth mappings. Writing the system relative degree as $\rho = \begin{bmatrix} \rho_1, \dots, \rho_m \end{bmatrix}^T$, the output dynamics can be described as (Wang et al., 2019):

$$\begin{aligned} y^{(\rho)} &= \begin{bmatrix} \mathcal{L}_f^{\rho_1} h_1(x) \\ \vdots \\ \mathcal{L}_f^{\rho_m} h_m(x) \end{bmatrix} + \begin{bmatrix} \mathcal{L}_{g_1}\mathcal{L}_f^{\rho_1 - 1} h_1(x) & \dots & \mathcal{L}_{g_m}\mathcal{L}_f^{\rho_1 - 1} h_1(x) \\ \vdots & \ddots & \vdots \\ \mathcal{L}_{g_1}\mathcal{L}_f^{\rho_m - 1} h_m(x) & \dots & \mathcal{L}_{g_m}\mathcal{L}_f^{\rho_m - 1} h_m(x) \end{bmatrix} u \\ &\triangleq \alpha(x) + \mathcal{B}(x)u \end{aligned} \tag{3.66}$$

where $\mathcal{L}_f^k h_i$ and $\mathcal{L}_{g_i}\mathcal{L}_f^k h_i$ represent repeated Lie derivatives of the function h_i along the vector fields f and g_i, with g_i being a column vector of the matrix G (Khalil, 2002).

The first-order Lie derivative is defined as (Khalil, 2002):

$$\mathcal{L}_f h(x) = \nabla h(x)\, f(x) = \sum_{i=1}^{n} \frac{\partial h(x)}{\partial x_i} f_i(x) \tag{3.67}$$

where ∇ is the gradient operator and $\nabla h(x)$ is defined as:

$$\nabla h(x) = \frac{\partial h(x)}{\partial x} = \begin{bmatrix} \frac{\partial h(x)}{\partial x_1} & \frac{\partial h(x)}{\partial x_2} & \cdots & \frac{\partial h(x)}{\partial x_n} \end{bmatrix} \tag{3.68}$$

Repeated (nested) use of the operation is possible:

$$\mathcal{L}_f^k h(x) = \mathcal{L}_f \left[\mathcal{L}_f^{k-1} h(x)\right] = \nabla \left[\mathcal{L}_f^{k-1} h(x)\right] f(x) = \frac{\partial}{\partial x} \left[\mathcal{L}_f^{k-1} h(x)\right] f(x) \tag{3.69}$$

which completes the definition of the Lie derivative terms in Equation 3.66.

For model-based NDI, Equation 3.66 can be used directly to construct a control law that linearises the input-output dynamics to a set of $\sum_{i=1}^{m} \rho_i$ integrators. Assuming that the control effectiveness matrix $\mathcal{B}(x)$ is invertible, the following control law is obtained:

$$u = \hat{\mathcal{B}}^{-1}(x)[v - \hat{\alpha}(x)] \tag{3.70}$$

where $\hat{\alpha}(x)$ and $\hat{\mathcal{B}}(x)$ represent on-board model estimates of $\alpha(x)$ and $\mathcal{B}(x)$, respectively, and $v \in \mathbb{R}^m$ is the pseudo-control vector generated by an auxiliary control law that is designed to meet the control objectives.

Considering a system with relative degree $\rho = 1$, such that $y^{(\rho)} = \dot{y}$, the control scheme can be represented as done in Figure 3.16.

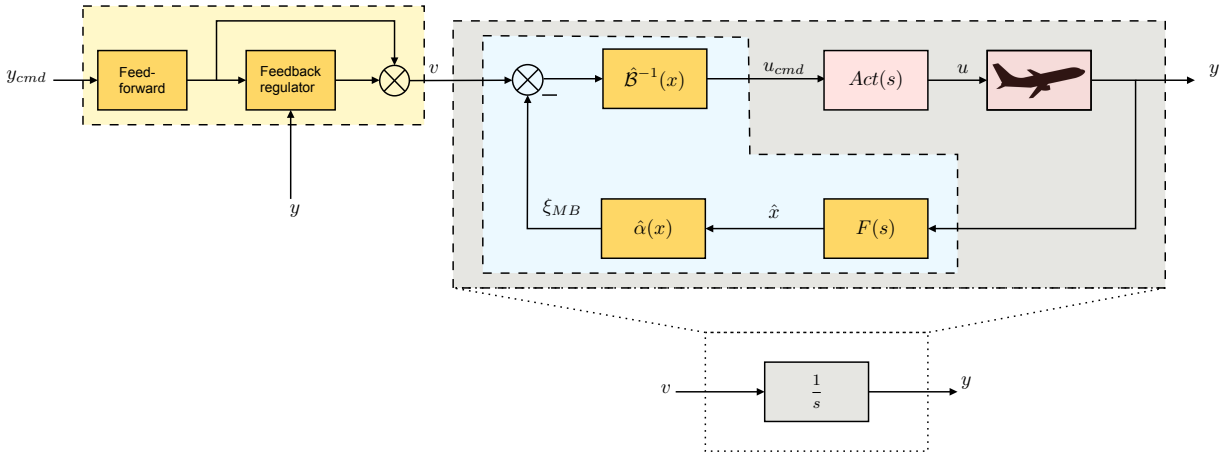

Figure 3.16: Model-based NDI implementation.

In the absence of On-Board model (OBM) uncertainty and disturbances, the closed-loop relationship between v and y is given by $\frac{1}{s}$. Dynamic inversion allows the design of the feedback regulator and the feedforward element using linear design methods to shape the desired dynamics. In principle, a single virtual control law can achieve the desired dynamics globally, i.e. over the full operating domain. Nevertheless, there are several aspects that jeopardize this principle (Pollack, 2024):

- The basic NDI framework does not provide a priori robustness guarantees. This does not mean that NDI-based control laws cannot be made robust; however, it does imply that robust design analysis and clearance must still be performed using additional insights based on, for example, local evaluations for individual flight conditions. Effectively, this results in shifting the burden of the control laws design from the synthesis to the analysis part. The lack of integration of a priori robustness in the synthesis can result in the need to redesign the control laws, thus making the overall design process longer.
- The presence of high-order dynamics associated with, e.g. finite bandwidth actuators, sensor processing, and filtering, and the limitation of control authority, leads to non-ideal inversion. Consequently, if the virtual control law is designed to shape the desired dynamics based on an incorrect closed-loop relationship between v and y, then the inversion error will propagate through the actual dynamic behaviour of the closed-loop system.
- The controlled variable (CV) that is inverted must be selected such that the zero dynamics are stable and behave satisfactorily. These zero dynamics refer to the part of the system that cannot be controlled through the selected CV (Khalil, 2002).
- The requirement for full-state feedback for the inversion often requires the use of observer-based solutions or other estimation techniques, as not all vehicle states may be available. A typical flight control example is air data variables such as angle-of-attack and angle-of-sideslip. The state estimator is represented in Figure 3.16 via the $F(s)$ block and its design may have a significant impact on the inversion path - factors not explicitly accounted for during the derivation of the inversion - and therefore have a significant impact on overall robustness and performance of the control law.

3.5.3. Sensor-Based INDI

Incremental Nonlinear Dynamic Inversion (INDI), referred to here as sensor-based (SB) INDI, seeks to address the robustness limitations of model-based NDI in an alternative way. Instead of relying on robust virtual control design, it circumvents the dependence on detailed and accurate airframe models by utilising direct sensor measurements of state derivatives instead. Consequently, control allocation is the only part of the control law that requires On-Board Model information, thus making it less expensive to implement than model-based NDI.

In order to obtain the incremental form (INDI), a common procedure is to make use of the Taylor expansion of the output dynamics around the system state at time $t - \Delta t$, where Δt represents the sampling time interval (Wang et al., 2019). Denoting this condition by the subscript 0 for ease of notation yields the following expression:

$$y^{(\rho)} = y_0^{(\rho)} + \left.\frac{\partial\big[\alpha(x) + \mathcal{B}(x)u\big]}{\partial x}\right|_0 \left(\underbrace{x - x_0}_{\Delta x}\right) + \mathcal{B}(x_0)\underbrace{(u - u_0)}_{\Delta u} + R_1 \tag{3.71}$$

where $\mathcal{R}_1$ represents the expansion remainder. Consequently, the time-scale separation assumption can be leveraged to design the incremental control input Δu, which assumes that all state-dependent and residual terms can be neglected (Grondman et al., 2018). This is often argued as justified in the case of high sampling rates and high-bandwidth actuators. The control law is completed by adding the control vector u_0 to the resulting incremental term:

$$u = u_0 + \hat{\mathcal{B}}^{-1}(x_0)\big[\nu - y_0^{(\rho)}\big]. \tag{3.72}$$

Thus, it is demonstrated that model information about $\alpha(x)$ is not required, as the control law is computed using sensor feedback from the previous control vector and the derivative of the control variable.

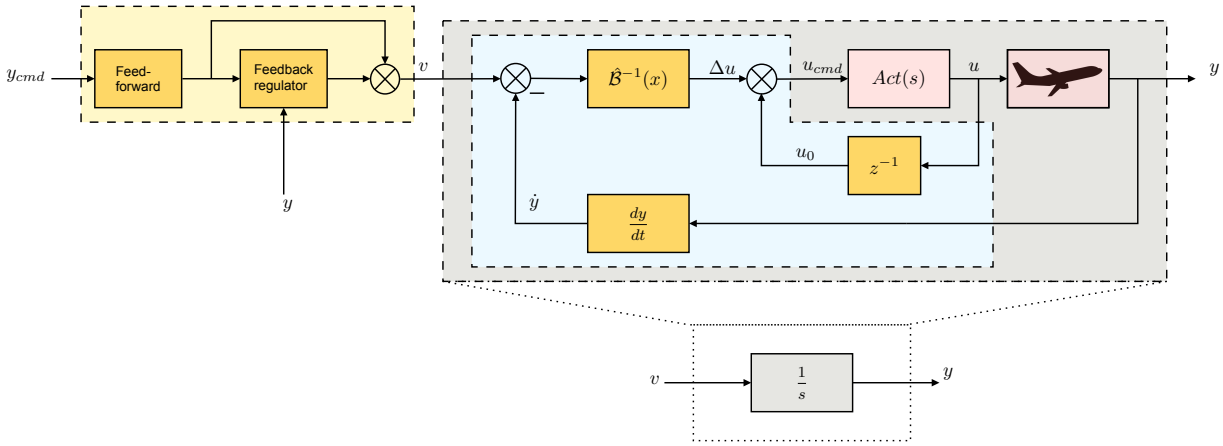

Figure 3.17: Sensor-based INDI implementation.

Defining $\xi^{MB} = \hat{\alpha}(x)$ and $\xi^{SB} = \dot{y} - \hat{\mathcal{B}}(x)u_0$ as the signals in the inversion path for model-based nonlinear dynamic inversion (MB) and sensor-based incremental nonlinear dynamic inversion (SB), it is demonstrated that in the absence of uncertainty and disturbances of the On-Board model (OBM) (Pollack, 2024):

$$\underbrace{\hat{\alpha}(x)}_{\xi^{MB}} = \underbrace{\dot{y} - \hat{\mathcal{B}}(x)u_0}_{\xi^{SB}} \tag{3.73}$$

However, the fact that the above equality only holds in the absence of uncertainty and disturbances is a direct indication that the different inversion schemes possess different robustness properties. Pollack and Van Kampen, 2023 highlights that these different robustness properties can clearly be understood by analysing the open-loop response of an MB and an SB dynamic inversion schematic with the same virtual control law controller. The authors demonstrate that SB inversion schemes result in an elevated open-loop response compared to MB ones, resulting in higher gain at low frequencies but also higher and often prohibitive crossover frequencies. The higher gain at low frequencies helps to understand the reason why Incremental Nonlinear Dynamic Inversion (INDI) is often claimed as more robust against aerodynamic model uncertainties. However, robustness at higher frequencies is penalised due to the insufficient roll-off. This helps to explain why unfiltered implementations of INDI control are often unsuccessful (Grondman et al., 2018). Practical implementations of SB INDI control require filtering via synchronised low-pass filters H_c as described in Figure 3.18 (Pollack, 2024):

$$\xi^{SB}_{filtered}(s) = H_c(s)[sy(s) - \hat{\mathcal{B}}(x)u(s)] \tag{3.74}$$

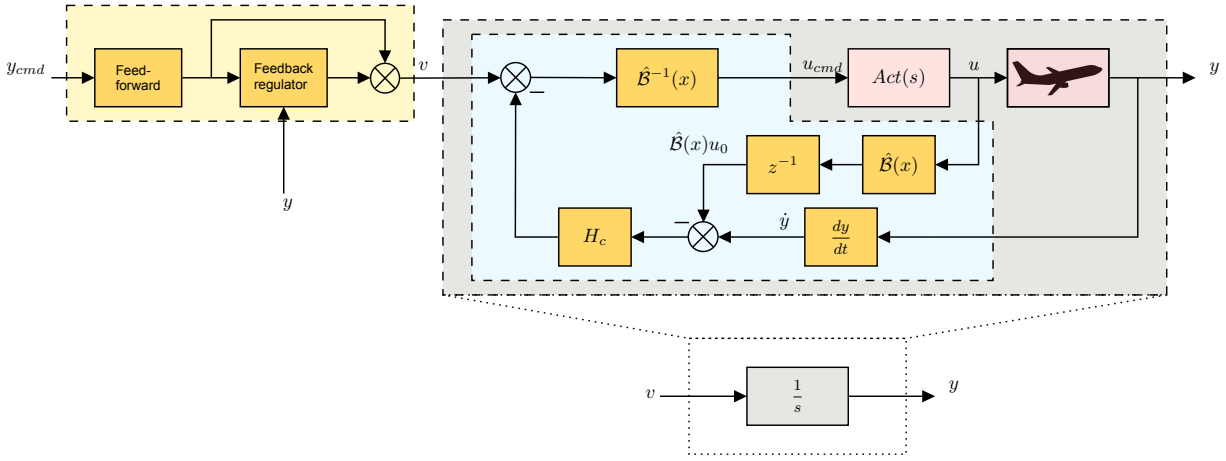

Figure 3.18: Sensor-based INDI implementation with synchronised low-pass filter H_c.

3.5.4. Hybrid-Based INDI

Pollack, 2024 suggests that reintroducing model information of the bare airframe dynamics in the form of a complementary augmentation element can lead to increased robustness of the control law. This concept was originally proposed in Jiali and Jihong, 2016 and formulated in Kim et al., 2021 and Kumtepe et al., 2022 as a Hybrid INDI approach. However, the formal definition of a scaled complementary filter (SCF) hybrid scheme, which will be used throughout the present study, originates from Pollack, 2024; Pollack et al., 2024 and can be defined as:

$$\begin{aligned}\xi^{HB}(s) &= (1 - K_c H_c(s))\xi^{MB}(s) + K_c H_c(s)\xi^{SB}(s) \\ &= \xi^{MB}(s) + K_c H_c(s)(\xi^{MB}(s) - \xi^{SB}(s)) \\ &= \xi^{MB}(s) + \underbrace{K_c H_c(s) e_\xi(s)}_{e^*_\xi}\end{aligned} \tag{3.75}$$

where K_c is a scaling gain $\in [0, I] \subset \mathcal{R}^{n_u \times n_u}$ and H_c is a filter element $\in \mathcal{R}H_\infty^{n_u \times n_u}$.

Given the relationships in Equation 3.75, Hybrid INDI can be interpreted as a standard model-based NDI with additional error compensation. Correspondingly, Hybrid INDI collapses to purely model-based NDI design when $K_c = 0$ and to a sensor-based INDI if $K_c = 1$ and the bandwidth of $H_c(s)$ is made sufficiently large since $\xi^{HB}(s) \to \xi^{SB}(s)$.

Figure 3.19 illustrates the Hybrid INDI control scheme. In INDI control schemes, feedback on the input signal $\mathcal{B}(x)u$ is necessary. There are two possible sources for input feedback: using feedback on the actuator position u_0 or internal Control Command (CC) feedback. The principle behind CC feedback relies on the approximation $CC \approx \hat{\mathcal{B}}(x)u_0$, which can be explained as follows:

$$u = G_{act}(s)\hat{\mathcal{B}}^{-1}(x)CC \Rightarrow CC \approx \hat{\mathcal{B}}(x)u \tag{3.76}$$

assuming ideal or close to ideal actuators such that $G_{act} \approx I$. Therefore, measuring the input u_0 and using input feedback $\hat{\mathcal{B}}(x)u_0$ for the control law is equal to simply using CC feedback. In reality, actuators have internal dynamics, and therefore, CC is merely an approximation of $\hat{\mathcal{B}}(x)u_0$. However, the fact that CC feedback is a location inside the control law itself makes it conceptually simpler to implement and does not require additional sensor measurements. Furthermore, the work of Pollack, 2024 suggests that little

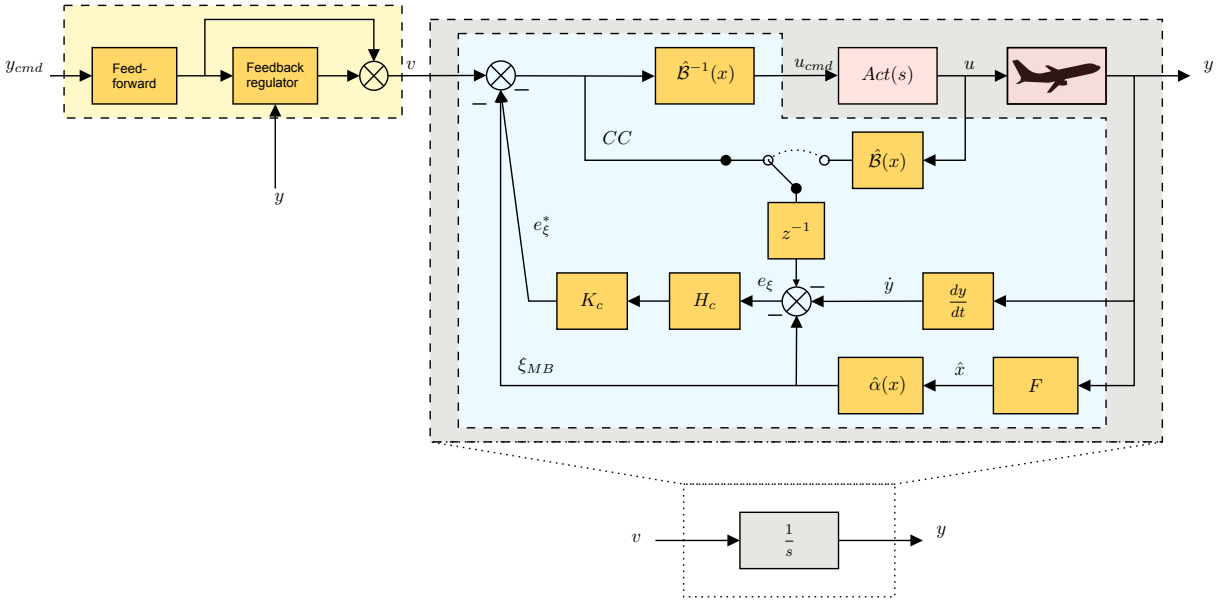

Figure 3.19: Hybrid-based INDI implementation.

limitations in terms of robustness of the control law result from this choice of input feedback signal. Thus, either choice for input feedback is valid.

Since the hybrid design allows navigating between model-based and sensor-based inversion schematics, it can be interpreted as an extra degree of freedom to modify the open-loop gain response and, thus, help address robustness. The main challenge with NDI-based control laws is how to tune the different elements in order to achieve a robust design given the inexistence of a priori robustness guarantees.

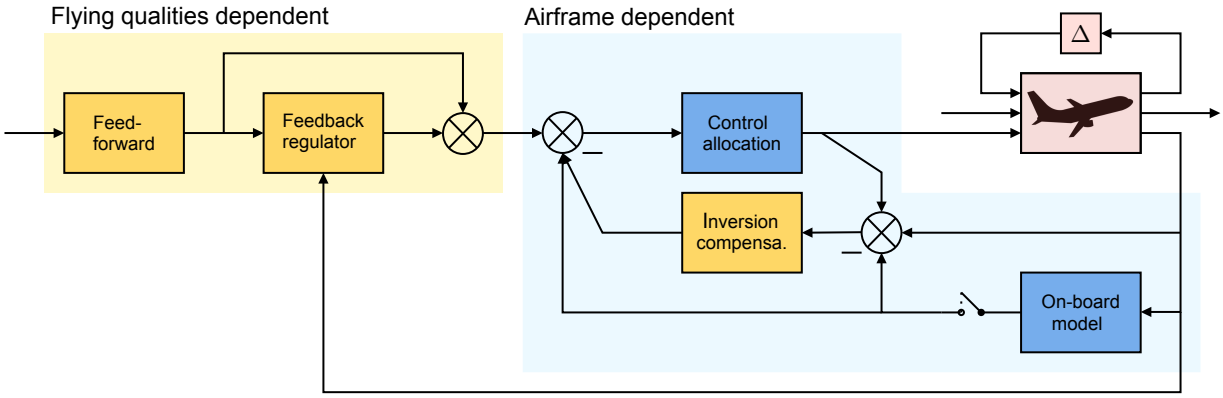

Figure 3.20: General description of a Hybrid INDI control schematic. The yellow blocks refer to the tunable elements to achieve a robust design, while the blue blocks refer to airframe-dependent blocks.

4

$\mathcal{H}_\infty$ Loop-Shaping for Dynamic Inversion Control: A Critical Perspective

This chapter is devoted to the analysis of the work of G. Papageorgiou and Polansky, 2009 since it proposes a methodology to combine Dynamic Inversion with $\mathcal{H}_\infty$ Loop-Shaping Control, thus making it a relevant and rare article. First of all, the proposed methodology will be presented, followed by an analysis of the multiple assumptions made. It will be shown that these assumptions lead to the proposed methodology not having the guarantees of the $\mathcal{H}_\infty$ LSDP. As a result, this chapter aims to demonstrate that combining Dynamic Inversion with $\mathcal{H}_\infty$ LSDP is still an open question.

4.1. Analysis of the Proposed Method

This section aims to explain the procedure proposed by G. Papageorgiou and Polansky, 2009 to combine Dynamic Inversion with $\mathcal{H}_\infty$ Loop-Shaping (or McFarlane-Glover Loop-Shaping). Some block diagrams are recreated from the ones presented in the paper.

4.1.1. DI Control

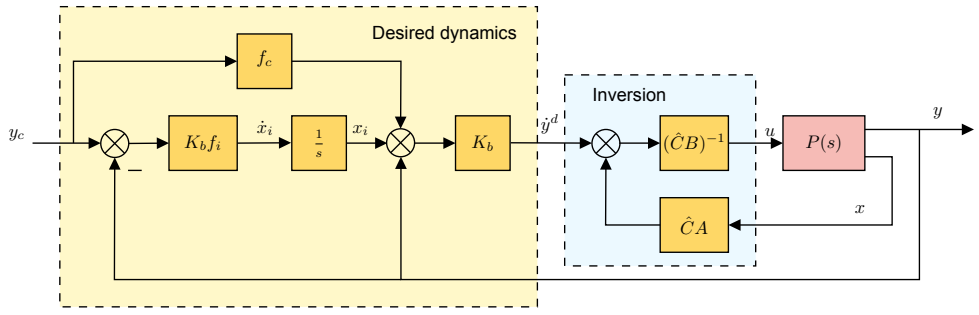

Figure 4.1: DI controller with a novel approach to the model used to derive the inversion. Block diagram recreated from G. Papageorgiou and Polansky, 2009.

As outlined in Chapter 3, the architecture of a Dynamic Inversion (DI) controller is structured with two loops: an inner inversion loop and an outer loop that imposes the desired dynamics. The inner loop attempts to invert the system dynamics so that the resulting closed-loop behaviour equates to an integrator (assuming a plant with relative degree one) or a chain of integrators (for a system with a relative degree greater than one). Meanwhile, the outer loop aims to define the desired closed-loop performance and

address robustness. For the remainder of the chapter, a plant model with a relative degree of one will be considered. In order to derive the control law for DI architectures, the first step is to define the controlled variables and differentiate them once. Consider, for instance, a SISO LTI plant model $P = [A, B, C, 0]$ of relative degree one. The initial step in deriving the inversion involves (G. Papageorgiou & Polansky, 2009):

$$\begin{aligned} &y = Cx \Rightarrow \\ &\dot{y} = C\dot{x} = CAx + CBu \end{aligned} \tag{4.1}$$

where $x \in \mathbb{R}^n$ represents the state of the system, $y \in \mathbb{R}$ corresponds to the output and controlled variable, and $u \in \mathbb{R}$ denotes the control input. Following this, the next step is to equate the derivative of the controlled variable $\dot{y}$ to the output of the desired dynamics $\dot{y}^d$ and subsequently solve for the control input u (G. Papageorgiou & Polansky, 2009):

$$\begin{aligned} &\dot{y} = \dot{y}^d \Rightarrow \\ &CAx + CBu = \dot{y}^d \Rightarrow \\ &u = (CB)^{-1}\left(\dot{y}^d - CAx\right) \end{aligned} \tag{4.2}$$

where $\dot{y}^d \in \mathbb{R}$ represents the output of the desired dynamics, typically referred to in other studies as the virtual control v. In the ideal scenario where no plant model uncertainty is present and assuming zero initial conditions, the inversion process effectively cancels out the system dynamics, ensuring the integrator-like response of the closed-loop transfer function from $\dot{y}^d$ to y. For SISO systems, CB is a non-zero scalar, but for multi-input multi-output (MIMO) systems, CB is not necessarily a square or invertible matrix. Addressing the inversion of CB falls under the domain of control allocation, which is out of the scope of the present thesis. Several options are available for selecting the desired dynamics; however, to maintain consistency with the work of G. Papageorgiou and Polansky, 2009, the one shown in Figure 4.1 is used:

$$\begin{aligned} \dot{x}_i &= K_b f_i \left(y^c - y\right) \\ \dot{y}^d &= K_b \left(x_i + f_c y^c - y\right) \end{aligned} \tag{4.3}$$

where x_i is the integrator state and y^c is the commanded output. Thus, assuming no plant model uncertainty and zero initial conditions (G. Papageorgiou & Polansky, 2009):

$$\begin{aligned} &\dot{y} = K_b \left(x_i + f_c y^c - y\right) \Rightarrow \\ &\ddot{y} + K_b \dot{y} = K_b \dot{x}_i + K_b f_c \dot{y}^c \Rightarrow \\ &\ddot{y} + K_b \dot{y} + K_b^2 f_i y = K_b f_c \dot{y}^c + K_b^2 f_i y^c \Rightarrow \\ &\left(\text{ choose } f_c + \frac{f_i}{f_c} = 1\right) \\ &\left(s + K_b f_c\right)\left(s + K_b \frac{f_i}{f_c}\right) y = K_b f_c \left(s + K_b \frac{f_i}{f_c}\right) y^c \Rightarrow \\ &\frac{y}{y^c} = \frac{K_b f_c}{s + K_b f_c} \end{aligned} \tag{4.4}$$

where s is the Laplace operator.

Conventionally, with DI controllers, the designer enjoys some degrees of freedom to address robustness and meet closed-loop requirements, such as the virtual control law and the command path, which together

constitute the desired dynamics, and the choice of the controlled variable y. In the work of G. Papageorgiou and Polansky, 2009, another degree of freedom is introduced: the inverted model. In aerospace contexts, this is referred to as the On-Board AirCraft (OBAC) model. While it shares the same state-space dynamics as the nominal plant model P (i.e. identical A and B matrices), its output matrix C might differ from the plant one. This output matrix is introduced in G. Papageorgiou and Polansky, 2009 as a non-standard tuning parameter which can be used to improve robustness and is depicted as $\hat{C}$. Naturally, it is often sufficient to choose $\hat{C} = C$, but in case a different $\hat{C}$ is selected, the nominal transfer function from $\dot{y}^d$ to y will not be equal to an integrator, as, effectively, it results in a different controlled variable (CV).

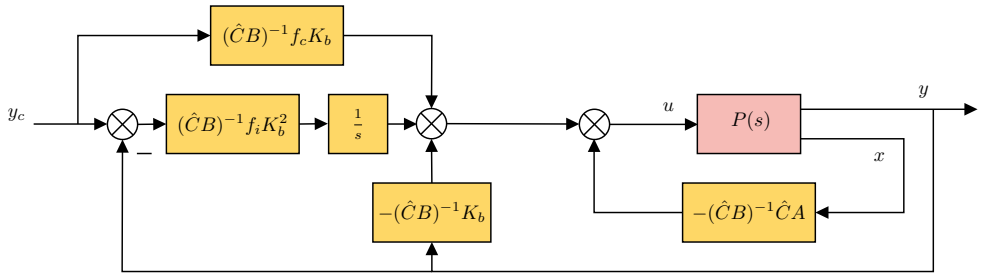

Figure 4.2: The four 'gains' of a typical DI controller. Block diagram recreated from G. Papageorgiou and Polansky, 2009.

4.1.2. $\mathcal{H}_\infty$ Loop-Shaping

In this section, a brief review of $\mathcal{H}_\infty$ Loop-Shaping will be given in order to make the following discussion clearer. The procedure boils down to two different steps: first the plant P is augmented with frequency-dependent weights W_1 and W_2, such that the shaped plant $P_s = W_2 P W_1$ has a desired open-loop response; in the second step, an $\mathcal{H}_\infty$ controller is synthesized, which minimizes the following $\mathcal{H}_\infty$ norm:

$$\min_{\text{stab } K_\infty} \left\| \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} \rightarrow \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} \right\|_\infty = \min_{\text{stab } K_\infty} \left\| \begin{bmatrix} I_p \\ -K_\infty \end{bmatrix} (I_p + P_s K_\infty)^{-1} \begin{bmatrix} I_p & P_s \end{bmatrix} \right\|_\infty \leq \gamma \triangleq \epsilon^{-1} \tag{4.5}$$

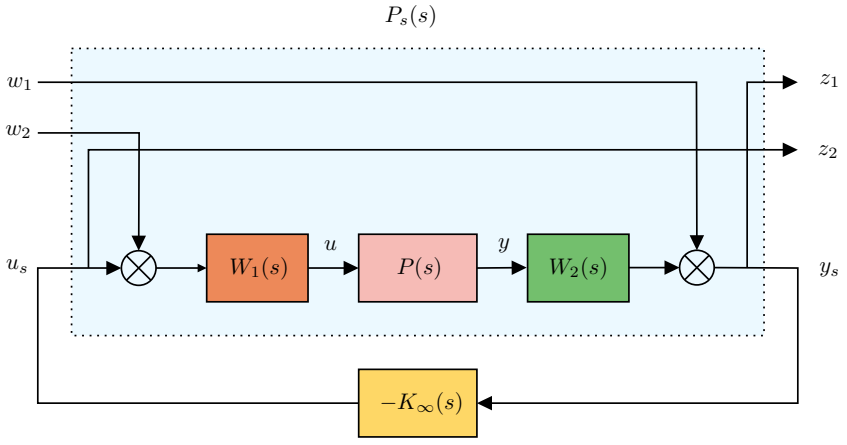

Figure 4.3: NCF robust stabilization problem written as mixed sensitivity 4-block problem.

According to G. Papageorgiou and Polansky, 2009, while standard $\mathcal{H}_\infty$ controller synthesis typically involves solving a pair of coupled Riccati equations, the $\mathcal{H}_\infty$ Loop-Shaping can be divided into three stages.

The first stage involves solving a normalised LQ problem, which takes the shaped plant P_s as input and yields the state feedback matrix F solution, as shown in Figure 4.4. The LQ problem is formulated as:

$$\min_{z_2 \in \mathscr{L}_2[0,\infty)} \int_0^\infty \left(z_1^2 + z_2^2\right) \mathrm{d}t \tag{4.6}$$

where $z_1 = P_s z_2$. In the second stage, the ε associated with the shaped plant P_s is computed. In the final stage, the values of P_s, F, and ε are used to compute the observer gain matrix L depicted in Figure 4.4.

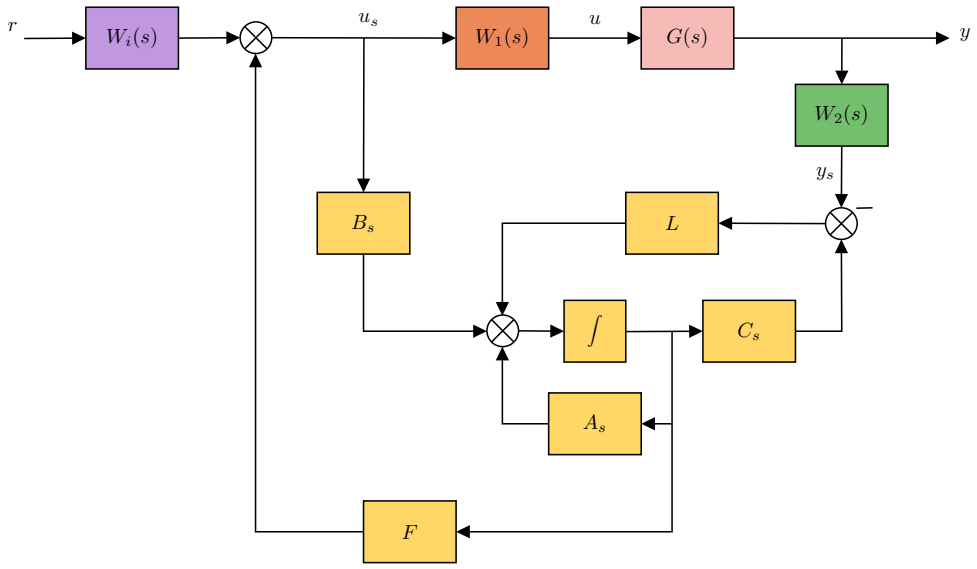

Figure 4.4: Observer-based structure for the $\mathcal{H}_\infty$ LSDP controller.

This odd description of the $\mathcal{H}_\infty$ Loop-Shaping solution deserves special attention since it is the genesis of the limitations of the proposed procedure. If the integral in Equation 4.6 is expanded as done in Equation 4.7:

$$\begin{aligned}
\min_{z_2 \in \mathscr{L}_2[0,\infty)} \int_0^\infty \left(z_1^2 + z_2^2\right) dt &= \min_{u_s \in \mathscr{L}_2[0,\infty)} \int_0^\infty \left(y_s^2 + u_s^2\right) dt \\
&= \min_{u_s \in \mathscr{L}_2[0,\infty)} \int_0^\infty \left((C_s x_s)^2 + u_s^2\right) dt \\
&= \min_{u_s \in \mathscr{L}_2[0,\infty)} \int_0^\infty \left(x_s^T C_s^T C_s x_s + u_s^T u_s\right) dt \\
&= \min_{u_s \in \mathscr{L}_2[0,\infty)} \int_0^\infty \left(x_s^T Q x_s + u_s^T R u_s\right) dt \quad (\text{with } Q = C_s^T C_s \text{ and } R = I),
\end{aligned} \tag{4.7}$$

then the final form can be recognised as a standard LQR problem with an optimal controller F given by the solution to the following Algebraic Riccati Equation (ARE):

$$\begin{aligned}
&A_s^T X + X A_s - (X B_s) R^{-1} (B_s^T X) + Q = 0 \\
&F = R^{-1} (B_s^T X) = B_s^T X
\end{aligned} \tag{4.8}$$

While it is true that parallelisms can be established between LQ control and $\mathcal{H}_\infty$ Loop-Shaping control, they possess obvious differences. This motivates a thorough analysis of the equivalences between $\mathcal{H}_\infty$

Loop-Shaping and a special case $\mathcal{H}_2$ problem in Section 4.1.3 to better understand the described $\mathcal{H}_\infty$ solution.

4.1.3. Equivalence between $\mathcal{H}_\infty$ LSDP and an equivalent Loop-Shaping $\mathcal{H}_2$ control

The analysis of the work of G. Papageorgiou and Polansky, 2009 raised the discussion about a possible equivalence between $\mathcal{H}_\infty$ LSDP and an equivalent Loop-Shaping $\mathcal{H}_2$ control, or, in other words, a special case LQG.

For a shaped plant P_s given by:

$$P_s = W_2 G W_1 = \left[\begin{array}{c|c} A_s & B_s \\ \hline C_s & D_s = 0 \end{array}\right] \tag{4.9}$$

the Algebraic Riccati equation (ARE's) of $\mathcal{H}_\infty$ Loop-Shaping control are given by (McFarlane & Glover, 1990; Sefton & Glover, 1990):

$$\begin{cases} A_s^T X + X A_s - (X B_s)(X B_s)^T + C_s^T C_s = 0 \\ A_s Z + Z A_s^T - (Z C_s^T)(Z C_s^T)^T + B_s B_s^T = 0 \end{cases} \tag{4.10}$$

For the LQG case, the Riccati equations are given by (Zhou & Doyle, 1998):

$$\begin{cases} A^T X + (A^T X)^T - (XB) R^{-1} (XB)^T + Q = 0 \\ AZ + (A^T Z)^T - (ZC^T) V^{-1} (ZC^T)^T + W = 0 \end{cases} \tag{4.11}$$

In the case where LQG is applied to a shaped plant P_s, and setting $R = V = I$, $Q = C_s^T C_s$ and $W = B_s B_s^T$, the LGQ Riccati equations coincide with the $\mathcal{H}_\infty$ Loop-Shaping ones in 4.10. This specific choice of the R, V, Q and W matrices and considering the plant $P(s)$ as being the same as the shaped plant $P_s(s) = W_2(s)G(s)W_1(s)$, boils down to an equivalent representation of the closed-loop system $T_{wz}(s)$ of the $\mathcal{H}_\infty$ Loop-Shaping problem and of the special case $\mathcal{H}_2$ problem (see Figures 4.5, 4.6 and 4.7).

However, since one is the solution of an $\mathcal{H}_2$ problem (minimization of the $\|T_{wz}(s)\|_2$ norm) and the other the solution of an $\mathcal{H}_\infty$ problem (minimization of the $\|T_{wz}(s)\|_\infty$ norm), control differences appear in the way the controller is constructed.

By its nature, the LQG controller has an observer-based structure, i.e. the internal states of the controller are estimates of the plant states. Assuming that the controller is proper and has the following state-space structure:

$$K = \left[\begin{array}{c|c} A_k & B_k \\ \hline C_k & D_k = 0 \end{array}\right] \tag{4.12}$$

then the LQG controller can be constructed as follows:

$$K_{LQG} = \begin{bmatrix} A - B_s B_s^* X - Z C_s^* C_s & Z C_s^* \\ -B_s^* X & 0 \end{bmatrix} \tag{4.13}$$

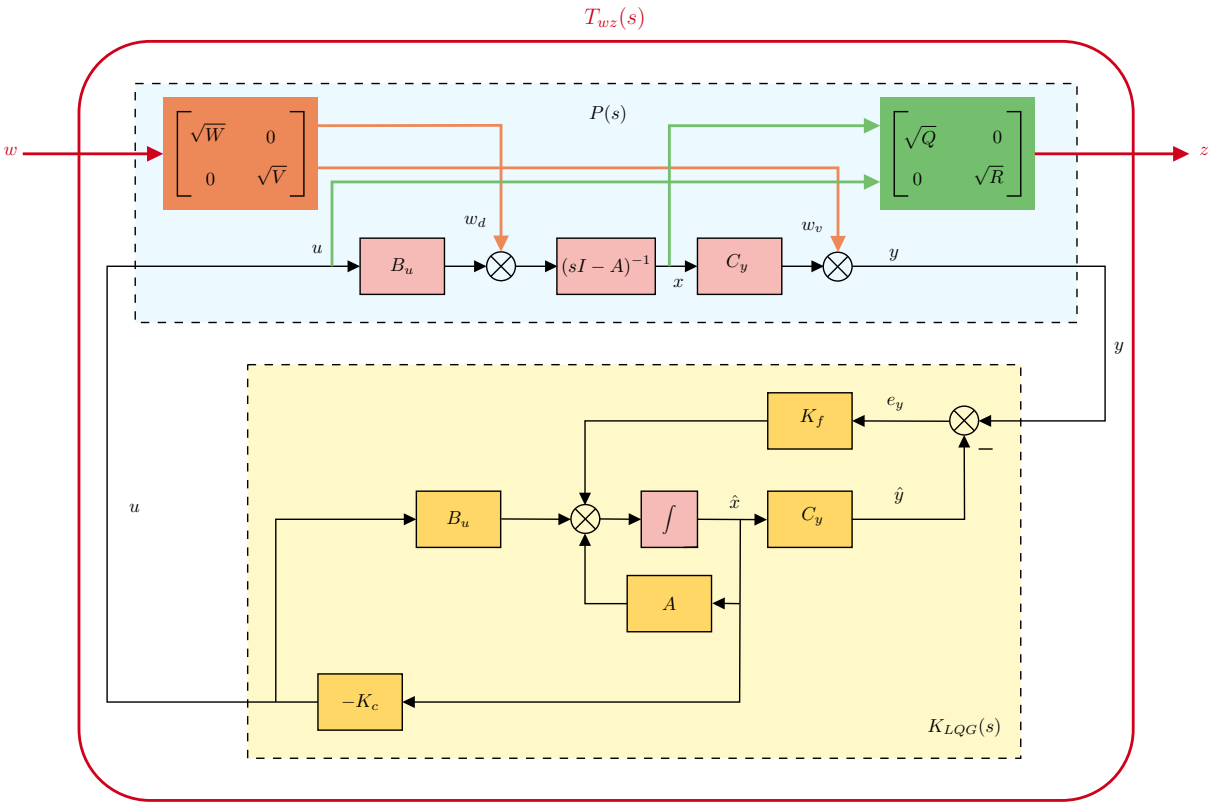

Figure 4.5: Generic structure for the $\mathcal{H}_2$ problem associated with the LQG controller.

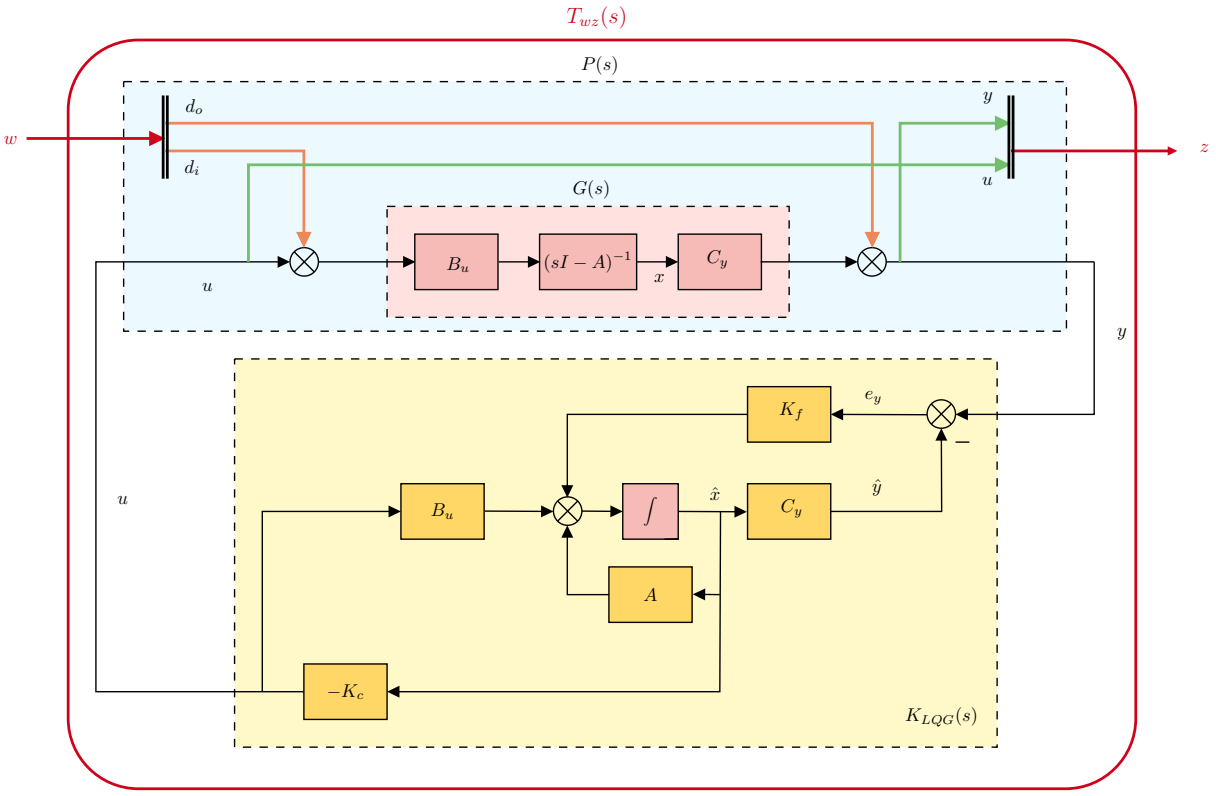

Figure 4.6: Special case LQG which results in the same T_{wz} as the one in $\mathcal{H}_\infty$ Loop-Shaping.

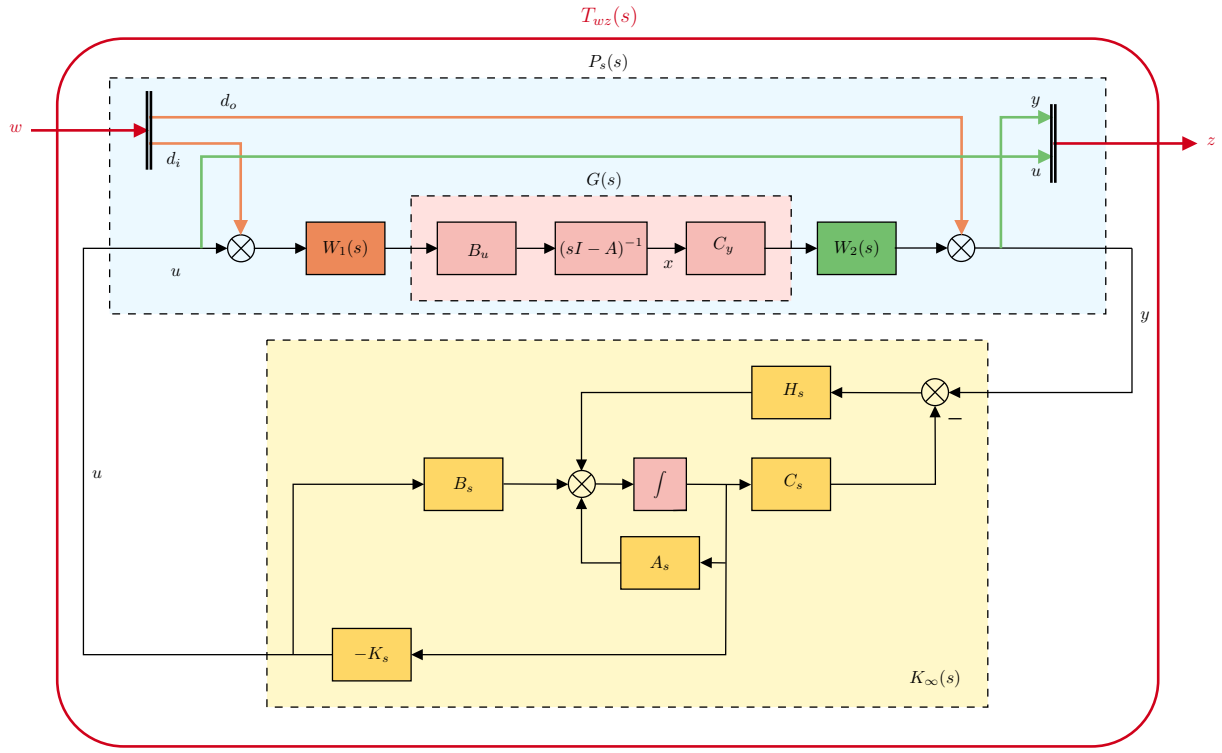

Figure 4.7: $\mathcal{H}_\infty$ Loop-Shaping with controller with observer-based structure.

The $\mathcal{H}_\infty$ controller, on the other hand, can be parameterised in infinitely different ways, according to the choice of the Youla Parameterisation (Zhou et al., 1996). If the Youla Parameterisation Q is set equal to 0, the solution to the central controller is obtained:

$$K_{H_\infty central} = \begin{bmatrix} A - B_s B_s^* X + (I - \gamma^{-2}(I + XZ))^{-1} Z C_s^* C_s & \gamma^2(((1-\gamma^2)I + XZ)^T)^{-1} Z C_s^* \\ -B_s^* X & 0 \end{bmatrix} \tag{4.14}$$

This controller is referred to as the central one. This controller is unstructured, which means that the internal states of the controller are not estimates of the states of the plant, as in a controller with an observer-based structure.

However, a particular advantage of $\mathcal{H}_\infty$ Loop-Shaping control is that it can be written in an observer-based form, similar to the LQG one:

$$K_{H_\infty observer} = \begin{bmatrix} A_s - B_s B_s^* (I - \gamma^{-2}(I + XZ))^{-1} X - Z C_s^* C_s & Z C_s^* \\ -B_s^* (I - \gamma^{-2}(I + XZ))^{-1} X & 0 \end{bmatrix} \tag{4.15}$$

The terms highlighted in blue evidence the commonalities between the LQG control solution and the $\mathcal{H}_\infty$ Loop-Shaping one. On one hand, the controller state feedback matrix C_k of the LQG solution is equal to the one from $K_{H_\infty central}$. However, in the LQG case, the internal states of the controller x_k would correspond to the estimates of the states x_s of the shaped plant P_s. For the unstructured $\mathcal{H}_\infty$ case, however, it would correspond to the internal states of the controller, where $x_k \neq x_s$. The differences between the two controllers are clear.

On the other hand, when the $\mathcal{H}_\infty$ Loop-Shaping controller is written in its observer-form, it can be clearly seen that "the controller is a normalized LQG observer with optimal state-feedback $F_{NFI} = -B_s^*(I - \gamma^{-2}(I + XZ))^{-1}X$" (Sefton & Glover, 1990), where X and Z are the solutions to the Riccati equations in 4.10 and F_{NFI} corresponds to C_k in 4.12.

As a result, when Papageorgiou and Polansky establish the parallelism between the LQ control (in 4.6) and $\mathcal{H}_\infty$ Loop-Shaping, in terms of OF control, the equivalence is present in the way the C_k matrix of the controller K is constructed for the LQG and the $\mathcal{H}_\infty$ LS central controller (which is unstructured). However, the procedure requires an observer-based structure, which the $\mathcal{H}_\infty$ Loop-Shaping central controller does not exhibit. On the other hand, if the $\mathcal{H}_\infty$ Loop-Shaping is constructed in its observer-based form, then the C_k matrix is computed differently than in the LQG case. Therefore, it is difficult to explain the characterisation of the $\mathcal{H}_\infty$ controller synthesis provided by G. Papageorgiou and Polansky, 2009, particularly the first step in Equation 4.6.

The previous discussion on the output feedback solution of $\mathcal{H}_2$ and $\mathcal{H}_\infty$ LS control is relevant since many flight control applications are output feedback problems. For instance, even the solution proposed in G. Papageorgiou and Polansky, 2009 introduces an AoA estimator in the pitch-rate DI controller, shifting the nature of the problem from state feedback to output feedback. Nevertheless, conceptually, DI control is, in its nature, a state feedback approach. According to Khargonekar and Shim, 1994, the state feedback controller which maximises the stability of a plant subject to uncertainty in its normalised right coprime factors is a LQR gain. It is of interest to explore the ramifications of this state-feedback solution in future studies.

This discussion highlights the differences between the $\mathcal{H}_2$ and $\mathcal{H}_\infty$ OF control solutions and provides the necessary context to understand the limitations of the approach proposed in G. Papageorgiou and Polansky, 2009 and described below.

4.1.4. Tuning a DI Controller using McFarlane-Glover Loop-Shaping

This section is devoted to explaining the proposed method from G. Papageorgiou and Polansky, 2009 to tune a DI controller using $\mathcal{H}_\infty$ LS. The sequential steps are as follows:

Procedure by G. Papageorgiou and Polansky, 2009 to tune DI controllers using $\mathcal{H}_\infty$ LS. The enumerated steps are quoted verbatim from G. Papageorgiou and Polansky, 2009:

1. With respect to Figure 4.3, set W_1 = 1 and shape the plant model P with W_2 = K_p + $\frac{K_i}{s}$. The proportional gain is used to specify the desired rise time, and the ratio of the two gains is used to specify the desired settling time.
2. Solve the normalized LQ problem 4.6. Denote the resulting state feedback matrix by $F = [F_i \; F_x]$, where F_i multiplies the integrator state of W_2 and F_x multiplies the state or state estimate of P.
3. Choose K_b, and compute $\hat{C}$, f_i, and f_c (the remaining parameters in Figure 4.2) from F_x, K_i, and K_p. Since K_b does not affect the achieved feedback properties, it is chosen to normalise one of the elements of $\hat{C}$.
4. If a state estimator is required, i.e. a complete state measurement is not available, then compute ε. If $\varepsilon < 0.4$, then go back to Step 1 and modify the integral and/or proportional gains. First, try reducing the ratio of the two gains to reduce the phase lag around cross-over. If this is not sufficient, then reduce the proportional gain. If $\varepsilon \geq 0.4$, then compute the

McFarlane–Glover Loop-Shaping observer gain matrix or design a state estimator using some other method.

The described steps, together with block diagram manipulations, allow mapping F_x and the gains of W_2 to the four gains of the DI controller presented in Figure 4.2. From Figure 4.4, by defining W_1 = 1, W_2 equal to a PI and assuming full state feedback (removing the state estimator), it is possible to rewrite the block diagram as the one in Figure 4.8. By noting that the proportional path of the W_2 weight is not part of the feedback signal to the controller and expanding the command pre-filter, we arrive at Figure 4.9. To complete the block manipulation, the following lemma is required:

Lemma stated without proof in G. Papageorgiou and Polansky, 2009:

Lemma 1 *Given a SISO plant model P of relative degree one and with $P(j0) \neq 0$. Solve the normalised LQ problem 4.6 with P_s = $(K_p + K_i/s)\cdot P$. If a solution exists, then*

$$F_i^2 = 1, \qquad T_{y^c \to y}(j0) F_i = -1, \qquad det(A + BF_x) \neq 0$$

Under the assumption that $F_i = -1$, then Figure 4.9 can be reworked to Figure 4.10. Finally, it is possible to compute the desired DI parameters by equating the gains of Figures 4.10 and 4.2 as follows:

$$\begin{aligned} F_x =& (\hat{C}B)^{-1}\left(-\hat{C}A - K_b C\right) \Rightarrow \hat{C}BF_x = -\hat{C}A - K_b C \Rightarrow \\ \hat{C} =& -K_b C\left(A + BF_x\right)^{-1}, \quad [\det\left(A + BF_x\right) \neq 0 \text{ from Lemma 1}] \end{aligned} \tag{4.16}$$

$$f_c = \frac{\hat{C}BK_p}{K_b} \tag{4.17}$$

$$f_i = \frac{\hat{C}BK_i}{K_b^2} \tag{4.18}$$

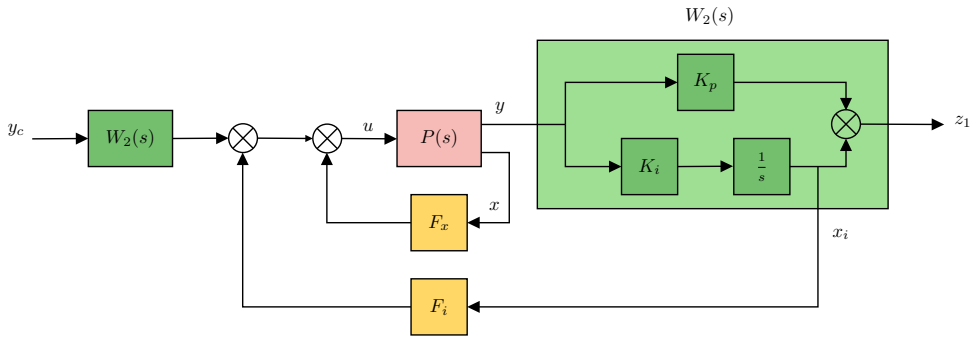

Figure 4.8: Step 1 of the block diagram manipulations. Block representation from G. Papageorgiou and Polansky, 2009.

In the original work from G. Papageorgiou and Polansky, 2009, Lemma 1 is stated without proof, but given its importance to the derivations, Section 4.3 provides further insights into its implication in the procedure. In the steps of the procedure, Papageorgiou and Polansky mention that, according to Step 3,

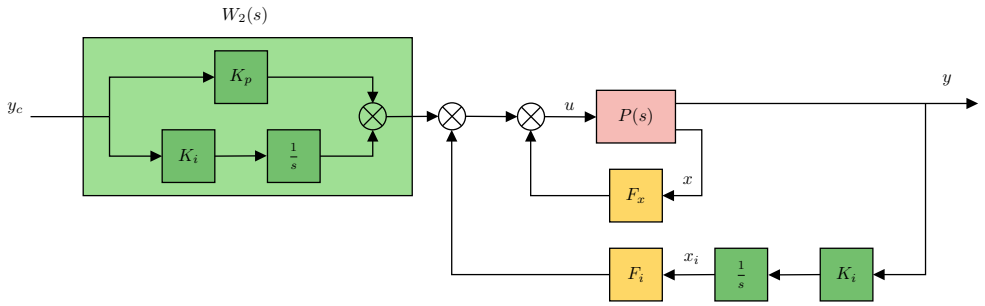

Figure 4.9: Step 2 of the block diagram manipulations. Block diagram recreated from G. Papageorgiou and Polansky, 2009.

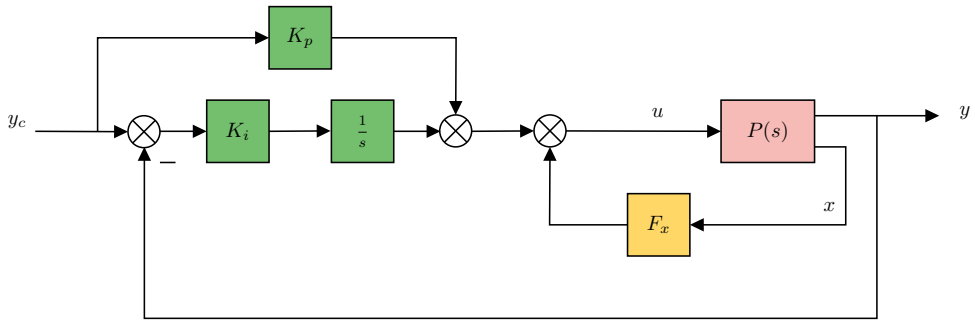

Figure 4.10: Final step of the block diagram manipulations. Block diagram recreated from G. Papageorgiou and Polansky, 2009.

LQ theory provides the stability margins at the plant input of LQR control. Furthermore, they mention that in the case a state estimator is needed, the degradation of the stability margins can be arbitrarily bad, as demonstrated in the landmark study about the absence of robustness guarantees with LQG regulators in Doyle, 1978. In that case, they recommend the use of the McFarlane–Glover ($\mathcal{H}_\infty$) Loop-Shaping state estimator or an estimator 'close' to the McFarlane–Glover Loop-Shaping estimator so that robustness margins 'close' to the $\mathcal{H}_\infty$ Loop-shaping control are expected.

First of all, it is not detailed what is meant by an estimator 'close' to the McFarlane–Glover Loop-Shaping estimator. Moreover, this description of the controller obtained in Step 3 reveals that the state feedback controller is computed based on LQ theory, i.e. $\mathcal{H}_2$ theory. In light of the discussion conducted in the previous section, the use of the state feedback controller from LQ theory with an $\mathcal{H}_\infty$ estimator computed posteriorly based on ε leads to the $\mathcal{H}_\infty$ LS central controller solution. However, as this is an unstructured controller which does not exhibit an observer-based structure, it does not make sense to use it in the context of DI control. Curiously, in the article, Papageorgiou and Polansky use the state feedback matrix from LQ optimisation (assuming full state feedback) but introduce an arbitrary AoA estimator, which results in the absence of any a priori robustness guarantees, as commonly found with output-feedback $\mathcal{H}_2$ control. Effectively, the only commonality with $\mathcal{H}_\infty$ LS control appears to be the first step, where the plant is shaped with weighting filters.

4.2. Verification of the Assumptions Made

In this section, the discussion provided above will be verified by recreating some of the article's results. From the 12 LTI short-period models of a Boeing 747 listed in the original source, the plant which yielded the best-performing DI controller was model number 8, and consequently, that was the model chosen for a

more thorough analysis in the article. Therefore, the same plant model is chosen in this work, which has the following dynamics:

$$\begin{bmatrix} \dot{\alpha} \\ \dot{q} \end{bmatrix} = \begin{bmatrix} Z_\alpha & 1 \\ M_\alpha & M_q \end{bmatrix} \begin{bmatrix} \alpha \\ q \end{bmatrix} + \begin{bmatrix} 0 \\ M_\delta \end{bmatrix} \delta = \begin{bmatrix} -0.666 & 1 \\ -0.947 & -0.653 \end{bmatrix} \begin{bmatrix} \alpha \\ q \end{bmatrix} + \begin{bmatrix} 0 \\ 1.349 \end{bmatrix} \delta \tag{11}$$

The W_2 filter, a PI controller, was selected following the guidelines in the paper, such that the zero of W_2 is set to 2 rad/s and $|W_2P|$ = 1 at 3 rad/s. The obtained gains for the W_2 filter and the gain of the shaped plant (P_s = W_2P) are presented below.

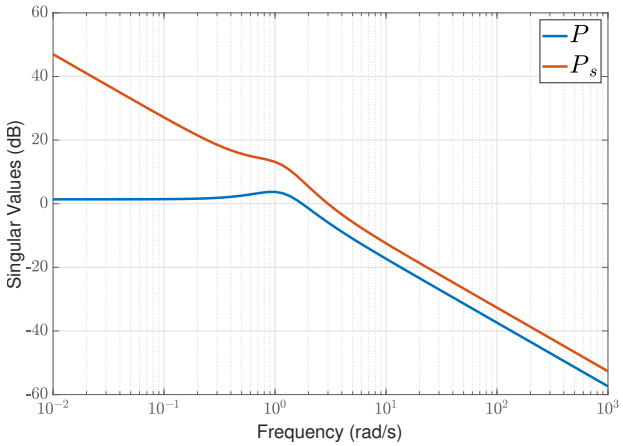

Figure 4.11: Open-Loop gain of G_s evidencing the 0 dB gain at 3 rad/s.

Table 4.1: Values of the PI W_2 filter

Parameter	K_p	K_i
Value	1.723	3.446

Furthermore, the special case LQG controller, the unstructured and the observer-based $\mathcal{H}_\infty$ controller are presented below:

$$K_{LQG} = \left[\begin{array}{ccc|c} -1.255 & 0 & 1.284 & 1.255 \\ -0.3074 & -0.666 & 0.4703 & 0.3074 \\ -2.38 & -0.1678 & -5.459 & 1.031 \\ \hline -1 & 0.5776 & -2.246 & 0 \end{array}\right] \tag{4.19}$$

$$K_{H_\infty unstrct} = \left[\begin{array}{ccc|c} -568.7 & 0 & -976.4 & 568.7 \\ -143.1 & -0.666 & -245.5 & 143.1 \\ -542 & -0.1678 & -935.2 & 540.6 \\ \hline -1 & 0.5776 & -2.246 & 0 \end{array}\right] \tag{4.20}$$

$$K_{H_\infty obs} = \left[\begin{array}{ccc|c} -1.255 & 0 & 1.284 & 1.255 \\ -0.3074 & -0.666 & 0.4703 & 0.3074 \\ -680.7 & 340.1 & -1503 & 1.031 \\ \hline -503.8 & 252.8 & -1112 & 0 \end{array}\right] \tag{4.21}$$

The inclusion of the LQG and $\mathcal{H}_\infty$ LS controllers' solution is done to maintain the generality of the discussion of the previous section about the commonalities and differences between the controllers; nevertheless, as Papageorgiou and Polanksy assume the full-state feedback solution, only the matrix C_k from the controllers is required to recreate the results. If the value of C_k from the LQG or the $\mathcal{H}_\infty$ unstructured controller (notice that C_k is the same even though the controllers are not the same) is used to derive the different DI parameters (this is referred to as the Tuned DISE controller) according to formulae 4.16, 4.17, 4.18, then the values presented in Table 4.2 are obtained.

Table 4.2: Verification outcome of the procedure outlined in G. Papageorgiou and Polansky, 2009.

	Obtained Tuned DISE	From G. Papageorgiou and Polansky, 2009
K_b	3.9354	3.93
f_i	0.3002	0.297
f_c	0.5906	0.59
$\hat{C}$	[-0.2520 1]	[-0.248 1]

From the table above, it can be observed that the results were successfully recreated and that the deviation in the results can be attributed to rounding differences in intermediate computations. Having done this, it is important to note that the state feedback matrix from LQ optimization, $F = [F_I \ F_\alpha \ F_q]$ corresponds to the C_k matrix of the K_{LQG}, since it is assumed it multiplies the states of the plant. The pitch rate is assumed to be measured, and the angle of attack is estimated based on model information, which introduces an estimator in the controller. If one would recreate the procedure using the C_k matrix of the $\mathcal{H}_\infty$ Loop-Shaping controller written in an observed-based way, then:

$$F = C_k = [-503.8 \ 252.8 \ -1112] \tag{4.22}$$

and following the equations in 4.16, 4.17 and 4.18, it would result in the DI gains illustrated in Table 4.3, which confirms that the equivalence between both methods does not hold when the $\mathcal{H}_\infty$ Loop-Shaping control solution is used.

Table 4.3: Discrepancy in the DI gains using the C_k matrix of the observer-based $\mathcal{H}_\infty$ Loop-Shaping.

	$\mathcal{H}_\infty$ equivalent	Obtained Tuned DISE
K_b	990.1599	3.9354
f_i	0.0023	0.3002
f_c	4.7415e-06	0.5906
$\hat{C}$	[510.6672 1]	[-0.2520 1]

4.3. Further Insights to Lemma 1

According to Lemma 1, the gain associated with the integrator state always equals -1. Since the fact that this gain is -1 is central to the derivations at hand, it is of interest to understand the implications of this lemma.

In order to understand the principles behind this lemma, it is important to note that the post-compensator filter W_2 is a PI with a transfer function equal to:

$$W_2(s) = \frac{K_p s + K_i}{s} \tag{4.23}$$

Thus, the original system P, given by:

$$\begin{bmatrix} \dot{\alpha} \\ \dot{q} \end{bmatrix} = \underbrace{\begin{bmatrix} Z_\alpha & Z_q \\ M_\alpha & M_q \end{bmatrix}}_{A} \begin{bmatrix} \alpha \\ q \end{bmatrix} + \underbrace{\begin{bmatrix} 0 \\ M_\alpha \end{bmatrix}}_{B} \delta \tag{4.24}$$

is augmented by W_2, such that P_s = $W_2 P$, where the state-space of W_2 can be formulated as:

$$\begin{aligned} \dot{x}_i &= A_{W_2} x_i + B_{W_2} y \\ z_1 &= C_{W_2} x_i + D_{W_2} y \end{aligned} \tag{4.25}$$

and the state-space of P_s as:

$$\begin{aligned} \begin{bmatrix} \dot{x_i} \\ \dot{\alpha} \\ \dot{q} \end{bmatrix} &= \underbrace{\begin{bmatrix} A_{W_2} & 0 & B_{W_2} \\ 0 & Z_\alpha & Z_q \\ 0 & M_\alpha & M_q \end{bmatrix}}_{A_s} \begin{bmatrix} x_i \\ \alpha \\ q \end{bmatrix} + \underbrace{\begin{bmatrix} 0 \\ 0 \\ M_\delta \end{bmatrix}}_{B_s} \delta \\ z_1 &= \underbrace{\begin{bmatrix} C_{W_2} & 0 & D_{W_2} \end{bmatrix}}_{C_s} \begin{bmatrix} x_i \\ \alpha \\ q \end{bmatrix} \end{aligned} \tag{4.26}$$

A single transfer function has infinitely many state-space formulations that all describe the same input-output relationship, as the transfer function does not uniquely determine internal states in a state-space model. Thus, the choice of the internal state of the state-space model of the W_2 can lead to a different representation of W_2 and, consequently, P_s. The reason why this is relevant will be made clear. However, for the remainder of this derivation, it is assumed that $A_{W_2} = 0$.

Choosing the integrator state x_i to be equal to $K_i \cdot \frac{1}{s} \cdot q$, the following state-space representation of P_s is obtained:

$$\begin{aligned}\begin{bmatrix}\dot{x_i}\\ \dot{\alpha}\\ \dot{q}\end{bmatrix} &= \underbrace{\begin{bmatrix}0 & 0 & K_i\\ 0 & Z_\alpha & Z_q\\ 0 & M_\alpha & M_q\end{bmatrix}}_{A_s}\begin{bmatrix}x_i\\ \alpha\\ q\end{bmatrix} + \underbrace{\begin{bmatrix}0\\ 0\\ M_\delta\end{bmatrix}}_{B_s}\delta\\ z_1 &= \underbrace{\begin{bmatrix}1 & 0 & K_p\end{bmatrix}}_{C_s}\begin{bmatrix}x_i\\ \alpha\\ q\end{bmatrix}\end{aligned} \tag{4.27}$$

As mentioned in Equation 4.10, to compute the feedback controller, one needs to compute the solution to the following ARE:

$$A_s^\top X + XA_s - (XB_s)(XB_s)^\top + C_s^\top C_s = 0 \tag{4.28}$$

The feedback controller can then be computed from:

$$F_x = \begin{bmatrix} F_{x_i} & F_\alpha & F_q \end{bmatrix} = -B_s^\top X \tag{4.29}$$

By noting that B_s has the structure as shown in system 4.27, and that the solution for the Riccati matrix $X > 0$ has a full-block structure:

$$X = \begin{bmatrix} X_{11} & X_{12} & X_{13}\\ X_{21} & X_{22} & X_{23}\\ X_{31} & X_{32} & X_{33}\end{bmatrix} \tag{4.30}$$

then, the feedback gain on the integrator state F_{x_i} can be computed as:

$$F_x = \begin{bmatrix} -M_\delta X_{31} & -M_\delta X_{32} & -M_\delta X_{33} \end{bmatrix} \Rightarrow F_{x_i} = -M_\delta X_{31} \tag{4.31}$$

Solving the ARE for the system 4.27, considering X symmetric and real, yields the following solution to X_{31} and, consequently, to F_{x_i}:

$$X_{31} = \frac{1}{M_\delta} \Rightarrow F_{x_i} = -M_\delta \frac{1}{M_\delta} = -1 \tag{4.32}$$

Thus, it is shown how the integrator state's -1 gain described in Lemma 1 arises. The question that arises is whether this lemma holds for a different state-space realisation of the same transfer function of the PI structure of W_2. Considering the generic state-space system 4.26 with $A_{W_2} = 0$ and solving again the ARE, the solution is:

$$X_{31} = \frac{C_{W_2}}{M_\delta}, \text{ and thus } F_{x_i} = -M_\delta \frac{C_{W_2}}{M_\delta} = -C_{W_2} \tag{4.33}$$

This means that the same transfer function of W_2, under a different state-space realisation, leads to a different value for F_{x_i}. As a result, if one wants to have $F_{x_i} = -1$, one needs to define the integrator

state exactly as done in 4.27. However, this is not necessarily problematic. If Figure 4.10 is inspected, it is possible to observe that the necessary assumption is that the gain on the integrator path of q (referred to as y in the figure) needs to be the same as the one in the path of the integrator of q_{cmd} (referred to as y_c in the figure). Thus, by defining x_i = q, the following realization is obtained:

$$A_s = \begin{bmatrix} 0 & 0 & 1 \\ 0 & Z_\alpha & Z_q \\ 0 & M_\alpha & M_q \end{bmatrix}, \quad B_s = \begin{bmatrix} 0 \\ 0 \\ M_\delta \end{bmatrix}, \quad C_s = \begin{bmatrix} K_i & 0 & K_p \end{bmatrix}. \tag{4.34}$$

For this realization of W_2, the solution to X_{31} will be equal to $\frac{K_i}{M_\delta}$ and thus $F_{x_i} = -K_i$. In this case, F_{x_i} is already the gain on the entire path of q and looking back to how G. Papageorgiou and Polansky, 2009 defines the integrator state x_i, it is possible to observe that although $F_{x_i} = -1$, since $x_i = \frac{K_i \cdot q}{s}$, the actual gain on the integrator path of q is $K_i \cdot F_{x_i}$ = $K_i \cdot -1$ = $-K_i$. Thus, choosing a different realisation of W_2, which leads to a different definition for the integrator state x_i, has no implications in the procedure.

4.4. Discussion

The current chapter serves as a motivation for the path chosen throughout the rest of the thesis, from the various possible approaches to combining Dynamic Inversion with $\mathcal{H}_\infty$ Loop-Shaping control identified in the literature study in Section 2.1. It has been demonstrated that in the OF case, the results from G. Papageorgiou and Polansky, 2009 would fall into the same pitfalls of $\mathcal{H}_2$ control and, consequently, LQG control, which, as demonstrated by Doyle, 1978, come with no robustness guarantees. In fact, with this procedure, a designer must test their controller margins after the design, like one would need to do for either LQG or Dynamic Inversion control. The only equivalence to $\mathcal{H}_\infty$ Loop-Shaping control is in the first step, where the plant is augmented with shaping filters. Considering an observer-based structure, if the $\mathcal{H}_\infty$ Loop-Shaping controller is used, then the equivalence with DI control does not hold, as shown in Table 4.3.

In an attempt to find a possible equivalence between both approaches that could leverage the advantages of each method while mitigating its limitations, it is concluded that an alternative procedure should be considered. This does not mean that this procedure does not offer a lot of potential, since the idea of making the choice of the controlled variable free by making the matrix C a tunable parameter appears to give an extra degree of freedom to the DI controller structure, and thus, an extra degree of freedom to attain a robust design.

However, for this thesis's scope, the DI controller's structure will be restricted to cases where $C = \hat{C}$ to maintain the physical insight of the integrator behaviour of the transfer function from $\dot{y}^d$ to y. These are model-based inversion, sensor-based inversion, and a relatively new hybrid inversion scheme configuration approach. Similarly, a hybrid inversion scheme introduces more degrees of freedom to the inversion to address robustness, just like making $\hat{C}$ not necessarily equal to C does; nevertheless, unlike the $\hat{C}$ tunable parameter, where robustness and nominal performance are closely tied, the extra DoF introduced in the hybrid inversion approach to address robustness is independent from closed-loop nominal behaviour. The main goal consists of replicating the design philosophy of Pollack, 2024, where the various elements of a Hybrid IDI controller (virtual law controller and inversion gains) are left as tunable elements and are then optimised based on a linear robust control synthesis method: μ-synthesis. In the present thesis, the goal is to use the $\mathcal{H}_\infty$ Loop-Shaping Design Procedure instead to optimise the Hybrid IDI controller.

5

Combining $\mathcal{H}_\infty$ LSDP with Dynamic Inversion Control

This chapter is devoted to a study of the fundamental robustness properties of linear dynamic inversion approaches, such as Model-Based Dynamic Inversion, Sensor-Based (Incremental) Dynamic Inversion and Hybrid-Based (Incremental) Dynamic Inversion. The goal is to obtain a better understanding of the robustness implications of the different inversion strategies and analyse what connections can be made to $\mathcal{H}_\infty$ Loop-Shaping Control. The scope of the analysis is limited to the short-period dynamics of the open-access simulation model (low-fidelity version) of the General Dynamics F-16 (Russell, 2003). Following this, the acquired insights are used for the development of a novel approach to combine $\mathcal{H}_\infty$ Loop-Shaping with a DI structured controller. This approach employs the use of $\mathcal{H}_\infty$ LSDP and non-smooth optimization techniques to solve the $\mathcal{H}_\infty$ synthesis problem under structural constraints to optimize the controller.

5.1. Robustness Implications of Different Inversion Strategies

This section will analyse the broken-loop gain of different inversion strategies (Model-Based, Sensor-Based and Hybrid-Based) at the plant I/O. This analysis is often disregarded in the DI framework, but the work of Pollack and Van Kampen, 2023 suggests its usefulness in understanding the fundamental robustness trade-offs of the different Dynamic Inversion strategies (refer to the discussion in Section 3.5). One of the most interesting takeaways is that the information provided by the μ-analysis can also be extracted from simply analysing the broken-loop gain at the plant input (Pollack & Van Kampen, 2023). This is one of the main reasons that motivated the research on the combination of the $\mathcal{H}_\infty$ LSDP with DI-based controllers.

However, to isolate the effect of the different inversion strategies on the broken-loop gain and to evidence that robustness properties of (I)NDI remain dependent on the nature of the open-loop plant, a similar analysis is performed considering only the inversion loop and multiple flight conditions. For the study in question, the flight envelope displayed in Figure 5.1 is considered, and the actuator model used is a first-order model with $w_{act} = 20.2$ rad/s.

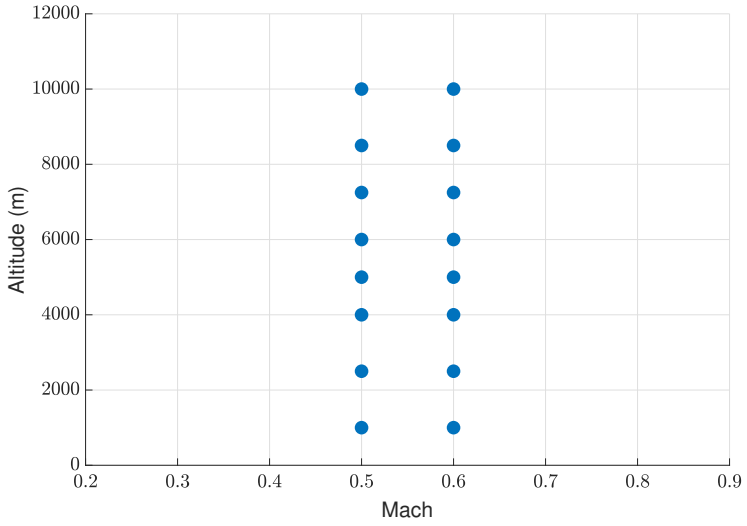

Figure 5.1: Flight-envelope considered.

5.1.1. Model-Based Inversion

Considering an ideal model-based inversion scenario, where it is assumed that there is access to ideal measurements of both the pitch-rate q and the angle-of-attack (AoA) α as displayed in Figure 5.2, it is of interest to analyse the broken-loop gain at the plant I/O (the locations represented with the vertical bars) and the closed-loop relationship $v \rightarrow y$ for each flight condition. The results are displayed in Figure 5.3.

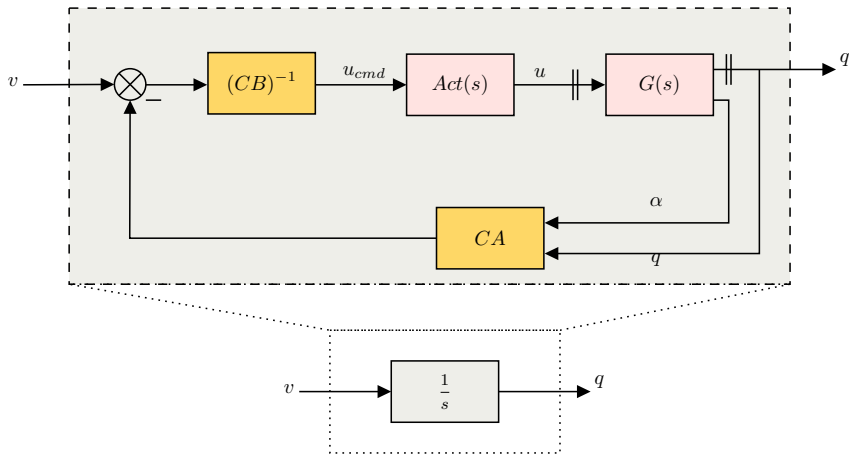

Figure 5.2: MB inversion schematic considering that measurements of all states are available.

As expected, at the virtual control location v, the system exhibits the $\frac{1}{s}$ behaviour at low to intermediate frequencies (see Figure 5.3), with an additional roll-off at high frequencies due to the actuator dynamics (unaccounted for in the inversion). However, the broken-loop gain at the plant I/O varies with flight condition, which confirms the fact that the robustness properties are local properties that are flight condition-dependent. Moreover, inversion alone is not sufficient to warrant desired broken-loop responses according to classical loop-shaping guidelines (Skogestad & Postlethwaite, 2005), which highlights the importance of the virtual law controller in guaranteeing a robust design.

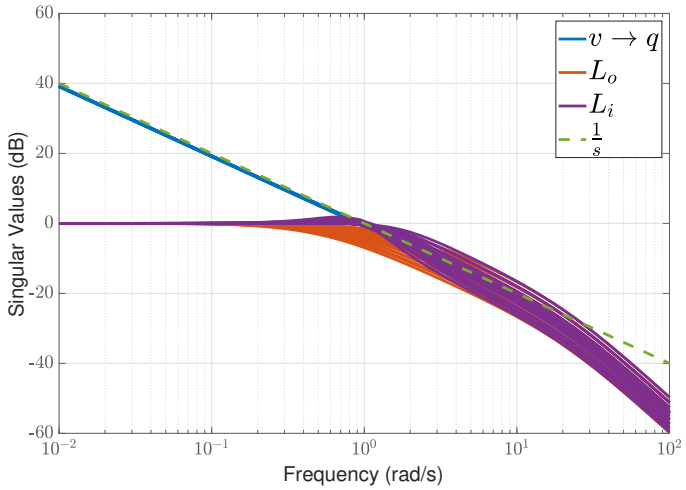

Figure 5.3: Broken-loop response at the plant I/O and closed-loop relationship between v and q for different flight condition points using MB DI.

Accurate measurements of air data angles are often difficult to obtain, especially in high-speed conditions (Honeywell Technology Center, Lockheed Martin Skunk Works and Lockheed Martin Tactical Aircraft Systems, 1996). Therefore, it is preferred to make use of inertially derived estimations of the angle-of-attack (AoA) α. The α estimator introduced is based on a combination of measured pitch rate with the short-period force equation (G. Papageorgiou & Polansky, 2009), resulting in the following definition of the estimated states $\hat{x}$ (see the short-period state-space definition in Equation 6.1):

$$\hat{x}(s) = \left[\hat{\alpha}(s)\ q(s)\right]^T = F(s)\ q(s) = \left[\frac{1+Z_q/V_o}{s-Z_\alpha} \quad 1\right]^T q(s) \tag{5.1}$$

The introduction of the estimator effectively makes the nature of the problem SISO (see Figure 5.4), resulting in equal broken-loop responses at the plant input and output.

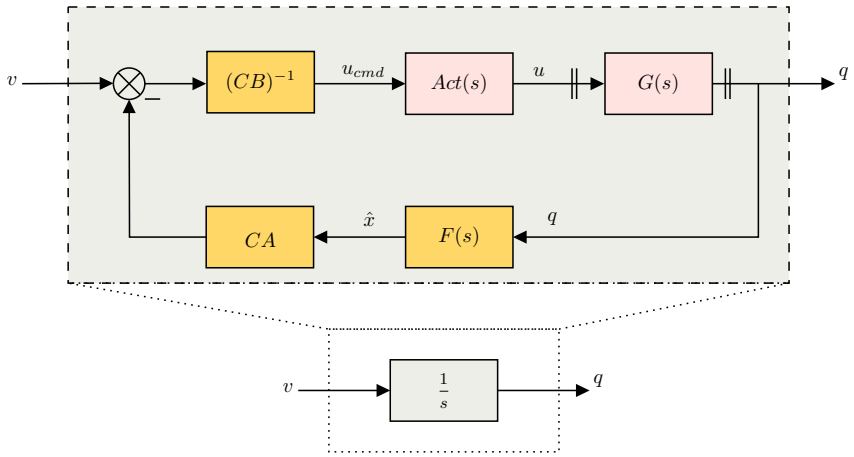

Figure 5.4: MB inversion schematic with AoA α estimator.

In this scenario, at the virtual control law location v, the transfer function from v to q is still close to $\frac{1}{s}$, even though there is a distortion of the gain at low frequencies due to the introduction of the estimator. The response in terms of the broken-loop gain at plant I/O in Figure 5.5 is similar to the one in Figure 5.3, and as such, similar conclusions are retrieved.

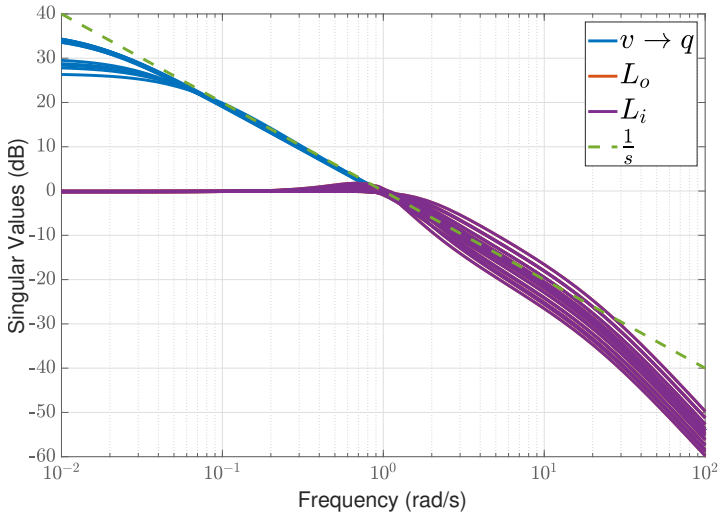

Figure 5.5: Broken-loop response at the plant I/O and closed-loop relationship between v and q for different flight condition points using MB DI with AoA α estimator.

5.1.2. Sensor-Based Inversion

The same analysis is conducted for a Sensor-Based inversion scheme, but the output location is disregarded for simplicity. Since this approach makes use of output differentiation and input feedback, estimation of the angle of attack is not necessary.

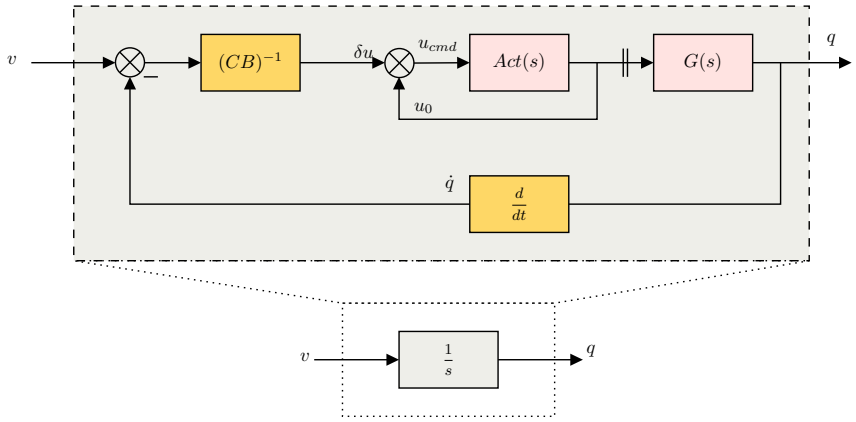

Figure 5.6: SB IDI inversion schematic.

Comparing Figures 5.5 and 5.7, it is possible to observe that SB IDI results in an elevated open-loop response compared to MB DI, resulting in a significantly higher gain at low frequencies but also a higher crossover frequency. This simple analysis confirms the main takeaways from Pollack and Van Kampen, 2023 that SB IDI is not necessarily more robust than MB IDI. Namely, it demonstrates how unfiltered IDI approaches could warrant robustness issues against high-frequency uncertainty and the importance of the synchronisation low-pass filter H_c in ensuring a robust IDI design. Moreover, it adds clarity to the typical guidelines stating that PI gains can be decreased in the case of SB INDI control.

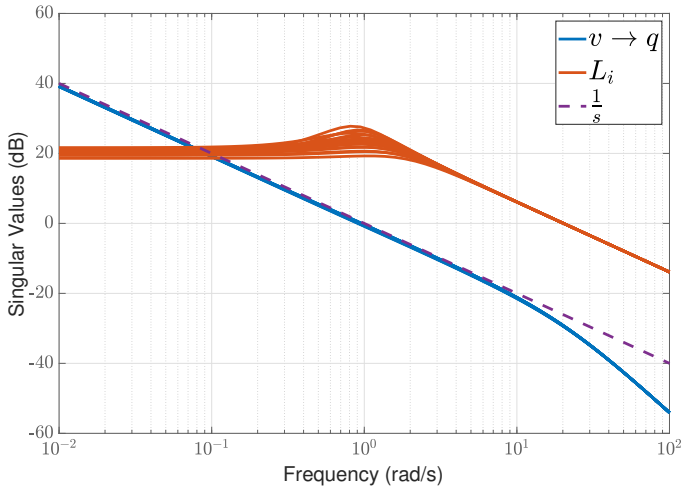

Figure 5.7: Broken-loop response at the plant I/O and closed-loop relationship between v and q for different flight condition points using SB IDI.

5.1.3. Hybrid-Based Inversion

The previous analysis has demonstrated that the robustness characteristics of a fully sensor-based IDI are complementary to those of the classical model-based DI variant. Therefore, balanced robustness properties can be obtained by blending both inversion techniques. The following analysis will focus on one single point of the flight envelope ($Mach = 0.6$ and $h = 6000\ m$) to illustrate that Hybrid IDI can be thought of as a way to navigate between MB DI and SB IDI robustness properties by manipulating the scaling gain K_c and the bandwidth of the synchronizing filter H_c via w_{H_c}. The reason to focus on one single point is to visually demonstrate in a clear and simple way how the hybrid scheme offers this extra design flexibility to address robustness.

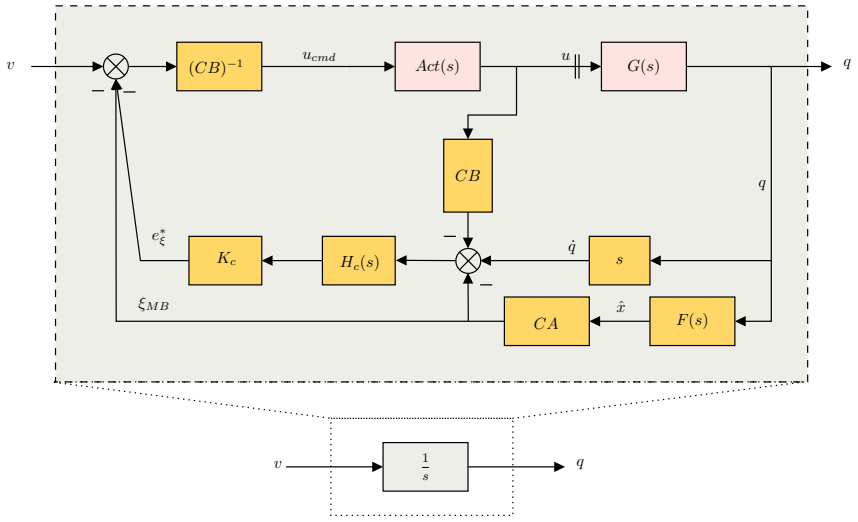

Figure 5.8: HB IDI inversion schematic.

For the Hybrid IDI schematic in Figure 5.8, it was assumed that ideal measurements of the actual input u were accessible and the synchronising filter H_c was set as a second-order Butterworth low-pass filter with cutoff angular frequency equal to w_{H_c}. From Figures 5.9, 5.10, and 5.11, it is possible to visualise how the Hybrid IDI inversion scheme offers extra degrees of freedom to adjust the broken-loop gain of the system

and, hence, address robustness. As expected, by setting $K_c = 0$, the Hybrid IDI schematic collapses to the MB DI, while setting $K_c = 1$ and making the bandwidth of the synchronisation filter large enough $w_{H_c} \to +\infty$ makes it collapse to an SB IDI scheme. Setting intermediate values for the hybrid inversion scheme ($K_c = 0.5$ and $w_{H_c} = 40$ rad/s) results in an intermediate broken-loop response between the MB DI and the SB IDI. Therefore, the advantages of using a Hybrid IDI scheme are clearly demonstrated.

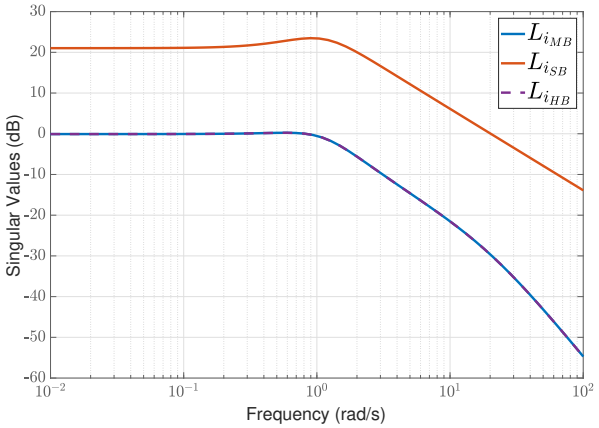

Figure 5.9: Broken-loop gain at the plant input for MB DI, SB IDI, and HB IDI with $K_c = 0$.

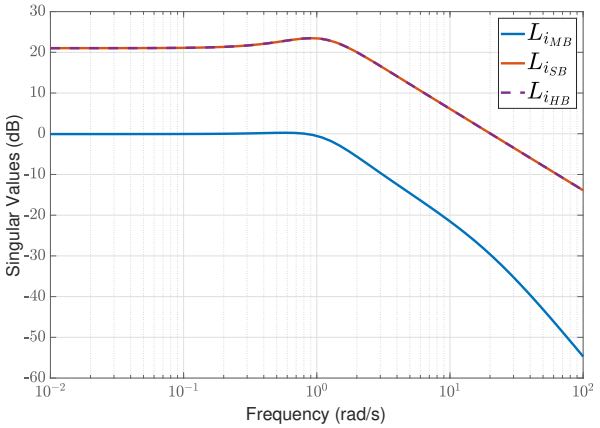

Figure 5.10: Broken-loop gain at the plant input for MB DI, SB IDI, and HB IDI with $K_c = 1$ and $w_{H_c} \to +\infty$.

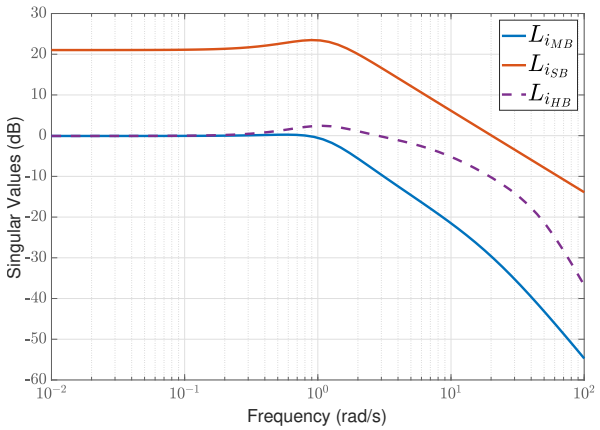

Figure 5.11: Broken-loop gain at the plant input for HB IDI with $K_c = 0.5$ and $w_{H_c} = 40$ rad/s.

5.1.4. Discussion

While the results in this chapter do not add much to the discussion conducted in Section 3.5, they do validate in a simple yet practical way all the important takeaways about the robustness implications of different inversion strategies.

This simple analysis raises concerns about the often-claimed “automatic gain-scheduled nature” of (I)NDI control laws. While it is true that under ideal inversion, the closed-loop relationship between v and y exhibits a consistent behaviour across the flight envelope, such consistency is not present where it actually matters from a robustness perspective: at the plant I/O. An analogy can be established with the LQG controller, which has guaranteed excellent (input) margins at a location internal to the controller, but not at the actual plant I/O, resulting in the absence of guaranteed margins for LQG regulators (Doyle, 1978). Furthermore, from a robust control perspective, the v signal has little relevance as it concerns a signal

internal to the controller itself. Attempting to design a robust control based on this v signal, as done in many (I)NDI applications, fundamentally "blinds" the designer to robustness. Therefore, the robust design of (I)NDI laws still requires some form of gain-scheduling to achieve consistent robust properties, either by employing linear $\mathcal{H}_\infty$-synthesis tools within the divide-and-conquer paradigm (Pollack, 2024) or by making use of LPV theory (Pollack et al., 2025).

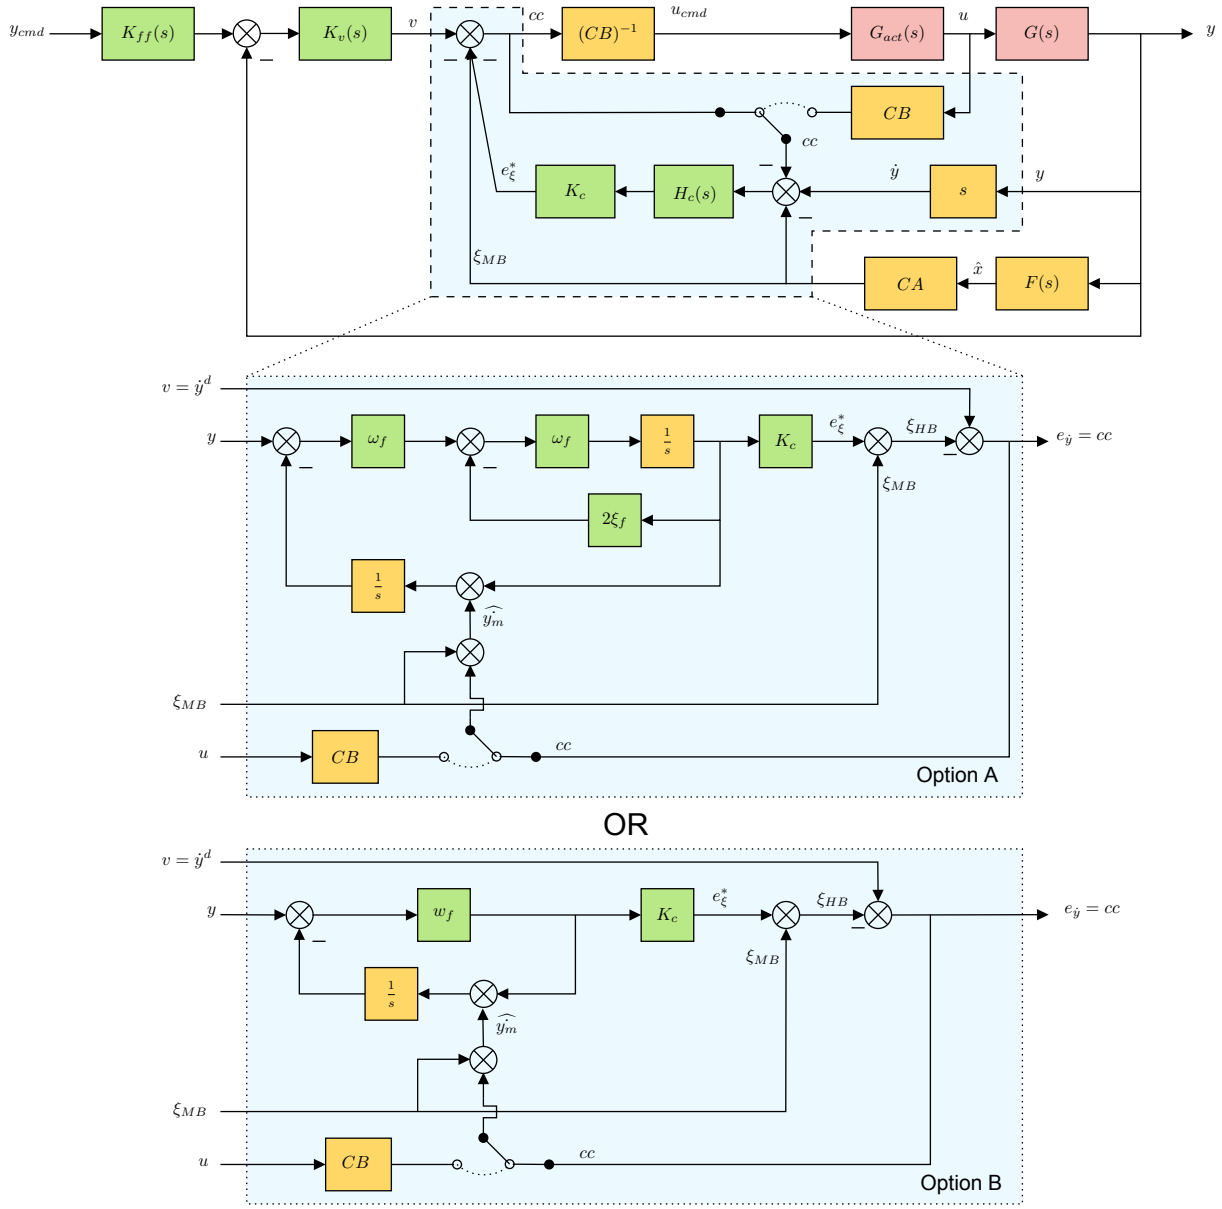

Figure 5.12: Hybrid IDI structure, where the green blocks refer to the tunable elements of the controller and the yellow blocks to the plant-dependent/physical blocks of the controller; the insert shows two options for the implementation of a scaled complementary filter (SCF) signal blending.

As a result, the discussion above justifies the use of a linear Hybrid IDI control law, which can be combined with $\mathcal{H}_\infty$-synthesis tools. The goal is to use the $\mathcal{H}_\infty$ LSDP to optimise the different hybrid gains to design a controller with robustness guarantees. The Hybrid IDI structure is displayed in Figure 5.12. For the sake of completeness, the input feedback case is also included in the schematic, even though the remainder of the thesis considers only the CC feedback case. The insert highlights that the hybrid

inversion can be implemented as a scaling complementary filter (SCF), which is advantageous since it does not require explicit signal differentiation. Two options are displayed, with option A referring to a second-order low-pass filter H_c and option B referring to a first-order one. As previously described, the analysis conducted in this section considered a second-order low-pass filter H_c, but for the following chapters, a first-order low-pass filter H_c is used to maintain commonality with the work of Pollack et al., 2024 and since it results in a lower complexity of the control law. The virtual law controller K_v is defined as a PI controller in series with a low-pass filter, and K_{ff} is a feedforward controller used to improve model-following performance. More details on the structure of K_{ff} are provided in the following chapters.

Despite the fact that this linear synthesis approach drifts away from the typical global approach to the design of (I)NDI control laws, it is important to notice that the architectural value of these incremental laws is maintained. Nevertheless, combining DI-based architecture with the $\mathcal{H}_\infty$ Loop-Shaping have its particular challenges:

The specific architecture of DI-based controllers imposes a certain structure on the inversion path and closes the loop at the virtual law v location. Adding weighting filters to the open-loop plant to adjust the open-loop response, as done in the $\mathcal{H}_\infty$ LSDP, may distort the inversion path and compromise the closed-loop relationship between v and y. On the other hand, trying to use the $\mathcal{H}_\infty$ LSDP from the virtual law location and relying on the closed-loop relationship between v and y would fail to consider the relevant location for robustness: the plant I/O. The fact that DI-based controllers have such a specific architecture and that the $\mathcal{H}_\infty$ LSDP also exhibits less flexibility than it would be desired (requirement to shape the open-loop response with weighting filters) requires an alternative setup of the $\mathcal{H}_\infty$ LSDP.

5.2. Equivalence between Classic $\mathcal{H}_\infty$ Loop-Shaping and an Alternative Schematic

In Zhou and Doyle, 1998, the authors demonstrate that the classical $\mathcal{H}_\infty$ LSDP can equivalently be interpreted as the more standard $\mathcal{H}_\infty$ problem formulation of minimising the $\mathcal{H}_\infty$ norm of the frequency-weighted gain from disturbances on the plant input and output to the controller input and output.

In order to demonstrate this equivalence, the steps of the derivation from Zhou and Doyle, 1998 are recreated. Recall that the NCF robust stabilisation problem can be written as a mixed sensitivity 4-block problem, as presented in Figure 5.13 and Equation 5.2.

$$\min_{stab\ K_\infty} \left\| \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} \rightarrow \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} \right\|_\infty = \min_{stab\ K_\infty} \left\| \begin{bmatrix} I_p \\ -K_\infty \end{bmatrix} (I_p + G_s K_\infty)^{-1} \begin{bmatrix} I_p & G_s \end{bmatrix} \right\|_\infty \leq \gamma \triangleq \epsilon^{-1} \tag{5.2}$$

Using the fact that $G_s = W_2 G W_1$ and defining as the total controller $K_s = W_1 K_\infty W_2$, and assuming that the weighting filters are square and invertible, it is possible to rewrite Equation 5.2 as:

$$\left\| \begin{bmatrix} I_p \\ -K_\infty \end{bmatrix} (I_p + G_s K_\infty)^{-1} \begin{bmatrix} I_p & G_s \end{bmatrix} \right\|_\infty = \left\| \begin{bmatrix} W_2 \\ -W_1^{-1} K_s \end{bmatrix} W_2^{-1} (I_p + G_s K_\infty)^{-1} W_2 \begin{bmatrix} W_2^{-1} G W_1 \end{bmatrix} \right\|_\infty \tag{5.3}$$

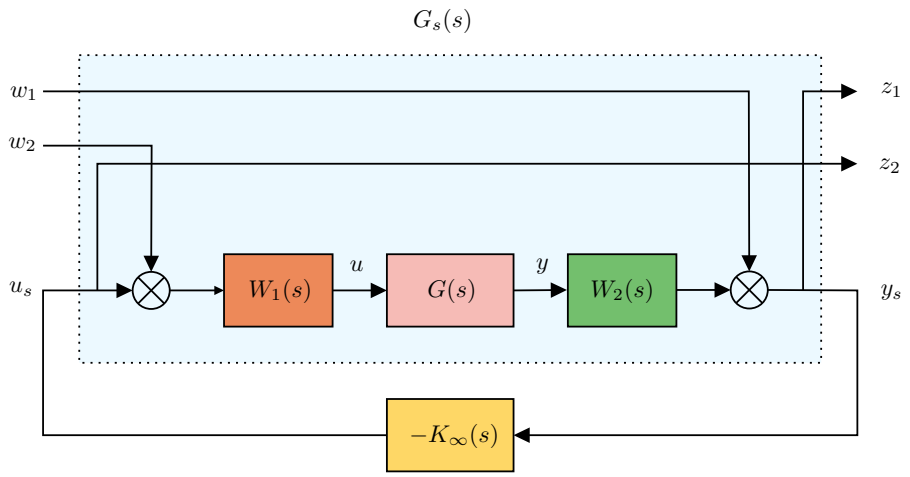

Figure 5.13: NCF robust stabilisation problem written as mixed sensitivity 4-block problem.

Using the fact that the inner part $W_2^{-1}\left(I_p + G_s K\right)^{-1} W_2$ can be written as:

$$\left[W_2^{-1}\left(I_p + W_2 G W_1 K_\infty\right) W_2\right]^{-1} = \left(I_p + G K_s\right)^{-1} \tag{5.4}$$

and the following equivalence holds:

$$\begin{aligned} &\left\| \begin{bmatrix} I_p \\ -K_\infty \end{bmatrix} \left(I_p + G_s K_\infty\right)^{-1} \begin{bmatrix} I_p & G_s \end{bmatrix} \right\|_\infty \leq \epsilon^{-1} \Leftrightarrow \\ &\left\| \begin{bmatrix} W_2 \\ -W_1^{-1} K_s \end{bmatrix} \left(I_p + G K_s\right)^{-1} \begin{bmatrix} W_2^{-1} & G W_1 \end{bmatrix} \right\|_\infty \leq \epsilon^{-1} \end{aligned} \tag{5.5}$$

If the multiplications in Equation 5.5 are expanded, the following particular weighted mixed sensitivity problem is achieved:

$$\left\| \begin{bmatrix} W_2(I_p + G K_s)^{-1} W_2^{-1} & W_2(I_p + G K_s)^{-1} G W_1 \\ -W_1^{-1} K_s (I_p + G K_s)^{-1} W_2^{-1} & -W_1^{-1} K_s (I_p + G K_s)^{-1} G W_1 \end{bmatrix} \right\|_\infty \leq \epsilon^{-1} \tag{5.6}$$

The previous result can be represented as a 4-block mixed sensitivity problem with weights on the disturbances and plant input and output, as shown in Figure 5.14 below.

$$\begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} W_2\left(I_p + G K_s\right)^{-1} W_2^{-1} & W_2\left(I_p + G K_s\right)^{-1} G W_1 \\ -W_1^{-1} K_s\left(I_p + G K_s\right)^{-1} W_2^{-1} & -W_1^{-1} K_s\left(I_p + G K_s\right)^{-1} G W_1 \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} \tag{5.7}$$

The advantage of this equivalent formulation is that it is possible to consider the inputs and outputs in the actual plant I/O instead of the shaped plant I/O's.

As described in the previous chapter, one of the challenges of combining $\mathcal{H}_\infty$ LSDP with Dynamic Inversion is that the fundamental robustness properties are tied to the open loop at the plant I/O and not to

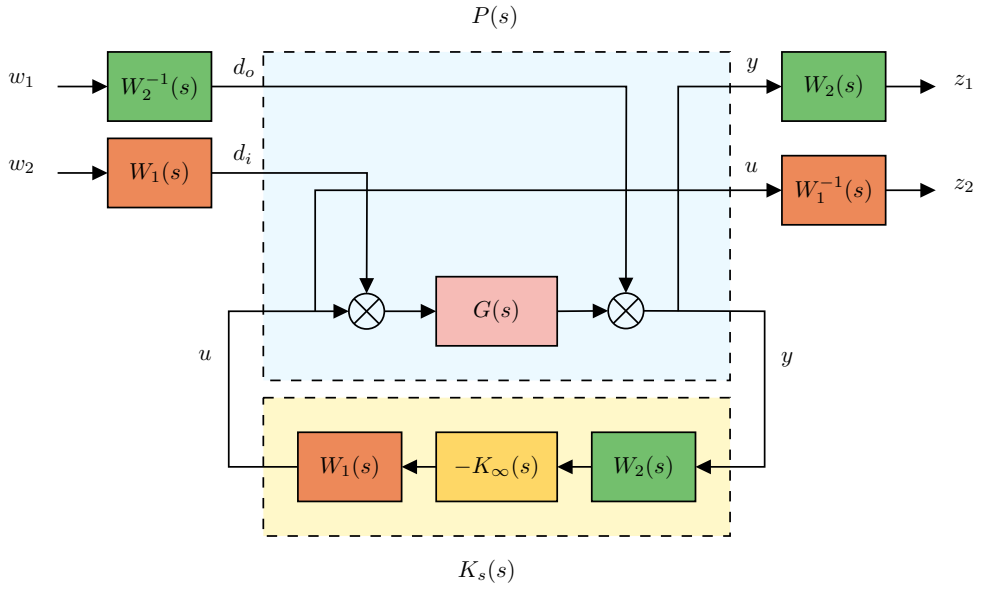

Figure 5.14: Alternative structure for the NCF LSDP robust stabilisation.

the virtual control law location. Since DI techniques close the loop at this virtual location, combining it with $\mathcal{H}_\infty$ LSDP raises concerns, as shaping the plant's open-loop response with the filters W_1 and W_2 may distort the inversion path.

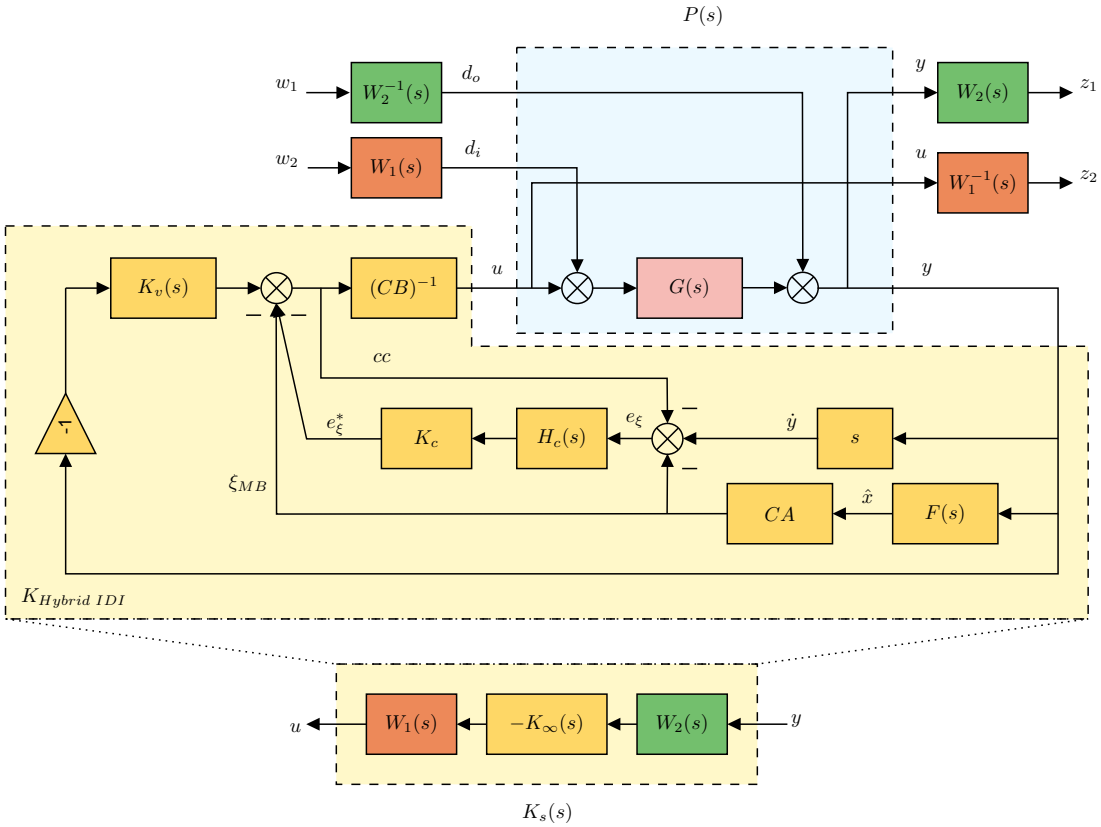

Figure 5.15: Leveraging alternative structure to tune Hybrid IDI controller using the $\mathcal{H}_\infty$ LSDP.

However, since this equivalent formulation allows to consider the disturbances and the outputs at the actual plant I/O, that difficulty is overcome. Nevertheless, the fact that the feedback controller K_s is constructed as $K_s = W_1 K_\infty W_2$ might be undesirable when highly structured control architectures are considered, like the one of Hybrid IDI. Despite that, it is hypothesised that if a given structured controller K_{struct} has compatible tunable elements to those of K_s, then K_{struct} can achieve similar results to K_s, by using non-smooth optimization techniques to solve the $\mathcal{H}_\infty$ synthesis problem under structural constraints (Apkarian & Noll, 2006a). This effectively makes the nature of the optimisation non-convex.

It has been established that there is a full-order solution for K_∞ which minimises the $\mathcal{H}_\infty$ norm from w to z. Alternatively, the $\mathcal{H}_\infty$ minimisation solution can be reimagined as: there is an optimal stabilisable controller K_s, subject to $W_1 K_\infty W_2$, which minimises the $\mathcal{H}_\infty$ norm from w to z. Reconsidering K_s as an appropriately structured controller K_{struct}, it is hypothesized that the minimisation of the $\mathcal{H}_\infty$ norm from w to z results in the tunable elements of K_{struct} being configured such that $K_{struct} \to K_s$ and $\gamma_{K_{struct}} \to \gamma_{K_s}$. Therefore, if the structure of Hybrid IDI K_{H-IDI} is compatible with K_s, it is expected that the multiple tunable elements in K_{H-IDI} are tuned such that $K_{H-IDI} \to K_s$ as displayed in Figure 5.15.

5.3. Extension to Address Flying Qualities in 2 DoF Setup

Specifications for flying qualities in manual flight control systems are typically defined using the modal parameters of Low-Order Equivalent System (LOES) models (Anonymous, 1990). These models represent the desired response of specific controlled variables to pilot inputs. LOES models are derived by simplifying high-order transfer function models that encompass the complete dynamics of the aircraft, including aerodynamics, actuator dynamics, controller dynamics, structural modes, and sensor dynamics. The LOES parameter values are generally determined based on military or other regulatory requirements to ensure the aircraft achieves satisfactory flight performance (Anonymous, 1990).

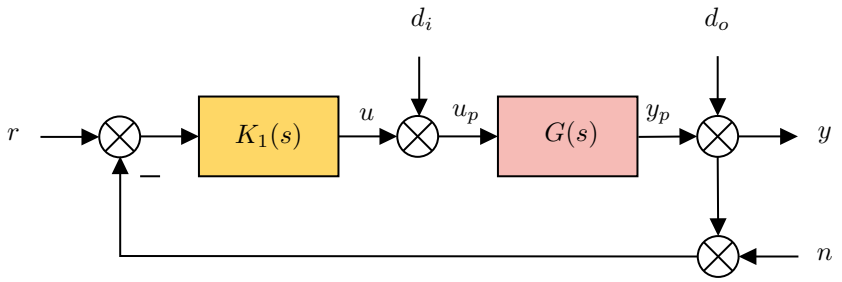

Figure 5.16: 1DoF Feedback Control Schematic.

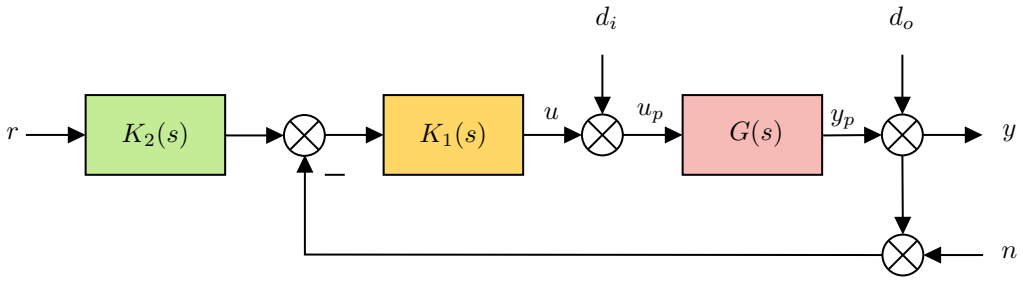

Figure 5.17: 2DoF Feedback Control Schematic.

Consider the two feedback configurations shown in Figures 5.16 and 5.17. The first is the standard 1DoF feedback, whereas the second is denominated the 2DoF. The feedback part K_1 is designed for nominal and/or robust stability (NS, RS) and disturbance rejection, whereas the feedforward part K_2 is for nominal and/or robust performance (NS, NP) (Limebeer et al., 1993). Concerning the former, the 'gang of six' (S_i, S_o, T_i, T_o, S_oG, KS_o) depend only on the feedback controller K_1 and determine RS with respect to all types of unstructured uncertainty (Skogestad & Postlethwaite, 2005). The feedforward controller K_2 is used to achieve better reference tracking, since $y = T_oK_2r$, where $T_o = GK_1(I + GK_1)^{-1}$ (Bates & Postlethwaite, 2002). Given a desired reference model $T_{ref}(s)$, one could argue that K_2 could be designed simply by setting $K_2 = T_o^{-1}T_{ref}$. This is not so simple, though, since T_o may contain unstable zeros that, when inverted, will give K_2 unstable poles. Also, K_2 may be required to have a predefined structure, typically a lead-lag structure (Bates & Postlethwaite, 2002).

In order to extend the procedure to a 2DoF setup, a tunable feedforward controller was added to the Hybrid IDI structure and model following performance requirements were introduced via minimization of the $\mathcal{H}_\infty$ norm from y_{cmd} to z_3 as shown in Figure 5.18.

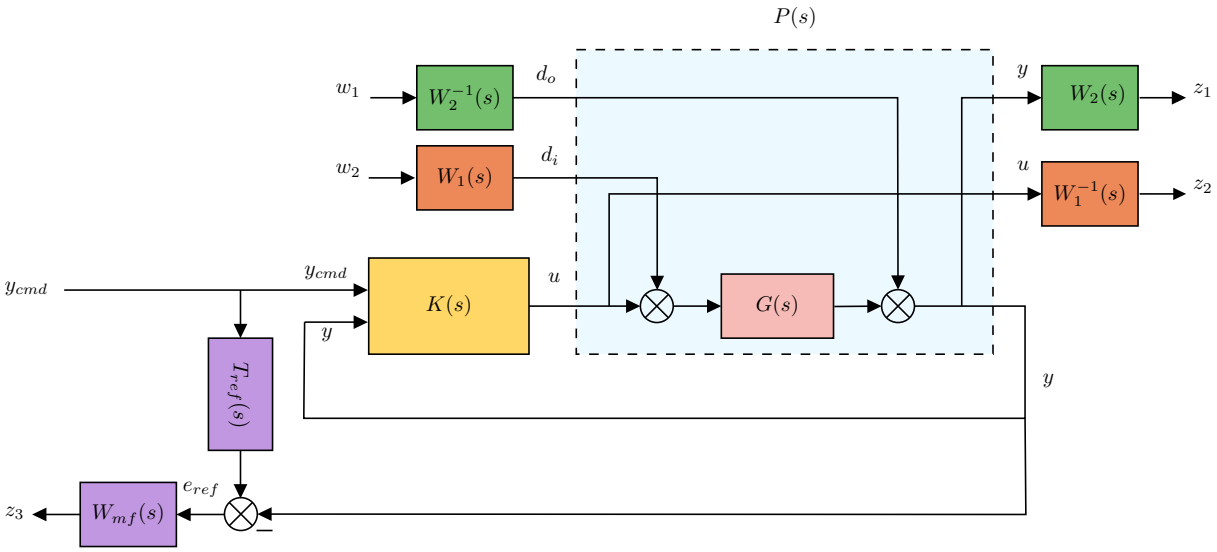

Figure 5.18: 2DoF Loop-Shaping Design Procedure.

For the specific case where the control structure is that of a Hybrid DI controller, Figure 5.18 equates to Figure 5.19. As a result, this procedure allows to systematically trade-off robustness with model following performance, as robustness is associated with the minimization of $\mathcal{H}_\infty$ norm from $[w_1\ w_2]^T \rightarrow [z_1\ z_2]^T$ and model following performance with the $\mathcal{H}_\infty$ norm from $y_{cmd} \rightarrow z_3$. As the model following performance is made more demanding, via the W_{mf} filter, the $\mathcal{H}_\infty$ norm from $[w_1\ w_2]^T \rightarrow [z_1\ z_2]^T$ is increased, decreasing the robustness of the overall system. Therefore, directly in the synthesis part, the designer can trade off these properties. More details on the working principles of designing W_{mf} are provided in the sections below.

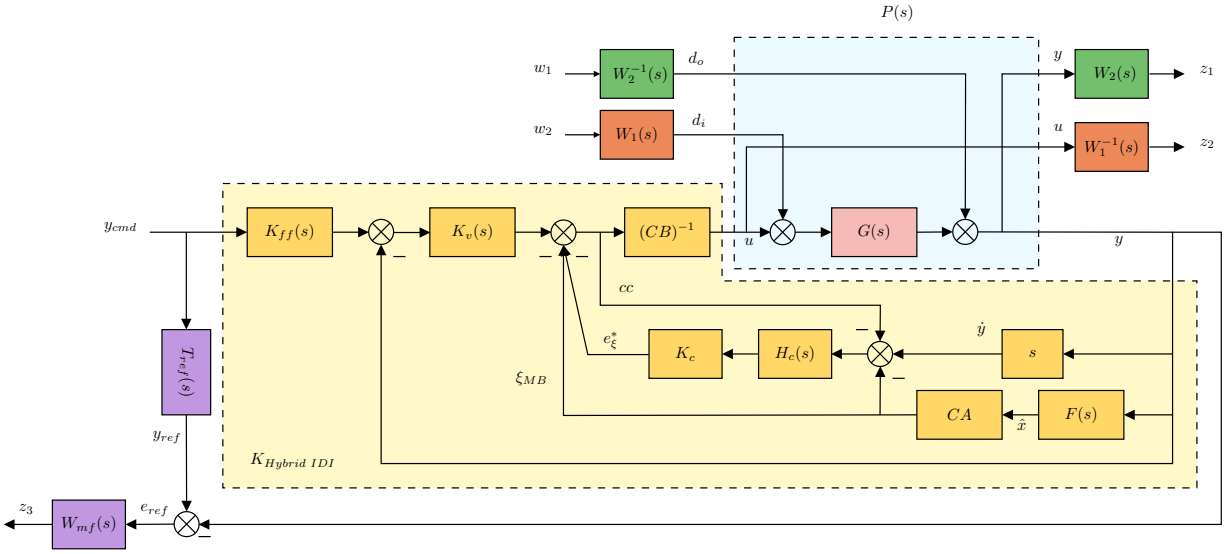

Figure 5.19: 2DoF Loop-Shaping Design Procedure for Hybrid IDI controller structure.

5.3.1. Choice of Reference Model based on Handling Qualities

The nature of the operational mission is directly linked to the formulation of the flight control law design goals. In this context, Anonymous, 1990 specifies aircraft class designation (I-IV) and identifies different flight phases (A-C). Under this criteria, the F-16 Fighting Falcon and the X-29, the two aircraft models used in the present work, fall within Class IV and the flight conditions used lie within Flight Phase Category A. Flying qualities are defined as “those characteristics of the complete air vehicle/system which allow the pilot/operator to perform to his/her satisfaction the flying tasks required to safely accomplish the mission, with an acceptable workload, while operating in the real world environment for which it is intended to operate” (Anonymous, 2002). The flying quality levels can be classified into three different levels (I-III), and it is desired to accomplish Level I satisfactory flying quality levels across the operational flight envelope.

Based on this, Anonymous, 1990 describes desired modal response characteristics. These modal parameters are given as Low Order Equivalent System (LOES) models of the desired response of certain controlled variables to pilot demands, as follows:

$$\frac{q(s)}{\delta_e(s)} = \frac{K_q \left(s + \frac{1}{T_{\theta_2}}\right) e^{-\tau_e s}}{s^2 + 2\xi\omega_n s + \omega_n^2} \tag{5.8}$$

Since the focus of this thesis is mainly on the combination of DI with $\mathcal{H}_\infty$ LSDP, a very simple approach was taken to include handling qualities into the choice of the reference model $T_{ref}(s)$. The CAP and the Gibson Dropback Criterion were used to define the multiple parameters that comprise $T_{ref}(s)$, similarly to what was done in Marques, 2024.

The Control Anticipation Parameter (CAP) serves as a physical measure of an aircraft’s manoeuvrability. It reflects the relationship between the initial pitching acceleration and the steady-state normal g-force experienced after a step input to the controls. This criterion highlights how a pilot’s ability to predict the aircraft’s flight path response is directly tied to this ratio (Gibson, 1999). The CAP can be described in terms of the variables of the LOES as described in Equation 5.9, and the flying quality levels are defined accordingly as depicted in Figure 6.8.

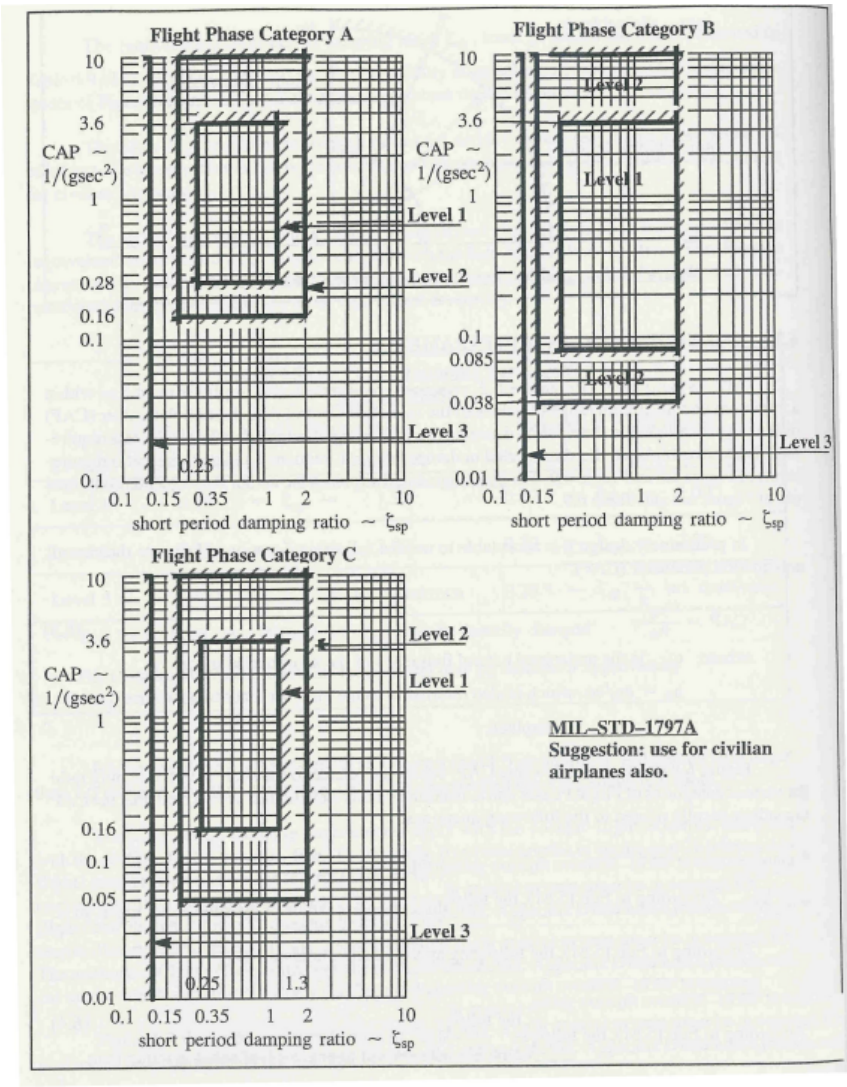

Figure 5.20: Control Anticipation Parameter and short-period damping ratio requirements. Retrieved from Roskam, 1998.

$$CAP = \frac{\dot{q}(0)}{n_z(\infty)} \approx \frac{g w_{sp}^2 T_{\theta_2}}{V_{TAS}} \tag{5.9}$$

The Gibson dropback criterion is defined by limiting two key parameters: the pitch rate overshoot ratio $\frac{q_m}{q_s}$ and the ratio of attitude dropback (or overshoot, depending on the direction of the transition when the step input is removed) to the steady-state pitch rate (Gibson, 1999). Satisfactory flying qualities according to the Gibson dropback criteria are specified in Figure 5.22, and the relevant parameters can be summarised as follows (Lombaerts & Looye, 2011):

1. q_m - maximum pitch rate;
2. q_s - steady-state value of pitch rate;
3. $\frac{q_m}{q_s}$ - pitch rate overshoot ratio;
4. DB: dropback, amount of negative transition towards final value after the step input has been removed;

5. OS: overshoot, amount of positive transition towards the final value after the step input has been removed.

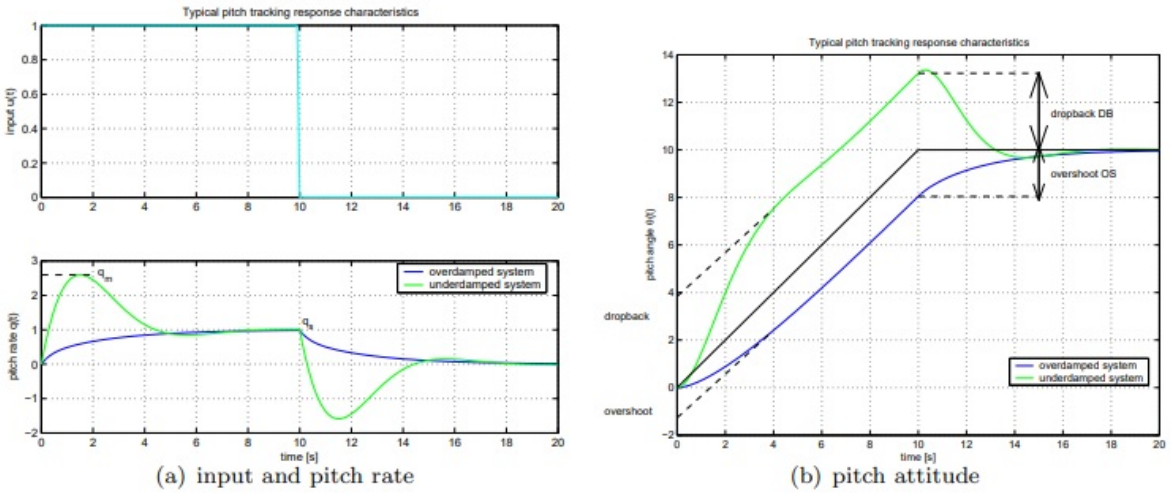

Figure 5.21: Typical pitch tracking response characteristics. Image retrieved from Lombaerts and Looye, 2011.

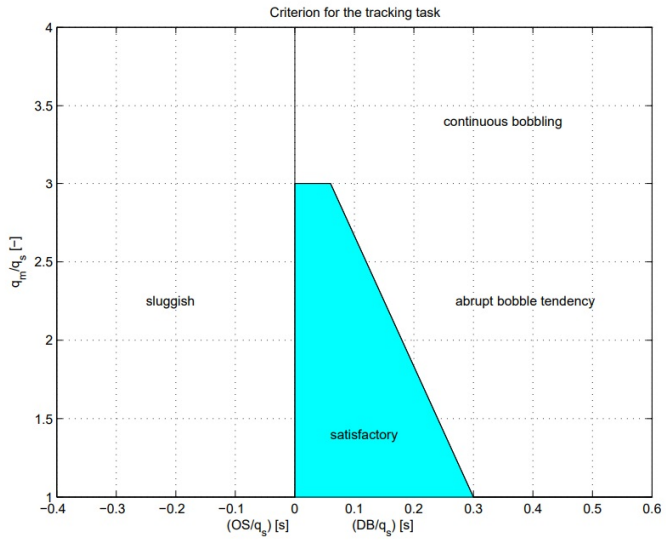

Figure 5.22: The criterion for the tracking task. Image retrieved from Lombaerts and Looye, 2011.

The Gibson can be described in terms of the variables of the LOES as follows:

$$\text{DB}/q_{ss} = T_{\theta_2} - \frac{2\xi}{\omega_n} \tag{5.10}$$

Thus, by choosing a point that adheres to Level I flying qualities in the CAP criteria in Figure 6.8, and a point that lies within the frame of satisfactory flying qualities in the Gibson criteria in Figure 5.22, the designer sets the desired CAP_{ref}, ξ_{ref}, $(DB/qss)_{ref}$. From this, it is possible to solve the system of equations in 5.11 and compute w_{ref} and $T_{\theta_{2_{ref}}}$. Regarding the remaining parameters, $K_{q_{ref}}$ is chosen such that T_{ref} has unitary static gain, and since the reference model is a pure second-order system, the equivalent time delay $\tau_{e_{ref}}$ is zero.

$$\begin{cases} \text{CAP} = \dfrac{g\,\omega_n\,T_{\theta_2}}{V_{TAS}} \\ \text{DB}/q_{ss} = T_{\theta_2} - \dfrac{2\xi}{\omega_n} \end{cases} \tag{5.11}$$

Consequently, T_{ref} is described as:

$$T_{ref}(s) = \frac{q(s)}{\delta_e(s)} = (w_n^2)(T_{\theta_2})\frac{s + \frac{1}{T_{\theta_2}}}{s^2 + 2\xi\omega_n s + \omega_n^2} \tag{5.12}$$

5.3.2. Co-Design

As described in the section above, the inclusion of model-following requirements is achieved via a model-following filter W_{mf} on the error signal e_{ref}. In this context, it is more intuitive to interpret the relevance of W_{mf}^{-1} as it is the filter that sets the bounds for the transfer function from y_{cmd} to e_{ref}, given that:

$$\left\| \frac{z_3(s)}{y_{cmd}(s)} \right\|_\infty \leq 1 \Leftrightarrow \left\| W_{mf}(s)\left(T_{ref}(s) - \frac{y(s)}{y_{cmd}(s)}\right) \right\|_\infty \leq 1 \Leftrightarrow \left\| T_{ref}(s) - \frac{y(s)}{y_{cmd}(s)} \right\|_\infty \leq \left\| W_{mf}^{-1}(s) \right\|_\infty \tag{5.13}$$

The inverse of the model following filter, W_{mf}^{-1}, is a high-pass filter. The designer has the freedom to set the DC gain, the filter bandwidth and the high-frequency gain. The lower the DC gain and the higher the bandwidth, the better the model-following performance. The high-frequency gain typically corresponds to a frequency region beyond the relevant range, so its role is not as critical in the model-following performance. The model following filter $W_{mf}(s)$ can be defined as follows:

$$W_{mf}(s) = \frac{s + w \cdot M}{M \cdot (s + A)} \quad \Rightarrow \quad W_{mf}^{-1}(s) = \frac{M \cdot (s + A)}{s + w \cdot M} \tag{5.14}$$

Choosing the different filter parameters often requires managing conflicting trade-offs and might require some iterations to strike a desired balance between performance and robustness. Inspired by the work of Pérez et al., 2022, this motivated the idea of using a co-design procedure, where the filter parameters are tunable parameters optimised in parallel together with the minimisation of the $\mathcal{H}_\infty$ norm in Equation 5.13.

In order to do so, the block diagram structure in Figures 5.23 and 5.24 is imposed in `MATLAB®` `Simulink`, with a complete overview depicted in Figure 5.25.

The DC gain of W_{mf}^{-1} is given by $\frac{A}{w}$, the peak gain is mainly associated with M and the bandwidth with the value w. Choosing the desired value for the DC gain is often trivial: -40 dB is often sufficient, thus, any value below that yields good steady-state tracking. Therefore, the choice of the parameter A is set by the designer and is not part of the co-design optimisation routine. Nevertheless, since the DC gain is given by $\frac{A}{w}$ and not merely A, the A parameter should be set with an idea of the lower bound for w. The w and M parameters are optimised such that the value of w is maximised and the value of M is minimised, and hence, the bandwidth of W_{mf}^{-1} is maximised and the peak value is minimised. The minimisation of the $\mathcal{H}_\infty$ norm of the model-following constraint in Figure 5.23 is set as a `HardGoal` in `MATLAB®` `systune`, while the minimisation of $i_w \to o_w$ and $i_M \to o_M$, displayed in Figure 5.24, are set as `SoftGoals` (Balas et al., 2022). Notice that the value of w and M of the soft constraints is the same as the one that integrates the model-following filter W_{mf}. Therefore, the non-smooth $\mathcal{H}_\infty$ optimization aims to maximize the bandwidth

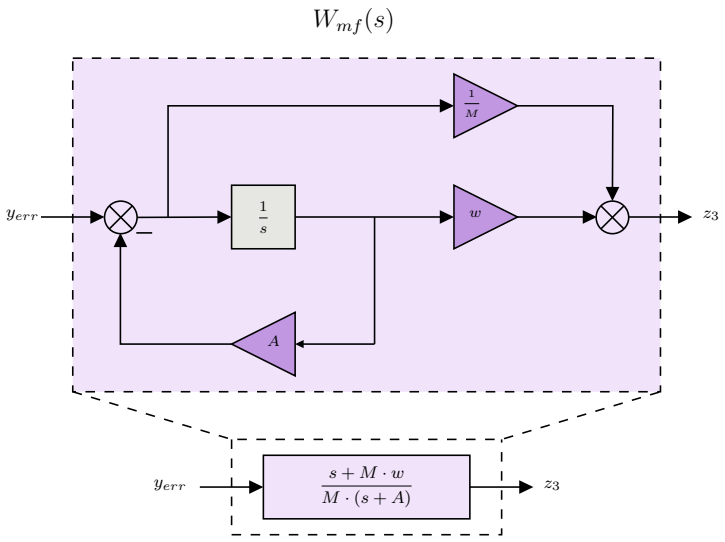

Figure 5.23: Implementation of the model-following filter $W_{mf}(s)$ in MATLAB® Simulink.

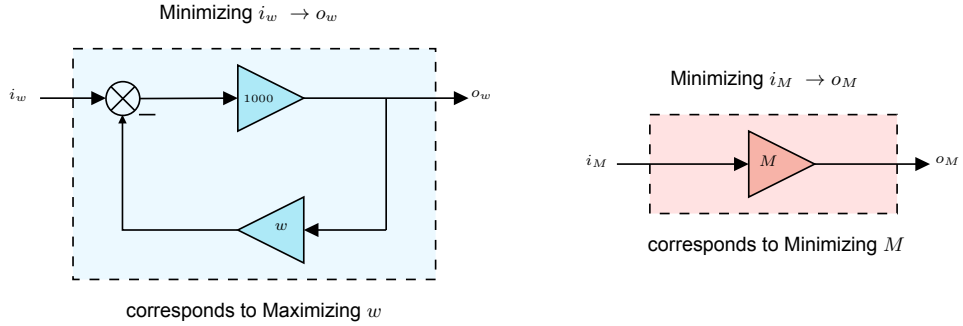

Figure 5.24: Implementation of the co-design parameters W and M in MATLAB® Simulink.

and minimize the peak value of W_{mf}^{-1}, thus making the model-following requirements tighter, while still ensuring that the transfer function $y_{cmd} \rightarrow e_{ref}$ is bounded by W_{mf}^{-1}.

Nevertheless, some additional tricks are necessary to make the co-design procedure work. Given the non-convex nature of the optimisation in MATLAB® systune, it is necessary to include bounds on the values of the model-following filter tunable parameters to limit the subset of the solution space (Pérez et al., 2022). Consequently, a lower bound is defined for the value of w and a certain range for the admissible values of M. Moreover, to make the co-design SoftGoals "competitive" with the one associated with the NCF robust stabilisation (SoftGoal with value γ) an ideal value is set for w and M, such that the minimisation of the $\mathcal{H}_\infty$ norm is as follows:

$$\|i_w \rightarrow o_w\|_\infty \leq \frac{1}{w_{ideal}} \tag{5.15}$$

$$\|i_M \rightarrow o_M\|_\infty \leq M_{ideal} \tag{5.16}$$

By implementing the SoftGoals as described above, not only w is maximized, and M minimized, but deviations from the ideal values are penalized in such a way that results in SoftGoals > 1 if $M > M_{ideal}$ and $w < w_{ideal}$. This can be seen as a trick to scale the co-design constraints and is necessary since

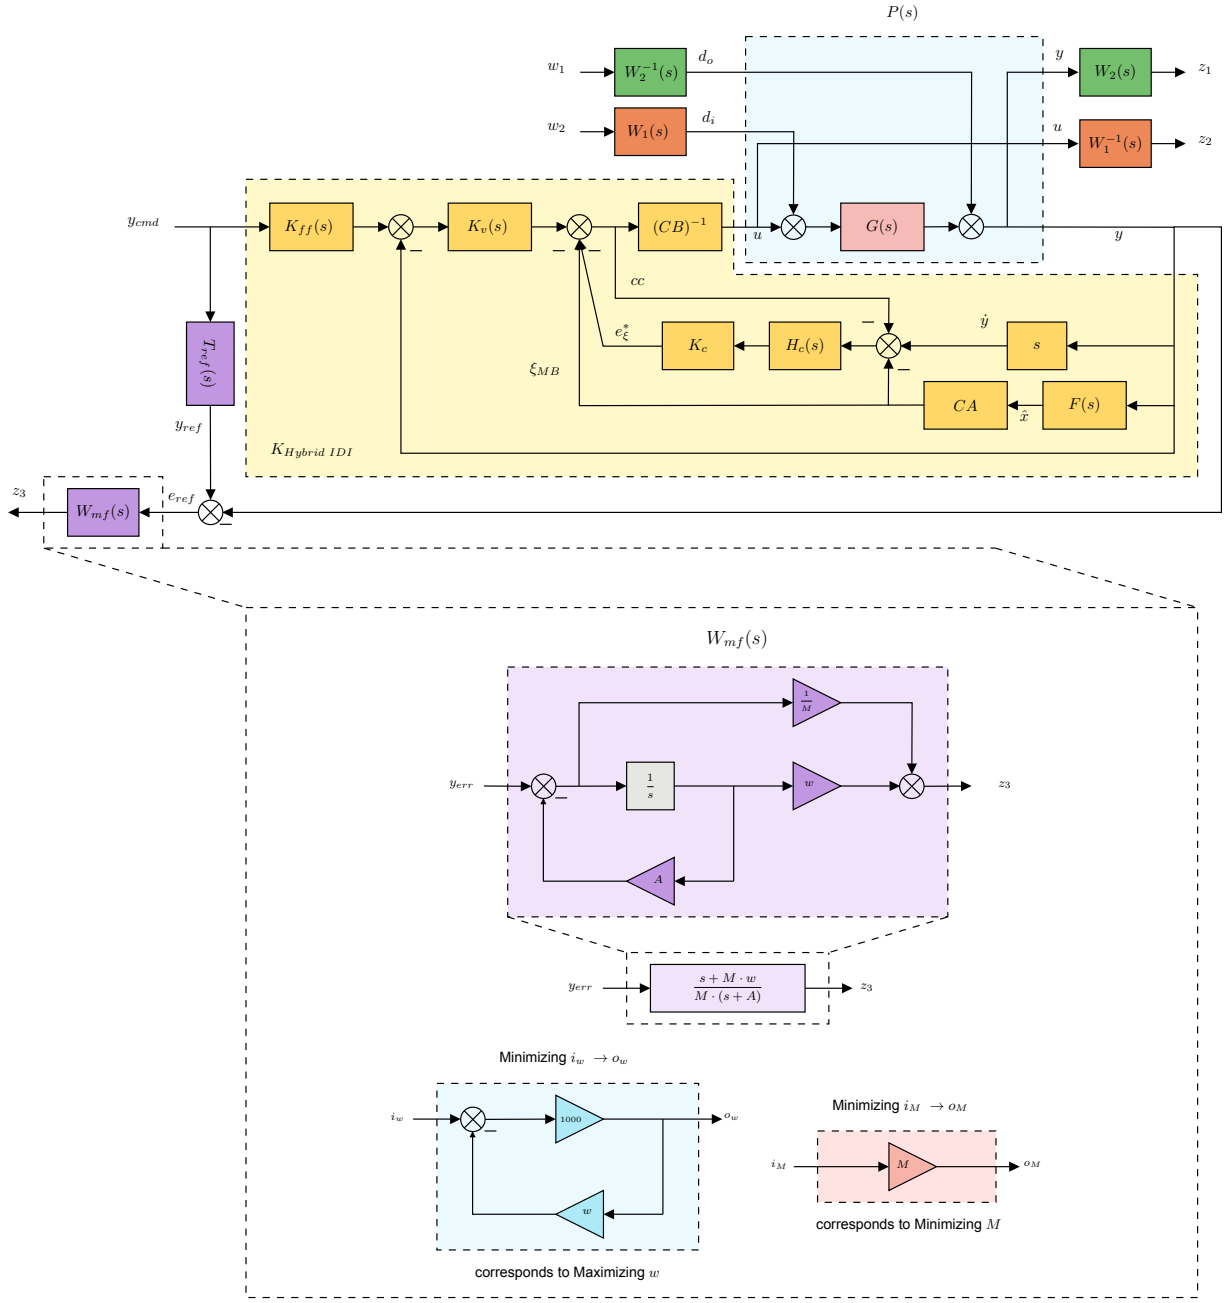

Figure 5.25: Overview of the co-design procedure to tune a Hybrid IDI using $\mathcal{H}_\infty$ LSDP and including model-following requirements.

γ, the value for the other `SoftGoal`, typically ranges from 2 to 4. In a scenario where the trade-offs between performance and robustness come into play, the value of the three `SoftGoals` (one for the NCF robust stabilization and two for the co-design) should yield similar values, as the optimization attempts to maximize the robustness of the system and hence minimize γ, but at the same time seeks to stretch the model following parameters the maximum it can for the associated robustness level. From this point on, demanding more performance via increasing the w_{ideal} and decreasing M_{ideal} will also increase γ and, thus, decrease the overall robustness of the system.

5.3.3. Remarks on Controller Structure

In Chapter 3, it was demonstrated that besides an adequate open-loop gain response, it was also necessary to have certain controller gain properties, namely, $\underline{\sigma}[K] \gg 1$ at low frequencies and $\bar{\sigma}[K] \ll K_{max}$ at high frequencies (under the assumption that K_{max} is not too large). Having $\bar{\sigma}[K] \ll K_{max}$ at high frequencies is particularly important, given that in this frequency range, $\bar{\sigma}[KS_0] \leq \bar{\sigma}[K]$, and ensuring a low $\bar{\sigma}[KS_0]$ is critical for control signal reduction and noise attenuation.

These relationships point out the fact that choosing an adequate controller structure is central to ensuring a successful design. More specifically, the inclusion of low-pass filters is of the utmost importance since it contributes to reducing the controller gain at high frequencies. However, many NDI/INDI control implementations seem to disregard this aspect. If a PI structure is considered for the virtual law, it is possible to infer that the gain of this controller at high frequency is given by:

$$\lim_{s \to +\infty} K_p + \frac{K_i}{s} = K_p \tag{5.17}$$

Consequently, a PI alone does not ensure $\bar{\sigma}[K] \ll K_{max}$, given that the value of K_p will often be significantly higher than the desired K_{max}. For that reason, the virtual law controller K_v was chosen to be a PI in series with a low-pass filter, as follows:

$$K_v(s) = \frac{w_{lp}}{s + w_{lp}} \frac{K_p s + K_i}{s} \tag{5.18}$$

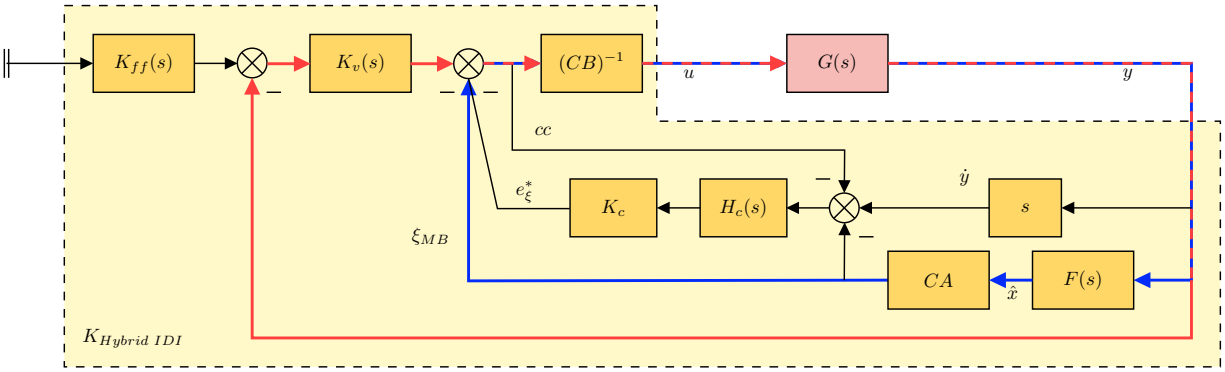

Figure 5.26: Parallel nature of the control structure of DI-based control schemes.

Nevertheless, given the parallel architecture of DI-based control schemes, this alone is not sufficient to guarantee $\bar{\sigma}[K] \ll K_{max}$. Consider, for the sake of simplicity, an MB scheme ($K_c = 0$), such that only the blue path in the inversion schematic in Figure 5.26 is relevant. As can be seen, the control signal u is both a function of the virtual control law path (in red) and of the inversion path (in blue). Therefore, even if the structure of K_v is such that for high frequencies, $K_v(+\infty) \to 0$, if the gain in the inversion path $K_{inv}(s)$ is not low in this frequency range, then $\bar{\sigma}[K]$ will also not be, given that:

$$\lim_{s \to +\infty} |K(s)| = \lim_{s \to +\infty} |-(CB)^{-1}(K_v(s) + K_{inv}(s))| \approx \lim_{s \to +\infty} |-(CB)^{-1} K_{inv}(s)| \tag{5.19}$$

where $K_{inv}(s)$ is given by the gain in the inversion path $CAF(s)$.

Notice that given the SISO nature of the problem, the discussion greatly simplifies, since $|K(s)|$ corresponds to $\bar{\sigma}[K(s)]$, and that is why Equation 5.19 makes use of the gain of the transfer function

instead of the singular values. For a MIMO case, the relationships are more complicated, but a similar conclusion is to be expected.

Accordingly, the inversion path also plays a key role in ensuring low controller gain at high frequencies. However, most of the elements of the inversion path are model-related and, therefore, fixed, with the exception of the estimator $F(s)$. This demonstrates that the designer can exploit this degree of freedom to help address the issue of $\bar{\sigma}[K]$ at high frequencies. This was not done for the present thesis, but it is highly recommended that the implications of the estimator on the controller gain be explored further in future studies. Instead, a simple estimator was used, which uses the short-period dynamics model information to estimate the angle of attack α, as described in Equation 5.1.

It is, therefore, possible to infer that the low-pass on K_v is not sufficient to guarantee low controller gain at high frequency (despite the fact that it contributes positively to the latter fact). However, this issue has limited solutions due to the typical control structure of DI-based control laws. One solution would be, for instance, to move the low-pass filter to after the virtual control sum signal, such that it affects the entire control signal path. However, in doing so, the inversion path would be distorted. As a result, the choice was made to maintain the low-pass filter in the virtual control law. Moreover, in digital design implementations, the use of anti-aliasing filters at the output of the plant actually contributes to solving the issues described above since it essentially acts as a low-pass filter in the control signal path.

5.3.4. Standard EMF (Explicit Model Following) DI control architectures

The Hybrid IDI feedforward term proposed in the present research deviates from what is commonly found in literature for NDI and INDI controllers, which normally make use of an Explicit Model Following (EMF) architecture, as shown in Figure 5.27 (Grondman et al., 2018; Kolesnikov, 2005; Miller, 2011). However, the concept of EMF heavily relies on the assumption that $\frac{y}{v} = \frac{1}{s}$, as it is demonstrated below. For a moment, consider the following closed-loop relationship $\frac{y}{v} = \eta(s)$. It can be proven that:

$$\frac{y(s)}{y_{cmd}(s)} = T_{ref}(s) \cdot \frac{\eta(s) \cdot (s + K_v(s))}{1 + K_v(s) \cdot \eta(s)} \tag{5.20}$$

which implies:

$$\frac{y(s)}{y_{cmd}(s)} = T_{ref}(s) \Rightarrow \frac{\eta(s) \cdot (s + K_v(s))}{1 + K_v(s) \cdot \eta(s)} = 1 \tag{5.21}$$

This equality only holds if:

$$\eta(s) \cdot (s + K_v(s)) = 1 + K_v(s) \cdot \eta(s) \Leftrightarrow \eta(s) = \frac{1}{s} \tag{5.22}$$

which means that model matching is only guaranteed under ideal inversion, in which case, ideal model following performance is independent of the virtual control law. However, inversion distortions are introduced either by the estimator F, by the fact that the actuator model is not taken into consideration in the MB inversion, or due to the synchronisation LPF H_c (Pollack, 2024). Therefore, either a purely model-based setup (K_c = 0), a purely sensor-based one (K_c = 1) or a hybrid one will always have elements in its inversion path contributing to inversion distortion.

Nevertheless, the work of Pollack, 2024 derived explicitly how the error signal transfer function $\frac{e_{ref}(s)}{y_{cmd}(s)}$ can be made arbitrarily zero by using the inversion compensation, increasing the open-loop gain by tuning

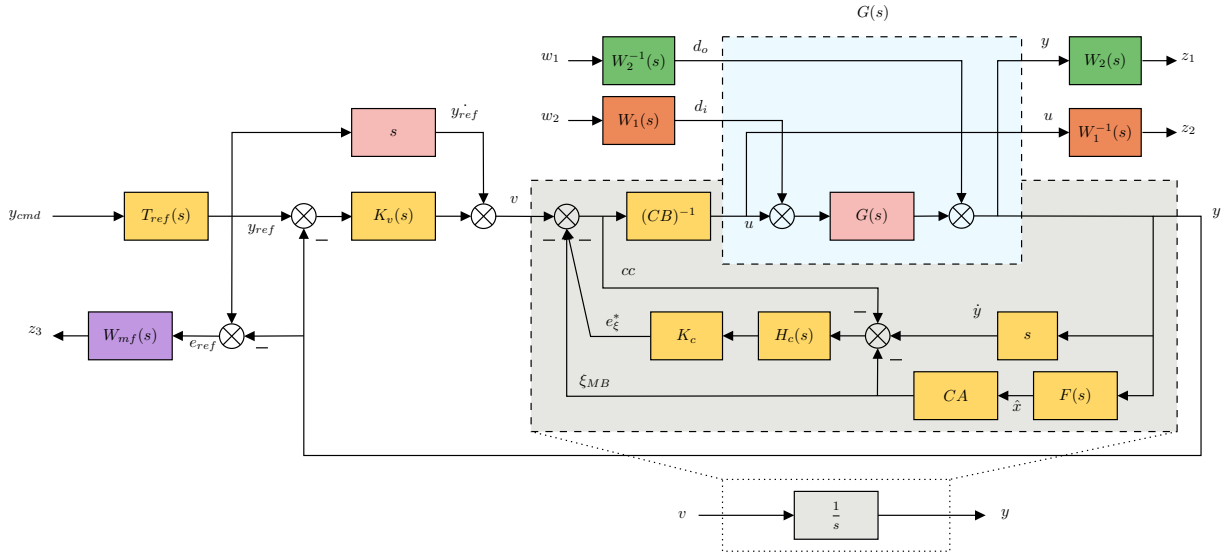

Figure 5.27: Explicit Model Following DI control architecture.

the virtual control law or by a combination of both. This dual effect demonstrates the multi-loop robustness functionality of INDI control laws and justifies the success of Explicit Model-Following architectures.

Part of the motivation for not using an EMF setup is related to the fact that it is desired to stay more in line with the common 2DoF $\mathcal{H}_\infty$ LSDP approach. Moreover, due to the uncertainty faced about the possibility of combining a DI control structure with the $\mathcal{H}_\infty$ LSDP, the addition of the feedforward element makes it possible to decouple the robustness of the system from the model following requirements. The designer can first consider only robustness (minimization of the $\mathcal{H}_\infty$ norm from $[w_1\ w_2]^T \rightarrow [z_1\ z_2]^T$) and tune K_v, K_c and H_c. Then, in the second step, tune K_{ff} for the model following the performance. As such, more degrees of freedom are given to the controller to address robustness, and since the entire procedure is built upon the hypothesis that the feedback part of $K_{H-IDI} \rightarrow K_s$, then, it is of the utmost importance to guarantee that these degrees of freedom are present. Nevertheless, when the method is implemented in a single stage, as proposed above, it also plays a role in achieving the desired model following performance since the feedback elements appear in the forward path of the controller. However, on a conceptual level, if to achieve a certain model following performance, the robustness is largely compromised, the designer always has the option to increase the order of the feedforward controller **(which only influences performance, and not robustness)**. This makes the tuning of the feedback elements K_v, K_c and H_c less constrained and, consequently, it is expected that robustness improves, as these parameters can more freely be configured such that $K_{H-IDI} \rightarrow K_s$.

Consequently, the addition of a feedforward element for model performance stays more within the realm of robust control applications. In addition, given the novel nature of the proposed procedure, it was considered that the approach which yields more degrees of freedom to achieve a successful design was more appropriate for this initial study. It is, however, highly recommended that combinations of EMF architectures be combined with the procedure proposed in this work in future studies.

5.4. Framework for Tuning DI Controllers Using the LSDP

As a result of the discussion and deep insight obtained conducted in the previous sections, the current section provides a framework to tune Dynamic Inversion Controllers (Model-Based, Incremental Sensor-Based and Hybrid schematics) using the $\mathcal{H}_\infty$ Loop Shaping Design Procedure.

The first step of the $\mathcal{H}_\infty$ LSDP is to translate all the design requirements (often time-domain requirements) to frequency-domain ones. This first step involves deciding on the LOES for flying qualities (in this case, T_{ref}) according to handling qualities guidelines. The designer can then set the desired crossover frequency based both on the closed-loop bandwidth of T_{ref} and on the limitations imposed by RHP poles and zeros (Skogestad & Postlethwaite, 2005).

After deciding on the crossover frequency for G_s, the general guidelines of LSDP should be followed to compute W_1 and W_2, i.e., W_1 should comprise of PI controllers to boost the low-frequency gain (integrator) and reduce the slope around cross-over (through the PI zero at $\frac{K_i}{K_p}$), and W_2 should contain low-pass filters for gain attenuation at high frequencies. It is important to highlight that W_1 and W_2 will not incorporate the final controller, and so their complexity can be increased without increasing the complexity of the final controller, whose structure is fixed to one of the selected Dynamic Inversion strategies. On one hand, this allows for a better shape of the open-loop plant, but on the other hand, it is important to keep in mind that the procedure builds on the hypothesis that $K_{struct} \rightarrow K_s$, which means that unnecessary inflating the order of the filters can lead to a design which is incompatible with the tunable elements in the structured DI controller.

After shaping the plant, the theoretical γ_{min} associated with the shaped plant $G_s = W_2 G W_1$ can be computed. According to literature, an adequate value for γ_{min} lies between 1 and 4 (Hyde, 1995). If an acceptable value for γ is obtained, then, under the assumption that K_{struct} is compatible with K_s, a non-smooth $\mathcal{H}_\infty$ minimization together with the schematic presented in Figure 5.15 can be leveraged to optimise the different controller elements in `systune`. As an initial verification, the designer should only use a 1DoF controller (considering only robustness) to check whether the structure of the DI is compatible with the standard full-order solution of $\mathcal{H}_\infty$ LSDP given by K_s. If, for instance, γ_{min} is 2.5, and the result of the optimization gives 2.7, then it can be said that K_{DI} is close to K_s (in general, the γ produced by the structured DI controller will be higher than the sub-optimal γ associated with the full-order solution K_s). If an unreasonable value for γ is obtained, then either the filters W_1 and W_2 need to be redesigned or more flexibility needs to be added to the DI controller, as, for instance, increasing the order of the virtual controller K_v or taking advantage of a hybrid inversion structure.

Under the assumption that the 1DoF design achieved satisfactory results, the 2DoF scheme can be used to design a DI controller which is not only nominally and robustly stable but also capable of achieving nominal and robust performance according to handling qualities requirements. This step requires that the designer set the desired bandwidth w_{ideal} and peak value M_{ideal} of W_{mf}^{-1} and employ the co-design procedure described in Section 5.3.2.

The above discussion can be summarised as follows:

Benchmark for tuning DI-based controllers using the $\mathcal{H}_\infty$ LSDP

1. Choose the reference model T_{ref} based on handling qualities requirements;
2. Choose the desired crossover frequency w_c according to T_{ref} and eventual RHP poles and zeros of the open-loop plant G;
3. Select W_1 and W_2 to achieve a desired open-loop response $G_s = W_2 G W_1$[a];
4. Choose the DI structure (MB, SB or HB)[b] and employ the structure in Figure 5.15 to tune the controller using the 1DoF $\mathcal{H}_\infty$ LSDP[c].
5. Implement the 2DoF architecture in Figure 5.19 and employ the co-design procedure to optimize the model-following filter W_{mf} parameters in parallel with the rest of the optimization[d].

---

[a] At the end of this step, compute γ_{opt} associated with the shaped plant $G_s = W_2 G W_1$. It should be below 4 according to design guidelines (Hyde, 1995). If not, redesign W_1 and W_2.

[b] It is expected that the Hybrid-based architecture yields better results, given the additional degrees of freedom to adjust the open-loop response.

[c] For a successful design, the obtained γ should be close to γ_{opt}, as $\gamma \to \gamma_{opt}$ results in $K_{DI/IDI} \to W_1 K_\infty W_2$.

[d] The designer sets a fixed A, related to the low-frequency gain of W_{mf}^{-1}, and sets an ideal bandwidth w_{ideal} and a maximum peak value M_{ideal}. The co-design procedure returns an optimised value for w and M.

6

X-29 Pitch Rate Controller Design

The previous chapter presented a methodology to combine $\mathcal{H}_\infty$ Loop-Shaping with a Dynamic Inversion schematic. In this chapter, this methodology will be used to develop a pitch-rate controller for NASA's X-29 experimental aircraft. The X-29 constitutes a particularly interesting case study since its highly unstable configuration poses a great challenge to the design of flight control laws. For this reason, the X-29 aircraft was one of the examples provided by Gunter Stein in his "Respect the Unstable" Bode Prize lecture (Stein, 2003). Thus, if the proposed procedure enables the designer to systematically address robustness and performance trade-offs for a highly unstable plant with limited stability margins, then this demonstrates the potential applicability of the developed method for tuning DI controllers with guaranteed robustness margins.

6.1. X-29 Model

The X-29 is a former NASA research aircraft featuring a forward-swept wing design, developed to explore potential aerodynamic benefits made possible by advanced composite materials engineered for aerodynamic efficiency. The wings are swept forward instead of aft, and a large canard surface was positioned close to the main wing to leverage beneficial aerodynamic interactions (Webster & Purifoy, 1991).

Since the centre of pressure (CoP) is located ahead of its centre of gravity (CoG), the airplane is open-loop unstable. A positive pitch motion increases the generated lift, which acts through the CoP – CoG offset to produce a positive pitch moment, further increasing the angle of attack and leading to a divergent (unstable) pitch attitude. While designed to be modestly unstable in the transonic and supersonic flight regimes, the X-29 exhibits a more accentuated instability at subsonic speeds. This is the result of a basic aerodynamic phenomenon that moves the centre of pressure of a lifting surface sharply aft as speeds go from sub- to supersonic. The associated short-period linearised dynamics of the X-29 exhibit two roots, one stable and one unstable, with the unstable real pole reaching values as large as +6 rad/s (Stein, 2003).

6.1.1. Open-Loop Analysis

The flight condition (M = 0.9 and altitude = 8000 ft) chosen for the controller's design is extremely challenging due to the plane's violent instability and the aforementioned limitations. The linearised state-space models for different flight conditions can be retrieved from Bosworth, 1992. However, given that there are 3 different longitudinal inputs - a canard, a symmetric flap, and a strake flap (see Figure 6.1) - the

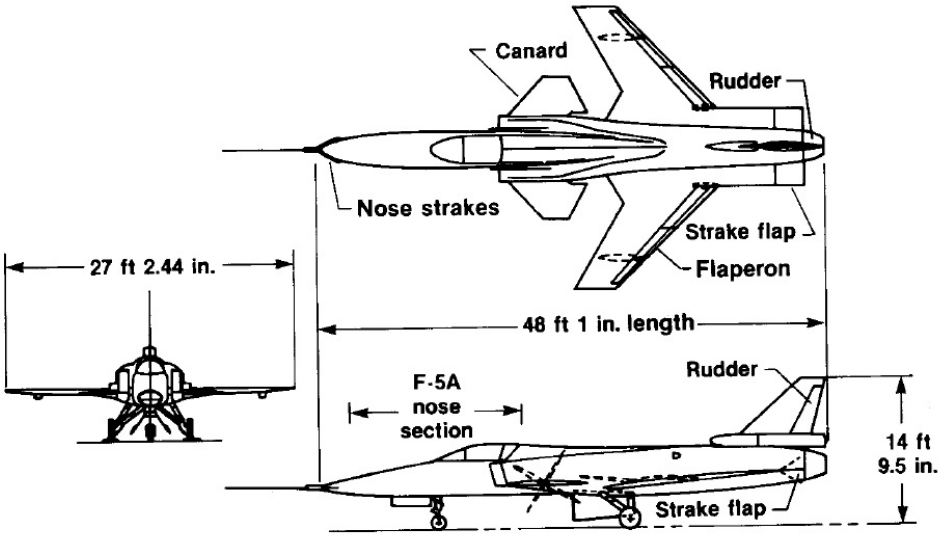

Figure 6.1: Three-view drawing of the X-29A airplane depicting the various control surfaces. Image retrieved from Kehoe et al., 1990.

Figure 6.2: Grumman X-29 flight tests [Credit NASA].

control allocation problem adds an extra degree of complexity to the problem. Therefore, the short-period dynamics model from Lavretsky and Wise, 2024 was used, where a linear combination of the 3 inputs is used to achieve a single virtual input δ_e, which makes the nature of the problem SISO:

$$\begin{bmatrix} \dot{\alpha} \\ \dot{q} \end{bmatrix} = \begin{bmatrix} \frac{Z_\alpha}{V_0} & 1 + \frac{Z_q}{V_0} \\ M_\alpha & M_q \end{bmatrix} \begin{bmatrix} \alpha \\ q \end{bmatrix} + \begin{bmatrix} \frac{Z_{\delta_e}}{V_0} \\ M_{\delta_e} \end{bmatrix} \delta_e = \begin{bmatrix} -2.241 & 0.9897 \\ 44.74 & -0.9024 \end{bmatrix} \begin{bmatrix} \alpha \\ q \end{bmatrix} + \begin{bmatrix} -0.2331 \\ -45.93 \end{bmatrix} \delta_e \tag{6.1}$$

In order to model the actuator dynamics, a second-order transfer function was used (Lavretsky & Wise, 2024):

$$G_{act}(s) = \frac{\delta_e(s)}{\delta_{e_{cmd}}(s)} = \frac{w_n^2}{s^2 + 2\xi w_n s + w_n^2} = \frac{70^2}{s^2 + (2)(0.7)(70)s + 70^2} \tag{6.2}$$

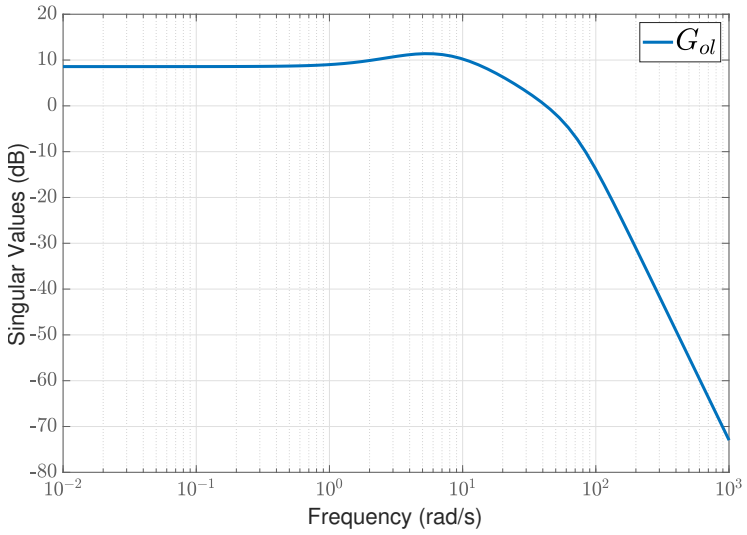

Figure 6.3: Open-loop response of the X-29 short period dynamics.

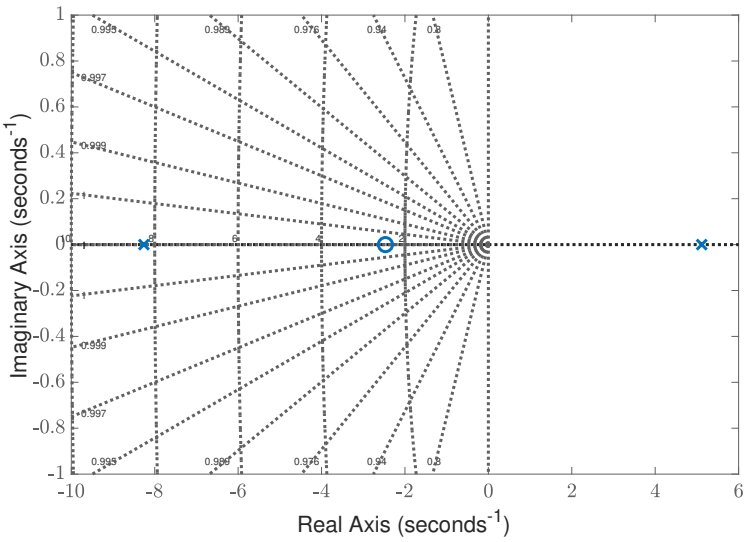

Figure 6.4: Pole-zero map of the X-29 short period dynamics.

As can be observed in Figure 6.4, the open-loop system has an unstable pole at $p = 5.12$ rad/s. This unstable RHP pole imposes a lower bound on the allowed bandwidth, with an approximate bound of $w_c > 2p$ (Skogestad & Postlethwaite, 2005). Therefore, the crossover frequency of the shaped plant G_s should be made to be at least $w_c > 2p = 10.24$ rad/s.

6.2. Continuous-Time Design

In this section, the procedure outlined in Chapter 5 will be used to optimise the gains of a Hybrid IDI continuous pitch-rate controller for the X-29. First, an initial verification of the procedure will be conducted without including any model-following requirements, and, upon the verification of its success, a 2DoF Hybrid IDI controller will be tuned, including model-following requirements using the co-design procedure described in Section 5.3.2.

6.2.1. Initial Verification

In order to boost the low-frequency gain of the open-loop response, have a -20 dB slope around crossover, and have the crossover frequency at the desired value of 10.2 rad/s, W_1 was set as a PI, while, for simplicity, W_2 was set to 1.

$$W_1 = K_p + \frac{K_i}{s} = \frac{K_p s + K_i}{s} = \frac{K_p(s + K_i/K_p)}{s} = \frac{0.24183(s + 8.166)}{s} \tag{6.3}$$

The obtained shaped plant, given by G_s = $G_{ol}W_1$ fulfils the desired shaping, as can be observed in Figure 6.5. In order to assess if the following shaping of the plant is compatible with a robust design, a suboptimal γ_{th} was computed with a tolerance of 1e-3 from the optimal γ_{min} associated with the shaped plant G_s. The obtained value for γ_{th} is equal to 2.3181, which corroborates that the shaped plant is compatible with a robust design since it is below 4 as suggested in the $\mathcal{H}_\infty$ LSDP guidelines (Hyde, 1995).

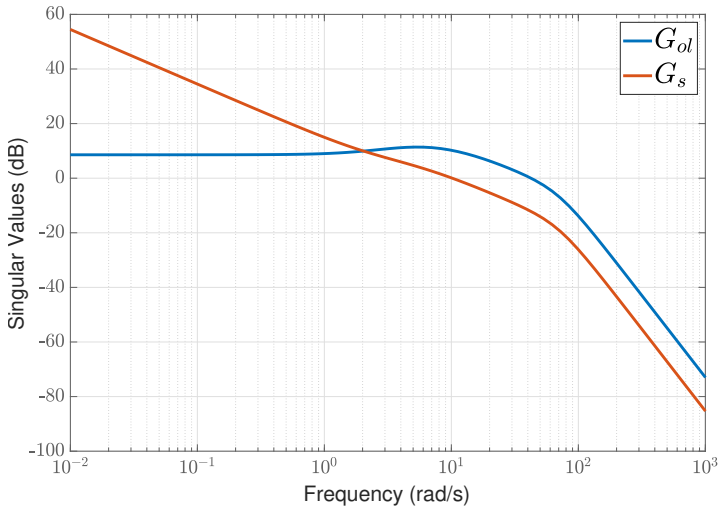

Figure 6.5: Shaping the open-loop plant G_{ol} with W_1 to achieve desired crossover frequency requirements.

As an initial verification step, the procedure is implemented without any model following requirements in order to assess if the structure of the Hybrid IDI is compatible with the standard $\mathcal{H}_\infty$ Loop-Shaping controller K_s. This boils down to the implementation described in Figure 5.15, and the results are presented in Table 6.1.

Table 6.1: Optimization gains of the Hybrid IDI using the $\mathcal{H}_\infty$ LSDP.

	Optimization		Virtual Control			Inversion Loop	
Parameter	γ_{th}	γ	K_p	K_i	w_{lp}	K_c	w_{H_c}
Value	2.3181	2.3622	12.7882	36.2545	52.12	0.3060	80.62

As can be observed, the obtained γ is similar to γ_{th}, which attests to the success of the optimization. If we first inspect the Bode plots of the full-order $\mathcal{H}_\infty$ LS controller and the tuned Hybrid IDI, it becomes evident that they are virtually identical, as demonstrated in Figure 6.6. Moreover, examining the obtained open-loop gain at the plant input in Figure 6.7 confirms that their open-loop responses are also nearly identical. Therefore, it is demonstrated that under the hypothesis that the $K_{H\text{-}IDI}$ has a compatible structure to that of K_s, the non-convex $\mathcal{H}_\infty$ synthesis displayed in Figure 5.15 results in $\gamma \to \gamma_{th}$ and $K_{H\text{-}IDI} \to K_s$.

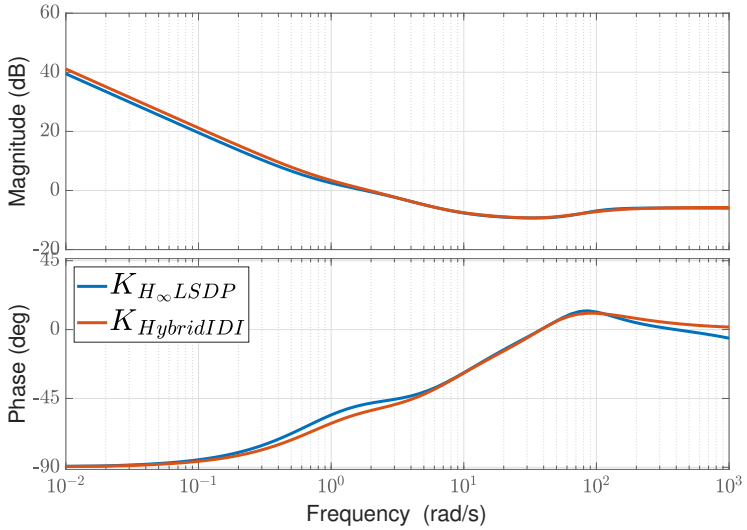

Figure 6.6: Gain and Phase comparison of standard full-order $\mathcal{H}_\infty$ controller and tuned Hybrid IDI.

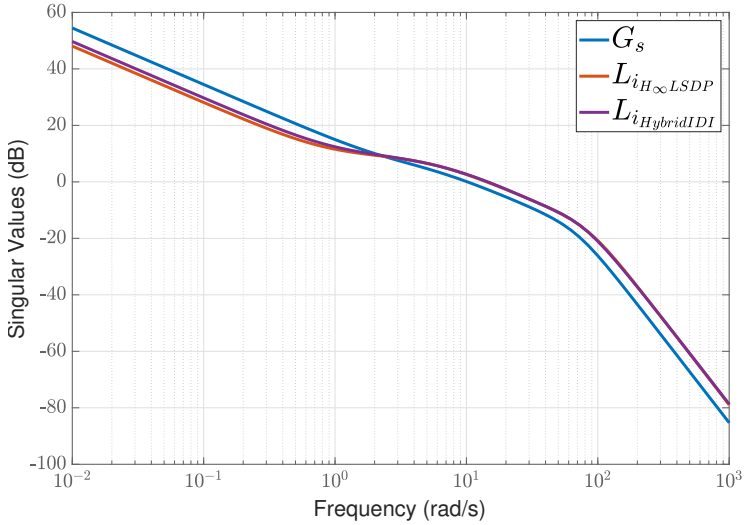

Figure 6.7: Open-Loop gain at plant input of standard full-order $\mathcal{H}_\infty$ controller and of the tuned Hybrid IDI.

6.2.2. 2DoF Continuous-Time Design

After verifying the compatibility of the procedure, the next step is to address the model following requirements. Given that the crossover frequency was set to 10.2 rad/s and that the reference model T_{ref} should have a compatible bandwidth (Skogestad & Postlethwaite, 2005), the following points in the CAP and Gibson criteria were chosen:

$$T_{ref} = \frac{15.2(s + 5.527)}{s^2 + (2)(0.6)(9.17)s + (9.17)^2} \tag{6.4}$$

The procedure outlined in Section 5.3.2 is then implemented to include the model following requirements via a co-design procedure, and the control architecture used is the one depicted in Figure 5.25. The results of the optimisation are presented in Tables 6.2 and 6.3.

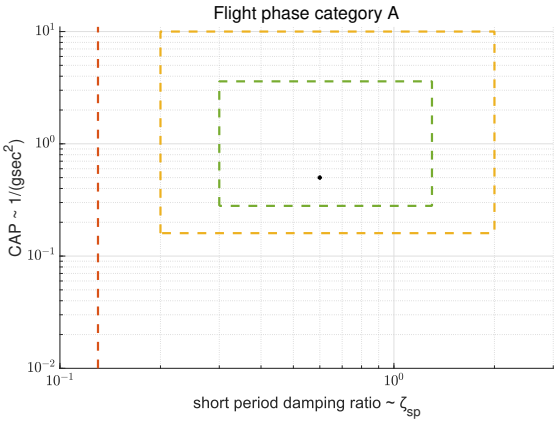

Figure 6.8: Chosen point for the CAP criteria that adheres to level 1 handling qualities.

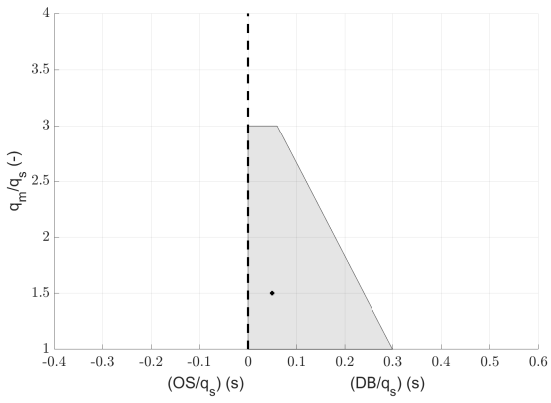

Figure 6.9: Chosen point for the Gibson criteria that adheres to level 1 handling qualities.

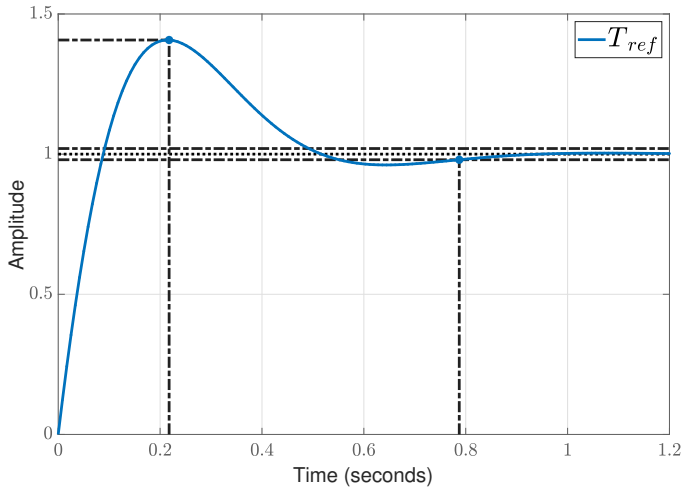

Figure 6.10: Reference model T_{ref} chosen according to the CAP and Gibson criteria.

Table 6.2: Filters setup and optimisation result of the Tuning Goals.

	Shaping Filters		NCF Rob. Soft Goal		Co-Design Soft Goals		MF Hard Goal
Parameter	W_1	W_2	γ_{th}	γ	w_e	M_e	W_{mf}
Value	$\frac{0.24183(s+8.166)}{s}$	1	2.3181	2.7452	2.7452	2.7452	0.9999

Table 6.3: Optimisation gains of the Hybrid IDI using the $\mathcal{H}_\infty$ LSDP.

	Virtual Control			Inversion Loop		Feedforward	MF Filter W_{mf}		
Parameter	K_p	K_i	w_{lp}	K_c	w_{H_c}	K_{ff}	A_e	w_e	M_e
Value	9.6493	39.3402	45	0.3963	67.24	$\frac{4.9898(s+15.8)}{s+78.54}$	0.3	54.6394	0.9608

From the optimisation results, it can be observed that the deterioration of the γ-value is higher than in the 1DoF case presented in Table 6.1. This is expected given the imposition of the model-following requirements. Nevertheless, the obtained γ is still reasonably close to γ_{th}, which hints at a successful and robust design.

The co-design `TuningGoals` have the same value as the γ, which corroborates that these constraints were adequately scaled to be competitive with the NCF robustification one (recall the discussion in Section 5.3.2). In order to achieve this scaling, the desired bandwidth value was set to $w_{e_{ideal}} = 150$ and the desired peak value was set to $M_{e_{ideal}} = 0.35$. The optimised values for w_e and M_e will, of course, never reach these desired values and the amount by which it violates this (soft) constraint is given precisely by the value of the co-design `TuningGoals`. If the designer were to increase $w_{e_{ideal}}$ and decrease $M_{e_{ideal}}$ even more, the values for the three soft constraints would increase, making γ larger, but the optimized value of w_e would increase and M_e would decrease. In other words, the designer would be demanding higher model-following performance from the system, which would come at the cost of robustness. Defining $w_{e_{ideal}}$ and $M_{e_{ideal}}$ (and even A_e) might require some iterations to strike the desired balance between robustness and performance. It is important to stress that inadequately scaling the co-design goals by setting a low $w_{e_{ideal}}$ and a high $M_{e_{ideal}}$ will lead to poor model-following performance, since from an optimization perspective, minimization of γ, and thus, robustness, is more easily achieved when the model-following requirements are less demanding.

The low-pass pole w_{lp} of the virtual control law K_v was manually set to 45 rad/s in an ad hoc manner, as this frequency lies above the crossover frequency but well within the controller bandwidth. This parameter was selected a priori, as it simplifies the optimisation process by reducing the number of parameters to be tuned. The PI gains are relatively high, which is justified by the severe instability of the open-loop system, necessitating large control signals for stabilisation. The optimisation takes full advantage of the Hybrid inversion structure, as can be seen by the intermediate value for the K_c gain. With respect to the feedforward element, the condition $\tau_{lead} > \tau_{lag}$ indicates that it speeds up the system response to satisfy the model-following requirements.

In order to get a more complete picture of the performance of the tuned controller, the famous “gang-of-six” closed-loop relationships, the open-loop response and how it compares to the desired shaped plant, the model-following performance, and the classical and disk-based margins will be analysed.

From the input and output broken-loop response (L_i and L_o) in Figure 6.11, it can be observed that they have a similar shape to the shaped plant G_s. This is a typical property of the $\mathcal{H}_\infty$ LSDP and once again hints at the success of the design procedure to tune a Hybrid IDI using the $\mathcal{H}_\infty$ LSDP. The cross-over frequency $w_c = 12$ rad/s is at a slightly higher frequency than the one set by G_s, which was $w_{c_{des}} = 10.2$ rad/s, but it is still at a quite reasonable value. All the CL transfer functions exhibit a desirable response, with the exception of KS_o. As described in Chapter 3, at low frequencies, KS_o depends only on the open-loop plant G. At high frequencies, on the other hand, it depends mostly on K, and it is required to have low controller gain, such that $\sigma[KS_o] \leq \sigma[K] \ll 1$. This motivated the inclusion of a low-pass filter in $K_v(s)$, but due to the parallel nature of the Hybrid IDI architecture, this alone is not sufficient, as explained in Section 5.3.3. This is the reason why KS_o does not roll off at high frequencies. When a more complete approach to control design is taken, with the inclusion of digital effects, the inclusion of an anti-aliasing filter will solve this issue, as it is shown in the following Section 6.3.

In terms of step-tracking performance, the pitch rate response to a step input closely matches the reference model T_{ref} response, as can be observed in Figure 6.12. The model following error is tightly bounded by W_{mf}^{-1}, as can be seen in Figure 6.12. This is expected, given the model-following hard constraint value of 0.9999. Given that T_{ref} was chosen based on level 1 handling qualities requirements, the fact that the pitch-rate response closely matches T_{ref} suggests that predicted level 1 handling qualities are achieved. However, the present work did not go beyond this analysis in terms of predicted handling qualities since it is not the primary scope of the research. It is therefore highly recommended that this

aspect is further studied in future research.

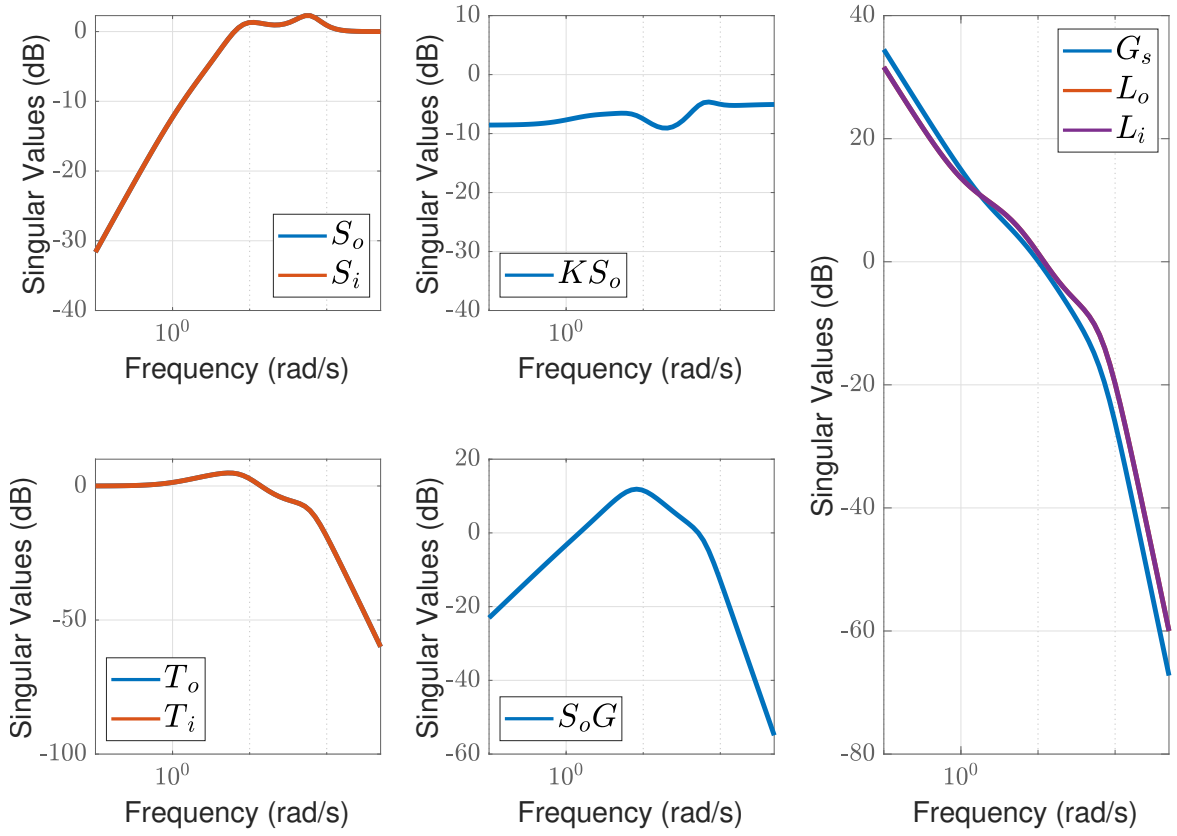

Figure 6.11: Gang-of-six and broken-loop results for the continuous-time Hybrid IDI controller.

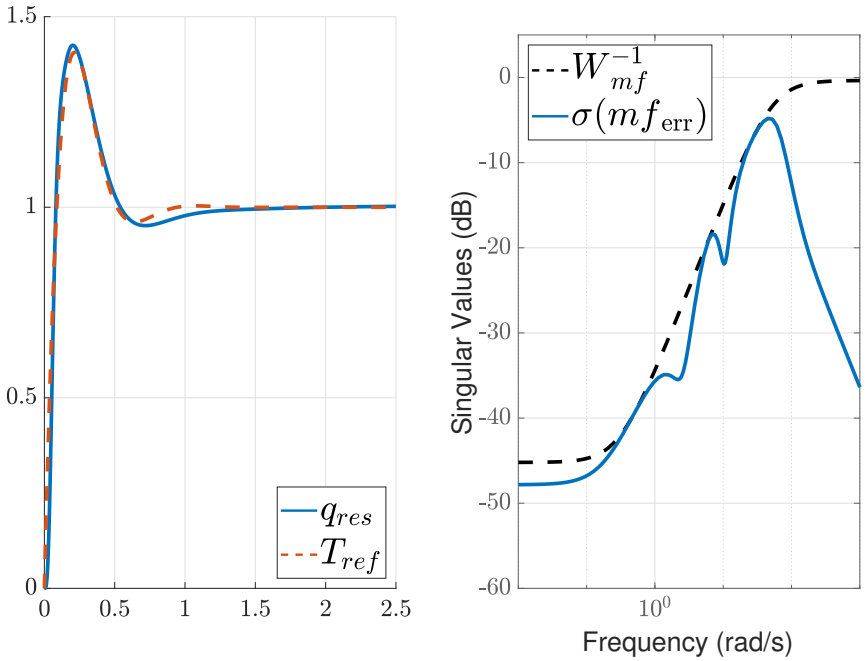

Figure 6.12: Step-tracking performance and model following error for the continuous-time tuned Hybrid IDI controller.

The classical gain and phase margins adhere to the golden rule for certification of a minimum Gain Margin of $\pm$ 6dB and minimum Phase Margin of $\pm$ 45° (Dobos-Bubno & Hartsook, 1977; Seiler et al., 2020). The disk-based margins also exhibit healthy values.

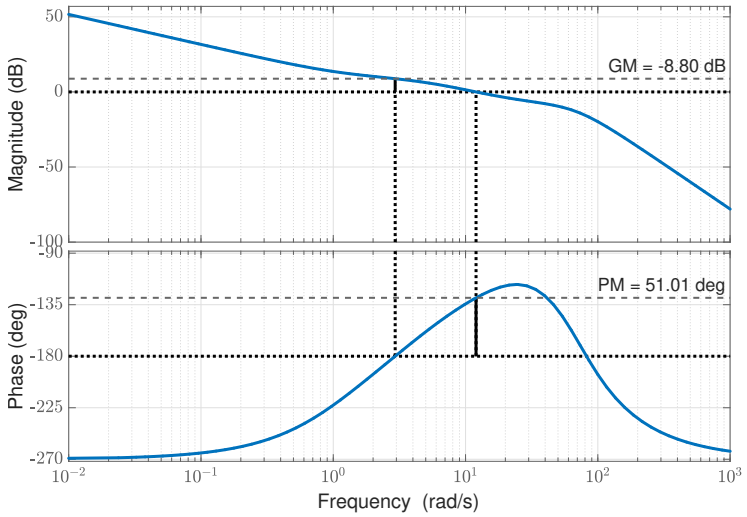

Figure 6.13: Classical margins for the continuous-time tuned Hybrid IDI controller.

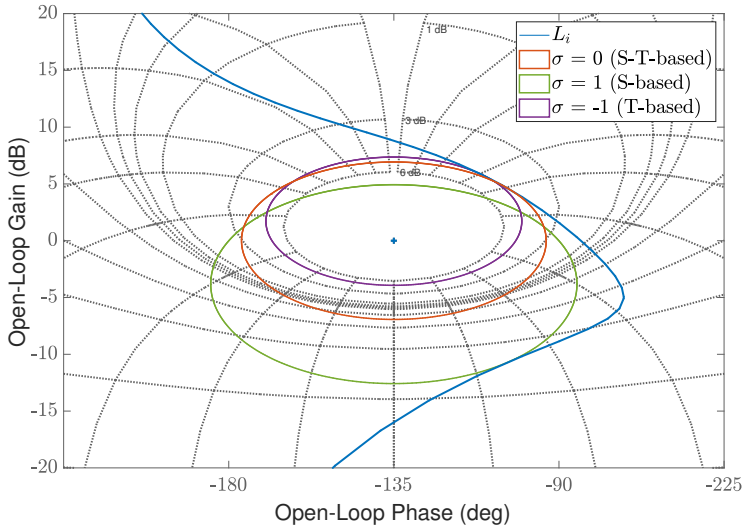

Figure 6.14: Nichols plot and Disk-based margins for the continuous-time tuned Hybrid IDI controller.

Table 6.4: Classical and Disk-based margins for the continuous-time Hybrid IDI design.

	Classical	Disk-based		
		S-T ($\sigma = 0$)	S ($\sigma = 1$)	T ($\sigma = -1$)
GM (dB)	-8.80	6.93	12.58	3.93
PM (°)	51.01	41.52	44.98	33.22

The discussion and the results presented above demonstrate a successful and robust Hybrid IDI controller. However, it will be demonstrated in the following section that when digital effects are included, the continuous-time results are overly optimistic.

6.3. Digital Design

In the previous section, it was demonstrated that a continuous-time design of the pitch-rate controller yielded acceptable results, given the system's limitations. However, control laws are usually implemented in digital form on modern aircraft. These digital, or discrete-time, controllers include discretisation effects that are not considered during a continuous-time design (Stevens et al., 2015).

The simplest way to approach digital control design is to transform a designed continuous-time controller to a discrete-time controller using continuous-to-discrete transformation methods, such as the bilinear transformation (BLT) or the matched pole-zero technique. This continuous controller redesign approach has the advantage of allowing the sampling period T to be chosen only after designing the continuous controller. However, controller discretisation methods, such as those based on the bilinear transform (BLT), are approximations. As a result, the sampling period T must be sufficiently small to ensure that the digital controller closely replicates the performance of the original continuous design (Åström & Wittenmark, 1997).

Another approach is to take into consideration the properties of the sampling process and also the computational delays by modifying the continuous-time controller to include continuous-time approximations of these digital effects. Given that in this procedure, the effects of discretisation are considered (to a certain extent), the discretisation of the modified continuous controllers results in a digital control with improved performance. This is the approach used throughout this section (Stevens et al., 2015).

In order to simulate a digital controller, the scheme shown in Figure 6.15 is used. There are multiple elements used to simulate the interaction between a continuous-time plant and a discrete-time controller, such as zero-order hold (ZOH), computational delays, and analogue-to-digital (A/D) converters. The idea behind the modified continuous-time controller design is to model these effects using continuous-time transfer function approximations.

The Zero-Order Hold is a D/A hold device, required to reconstruct the commanded signal u_{cmd} from a sampled u_{cmd_k}, and its effect, together with the sampler, can be approximated by:

$$G_{0s}(s) = \frac{1 - e^{-sT}}{sT} \tag{6.5}$$

Anti-aliasing filters are introduced using, usually, analogue circuitry in order to prevent aliasing effects of high-frequency measurement noise. These filters are introduced after the measuring devices and before the samplers, and take the form of simple low-pass filters given by:

$$H_a = \frac{a}{s + a} \tag{6.6}$$

The delay associated with a computational time of Δ has the following transfer function:

$$G_{comp}(s) = e^{-s\Delta}, \tag{6.7}$$

In order to design a continuous controller which takes into consideration these multiple discretisation and anti-aliasing effects, these are re-imagined as being part of the plant to be controlled. However, the approximations of the ZOH and of the computational delay are not rational; thus, Padé approximants of $e^{-s\Delta}$ and e^{-sT} are used to approximate these effects via rational transfer functions. The higher the order

of the transfer function used, the better the approximation. For this reason, higher order transfer functions, as opposed to the typically first order used in many applications, of the computational delay and the ZOH plus sampler, respectively, were used (Stevens et al., 2015):

$$G_{comp}(s) = \frac{(s\Delta)^2/20 - 2s\Delta/5 + 1}{(s\Delta)^3/60 + 3(s\Delta)^2/20 + 3s\Delta/5 + 1} \tag{6.8}$$

$$G_{0s}(s) = \frac{-(sT)^4/840 + 23(sT)^2/840 - sT/14 + 1}{(sT)^3/120 + (sT)^2/14 + 3sT/7 + 1} \tag{6.9}$$

Therefore, by including all of these effects directly into our continuous-time design, in the final step, when the tuned continuous-time controller is discretised using the BLT method (also referred to as Tustin's method), the final digital controller will have a better performance. It is important to notice that the anti-aliasing filter is usually implemented using analogue circuitry, thus making it part of the plant, while $G_{comp}(s)$ is not explicitly implemented since it is a model of the computation delay and $G_{0s}(s)$ is implemented by the ZOH and the sampler.

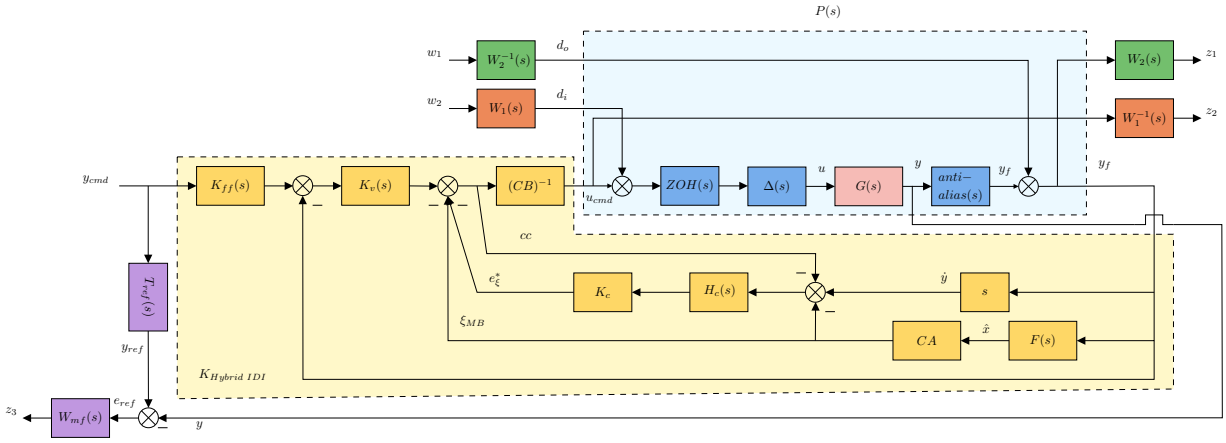

Figure 6.15: 2DoF $\mathcal{H}_\infty$ Loop-Shaping Design Procedure for Digital Hybrid IDI controller structure.

6.3.1. "Respect the Unstable": A Cruel Lesson on the Fundamental Limitations of Control Systems and Conservation of Dirt

In his Bode Prize lecture, Gunter Stein discussed the importance of the Bode Integrals in understanding the fundamental limitations associated with the control task. The integral in Equation 6.10 applies to stable plants, while the one in Equation 6.11 applies to unstable plants. These integrals can be interpreted as conservation laws. They explicitly indicate that a specific quantity, the integrated value of the logarithm of the sensitivity function's magnitude, remains constant under the influence of feedback. These relationships apply **to every controller, regardless of how it was designed**. Enhancing sensitivity in one frequency range comes at the cost of sensitivity deterioration in another frequency range, and the price to be paid is higher in case the plant is open-loop unstable (Stein, 2003).

$$\int_0^\infty \ln |S(j\omega)|\, d\omega = 0 \tag{6.10}$$

$$\int_0^\infty \ln |S(j\omega)|\, d\omega = \pi \cdot \sum_{i=1}^{N_p} \mathrm{Re}(p_i) \tag{6.11}$$

Stein proposes that this quantity being conserved, the integrated log of sensitivity magnitude, be called "dirt". As a result, a control engineer's task is, in Layman's terms, that of a ditch digger: to shove dirt from a certain frequency to another frequency to meet certain requirements, as illustrated in Figure 6.16. In this context, when a designer employs more abstract tools, such as linear quadratic Gaussian (LQG), $\mathcal{H}_\infty$, or another method, it can be seen as remotely operating a ditch-digging machine which moves the dirt according to the set of weights specified by the designer, as exemplified in Figure 6.17 (Stein, 2003).

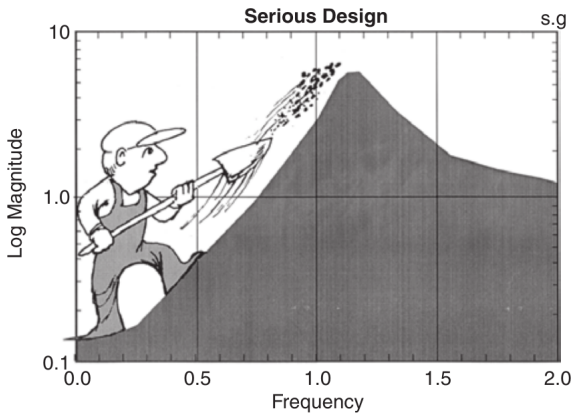

Figure 6.16: Sensitivity reduction at low frequency unavoidably leads to sensitivity increase at higher frequencies. Image retrieved from Stein, 2003.

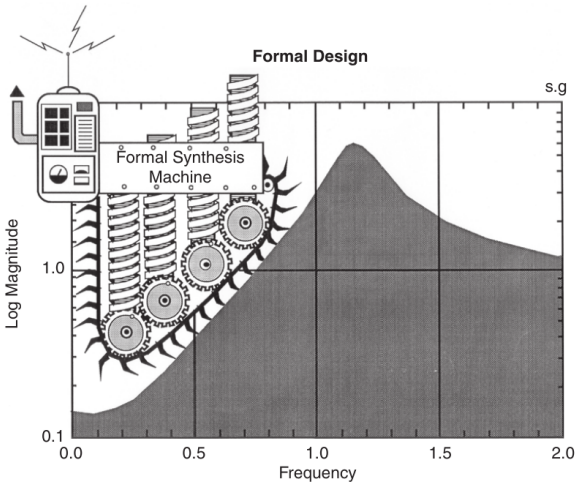

Figure 6.17: Sensitivity shaping automated by modern control tools. Image retrieved from Stein, 2003.

For unstable plants, the integral relationship in 6.11 dictates that the area above the 0 dB is at best equal to π times the sum of the real part of the unstable poles of the plant. Consequently, the more unstable the plant is, the higher the area above 0 dB. One has to wonder about the possibility of just spreading the "dirt" across the infinite frequency to make the layer extremely thin, thus making the peak of the Sensitivity function close to 1. Unfortunately, the harsh reality is that in control, there is no such thing as "free lunch". There is an upper bound frequency up to which the designer can pile dirt up, named available bandwidth Ω_a. The available bandwidth, as defined by Stein, 2003, is "the frequency up to which we can keep $GK(j\omega)$ close to a nominal design and beyond which we can only guarantee that the actual loop magnitude will attenuate rapidly enough". This is an a priori constraint imposed by the physical hardware.

As previously discussed, RHP poles dictate lower bounds on the crossover frequency, and, consequently, also on Ω_1. Furthermore, performance requirements might also impose a lower bound for Ω_1. Therefore, the region where the "dirt" can be piled up is limited according to the performance requirements and the physical hardware. The more the performance requirements and the worse the physical hardware, the shorter the frequency region. Given the integral relationship in 6.11, for a violently unstable plant, there is a minimum peak value M_s for the sensitivity function, as demonstrated in Figure 6.18.

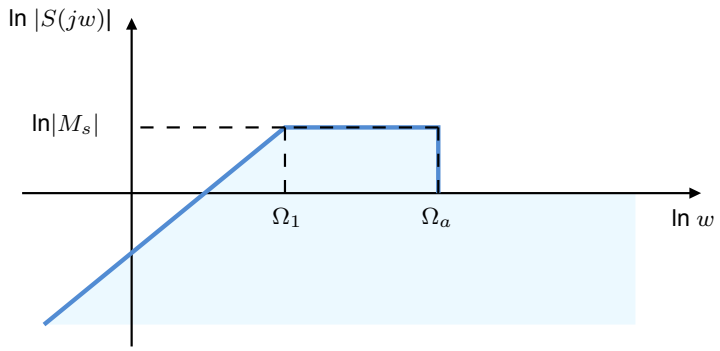

Figure 6.18: Generic ideal sensitivity function of an unstable plant.

The prototype sensitivity shape, idealised as a piece-wise continuous straight-line approximation, displayed on a log-log plot in Figure 6.18, can be defined as (Grauer, 2025):

$$|S(j\omega)| = \begin{cases} \frac{M_s \omega}{\Omega_1}, & \omega \leq \Omega_1 \\ M_s, & \Omega_1 < \omega \leq \Omega_a \\ 1, & \Omega_a < \omega \end{cases} \tag{6.12}$$

Defining $\sum \text{Re}(p) = p_u$ and substituting the straight-line approximations of the sensitivity function in Equation 6.12 into Equation 6.11 allows the Bode sensitivity integral to be written as:

$$\int_0^{\Omega_1} \ln \left| \frac{M_s(\omega)}{\Omega_1} \right| d\omega + \int_{\Omega_1}^{\Omega_a} \ln |M_s| d\omega + \int_{\Omega_a}^{\infty} \ln |1| d\omega = \pi p_u \tag{6.13}$$

These integrals evaluate to:

$$\Omega_1 \left[\ln(M_s) - 1\right] + (\Omega_a - \Omega_1) \ln(M_s) + 0 = \pi p_u, \tag{6.14}$$

from which the peak sensitivity M_s can then be solved for as:

$$M_s = \exp \left[\frac{\pi p_u + \Omega_1}{\Omega_a} \right] \tag{6.15}$$

The X-29 was one of the examples Stein, 2003 used to demonstrate the power of the Bode integral to understand the physical limitations of a system. He considered a flight condition where the plane was even more unstable, with a pole at 6 rad/s and then used the physical hardware information of the sensors, control processors, actuators, aerodynamics and airframe structures (fuselage bending modes) to establish Ω_a. The limiting factor was established as being the aircraft's mechanical structure and the sampling rate of its computers, which determined Ω_a to be approximately 40 rad/s. More recently, the work of Grauer, 2025 offers a new perspective on the computation of Ω_a. Stein then used the unstable nature of the plant and the performance requirements to choose Ω_1 equal to 3 rad/s. Based on this frequency region and using the relationship in 6.15, the resulting smallest sensitivity penalty M_s is approximately 1.75. It is well known that traditional phase margins are given in terms of the peak of the sensitivity function M_s by $PM = 2sin^{-1}(\frac{1}{2M_s})$. Therefore, a PM of 35° is the largest phase margin attainable for the flight condition at hand and given the hardware limitations. This value is below the golden rule of 45° of PM

and 6 dB of GM for certification (Dobos-Bubno & Hartsook, 1977), which led to the necessity of creating special certification requirements for this experimental aircraft to fly. The contractor decided to relax the requirements to 4.5 dB and 30°—assuming all known high-order dynamics were accounted for in the analysis — while the U.S. government flight test team opted to mandate flight-measured stability margins, setting minimum margins of 3 dB and 22.5° (Clarke et al., 1994).

This means that, irrespective of the chosen control technique to design a flight controller for X-29, regardless of whether it is LQG, $\mathcal{H}_\infty$ or some other control technique, the maximum that can be achieved is limited. It is noteworthy that this study did not entail the computation of a single controller. Nevertheless, the designer possesses a comprehensive understanding of the maximum achievable performance of any controller that may be developed for the system.

For the continuous design, good results were obtained for the classical gain and phase margins (-8.8 dB of GM and 51°). However, these were obtained by spreading the "dirt" of the sensitivity function across a large frequency interval, resulting in a smaller peak value of the sensitivity function, as can be seen in Figure 6.11. In reality, when the physical limitations of the hardware are considered, the discussion presented above evidences that the continuous-time results are overly optimistic, as it does not consider the fundamental limitations that impose a lower available bandwidth Ω_a. When considering a digital design, the limitations imposed by the sampling rate of the computers appear, which contribute to limiting the available bandwidth Ω_a. As such, the "dirt" can not be spread thin across a large frequency interval, resulting in a larger peak value for the sensitivity function and, hence, smaller guaranteed gain and phase margins. The limitations imposed by structural dynamics were not considered in the scope of this research.

6.3.2. 2DoF Digital Design

As mentioned before, the X-29 is a prime example of how hardware limitations impose limits on what the control system can achieve. The On-Board flight computer had an 80 Hz sampling rate, which was split in half for each control channel: longitudinal and lateral (Bosworth, 1992). As a result, the longitudinal controller had a sampling frequency of 40 Hz, or, in other words, a sampling time T = 0.025s. Additionally, NASA reports claim that the computational delay Δ was, on average, 10 ms (Bosworth, 1992). Given that the sampling frequency for the longitudinal controller channel is 40 Hz, the Nyquist frequency is equal to 20 Hz. Thus, the anti-aliasing cutoff frequency was selected to be 19 Hz, just below the Nyquist frequency. Furthermore, as standard guidelines point out, the maximum frequency of the controller poles should be well beneath the Nyquist frequency. The maximum pole frequency was set to 12.5 Hz, which is still relatively high given the Nyquist frequency, but it was considered that the challenge and the limitations of the problem at hand required so to achieve a controller with good performance.

Given the challenges that arose due to the reduced sampling rate available to control a highly unstable aircraft, the feedforward order was increased to a second-order transfer function, given by:

$$K_{ff}(s) = K\frac{(s+z_1)(s+z_2)}{s^2 + 2\xi w_n s + {w_n}^2} = K\frac{(s+z_1)(s+z_2)}{s^2 + (2)(1)(w_{max})s + w_{max}^2} \tag{6.16}$$

where the poles of the feedforward controller were fixed to the real axis at the maximum allowed frequency $w_{max} = 12.5$ Hz $= 12.5 \cdot 2\pi$ rad/s. Due to the high-frequency nature of these poles, which will appear as closed-loop poles of the system, their effect is negligible. This leaves 2 zeros, which can be used to cancel out undesired dynamics and achieve a better model following response.

Furthermore, the pre-compensator filter W_1 was slightly adjusted in order to increase the γ_{th} associated with the shaped plant G_s by moving the zero of W_1 to a slightly lower frequency. The new W_1 is given by:

$$W_1 = \frac{0.25519(s + 7.166)}{s} \tag{6.17}$$

The non-smooth $\mathcal{H}_\infty$ minimisation optimization tool in `systune` was then employed as displayed in Figure 6.15, and the results are presented in the tables below.

Table 6.5: Filters setup and optimisation result of the Tuning Goals for the digital design.

	Shaping Filters		NCF Rob. Soft Goal		co-design Soft Goals		MF Hard Goal
Parameter	W_1	W_2	γ_{th}	γ	w_e	M_e	W_{mf}
Value	$\frac{0.25519(s+7.166)}{s}$	1	2.9030	3.3501	3.3501	3.3501	1.0000

Table 6.6: Optimisation gains of the digital Hybrid IDI using the LSDP with model-following requirements.

	Virtual Control			Inversion Loop		Feedforward	MF Filter W_{mf}		
Parameter	K_p	K_i	w_{lp}	K_c	w_{H_c}	K_{ff}	A_e	w_e	M_e
Value	9.86	27.99	45.00	0.38	78.54	$\frac{10.85(s^2+28.67s+573.70)}{(s+78.54)^2}$	0.36	26.86	1.17

From the optimisation results, it can be seen that even with the adjustment of the W_1 filter, the inclusion of the digital effects led to a deterioration of the maximum achievable γ_{th}, with an increase from 2.3181 to 2.9030. The obtained γ is relatively higher than γ_{th} but still within the acceptable interval of 1 to 4.

Just as in the continuous-time optimisation, the co-design `TuningGoals` have the same value as γ. Given the robustness challenges that arose from the limited sampling time, the model-following requirements had to be relaxed. The desired bandwidth value was set to $w_{e_{ideal}} = 90$ and the desired peak value was set to $M_{e_{ideal}} = 0.35$, which reflects a lower value for $w_{e_{ideal}}$ compared to the continuous-time optimization. Making $w_{e_{ideal}}$ larger would result in an even higher value for γ, so a compromise was made between robustness and performance.

Similarly, the low-pass pole w_{lp} was manually set to 45 rad/s. The integrator gain K_i of the PI has been reduced, which comes at the cost of reduced performance to increase robustness. The proportional term K_p and the inversion gains K_c and w_{H_c} present values similar to the previous optimisation. As described above, the poles of the feedforward element K_{ff} were set at the maximum allowed frequency and thus have little impact on the response of the system. The motivation for this decision comes from the fact that the earlier optimisations, where the location of the poles was unconstrained, resulted in the placement of the stable poles at extremely high frequencies. The relevant elements to be analysed are the location of the zeros of K_{ff}, as they can be used to cancel the undesired dynamics of the closed-loop system and, thus, improve the model-following performance.

Figure 6.19 illustrates the Pole-Zero map of the closed-loop system with and without the feedforward controller. The closed-loop system without the feedforward element corresponds to the transfer function of an input signal r entering the system after the feedforward element, while the transfer function from $q_{cmd} \rightarrow q$ corresponds to the closed-loop transfer function of the closed-loop system with the feedforward element. Since both the Hybrid IDI feedback elements and the feedforward elements were tuned simultaneously

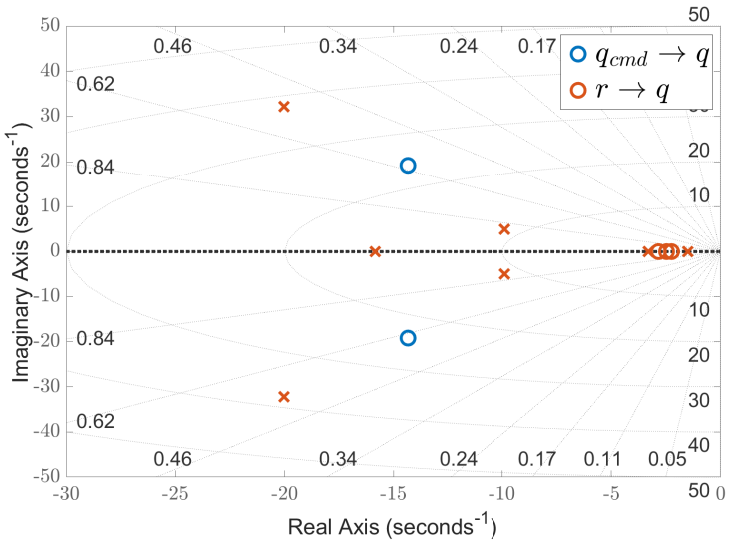

Figure 6.19: Pole-Zero map of the closed-loop system without the feedforward ($r \rightarrow q$) and with the feedforward element ($q_{cmd} \rightarrow q$).

for both robustness and model-following performance, it is challenging to isolate the contribution of the feedforward to performance. Nevertheless, from Figure 6.21, it is possible to attest to the success of the feedforward element in achieving good model-following performance.

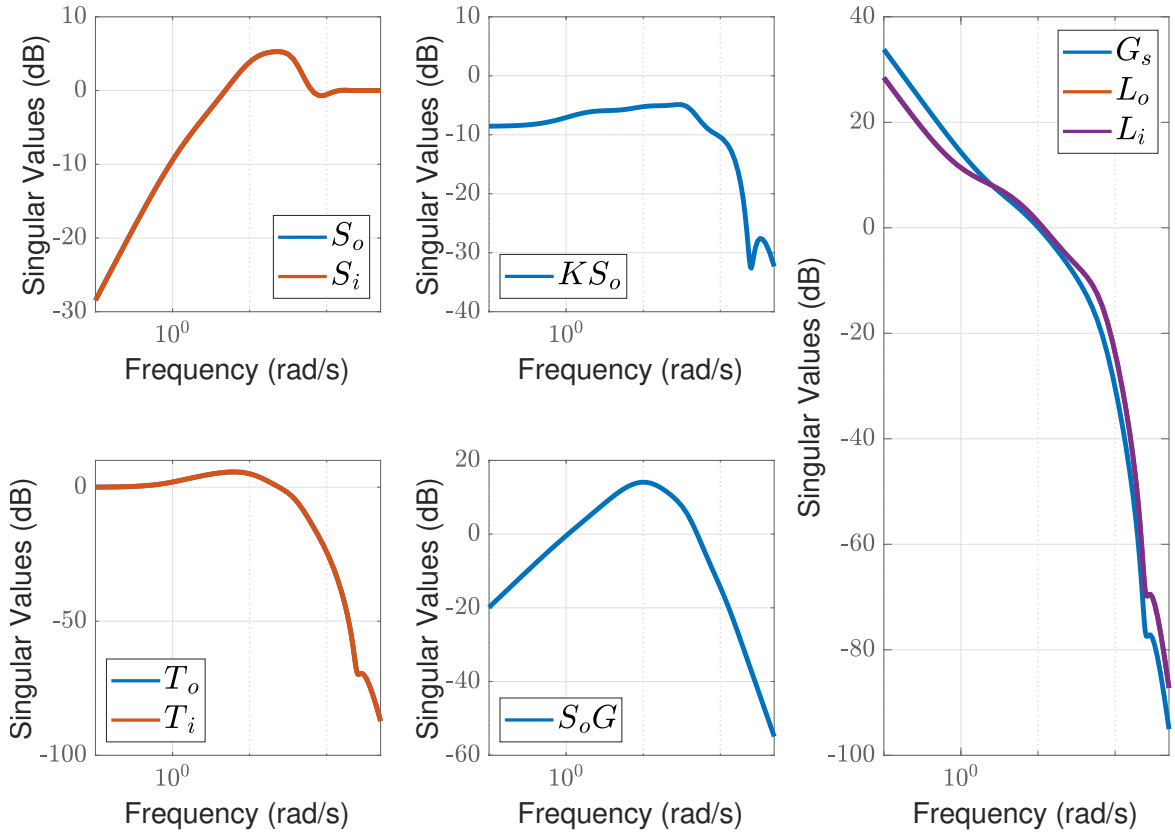

Figure 6.20: Gang of six and broken-loop results for the digital Hybrid IDI controller.

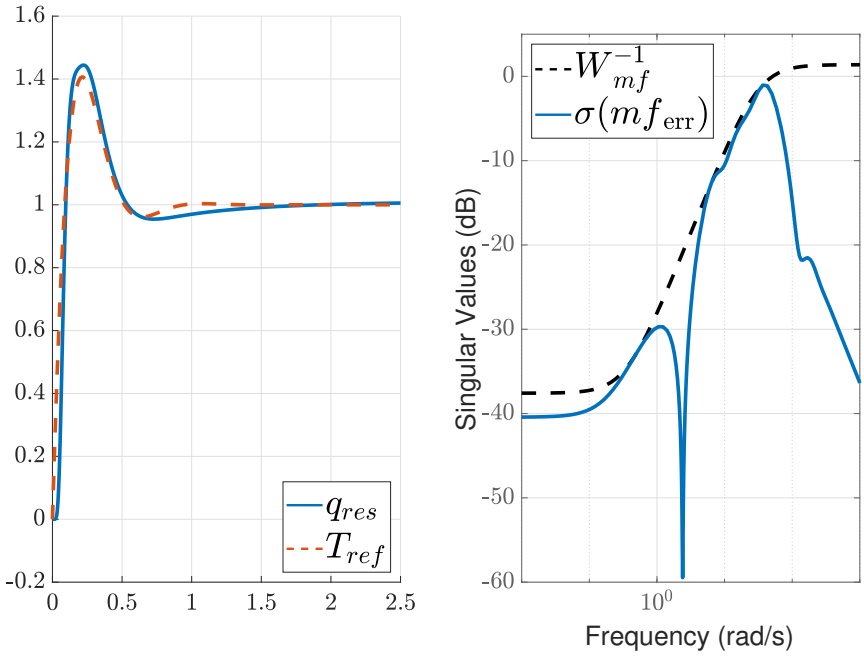

Figure 6.21: Step-tracking performance and model following error for the digital Hybrid IDI controller.

One aspect that deserves attention is the fact that the KS_o response now fulfils the high-frequency requirement $\sigma[KS_o] \ll 1$. This is the result of the anti-aliasing filter, which acts as a low-pass filter. In this context, the reader may question the relevance of the low-pass element in $K_v(s)$. However, the anti-aliasing LPF pole is placed at frequencies slightly below the Nyquist frequency (in this case, 19 Hz, which is approximately 119 rad/s), while the maximum controller bandwidth is usually set way below this value. The 45 rad/s value for w_{lp} is well within the controller bandwidth and is important to contribute to control signal attenuation at more intermediate frequencies.

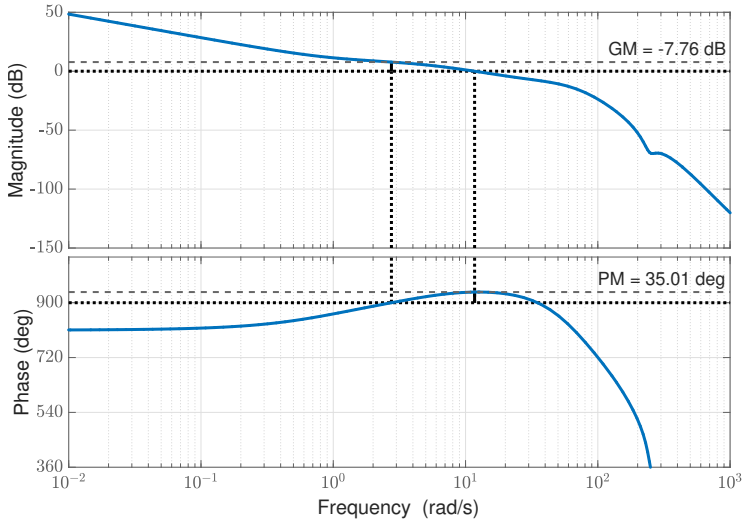

Figure 6.22: Classical margins for the digital Hybrid IDI controller.

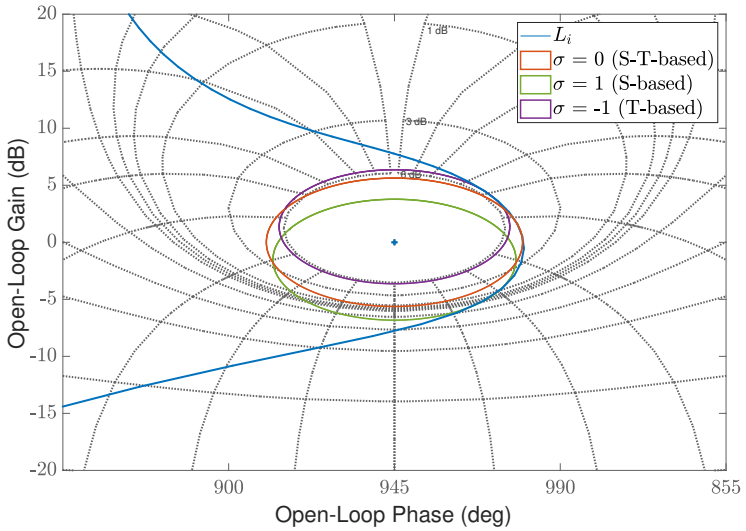

Figure 6.23: Nichols plot and Disk-based margins for the digital Hybrid IDI controller.

The performance deterioration of the system can be analysed through the input and output broken-loop response. Compared to the previous optimisation, this decrease is attributed to the reduction in K_i gain, which results in a lower gain at low frequencies in the digital design. The deterioration of the performance can also be attributed to the relaxation of the co-design goals, which resulted in W_{mf}^{-1} having a larger DC gain, smaller bandwidth and higher peak value.

Table 6.7: Classical and Disk-based margins for the digital Hybrid IDI design.

	Classical	Disk-based		
		S-T ($\sigma = 0$)	S ($\sigma = 1$)	T ($\sigma = -1$)
GM (dB)	-7.76	5.63	6.83	3.63
PM (°)	35.01	34.77	31.59	30.11

Even so, the fact that the digital effects and the hardware limitations constrain the available bandwidth Ω_a, means that there is a smaller region to pile up the “dirt”. This can be seen in the sensitivity plot, where the sensitivity rolls off at a lower frequency than before and exhibits a higher peak value of M_s. Consequently, the robustness of the system is severely impacted. This justifies the higher γ value and the decreased classical and disk-based gain and phase margins below the typical requirement of 45° of PM. The obtained value for the PM of 35° is in accordance with the X29’s stability margins measured explicitly in flight (Gera & Bosworth, 1987), as can be observed in Figure 6.24.

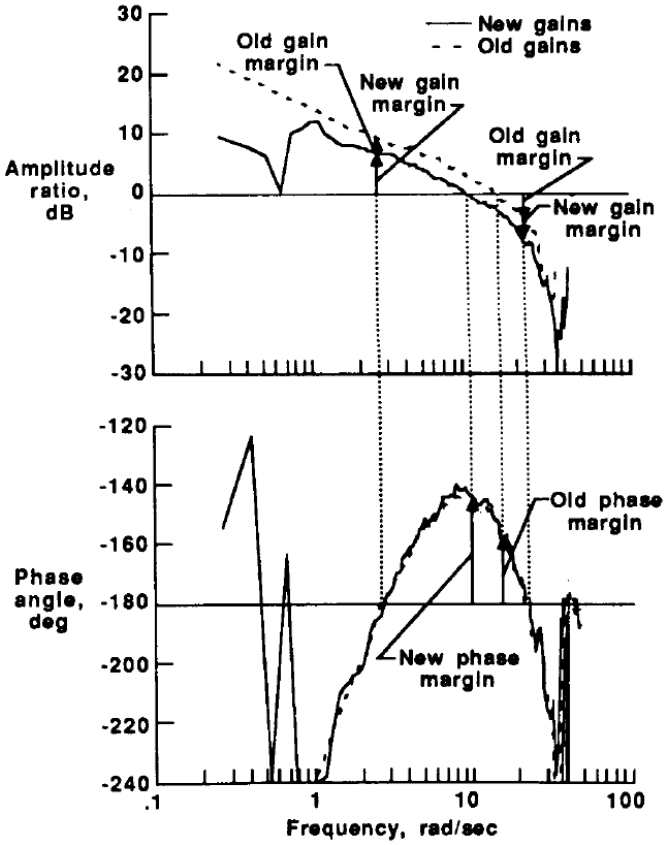

Figure 6.24: Flight-measured pitch-axis Bode plot of two different control law iterations for the X-29. Even the updated controller only had around 35° of PM. Image retrieved from Gera and Bosworth, 1987.

6.3.3. Discrete-Time Implementation of the Hybrid IDI controller

After the synthesis of the continuous-time Hybrid IDI controller with the inclusion of continuous-time models of the digital effects, the last step is to transform the continuous-time controller to a discrete-time controller. In order to do so, the bilinear transform, also known as Tustin's method, was used using the sampling time of $T_s = 25$ ms based on the X-29 sampling rate.

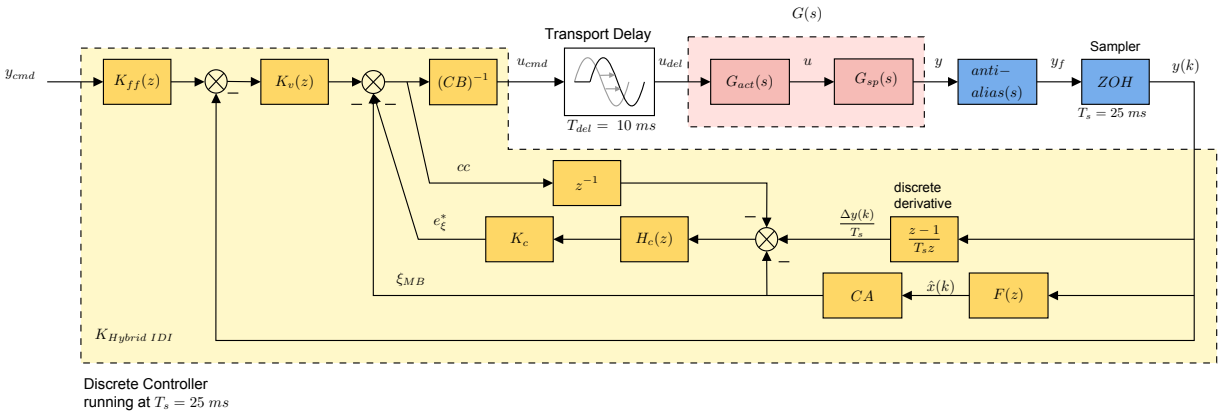

Figure 6.25: Implementation of the digital Hybrid IDI controller in MATLAB® Simulink.

$$K_v(s) = \frac{443.45(s+2.84)}{s(s+45)} \rightarrow K_v(z) = \frac{3.674 + 0.2519z^{-1} - 3.422z^{-2}}{1 - 1.28z^{-1} + 0.28z^{-2}} \tag{6.18}$$

$$K_{ff}(s) = \frac{10.85(s^2 + 28.67s + 573.70)}{(s+78.54)^2} \rightarrow K_{ff}(z) = \frac{4.002 - 5.033z^{-1} + 2.021z^{-2}}{1 - 0.01842z^{-1} + 8.483e^{-5}z^{-2}} \tag{6.19}$$

$$H_c(s) = \frac{78.54}{s+78.54} \rightarrow H_c(z) = \frac{0.49539(1+z^{-1})}{1 - 0.00921z^{-1}} \tag{6.20}$$

$$\frac{\hat{\alpha}(s)}{q(s)} = \frac{0.9897}{s+2.241} \rightarrow \frac{\hat{\alpha}(z)}{q(z)} = \frac{0.012034(1+z^{-1})}{1 - 0.9455z^{-1}} \tag{6.21}$$

Figure 6.25 illustrates how the discrete-time Hybrid IDI controller was implemented in `Simulink`. The plant dynamics $G(s)$ and the anti-aliasing filter are implemented as continuous-time blocks. In order to simulate the sampler, a ZOH block is implemented with a sampling time of $T_s = 25$ ms. The Hybrid IDI controller is a discrete subsystem with the same sampling time of $T_s = 25$ ms. In order to differentiate the output signal, the discrete derivative block is used, and given the internal Control Command (CC) feedback structure, a Delay block is used on the CC signal with a delay length equal to one time step.

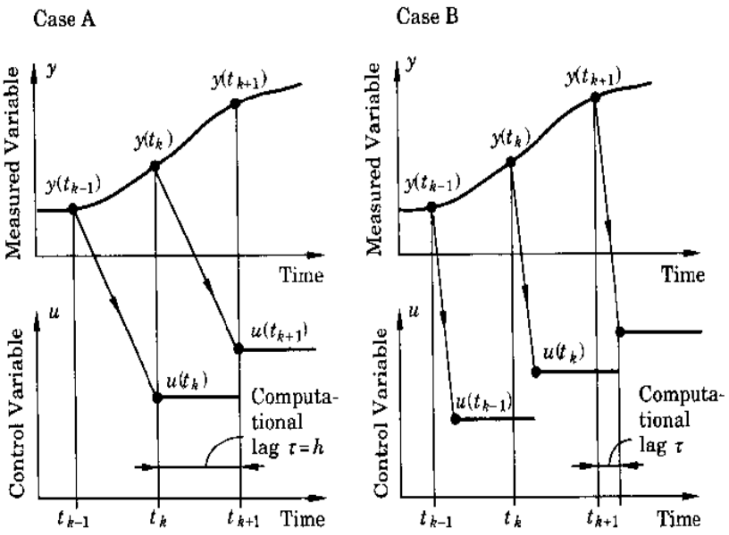

Figure 6.26: Two alternative ways of synchronising measured outputs with the computed input by the control system. In case A, the computed control signal to be applied at time t_{k+1} makes use of signals measured at time t_k. On the other hand, in case B, the control signals are computed and immediately applied. Image retrieved from Åström and Wittenmark, 1997.

As described in Åström and Wittenmark, 1997, to compute the control signal at a certain time t_{k+1} based on measured variables at a previous time t_k, there are two possible approaches. For the first one, case A, the signal measured at time t_k is used to compute the control signal to be applied at time t_{k+1}, which means that the computational lag is fixed to be exactly one time-step cycle T_s. The second approach, case B, consists of changing the control signal as soon as the computation is made, which effectively results in a computational lag equal to the computational delay T_{del}. Given the relatively high sampling time T_s of the discrete-time controller, the second option for the synchronisation of inputs and outputs was chosen. In order to simulate this in `Simulink`, a transport delay block was used with $T_{del} = 10$

ms, which results in u_{del} being delayed in comparison with u_{cmd}, which, in turn, is synchronized with $y(k)$. Consequently, measurements at a time t_k are used to produce an input command at time $t_k + T_{del}$, thus simulating the case B approach.

To analyse the nominal performance of the discrete-time Hybrid IDI controller, pitch-rate step-inputs with amplitude equal to 5°/s were provided as a command signal y_{cmd} resembling a doublet manoeuvre. The same exact y_{cmd} is provided to the modified continuous-time Hybrid IDI controller in Section 6.3.2, simulated in continuous-time. The pitch-rate response can be observed in Figure 6.27, and a zoomed-in portion of the response is displayed in Figure 6.28. The results show that the discrete-time implementation, despite the similar rise and settling time, results in a larger overshoot. This can be explained by the fact that the ZOH and the computational delay continuous-time models used for the synthesis were approximations that do not perfectly match the discrete-time implementation of these blocks. Furthermore, this specific implementation of the CC signal contains an internal delay (see the delay block in the CC signal path in Figure 6.25), which is unaccounted for in the design. This, together with the low sampling rate compared to other fighter flight control applications, typically between 50 and 80 Hz (Kim et al., 2024), contributes to the differences observed between the discrete-time controller and its continuous-time approximation.

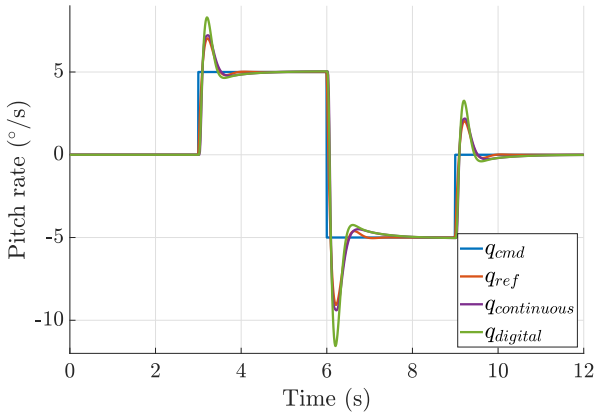

Figure 6.27: Pitch rate response of discrete-time Hybrid IDI controller and comparison with continuous-time controller with modelled digital effects.

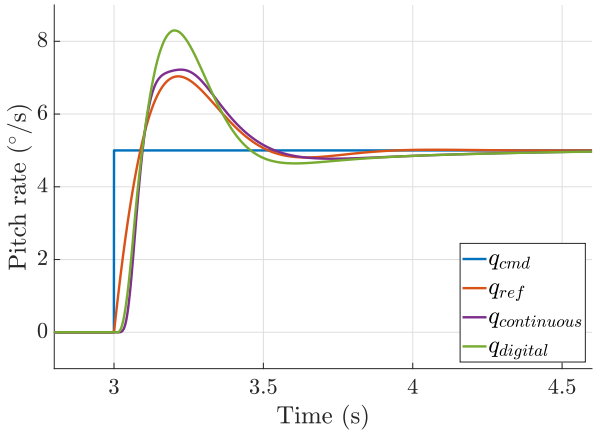

Figure 6.28: Zoomed in pitch rate response of discrete-time Hybrid IDI controller.

Nevertheless, if the input u response of the discrete-time Hybrid IDI is compared with the continuous-time one, it is possible to observe in Figure 6.29 that the differences are minimal. Figures 6.30 and 6.31 highlight the correct implementation of the synchronisation issue described above.

Overall, the discrete-time Hybrid IDI controller implementation evidences the success of the used approach to optimise the digital controller: modify the continuous-time controller to take into account discretisation effects, optimise it using the procedure described in Chapter 5 and then discretise the optimal continuous-time controller. The differences in the response of the discrete controller to the continuous counterpart can be attributed to the low sampling rate, the approximation nature of the ZOH and the computational delay of continuous-time transfer functions.

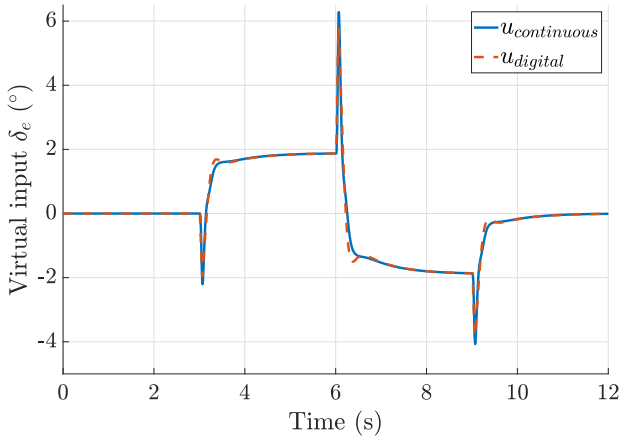

Figure 6.29: Comparison of input δ_e response of digital Hybrid IDI controller and continuous-time controller with modelled digital effects.

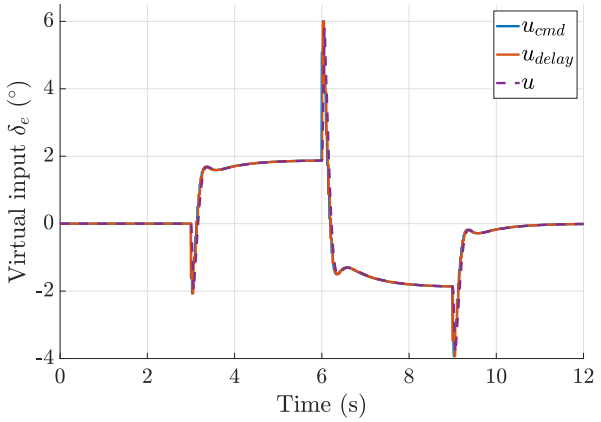

Figure 6.30: Input δ_e response of discrete-time Hybrid IDI controller.

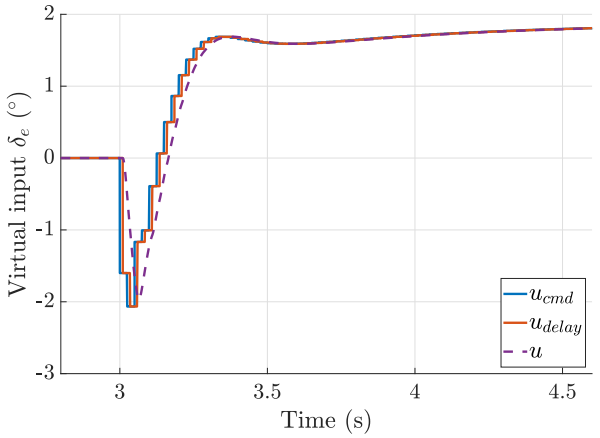

Figure 6.31: Zommed in input δ_e response of discrete-time Hybrid IDI controller.

Implementation-wise, multiple improvements can be made to decrease the controller complexity and better match the closed-loop response of the modified continuous-time design and the digital controller. The use of a SCF structure, as described in Figure 5.12, would avoid the explicit discrete differentiation of the output feedback, decreasing the controller complexity. The Hybrid IDI schematic can also be rewritten equivalently, removing the algebraic loop associated with the CC signal and, thus, the delay on that signal path in the discrete controller. This leads to the following control law (Pollack et al., 2024):

$$u_{cmd}(s) = (CB)^{-1}C_c(s)(v(s) - [1 - K_cH_c(s)]\xi_{MB} - [K_cH_c(s)]sy) \tag{6.22}$$

where $C_c(s) \triangleq (1 - K_cH_c(s))^{-1}$. Considering a first-order low-pass filter $H_c(s) = \frac{w_f}{s+w_f}$, and, consequently, that $C_c(s) = \frac{s+w_f}{s+(1-K_c)w_f}$, then it is possible to rewrite Equation 6.22 as:

$$\begin{aligned} u_{cmd}(s) &= (CB)^{-1}(C_c(s)v(s) - C_c(s)[K_cH_c(s)]sy - \xi_{MB}) \\ &= (CB)^{-1}\left(\frac{s+w_f}{s+(1-K_c)w_f}v(s) - \frac{K_cw_f}{s+(1-K_c)w_f}sy - \xi_{MB}\right) = (CB)^{-1}(v^* - \xi_{MB}) \end{aligned} \tag{6.23}$$

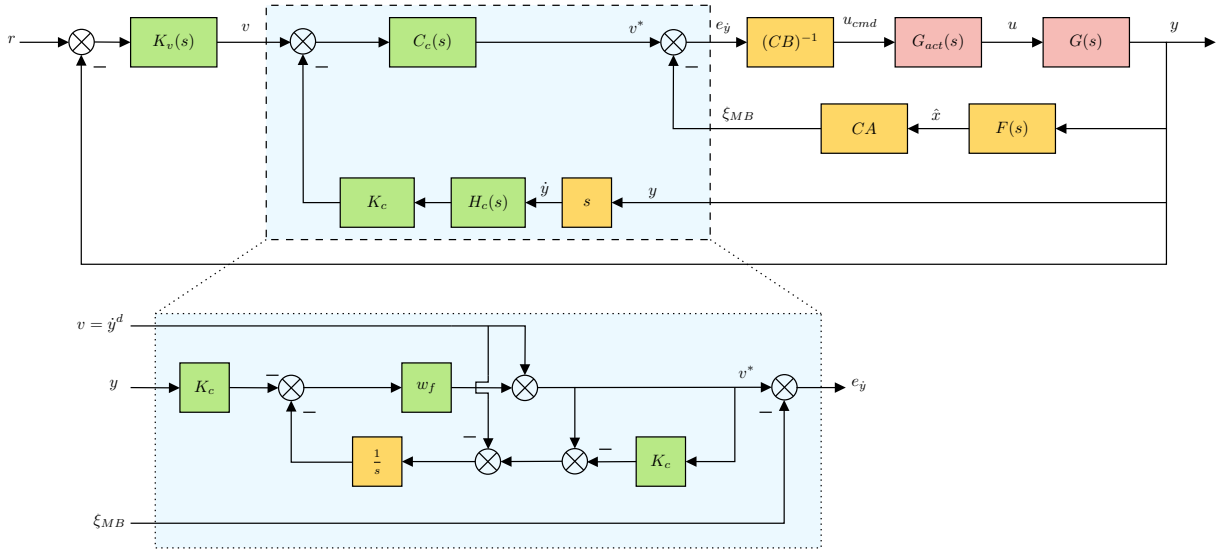

Figure 6.32: Equivalent Hybrid IDI structure without CC feedback, where the green blocks refer to the tunable elements of the controller and the yellow blocks to the plant-dependent/physical blocks of the controller; the insert shows a simplified implementation of the control law which does not require explicit differentiation of the output variable.

This boils down to the equivalent regulation formulation of the Hybrid IDI control law using CC feedback from Figure 5.12 as the one depicted in Figure 6.32. To avoid confusion, please note that, despite the similar notation, the auxiliary regulation term $C_c(s)$ and the Control Command feedback signal CC refer to different concepts. Similarly to the description of a scaled complementary filter (SCF) signal blending implementation in Figure 5.12, it is also possible to implement this control law in a simplified way, as described in the insert in Figure 6.32. This results in an equivalent control law that does not require explicit differentiation or CC feedback.

An aspect that deserves attention in future studies is the analysis of the performance of the discrete-time controller with a more complete nonlinear dynamics model. The current analysis is restricted to the use of a linear model for the actuators and the short-period dynamics, which can yield overly optimistic results, as it does not consider the actuators' position and rate limits and their potential interaction with the nonlinear aerodynamic forces and moments. Including these factors in the analysis could provide a more realistic evaluation of the controller's performance, particularly under extreme flight conditions or during aggressive manoeuvring, where nonlinear effects become significant.

6.4. Discussion and Conclusion

To further elaborate on the proposed procedure to tune Hybrid IDI controllers using the $\mathcal{H}_\infty$ LSDP, the block manipulation presented in Figure 6.33 is used (recall that throughout the procedure, it was assumed that W_1 and W_2 are invertible).

The closed-loop implementation reduces to simply having the controller K_{H-IDI} and the plant G, as the introduced filters perfectly cancel with their respective inverses. However, by considering the disturbances and outputs entering the specified locations, the standard $\mathcal{H}_\infty$ LSDP framework is recovered, wherein the controller K^*_∞ is computed to robustify the shaped plant given by W_2GW_1. Referring back to the proposed procedure, it was assumed that $K_{H-IDI} \to W_1K_\infty W_2$, implying that in the ideal scenario where $K_{H-IDI} = W_1K_\infty W_2$, then $K^*_\infty = W_1^{-1}W_1K_\infty W_2W_2^{-1} = K_\infty$. However, since K_{H-IDI} is highly

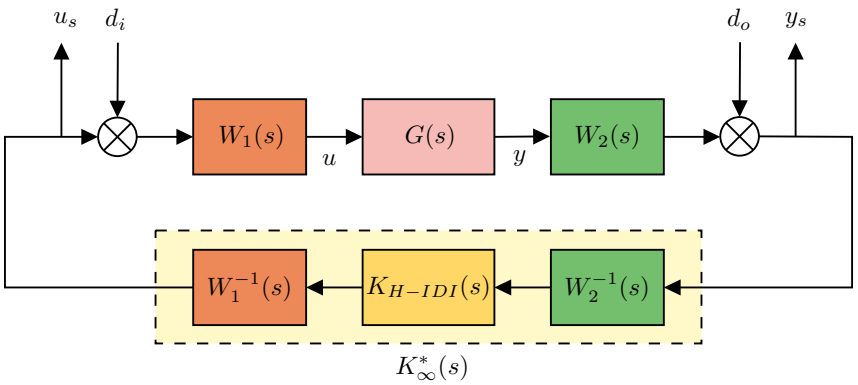

Figure 6.33: Equivalent closed-loop representation with addition of "dummy" filters and its respective inverses.

structured, it serves only as an approximation to $K_s = W_1 K_\infty W_2$, and thus it is necessary that the structure of K_{H-IDI} is sufficiently 'rich' such that $K_\infty^* \to K_\infty$ and, therefore, $\epsilon \to \epsilon_{opt}$. Nevertheless, when the $\mathcal{H}_\infty$ norm from $[d_o\, d_i]^T \to [y_s\, u_s]^T$ is computed using this equivalent closed-loop representation, the resulting γ-value matches exactly with that obtained during the synthesis procedure. This consistency underscores the validity of the proposed method for tuning structured controllers, specifically the Hybrid IDI controller, while maintaining the robustness guarantees of $\mathcal{H}_\infty$ LS.

Hyde and Papageorgiou, 2001 makes use of a similar concept to analyse the robustness of different DI controllers. The idea of the article is to generalise the analysis to any linear controller K, regardless of the design method, using the NCF stability test, which is inherently integrated into the synthesis process of $\mathcal{H}_\infty$ LSDP.

The approach involves transforming the controller and the plant into an equivalent normalised coprime factor controller and its corresponding weighted plant, as illustrated in Figure 6.33. The goal is to determine diagonal weights W_1, W_2 such that $W_1^{-1} K W_2^{-1}$ becomes the optimal normalised coprime factor controller for the weighted plant $W_2 G W_1$. This is achieved by solving (Hyde & Papageorgiou, 2001):

$$\min_{W_1,W_2} \|T_{d\to e}\|_\infty = \min_{W_1,W_2} \left\| \begin{bmatrix} W_2 & \\ & W_1^{-1} \end{bmatrix} \begin{bmatrix} I \\ K \end{bmatrix} (I - GK)^{-1} \begin{bmatrix} I & G \end{bmatrix} \begin{bmatrix} W_2^{-1} & \\ & W_1 \end{bmatrix} \right\|_\infty \triangleq \frac{1}{\epsilon_\mu}, \tag{6.24}$$

where d consists of $[d_o\, d_i]^T$, e consists of $[y_s\, u_s]^T$ and W_1, W_2 are restricted to being diagonal, stable and minimum phase. If the same values for W_1 and W_2 are used, then ϵ_μ yields the same ϵ value obtained from the structured $\mathcal{H}_\infty$ minimisation. Nevertheless, as the analysis tool searches for W_1 and W_2 weights such that the structured controller $K = K_{opt}$ for the associated shaped plant G_s, W_1 and W_2 are potentially very high order terms, but the resulting $\epsilon_\mu = \epsilon_{opt}$. This is something the synthesis procedure proposed in this study can not guarantee, as the synthesized controller K is computed from chosen W_1 and W_2 weights, and so one can only hope that the structure of K is 'rich' enough such that $K \to K_{opt}$ and, therefore, $\epsilon \to \epsilon_{opt}$.

The authors then make use of this tool to analyse the robustness of different dynamic inversion controllers: NDI with a P gain in the virtual law control, referred to as Standard NDI, NDI with a PI in the virtual law control, referred to as P+I NDI, and what the authors refer to as Robust Inverse Dynamics

Estimation (RIDE), which is essentially an unfiltered INDI controller. Their results demonstrate that both the Standard NDI and the RIDE controller give very poor results in terms of ϵ_μ, while the P+I NDI gives acceptable results. The fact that a simple proportional gain as the virtual control law gives poor robustness is in line with the discussion conducted in Chapter 5. Moreover, the fact that the RIDE controller also gives poor robustness is in accordance with the research results of Pollack, 2024, which demonstrated that INDI is not necessarily more robust than its model-based counterpart and that its success in achieving a robust design is dependent on the use of a synchronization filter H_c. It is, therefore, natural that an unfiltered implementation of INDI results in a controller with poor robustness.

Since its value remains primarily from an analysis standpoint, the designer must employ some method/guidelines to design the virtual control law and "hope" that the a-posteriori analysis using this tool confirms a successful robust design. The significance of the W_1 and W_2 filters is also less clear in this scenario, as they are computed only after the synthesis of the controller. On the other hand, the procedure described in this study **retains the core purpose of W_1 and W_2 in shaping the open-loop plant** and, given a successful design, bounds the broken loop gain at the plant I/O. Moreover, the synthesis results in a $\gamma = \epsilon^{-1}$ value, which ties directly to the robustness of the system to NCF uncertainty, and the closer the values of ϵ and ϵ_{opt}, the closer the obtained controller K is to the optimal LS controller K_{opt}. The gap $\epsilon_{gap} = \epsilon_{opt} - \epsilon$ is a direct indication of how 'optimal' the obtained controller is compared to the optimal full-order $\mathcal{H}_\infty$ LS controller. Overall, **the procedure effectively allows the designer to tune DI controllers with a priori robustness guarantees**.

7

X-29 Lateral-Directional MIMO Controller Design

The previous chapter presented a methodology to combine $\mathcal{H}_\infty$ Loop-Shaping with a Dynamic Inversion schematic. More precisely, the Hybrid IDI controller structure was tuned from a $\mathcal{H}_\infty$ Loop-Shaping synthesis point of view. Nevertheless, the design and tuning of pitch-rate controllers may successfully employ simpler classical tools, such as Bode plots, root locus, frequency domain analysis, and SISO loops (Skogestad & Postlethwaite, 2005). The limitations of these more classical methods become evident for MIMO applications, since the classical loop-at-the-time control comes to its limits when addressing systems with high internal coupling, as it limits full exploitation of the concepts of directionality associated with these systems. This motivated the origin of synthesis tools to address the multivariable nature of the problem, from which $\mathcal{H}_\infty$ control stands out. Therefore, it is only logical to assess the potential of the procedure with a MIMO system.

This chapter aims to explore how the procedure can be extended to a MIMO case and evaluate the potential of the design philosophy. Given the complexity of the task at hand, this chapter aims primarily to start and open up the discussion rather than close it, which means that the results in this section are predominantly preliminary and deserve further investigation. In order to demonstrate the expandability of the procedure to a MIMO Hybrid IDI control architecture, a lateral-directional controller for the X-29 was designed with the chosen control variables being the roll rate p and yaw rate r.

One of the main challenges of $\mathcal{H}_\infty$ Loop-Shaping is precisely how to design the shaping filters W_1 and W_2. For SISO cases, this part is somewhat trivial since shaping the singular value of a SISO system corresponds to adjusting the gain of a single transfer function and without explicit regard to the plant phase, as opposed to classical loop shaping. However, for MIMO cases, the task is more complex since the singular values (MIMO 'gain' of the system) are associated with the directionality (input and output directions) of the system. This often implies a process of trial and error to come up with satisfactory filters W_1 and W_2, such that the shaped plant G_s given by W_2GW_1 meets the desired shape (G. Papageorgiou, 1998). Hyde, 1995 provides a systematic method to design diagonal filters W_1 and W_2, which works particularly well for plants with dominant diagonal terms. However, as G. Papageorgiou, 1998 describes:

"Experience gained with designing $\mathcal{H}_\infty$ Loop-Shaping controllers indicates that diagonal weights do not work well with plants that have strong cross-coupling between channels. The task of choosing W_1 and W_2 for such plants can be difficult and time-consuming if done in an ad-hoc manner. Furthermore, there is

no guarantee that the desired loop shape will be achieved."

In this same work, G. Papageorgiou, 1998 proposes a systematic procedure to design non-diagonal filters W_1 and W_2, which allows the designer to specify the desired loop shape more accurately as he exploits more degrees of freedom. This work has caught little attention within the robust control community, despite the implications that it has on making the design of the shaping filters more straightforward and systematic. For this reason, the present chapter will also dive into this methodology and the implications it has when integrated with the proposed procedure, which combines the $\mathcal{H}_\infty$ LSDP synthesis with Hybrid IDI control structures.

Consequently, this chapter is structured as follows: Firstly, the methodology for the design of non-diagonal filters proposed by G. Papageorgiou, 1998 will be reviewed and re-created and used to systematically shape the plant for the lateral-directional control task. Secondly, the implications of using non-diagonal filters to shape the plant, which in the standard $\mathcal{H}_\infty$ LSDP integrates the final robust controller K_s, and to tune the Hybrid IDI lateral-directional controller, whose structure philosophy builds on decoupling the action on the multiple channels, will be discussed. Finally, diagonal weights W_1 and W_2 will be used to tune a Hybrid IDI lateral-directional controller for the X-29 for a certain flight condition, using the procedure described in Chapter 5.

7.1. Open-Loop Analysis of the X-29 Model

The linear state-space equations governing the lateral-directional dynamics of the X-29 at a flight condition of 20,000 ft altitude and Mach 0.7 were obtained from Bosworth, 1992. The two control inputs are the differential flap δ_a and the rudder δ_r (Bosworth, 1992). The controlled variables are the roll-rate p and yaw-rate signals r.

$$\begin{aligned}\begin{bmatrix}\dot{\beta}\\ \dot{p}\\ \dot{r}\end{bmatrix} &= \begin{bmatrix}\frac{Y_\beta}{V_0} & \frac{Y_p}{V_0} & \frac{Y_r}{V_0}-1\\ L_\beta & L_p & L_r\\ N_\beta & N_p & N_r\end{bmatrix}\begin{bmatrix}\beta\\ p\\ r\end{bmatrix} + \begin{bmatrix}\frac{Y_{\delta_a}}{V_0} & \frac{Y_{\delta_r}}{V_0}\\ L_{\delta_a} & L_{\delta_r}\\ N_{\delta_a} & N_{\delta_r}\end{bmatrix}\begin{bmatrix}\delta_a\\ \delta_r\end{bmatrix}\\ &= \begin{bmatrix}-0.1596 & 0.07150 & -0.9974\\ -15.20 & -2.602 & 1.106\\ 6.840 & -0.1026 & -0.06375\end{bmatrix}\begin{bmatrix}\beta\\ p\\ r\end{bmatrix} + \begin{bmatrix}-0.0005980 & 0.0006718\\ 1.343 & 0.2345\\ 0.08974 & -0.07097\end{bmatrix}\begin{bmatrix}\delta_a\\ \delta_r\end{bmatrix}\end{aligned} \tag{7.1}$$

Pole	Damping	Frequency (rad/s)	Time Constant (s)
-2.45×10^0	1.00×10^0	2.45×10^0	4.08×10^{-1}
$-1.87 \times 10^{-1} + 2.77i$	6.72×10^{-2}	2.77×10^0	5.36×10^0
$-1.87 \times 10^{-1} - 2.77i$	6.72×10^{-2}	2.77×10^0	5.36×10^0

Table 7.1: X-29 lateral-directional poles at 20000 ft and Mach 0.7.

From Table 7.1, the roll subsidence mode can be identified from the single real negative root, while the Dutch roll mode can be identified from the stable complex conjugate pair of roots.

The aileron and rudder actuator dynamics are modelled using fourth-order transfer functions retrieved from Bosworth, 1992:

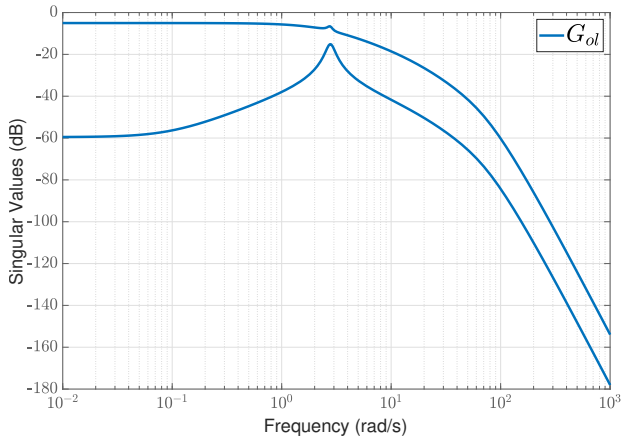

Figure 7.1: Singular Value Decomposition of the open-loop plant G_{ol} including the X-29 lateral-directional dynamics and actuators.

$$G_{\delta_a}(s) = G_{\delta_r}(s) = \frac{(54.1^2)(71.4^2)}{(s^2 + 2(1.53)(54.1)s + (54.1)^2)(s^2 + 2(0.735)(71.4)s + (71.4^2))} \tag{7.2}$$

The resulting open-loop plant G_{ol}, which incorporates the actuator and the lateral-directional dynamics, has the singular value decomposition illustrated in Figure 7.1.

7.2. Procedure to Compute Non-Diagonal Weights

For clarity, the steps taken to recreate the approach will be stated below, but it is important to stress that the procedure is the original work of G. Papageorgiou, 1998.

Procedure by G. Papageorgiou, 1998 to compute Non-Diagonal Weights for the $\mathcal{H}_\infty$ LSDP. The enumerated steps are quoted from G. Papageorgiou, 1998:

Consider a linear time-invariant scaled nominal plant model G with m inputs and p outputs. To make the following steps simpler, it is assumed that $m = p$, i.e. the plant G has the same number of inputs and outputs. The first step of the procedure is to determine the desired loop shape. Afterwards, the singular values need to be isolated in a way such that $\sigma_i(i = 1, \ldots, m)$ denotes the i^{th} singular value of G and varies smoothly with frequency. This means that σ_1 is not necessarily the maximum singular value across frequency, but rather the ordering is such that it tracks the SV continuous lines across frequency. The ordering of the singular values of G is according to the order at $\omega = 0$:

$$\Sigma_G = diag(\sigma_1, \ldots, \sigma_m) \tag{7.3}$$

Every σ_i is associated with a transfer function f_i. The desired i^{th} singular value of the weighted plant is $\sigma_i|f_i|$. Hence, the gain of f_i, denoted by $|f_i|$, is how the designer specifies the desired loop shape. It is important to note that f_i is restricted to be a unit in $\mathcal{RH}_\infty$, i.e. f_i cannot contain integrators. This translates to the designer choosing a diagonal transfer matrix $W = diag(f_1, \ldots, f_m)$, such that the desired singular values of the weighted plant are:

$$\Sigma_{G_s}^{des} = \Sigma_G W = diag(\sigma_1, \ldots, \sigma_m) \times diag(|f_1|, \ldots, |f_m|). \tag{7.4}$$

W is split into two diagonal transfer matrices Π, Γ such that W = $\Pi\Gamma$. Γ reflects the desired conditioning and singular values of W_1, and Π the desired conditioning and singular values of W_2. Note that the designer is free to choose Γ or Π equal to the identity. Besides, there should be no pole/zero cancellations between Π and Γ so as to not unnecessarily inflate the order of Π and Γ. Following this, stable and minimum phase weights, W_1 and W_2, are computed that shape the open-loop plant such that the singular values of the weighted plant are close to:

$$\Sigma_{G_s}^{des} = \Pi\Sigma_G\Gamma = diag(\sigma_1|f_1|, \ldots, \sigma_m|f_m|). \tag{7.5}$$

Computing W_1 and W_2:

In the first step of the proposed design procedure described above, two stable and minimum phase diagonal transfer matrices Π, Γ were chosen, reflecting the desired conditioning and singular values of W_2 and W_1, respectively. This section describes how to compute stable and minimum phase non-diagonal weights W_1 and W_2 so that the singular values of the weighted plant are close to $\Sigma_{G_s}^{des}$, i.e.

$$|\sigma_i(W_2GW_1) - \sigma_i|f_i|| \leq \beta, \quad \forall i = 1, \ldots, m, \tag{7.6}$$

where β is a function of frequency and $\beta > 0$, $\forall\omega$.

1. Choose a frequency range $[\omega_L, \omega_H]$ that contains the dynamics of G around the target closed-loop bandwidth and grid the range. At each frequency of the grid, perform a singular value decomposition of $G(j\omega)$
$$G(j\omega) = U_G(j\omega)\Sigma_G(j\omega)V(j\omega)^*, \tag{7.7}$$
such that the ordering of the singular values in Σ_G is uniform across the considered frequency range. This means that each singular value as a function of frequency varies smoothly with frequency. This ensures "continuous variation" of the elements of V across $[\omega_L, \omega_H]$ when the all-pass factors up to which V is unique are carefully selected.

   By "continuous variation" it is meant continuous variation when the grid is infinitely dense. It is assumed that all the dynamics of U_G and V are contained in the frequency range $[\omega_L, \omega_H]$; this can be easily verified. If not, the designer must increase the frequency range.

2. Fit a transfer function to each element V_{jl} of V $(j, k = 1, \ldots, m)$, without restricting the transfer functions to be stable and minimum phase. Denote the resulting transfer matrix by $\hat{V}$. The better the fit, the smaller the value of δ
$$\overline{\sigma}\big[\hat{V}(j\omega) - V(j\omega)\big] \leq \delta, \quad \forall\,\omega. \tag{7.8}$$
The above inequality ensures that $\hat{V}$ is close to V outside the frequency range $[\omega_L, \omega_H]$.

3. $W = \Pi\,\Gamma$, with $\Pi = \text{diag}(\pi_1, \ldots, \pi_m)$ and $\Gamma = \text{diag}(\gamma_1, \ldots, \gamma_m)$. If δ is small $\sigma_i(G\hat{V}\Gamma) \simeq \sigma_i|\gamma_i|$, but $\hat{V}\Gamma$ is not necessarily minimum phase and stable. Obtain a co-spectral factorisation of $\Psi = \hat{V}\Gamma(\hat{V}\Gamma)^{\sim}$:

$$\hat{V}\Gamma(\hat{V}\Gamma)^{\sim} = W_1 W_1^{\sim}, \tag{7.9}$$

where W_1 denotes the co-spectral factor of Ψ, $W_1^{\sim}(s) = W_1^T(-s)$ and $W_1, W_1^{-1} \in \mathcal{RH}_\infty$. The co-spectral factorisation can be obtained using the formulae given in Francis, 1987. Thus, W_1 is stable and minimum phase and $\sigma_i(GW_1) = \sigma_i(G\hat{V}\Gamma)$. Consequently, W_1 is the required weighting filter.

4. To compute W_2, perform an SVD of GW_1 at each grid point:

$$GW_1(j\omega) = U(j\omega)\Sigma_{GW_1}(j\omega)V_{GW_1}(j\omega)^*, \tag{7.10}$$

and fit transfer functions to the elements of U. Denote the resulting transfer matrix by $\hat{U}$. Then take a spectral factorisation of $U\Pi^{\sim}\Pi U^{\sim}$ to obtain the weighting filter W_2. Note that step 1 has to be repeated for GW_1, as $U(j\omega) \neq U_G(j\omega)$ because $\hat{V}$ is only an approximation of V.

As indicated in step 1, the designer must ensure that the elements of V (and U in step 4) vary "continuously" across the frequency range $[\omega_L, \omega_H]$, so that the fit obtained in step 2 is sensible. This requires that the ordering of the singular values is uniform across the frequency range considered in combination with careful selection frequency by frequency of the all-pass factors up to which V is unique. The `MATLAB®` `svd` tool arranges the singular values in decreasing order and, consequently, it cannot be used directly. However, `svd` can be modified to guarantee uniform ordering of the singular values across $[\omega_L, \omega_H]$ by looking at the singular vectors frequency by frequency. If $G(j\omega)$ has distinct singular values, then V is determined up to a right diagonal all-pass factor $\Phi = \text{diag}(e^{j\theta_1}, \ldots, e^{j\theta_m})$ with all $\theta_i \in \mathbb{R}$ (Horn & Johnson, 2012). If the designer ensures that the ordering of the singular values of G is uniform across the grid, and solves the 2-norm minimisation:

$$\alpha_{k,i}^{opt} = \arg\min_{\alpha} \|v_{\omega_k,i} - v_{\omega_{k+1},i}e^{j\alpha}\|_2 \tag{7.11}$$

at each pair of grid frequencies (ω_k, ω_{k+1}), where $v_{\omega_k,i}$ denotes the i^{th} column of $V(j\omega_k)$, then post-multiplying $V(j\omega_{k+1})$ by the all-pass factor $e^{jdiag(\alpha_{k,1}^{opt},\ldots,\alpha_{k,m}^{opt})}$ guarantees "continuous variation" of the elements of V. It is straightforward to show that minimization 7.11 has an analytic solution:

$$\alpha_{k,i}^{\text{opt}} = -\angle(v_{\omega_k,i}^* \cdot v_{\omega_{k+1},i}). \tag{7.12}$$

According to G. Papageorgiou, 1998, the procedure can be used in two different ways:

1. The designer directly designs non-diagonal weights that shape the singular values of the scaled open-loop plant. However, since the weights are fully populated, it is difficult to intuitively reduce the actuator usage. If the weights were diagonal instead, one could directly adjust the filter associated with the actuator with the large usage. Designing non-diagonal weights directly might require the engineer to modify the input scaling or change the shape of one or more of the singular values of the weighted plant. Thus, it removes some of the designer's insight and makes the tuning process less transparent.
2. The designer adds the proposed designed procedure after initially designing diagonal weights W_1 and W_2. The first step is to assess if the diagonal weights meet the control system design specifications.

In case they do not, the non-diagonal weights are designed using the procedure for the weighted plant W_2GW_1. According to Papageorgiou, this is a more intuitive approach as it allows for better trade-off robustness properties at the plant input and output. For instance, if robustness to input multiplicative uncertainty is poor and W_1 is ill-conditioned, then the condition number of W_1 - a diagonal weight - must be decreased and a non-diagonal W_2^+ designed to maintain similar nominal tracking performance.

More details about the procedure and the guarantees associated with it can be found in the original work of G. Papageorgiou, 1998.

7.2.1. Design of $\mathcal{H}_\infty$ Loop-Shaping Controller using Non-Diagonal Filters

This section is devoted to the use of the procedure described above to design a $\mathcal{H}_\infty$ Loop-Shaping controller using non-diagonal filters for a lateral-directional controller for the X-29 aircraft.

The first step is to choose the desired shape for the open loop. The most important decision is the choice for the crossover frequency w_c, which was chosen to be compatible with the reference models for the roll-rate and yaw-rate response. The reference model for the roll-rate response and for the yaw-rate response was set as second-order models due to the inherent second-order response of the closed-loop system with a PI and the $\frac{1}{s}$ closed-loop relationship between v and y from the dynamic inversion.

$$T_{ref_p} = \frac{w_n^2}{s^2 + 2\xi w_n s + w_n^2} = \frac{8^2}{s^2 + (2)(0.707)(5)s + 8^2} \tag{7.13}$$

$$T_{ref_r} = \frac{w_n^2}{s^2 + 2\xi w_n s + w_n^2} = \frac{5^2}{s^2 + (2)(0.707)(5)s + 5^2} \tag{7.14}$$

Based on these reference models, the desired crossover was set to lie between 6 and 10 rad/s. In order to leverage the procedure for the design of non-diagonal weights, the plant was first shaped by diagonal weights. In this preliminary shaping, W_1 was set as a diagonal block of PIs and W_2 was set as a diagonal block of first-order low-pass filters:

$$W_{1_{diag}} = \begin{bmatrix} \frac{K_{p_1}(s+K_{i_1}/K_{p_1})}{s} & 0 \\ 0 & \frac{K_{p_2}(s+K_{i_2}/K_{p_2})}{s} \end{bmatrix} = \begin{bmatrix} \frac{40(s+4)}{s} & 0 \\ 0 & \frac{40(s+4)}{s} \end{bmatrix} \tag{7.15}$$

$$W_{2_{diag}} = \begin{bmatrix} \frac{\omega_{lp_1}}{s+\omega_{lp_1}} & 0 \\ 0 & \frac{\omega_{lp_2}}{s+\omega_{lp_2}} \end{bmatrix} = \begin{bmatrix} \frac{45}{s+45} & 0 \\ 0 & \frac{45}{s+45} \end{bmatrix} \tag{7.16}$$

As can be seen in Figure 7.2, shaping the open-loop plant G_{ol} with the diagonal filters $W_{1_{diag}}$ and $W_{2_{diag}}$ did not produce the desired shape G_s. The frequency at which the highest and the lowest singular value cross the 0 dB line is starkly different, and the lowest singular value exhibits insufficient low-frequency gain and an undesirable flat response in the interval $\omega \in [0.1\ 1]$ rad/s. With this weighting filter structure, it was difficult to avoid these issues, even after iteratively adjusting the filter gains on a trial-and-error basis.

However, it suffices as an initial shape of the OL response. The procedure described above is then employed to design an additional non-diagonal W_1^+, such that the shaped plant exhibits the desired open-loop frequency response, where W_1 is given by:

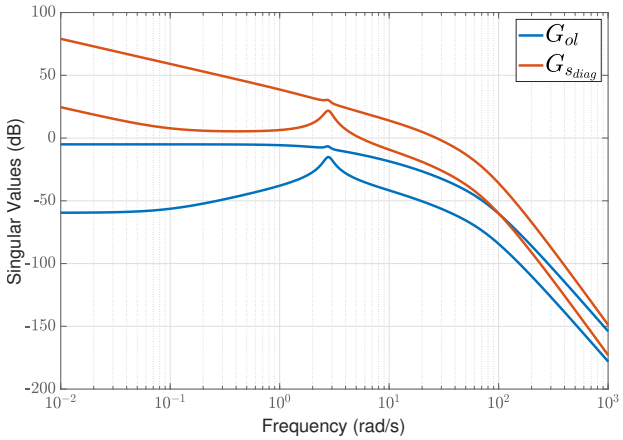

Figure 7.2: SVD of the open-loop plant G_{ol} and of the initially shaped plant $G_{s_{diag}}$ with diagonal weights.

$$W_1(s) = W_{1_{diag}}(s)W_1^+(s) \tag{7.17}$$

The first step of the procedure is to select individual transfer functions $f_i(s)$ that shape the OL-response $\sigma_i(G_s) = \sigma_i f_i$ of the plant to behave as $G_{s_{des}}$. Given that the plant was initially shaped by diagonal filters, the "new" open-loop plant is $G_{s_{diag}} = W_{2_{diag}} G_{ol} W_{1_{diag}}$. In order to address the issues of this initially shaped plant, the individual singular value lines are isolated, and then the individual transfer functions $f_i(s)$ are selected such that $\sigma_i f_i$ results in the desired shaping. Based on the SV of $G_{s_{diag}}$ in Figure 7.3, f_i was set as:

$$f_i = \begin{cases} f_1 = 0.2 \\ f_2 = \dfrac{s+4}{s+0.1} \end{cases} \tag{7.18}$$

which results in the desired SV lines illustrated in Figure 7.4.

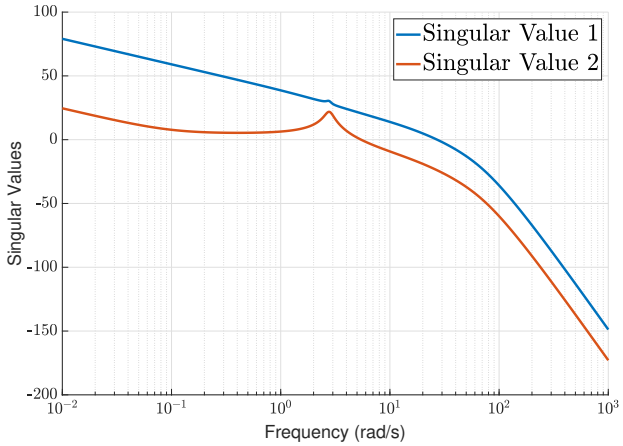

Figure 7.3: Isolating the continuous Singular Values of the initial shaped plant $G_{s_{diag}}$.

As demonstrated in Figure 7.4, it is different to multiply the plant G_{ol} with diagonal blocks containing the individual transfer functions f_i in its diagonals than it is to individually adjust the singular value lines with f_i. The goal of the procedure is now to compute a non-diagonal W_1^+ such that $G_{s_{initial}} W_1^+$ has an

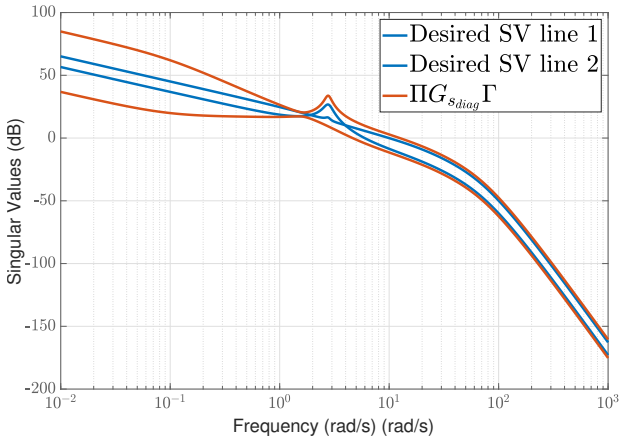

Figure 7.4: Desired Singular Values and difference to simply multiplying the plant G_{ol} with Π and Γ.

SVD equal to the desired SV lines depicted in blue in Figure 7.4.

Providing Π and Γ to the algorithm replicating the procedure in MATLAB®, the outcome is the non-diagonal filter W_1^+. Figure 7.4 illustrates the open-loop response of $W_{2_{diag}} G_{ol} W_{1_{diag}} W_1^+$. As can be seen in Figure 7.5, the open-loop response successfully matches the specified desired shaping.

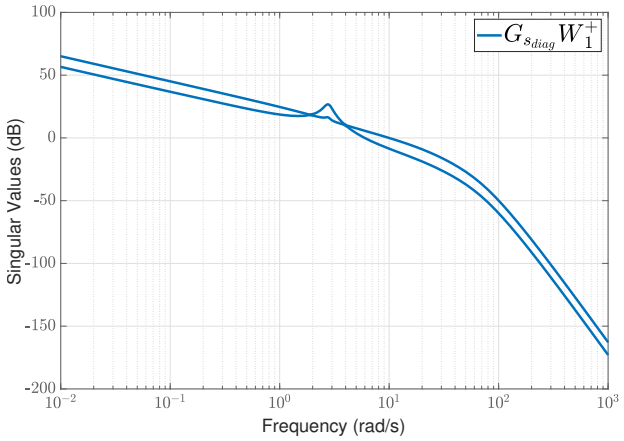

Figure 7.5: SVD of $G_{s_{diag}} W_1^+$ matching with the desired SV.

However, W_1^+ is typically a very high-order weighting filter (order 85 in this case) as the order depends on the order used to fit $\hat{V}(jw)$ and the higher the order, the better the fit. Since it is impractical to implement such a high-order weighting filter, $W_1^+(s)$ is model-order reduced a posteriori.

Using balanced truncation techniques, the minimum order $W_1^+(s)$ could be reduced while still yielding a good approximation in terms of $\sigma[G_{s_{diag}} W_{1_{red}}^+] \approx \sigma[G_{s_{diag}} W_1^+]$ was order 3, resulting in the following non-diagonal weighting filter state-space description:

$$W_{1_{red}}^{+} = \left[\begin{array}{c|c} A_k & B_k \\ \hline C_k & D_k \end{array}\right] = \left[\begin{array}{ccc|cc} -0.1388 & -0.3089 & 0.2352 & 0.007016 & -0.01504 \\ 0.1515 & 0.4096 & 1.837 & -0.0007371 & 0.002762 \\ -0.6941 & -4.288 & -0.8005 & -0.009274 & 0.01728 \\ \hline 14.33 & 28.99 & -20.28 & 0.2551 & -0.07017 \\ -210.7 & 14.74 & 6.961 & -0.4024 & 0.8896 \end{array}\right] \tag{7.19}$$

or, alternatively, the following transfer function block:

$$W_{1_{red}}^{+}(s) = \left[\begin{array}{cc} \frac{0.25512(s+0.5779)(s^2+s+5.323)}{(s+0.1)(s^2+0.4296s+7.771)} & \frac{-0.070165(s+8.499)(s^2-1.041s+2.201)}{(s+0.1)(s^2+0.4296s+7.771)} \\ \frac{-0.40242(s+4.157)(s^2+0.2329s+7.856)}{(s+0.1)(s^2+0.4296s+7.771)} & \frac{0.88959(s+4.011)(s^2+0.2608s+7.883)}{(s+0.1)(s^2+0.4296s+7.771)} \end{array}\right] \tag{7.20}$$

The SVD plot of the full-order and the reduced-order $W_1^+(s)$ in Figure 7.6 illustrates that they match accurately until around 1000 rad/s. The fact that they differ beyond this frequency is not problematic, given that it is a frequency well above the available system bandwidth. To further validate the reduced-order $W_1^+(s)$, the open-loop response of the plant shaped with the full-order and the reduced-order is compared in Figure 7.7. As can be seen, they are identical, which attests to the validity of the reduced-order model.

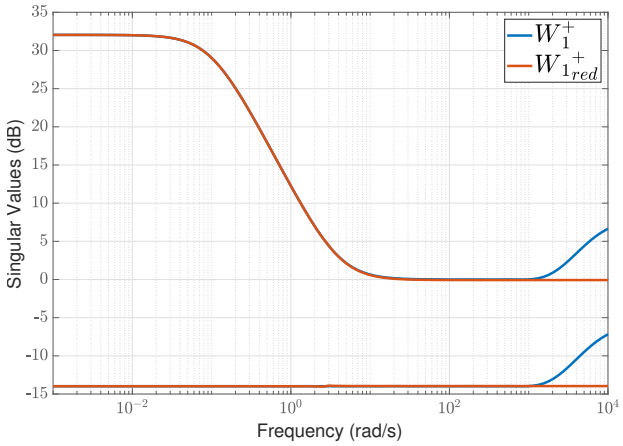

Figure 7.6: SVD of the full-order W_1^+ resulting from the procedure and a model-reduced $W_{1_{red}}^+$.

It has been demonstrated that the procedure was successfully understood and recreated to design a non-diagonal additional weighting filter W_1^+ to further improve the shaping of a previously shaped plant G_{ol} with diagonal filters $W_{1_{diag}}$ and $W_{2_{diag}}$. What remains to be seen is whether the final shaped plant is compatible with a robust control design using the $\mathcal{H}_\infty$ Loop-Shaping.

The obtained suboptimal $\gamma > \gamma_{min}$ associated with the shaped plant G_s, computed with a tolerance of 1e-3, was 2.5401. This value lies well within the practical guideline that states that γ should be below 4 (Hyde, 1995). The full-order $\mathcal{H}_\infty$ controller K_∞ was then computed by solving the 2 AREs. The resulting broken-loop response at the plant input L_i and at the plant output L_o of the controlled system is displayed in Figure 7.8.

The order of the open-loop plant together with the order of the weighting filters results in a full-order $\mathcal{H}_\infty$ controller of significant order, given by order$(K) = n_G + 2n_1 + 2n_2$; in this case, $n_G = 11$ (order 3

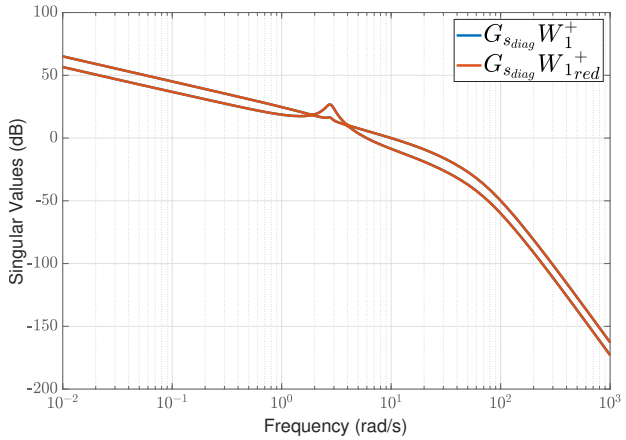

Figure 7.7: Comparison of the SVD of $G_{s_{diag}}W_1^+$ and $G_{s_{diag}}W_{1_{red}}^+$.

lateral-directional dynamics and order 8 actuator model), $n_1 = 5$ (order 3 non-diagonal weighting $W_{1_{red}}^+$ and order 2 $W_{1_{diag}}$) and $n_2 = 2$ (order 2 $W_{2_{diag}}$), resulting in the order of the controller being 25. Given that in this section, the main goal is not the practical implementation of the controller, its value will not be presented. The reader can, however, use the values presented above for G_{ol}, $W_{1_{diag}}$, $W_{2_{diag}}$ and $W_{1_{red}}^+$ to recreate it.

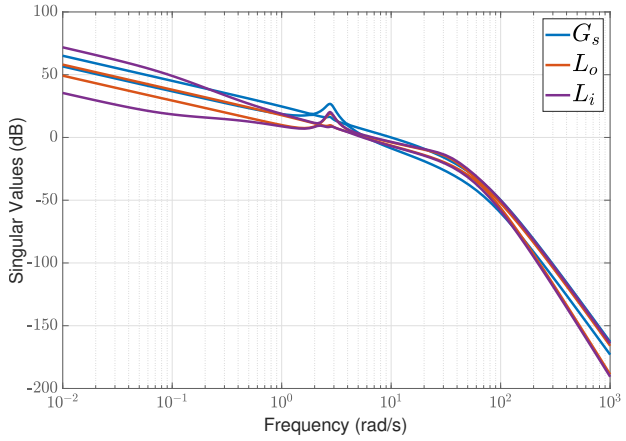

Figure 7.8: Broken-Loop response at the plant I/O of the full-order $\mathcal{H}_\infty$ Loop-Shaping controller associated with the shaped-plant $G_{s_{diag}}W_{1_{red}}^+$.

The results demonstrate that the broken-loop responses L_o and L_i closely match those of G_s, a key indicator of a successful and robust $\mathcal{H}_\infty$ Loop-Shaping Design Procedure (LSDP). This alignment reflects the effectiveness of the method, particularly given the small γ value achieved. Furthermore, these findings underscore not only the robustness of the approach but also its capability to simplify the design process for weighting filters, making it more systematic and accessible.

While NDI-like control laws philosophies rely on model and sensor information to invert the dynamics in a closed-loop fashion, this procedure for the design of non-diagonal filters makes use of the plant's model information as a way to modify the singular values in an open-loop fashion, since the weights contain an internal model of the singular vector matrices $V(jw)$ and $U(jw)$. Nevertheless, what is interesting about the non-diagonal filters' procedure is that it connects more directly with the robust control theory since robustness properties lie at the plant I/O.

7.3. Tuning a Hybrid IDI Controller Using the LSDP with Non-Diagonal Filters

Essentially, NDI-like architectures aim to decouple the system dynamics via inversion so that channels can be treated individually. In the case of the Hybrid IDI, the tunable elements which can be used to address robustness are the virtual controller and the scaled complementary filter in the inversion. Both these elements are diagonal blocks whose design philosophy often relies on treating the channels separately. However, the decoupling of the dynamics is achieved in terms of the closed-loop relationship between the virtual control signal v and the output y, and not in terms of the open-loop gain at the plant input and output. More precisely, even if the closed-loop decoupling makes it possible to design the virtual controller for the roll-rate p channel and the yaw-rate r channel separately, each controller will affect the singular values of the open-loop gain at the plant input and output, and hence affect robustness. This implies that even if the diagonal entries of the virtual control law controller are able to manipulate the decoupled closed-loop relationships between y_{cmd} and y, they might be insufficient to provide the desired open-loop response at the plant I/O. When dealing with coupled systems, such as lateral-directional flight dynamics, the concept of directionality associated with the input and output singular vectors plays an important role.

The non-diagonal filters resulting from the procedure contain an internal model of $\hat{V}(jw)$ and $\hat{U}(jw)$, which allows it to 'invert' the dynamics in an open-loop sense. For the Hybrid IDI to be able to replicate the $Ks = W_1 K_\infty W_2$, it would be necessary that the tunable elements of the Hybrid IDI controller can be tuned such that $K_{H-IDI} \to K_s$. As a result, choosing the virtual control law to be a diagonal block of PIs and the scaled complementary filter parameters fully decoupled, such that:

$$K_v(s) = \begin{bmatrix} K_{p_p} + \frac{K_{i_p}}{s} & 0 \\ 0 & K_{p_r} + \frac{K_{i_r}}{s} \end{bmatrix} \tag{7.21}$$

$$K_c = \begin{bmatrix} K_{c_p} & 0 \\ 0 & K_{c_r} \end{bmatrix} \tag{7.22}$$

$$H_c(s) = \begin{bmatrix} H_{c_p}(s) & 0 \\ 0 & H_{c_r}(s) \end{bmatrix} = \begin{bmatrix} \frac{w_{lp_p}}{s+w_{lp_p}} & 0 \\ 0 & \frac{w_{lp_r}}{s+w_{lp_r}} \end{bmatrix} \tag{7.23}$$

will impose structural constraints that may limit the possibility for K_{H-IDI} to be configured to mimic K_s and yield a satisfactory γ, that is, a robust design.

In order to assess this, the weighting filters W_1 and W_2 defined in Section 7.2.1 were used together with the Hybrid IDI structure displayed in Figure 5.15 to tune its elements using the $\mathcal{H}_\infty$-synthesis in `systune`.

Since the the controlled variables are roll-rate p and yaw-rate r, that is $y = \begin{bmatrix} p \\ r \end{bmatrix}$, and the inversion requires full-state feedback, a β estimator is introduced using the model information:

$$\hat{\beta}(s) = \frac{1}{s - Y_\beta/V_o} \begin{bmatrix} Y_p/V_o & Y_r/V_o - 1 \end{bmatrix} \underbrace{\begin{bmatrix} p \\ r \end{bmatrix}}_{y} \tag{7.24}$$

The estimated states $\hat{x}$ can now be defined as:

$$\hat{x}(s) = \begin{bmatrix} \hat{\beta}(s) & p(s) & r(s) \end{bmatrix}^T = F(s)\, y(s) = \begin{bmatrix} \frac{Y_p/V_o}{s - Y_\beta/V_o} & \frac{Y_r/V_o - 1}{s - Y_\beta/V_o} \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} p(s) \\ r(s) \end{bmatrix} \tag{7.25}$$

Consequently, the gain on the inversion path associated with model-based inversion ξ_{MB} can be defined as:

$$\xi_{MB}(s) = C\, A\, F(s) = \underbrace{\begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{C} \underbrace{\begin{bmatrix} Y_\beta/V_o & Y_p/V_o & Y_r/V_o - 1 \\ L_\beta & L_p & L_r \\ N_\beta & N_p & N_r \end{bmatrix}}_{A} \underbrace{\begin{bmatrix} \frac{Y_p/V_o}{s - Y_\beta/V_o} & \frac{Y_r/V_o - 1}{s - Y_\beta/V_o} \\ 1 & 0 \\ 0 & 1 \end{bmatrix}}_{F(s)} \tag{7.26}$$

while the error signal compensation e^*_ξ can be defined as:

$$e^*_\xi(s) = \begin{bmatrix} e^*_{\xi_p} \\ e^*_{\xi_r} \end{bmatrix} = K_c H_c(s) e_\xi(s) = \underbrace{\begin{bmatrix} K_{c_p} & 0 \\ 0 & K_{c_r} \end{bmatrix}}_{K_c} \underbrace{\begin{bmatrix} H_{c_p}(s) & 0 \\ 0 & H_{c_r}(s) \end{bmatrix}}_{H_c(s)} \underbrace{\left(\begin{bmatrix} s & 0 \\ 0 & s \end{bmatrix} y - CC - \xi_{MB}(s) \right)}_{e_\xi(s)} \tag{7.27}$$

where $CC \approx CBu$.

Table 7.2: Optimisation gains of the Hybrid IDI lateral-directional controller using the LSDP with non-diagonal W_1.

	Optimization		Virtual Control						Inversion Loop			
Parameter	γ_{th}	γ	K_{p_p}	K_{p_r}	K_{i_p}	K_{i_r}	w_{lp_p}	w_{lp_r}	K_{c_p}	K_{c_r}	$w_{H_{c_p}}$	$w_{H_{c_r}}$
Value	2.5401	5.5796	5.7313	3.9807	26.7594	20.6121	45.00	45.00	0.8639	0.9959	9.0438	1.357

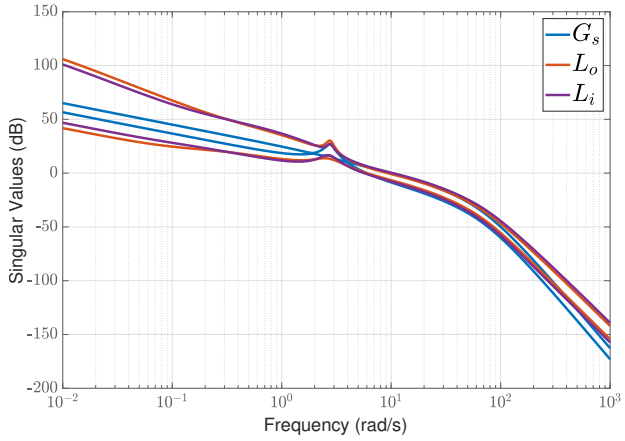

Figure 7.9: Broken-Loop response at the plant I/O of the Hybrid IDI controller tuned using the $\mathcal{H}_\infty$ LSDP with non-diagonal weights.

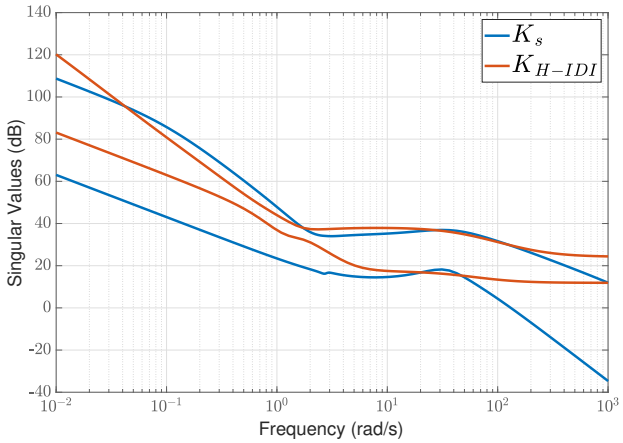

Figure 7.10: Comparison of the SVD of the full-order $\mathcal{H}_\infty$ controller K_s and the Hybrid IDI K_{H-IDI}.

The results of the optimisation are presented in Table 7.2. As can be seen, the obtained γ value is considerably above γ_{th} and is higher than the maximum recommended value of 4, which indicates poor robustness. The fact that W_1 requires significant integral action (W_1 is comprised of $W_{1_{diag}}$ which has an integrator and W_1^+ which further boosts the low-frequency gain of the shaped plant G_s) justifies the close to 1 values for K_c. In other words, given the need to achieve such high gain at low-frequencies the optimisation requires not only the integral action in the virtual control law but also maximising the incremental nature of the inversion scheme. This can be seen in Figure 7.10, where the $\mathcal{H}_\infty$ LS controller K_s has a higher slope (in magnitude) than -20dB between 0.01 and 1 rad/s and so, as K_{H-IDI} is tuned so that $K_{H-IDI} \rightarrow K_s$, it results in double integral action at low frequencies. In doing so, the H_c filter bandwidth is made sufficiently low in order to ensure adequate roll-off at high frequencies. However, a consequence of the relatively small bandwidth of H_c is that the stability margins around crossover are degraded, which helps justify the large γ value obtained (Pollack, 2024).

If the broken-loop response at the plant I/O is analysed in Figure 7.9, the higher gain at low frequencies compared to the full-order $\mathcal{H}_\infty$ solution for the same shaped plant G_s (recall Figure 7.8) is confirmed. Besides, compared to the $\mathcal{H}_\infty$ solution, the Hybrid IDI controller seems to result in a more aggressive transition around the crossover region, i.e. a steeper slope around crossover (compare L_o and L_i from Figure 7.8 and Figure 7.9). As mentioned above, this can be attributed to the relatively small bandwidth of H_c.

Overall, it is demonstrated that if the structure of the Hybrid IDI is not compatible with $K_s = W_1 K_\infty W_2$, then tuning a Hybrid IDI using the $\mathcal{H}_\infty$ LSDP is unfruitful. This limitation is not only attributed to the procedure but also to the $\mathcal{H}_\infty$ LSDP itself. The fact that the $\mathcal{H}_\infty$ LSDP has two separate steps, the design of weighting filters and the synthesis of the robust controller K_∞, makes it challenging to impose a specific a priori structure to K_s. As the designer increases the complexity of the weighting filters to achieve a better shape of the open-loop plant, so does the complexity of K_s. A possible solution is to increase the complexity of the Hybrid IDI controller, namely, increasing the order of the virtual controller K_v or by making K_c, H_c and K_v full-blocks instead of diagonal. In this case, by making the structure of the Hybrid IDI more 'rich', it increases the compatibility with $K_s = W_1 K_\infty W_2$. This, however, requires more research in future studies.

7.4. Tuning a Hybrid IDI Controller Using the LSDP with Compatible Diagonal Filters

As was demonstrated in the previous section, having compatible filters with the structure of Hybrid IDI is extremely important. This is due to the fact that the methodology of tuning a Hybrid IDI with the $\mathcal{H}_\infty$ LSDP has its foundation on the premise that the structured controller K_{H-IDI} has a compatible structure to $W_1 K_\infty W_2$ so that the tunable elements of K_{H-IDI} can be configured such that

$$K_{H-IDI} \rightarrow W_1 K_\infty W_2 \tag{7.28}$$

Thus, only diagonal W_1 consisting of PIs was chosen, and, for simplicity, W_2 was set to identity. In general, W_2 is chosen to be a low-pass filter (G. Papageorgiou, 1998), like it was done in section 7.2.1. However, since in this procedure, it is required that W_1 and W_2 be invertible, it often involves adding "dummy zeros", i.e. minimum-phase zeros at very high frequency, such that the blocks become invertible. However, it was found that these zeros on W_2 degraded the performance of the $\mathcal{H}_\infty$-synthesis, and so, it was decided to simply set W_2 to identity. This is one aspect that deserves more careful consideration in future research.

7.4.1. Choosing W_1

Unlike the SISO case addressed in Chapter 6, the choice of W_1 in a MIMO setup is particularly challenging. Given that the lateral dynamics are coupled, successfully following the guidelines proposed by Hyde, 1995 (which addressed a longitudinal model with dominant diagonal terms) proved to be challenging. Similar guidelines were followed to the ones in 7.2.1 to shape the open-loop plant, but the choice of W_1 involved a lot of trial and error and did not provide a way to accurately adjust the crossover frequency. The final choice for W_1 was:

$$W_1 = \begin{bmatrix} \frac{K_{p_1}(s+K_{i_1}/K_{p_1})}{s} & 0 \\ 0 & \frac{K_{p_2}(s+K_{i_2}/K_{p_2})}{s} \end{bmatrix} = \begin{bmatrix} \frac{10(s+7)}{s} & 0 \\ 0 & \frac{80(s+3)}{s} \end{bmatrix} \tag{7.29}$$

which resulted in the following shaped plant:

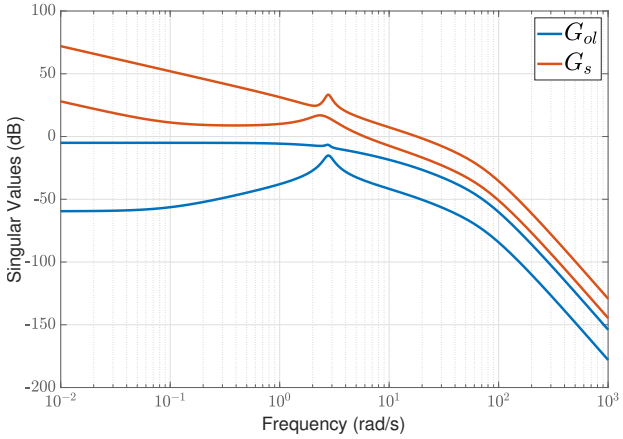

Figure 7.11: SVD of the shaped-plant G_s using diagonal weights.

From the resulting shaped plant, a few remarks can be made. The first one is that the lowest singular value $\underline{\sigma}$ does not have sufficiently high gain at frequencies between 10^{-1} and 2 rad/s. In order to address this, a low-frequency pole could be added to boost the gain at these frequencies, and another zero around crossover to compensate for the increase in the slope introduced by the additional pole. However, given the fixed structure of the Hybrid IDI, adding a low-frequency pole to W_1 (besides the already existing integrator pole) would essentially require the inversion compensation scheme to use as much incremental action as possible, i.e., setting K_c = 1, limiting the effect the inversion scheme could provide in addressing robustness, as was demonstrated in the previous section. Besides, it is not clear what element of this Hybrid IDI could mimic the effect of the second zero introduced in W_1, when most tunable elements are already constrained. In order to evaluate the compatibility of the shaped plant G_s with a robust control design, the suboptimal $\gamma_{th} < \gamma_{min}$ was computed with a tolerance of 1e-3. The resulting $\gamma_{th} = 2.3742$ demonstrates that the shaped plant is indeed compatible with a robust design as dictated by the $\mathcal{H}_\infty$ Loop-Shaping theory.

Table 7.3: Optimisation gains of the 1DoF Hybrid IDI lateral-directional controller using the LSDP with diagonal W_1.

	Optimization		Virtual Control						Inversion Loop			
Parameter	γ_{th}	γ	K_{p_p}	K_{p_r}	K_{i_p}	K_{i_r}	w_{lp_p}	w_{lp_r}	K_{c_p}	K_{c_r}	$w_{H_{c_p}}$	$w_{H_{c_r}}$
Value	2.3742	2.7139	8.6098	4.1224	23.7803	7.5605	45.00	45.00	0.4860	0.1142	78.54	54.65

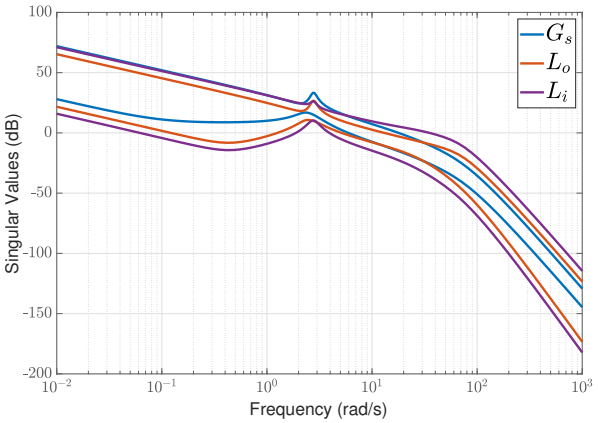

Figure 7.12: Broken-Loop response at the plant I/O of the full-order $\mathcal{H}_\infty$ controller K_s using diagonal weights.

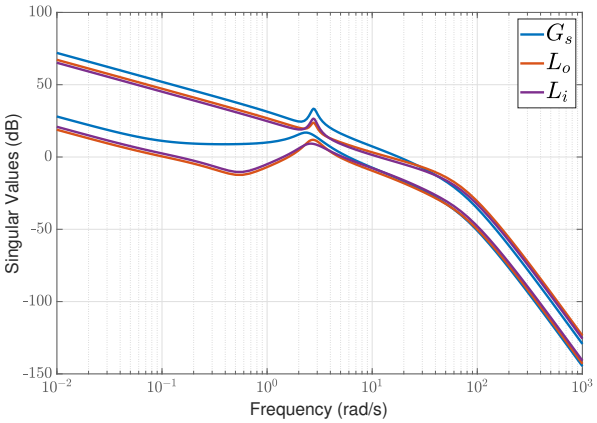

Figure 7.13: Broken-Loop response at the plant I/O of the 1DoF Hybrid IDI controller using diagonal weights.

However, analysing the broken-loop response at the plant I/O displayed in Figure 7.13, the low gain of the shaped plant G_s at low frequencies results in an even lower gain of the broken-loop gain. In fact, this even leads to the lowest singular values crossing the 0 dB line at around 10^{-1} rad/s, which anticipates a slow response of the yaw channel (the lowest singular value is primarily associated with the yaw channel).

To evaluate the preceding statement, the closed-loop response of the Hybrid IDI controller, tuned solely for robustness, is presented in Figure 7.14. This initial design excludes both model-following requirements and a feedforward element. The results confirm the anticipated significantly slow response of the yaw channel. While cross-coupling effects are reduced, they are not entirely eliminated due to inversion distortions. Furthermore, these effects are also influenced by the slow dynamics affecting the yaw channel.

The inclusion of a feedforward element and model-following performance requirements is expected to improve the closed-loop response of the controlled system. The $\mathcal{H}_\infty$-minimisation problem displayed in

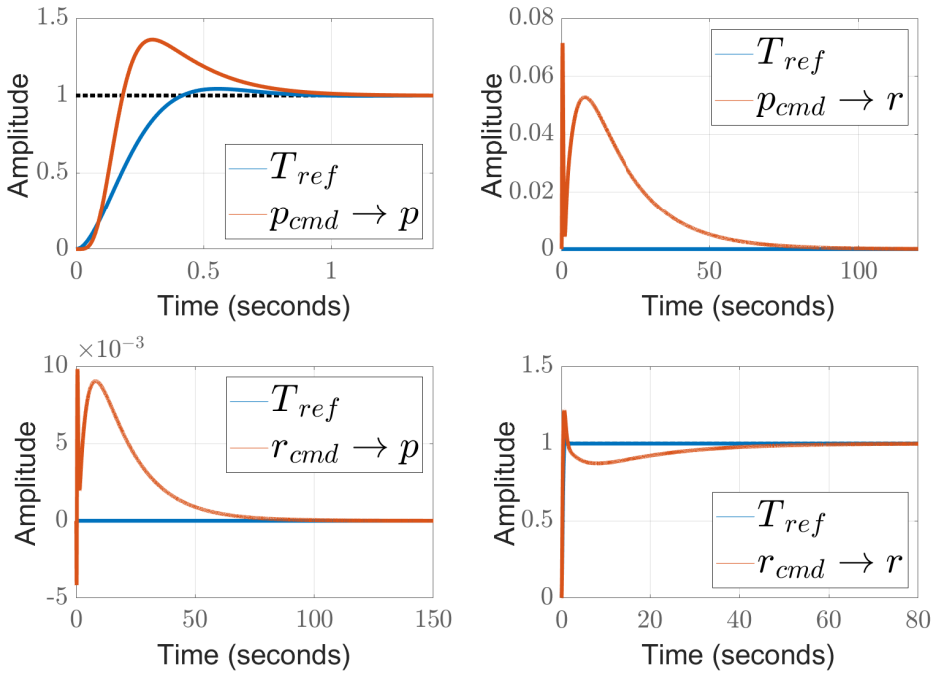

Figure 7.14: Model-following performance of the 1DoF Hybrid IDI controller using diagonal weights.

Figure 5.19 is used to tune the 2DoF Hybrid IDI. The results of the optimisation are displayed in Tables 7.4, 7.5 and 7.6.

Table 7.4: Filters setup and optimisation result of the Tuning Goals.

	Shaping Filters		NCF Rob. Soft Goal		Co-Design Soft Goals				MF Hard Goal	
Parameter	W_1	W_2	γ_{th}	γ	w_{e_p}	M_{e_p}	w_{e_r}	M_{e_r}	W_{mf_p}	W_{mf_r}
Value	$\begin{bmatrix} \frac{10(s+7)}{s} & 0 \\ 0 & \frac{80(s+3)}{s} \end{bmatrix}$	I_2	2.3181	3.1309	3.1309	3.1309	3.1309	3.1309	0.9999	1.0000

Table 7.5: Optimisation gains of the inversion and virtual control law of the Hybrid IDI using the $\mathcal{H}_\infty$ LSDP.

	Virtual Control						Inversion Loop			
Parameter	K_{p_p}	K_{p_r}	K_{i_p}	K_{i_r}	w_{lp_p}	w_{lp_r}	K_{c_p}	K_{c_r}	$w_{H_{c_p}}$	$w_{H_{c_r}}$
Value	6.9573	1.8360	27.3910	7.8324	45.00	45.00	0.5397	0.9872	78.54	4.988

Table 7.6: Optimisation gains of the feedforward element of the Hybrid IDI and of the model-following filter parameters using the $\mathcal{H}_\infty$ LSDP.

	Feedforward		MF Filter W_{mf}					
Parameter	K_{ff_p}	K_{ff_r}	A_{e_r}	A_{e_p}	w_{e_p}	w_{e_r}	M_{e_p}	M_{e_r}
Value	$\frac{1.8342(s^2+18.59s+143.7)}{(s+78.2)(s+3.364)}$	$\frac{1.1417(s^2+1.831s+7.563)}{s^2+4.378s+8.615}$	0.4111	0.2214	143.7295	79.8492	2.1916	2.1916

The obtained γ value adheres to the typical guidelines for $\mathcal{H}_\infty$ LSDP, and given all the model-following requirements, its deviation from γ_{th} is acceptable. The large K_{c_r} value can be explained by the model-following requirements requiring a larger open-loop gain at low frequencies. The relatively low value for the bandwidth of H_{c_r} is in accordance with what is expected from large (close to 1) K_c gains. The large

values for the bandwidth w_e of the model-following filters W_{mf_p} and W_{mf_r} suggest that desirable model-following performance was attained. This is confirmed by analysis of Figure 7.17, which demonstrates that the roll-rate and yaw-rate responses closely match the respective reference models. The performance improvement can be explained primarily by the action of the feedforward element, but the elevated broken-loop response at lower frequencies (compare Figures 7.13 and 7.15) also contributes to the latter fact. The cross-coupling effects are minimal, but undesired slow dynamics still appear in the $r_{cmd} \rightarrow p$ channel. In the tuning process, no consideration was given to these cross-coupling effects since the inversion scheme should inherently eliminate them. In reality, there will always be some residual inversion error, resulting in these cross-coupling effects, but the fact that they are quite reduced attests to the fact that they do not require further attention in the tuning phase.

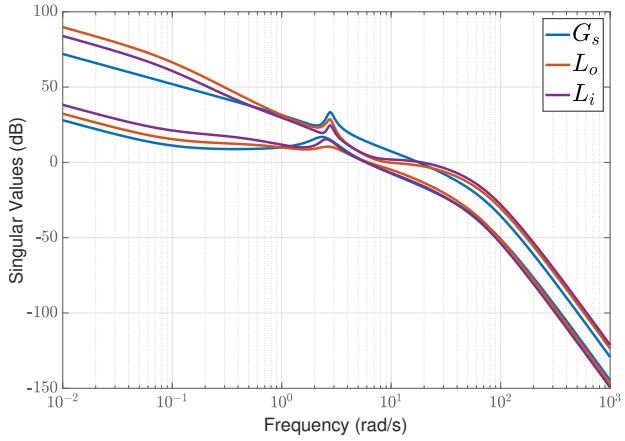

Figure 7.15: Broken-Loop response at the plant I/O of the 2DoF Hybrid IDI controller using diagonal weights.

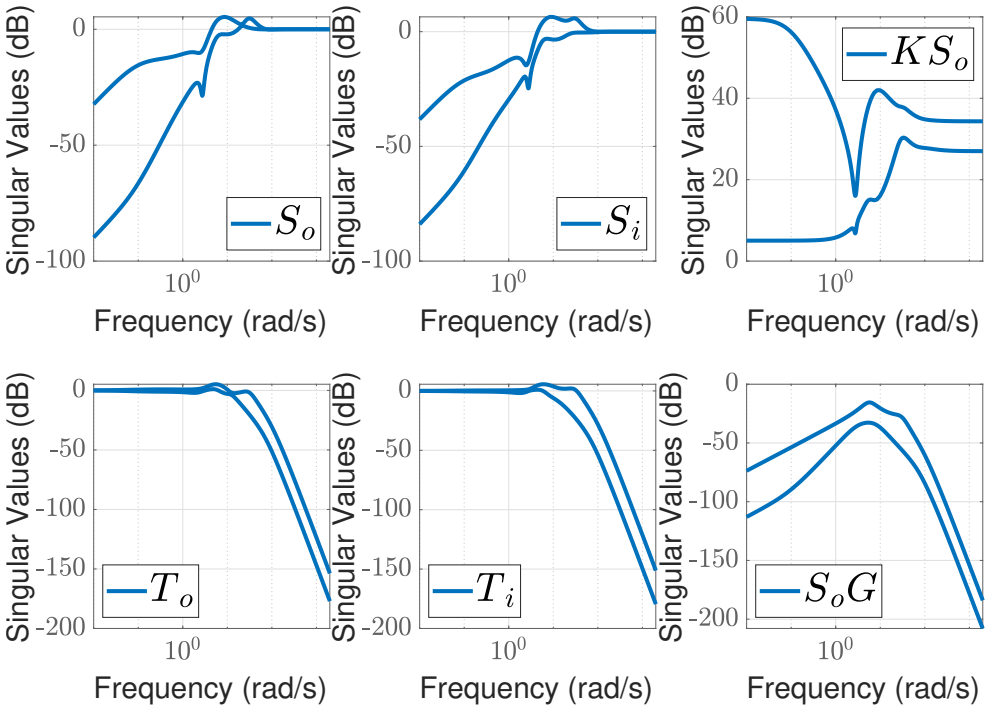

Figure 7.16: Gang-of-six closed-loop relationships of the 2DoF Hybrid IDI controller using diagonal weights.

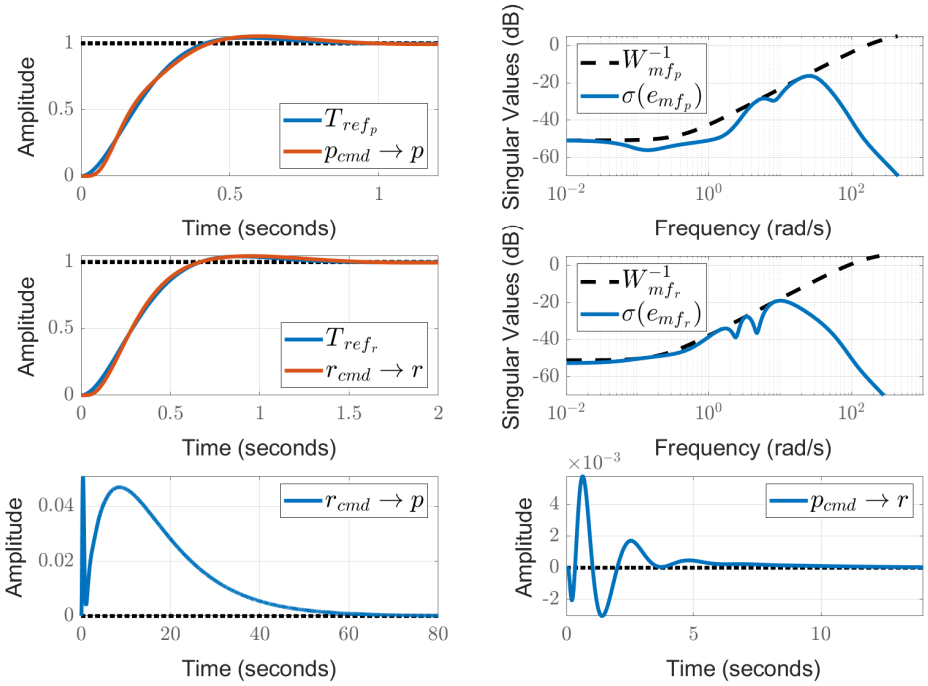

Figure 7.17: Model-following performance of the 2DoF Hybrid IDI controller using diagonal weights.

Analysing the "gang-of-six" closed-loop relationships in Figure 7.16 reveals that despite the apparent healthy shapes for most of them, KS_0 presents itself as extremely problematic. At low frequencies, it is required that $\overline{\sigma}[KS_0]$ is low (refer to the discussion in Section 3.3.2), and this essentially boils down to having large plant gain, i.e. sufficient control authority. However, what is observed is that $\overline{\sigma}[KS_0]$ is extremely high as a consequence of the lack of control authority on the yaw channel (observe the low-value for $\underline{\sigma}[G_{ol}]$ at low-frequencies in Figure 7.1). As a result, despite the apparent robustness of the system, the controller will produce large control signals, resulting in saturation of the actuators and possibly instability. To demonstrate this, linear time-domain simulations of pitch-rate and yaw-rate doublet manoeuvres are displayed in Figures 7.18, 7.19, 7.20 and 7.21.

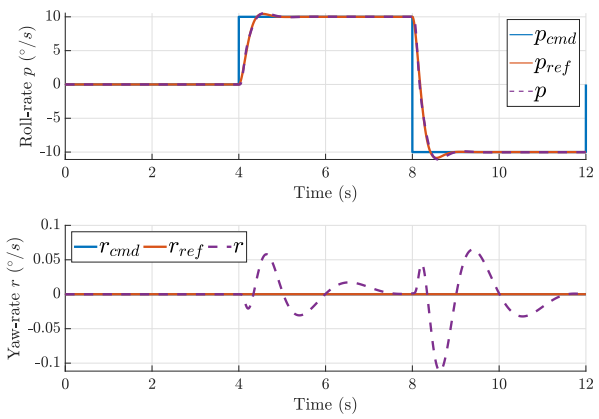

Figure 7.18: Time-domain simulation of the pitch-rate p and yaw-rate r response to a pitch-rate doublet manoeuvre.

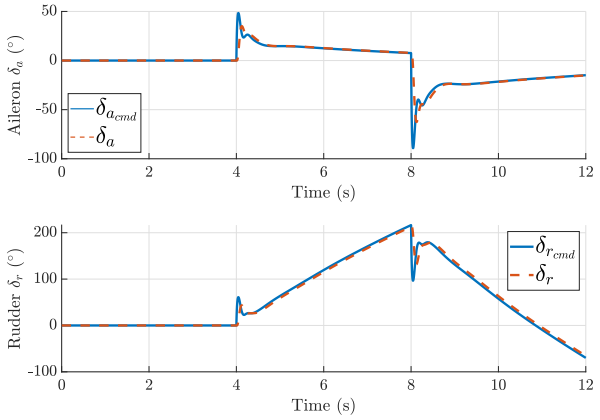

Figure 7.19: Time-domain simulation of the inputs δ_a and δ_r to a pitch-rate doublet manoeuvre.

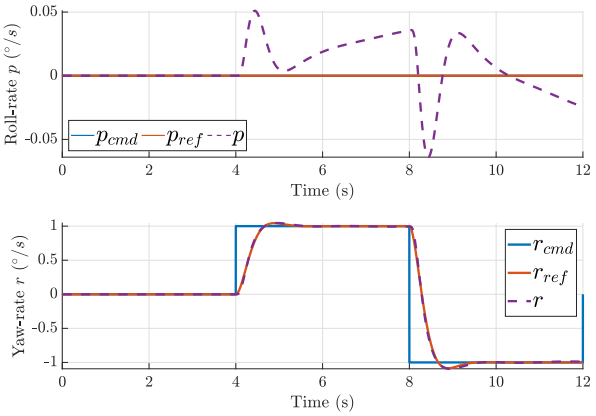

Figure 7.20: Time-domain simulation of the pitch-rate p and yaw-rate r response to a yaw-rate doublet manoeuvre.

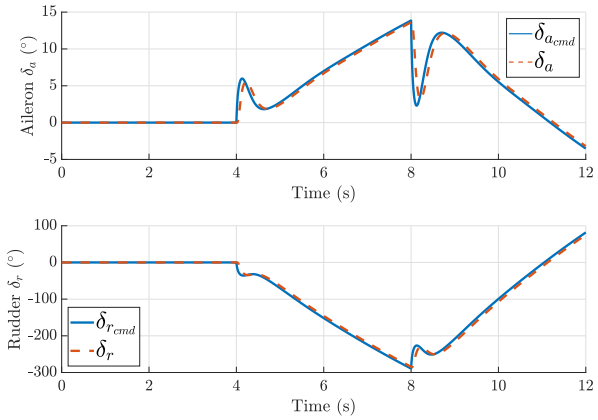

Figure 7.21: Time-domain simulation of the inputs δ_a and δ_r to a yaw-rate doublet manoeuvre.

From the time-domain simulations, it is possible to confirm the extremely large control signals, which are unrealistic in practical implementations. In Figure 7.18, it is possible to observe that the transient response on the yaw-rate channel r results in a deviation from the reference signal r_{ref}, which due to the integrator term in K_v and the incremental nature of the inversion, results in a continuously increasing desired acceleration v_r. This, associated with a plant with low control authority on the yaw-rate channel, will produce large rudder commands, as seen from the 4 to 8 seconds time interval in Figure 7.19. On the other hand, requiring tracking on the yaw-rate control channel will also produce extremely large rudder commands due to the low control authority on the yaw-rate channel, as seen in Figures 7.20 and 7.21.

The lack of control authority on the yaw-rate channel can be directly observed in the second row of the control effectiveness matrix CB:

$$CB = \begin{bmatrix} 1.3430 & 0.2345 \\ 0.0897 & -0.0710 \end{bmatrix} \tag{7.30}$$

which demonstrates that large control deflections are required to produce acceleration on the yaw-channel $\dot{r}$. Alternatively, if an SVD of the inverse of the control effectiveness matrix $(CB)^{-1}$ is considered, it is possible to infer that one singular value clearly dominates over the other, namely the one associated with the input $v_r - e^*_{\xi_r}$ and output δ_r, demonstrating, once more, the issues associated with the yaw-rate channel:

$$(CB)^{-1} = \begin{bmatrix} 0.6099 & 2.0154 \\ 0.7712 & -11.5421 \end{bmatrix} = U\Sigma V^T = \begin{bmatrix} -0.1686 & 0.9857 \\ 0.9857 & 0.1686 \end{bmatrix} \begin{bmatrix} 11.7351 & 0 \\ 0 & 0.7324 \end{bmatrix} \begin{bmatrix} 0.0560 & 0.9984 \\ -0.9984 & 0.0560 \end{bmatrix}^T \tag{7.31}$$

where (recall the definition of e^*_ξ from Equation 7.27):

$$\begin{bmatrix} \delta_{a_{cmd}} \\ \delta_{r_{cmd}} \end{bmatrix} = (CB)^{-1} \begin{bmatrix} v_p - e^*_{\xi_p} \\ v_r - e^*_{\xi_r} \end{bmatrix} \tag{7.32}$$

Saturation and rate limits of the actuators are aspects that are not considered in the presented linear $\mathcal{H}_\infty$ LSDP to tune a Hybrid IDI controller and, when dealing with plants with control authority deficiencies, can lead to unsuccessful control synthesis. If the time-domain simulations considered the actuator limits established in Bosworth, 1992, poor robustness and performance are expected. Once again, the lessons of Stein, 2003 on the importance of acknowledging the plant limitations for a successful control design task are demonstrated.

7.5. Discussion and Conclusion

The present chapter highlighted the potential of the use of the procedure presented in G. Papageorgiou, 1998 to systematically design non-diagonal weighting filters W_1 and W_2, which produce the desired shaping in the context of the $\mathcal{H}_\infty$ LSDP. The procedure was successfully recreated and used to design a $\mathcal{H}_\infty$ controller for the lateral-directional dynamics of the X-29 at a particular flight condition. Preliminary results of combining non-diagonal weights with the tuning procedure of the Hybrid IDI using the $\mathcal{H}_\infty$ LSDP did not yield satisfactory results.

Nevertheless, it has been demonstrated that the choice of the X-29 lateral-directional dynamics was an unfortunate choice due to the lack of control authority on the yaw channel. The low-plant gain at low frequencies required essentially double integrator action to adequately shape the open-loop plant, which, given the Hybrid IDI structure, imposed $K_{c_r} \approx 1$, resulting in robustness challenges, i.e. a large γ value.

The use of diagonal weights proved to be more compatible with the Hybrid IDI structure. The $\mathcal{H}_\infty$ LSDP tuning of the Hybrid IDI yielded an acceptable γ value, and the included model-following requirements were successfully met. Therefore, while the obtained γ value for the Hybrid IDI controller is higher than γ_{th} associated with the full-order $\mathcal{H}_\infty$ LS controller, the Hybrid IDI controller enjoys a good level of visibility and low complexity in the controller structure, favourable aspects to its implementation in a real aircraft system. However, the lack of control authority on the yaw channel resulted in large control signals unfeasible in practical applications due to actuator position and rate limits. The importance of understanding the lessons taught by Stein's Bode Prize lecture (Stein, 2003) on the fundamental limitations of control is then apparent. This highlights the importance of adequately analysing the open-loop plant at hand and its limitations, since most of these limitations can be immediately identified in this first stage, before even any control synthesis.

The current work did not, by any means, invalidate the idea of using non-diagonal weights with a Hybrid IDI structure. It still holds its potential, but it requires the selection of a more adequate MIMO plant in future studies to better assess its compatibility. Furthermore, the effects of the multiple Hybrid IDI elements in a MIMO context still need to be better understood. While the work of Pollack, 2024 explicitly derived the effects of the various elements of the Hybrid controller structure on the broken-loop response at the plant input L_i, that was done for a SISO plant. It is expected that in a MIMO scenario, the effects of directionality play an important role. Therefore, some asymmetry is expected in how the scaled complementary filter elements on each individual channel affect the overall broken-loop response at the plant I/O.

Overall, it was demonstrated how the procedure to tune Hybrid IDI control laws using the $\mathcal{H}_\infty$ LSDP can be extended to MIMO plants. In future research, it would be valuable to assess the procedure in a MIMO plant which does not exhibit the control authority issues that the chosen X-29 lateral-directional model does.

8

Conclusions & Recommendations for Future Work

8.1. Conclusion

The present study's research objective was to investigate how to combine $\mathcal{H}_\infty$ Loop-Shaping and Nonlinear Dynamic Inversion control with application to fighter aircraft flight control law design. The study proposed a novel framework leveraging the $\mathcal{H}_\infty$ Loop-Shaping Design Procedure to optimise a structured linear variant of Incremental Nonlinear Dynamic Inversion (INDI) control, a Hybrid IDI controller. Results demonstrate that the approach achieves robustness guarantees comparable to standard full-order $\mathcal{H}_\infty$ controllers while maintaining the simplicity and modular architecture of NDI-like structures, thereby combining the advantages of both techniques. Thus, the research objective has been accomplished.

8.1.1. Revisiting the Research Questions

1. **To what extent does the proposed procedure to combine $\mathcal{H}_\infty$ Loop-Shaping and a structured Dynamic Inversion controller preserve the robustness guarantees of $\mathcal{H}_\infty$ Loop-Shaping control?**

   The analysis of the work of G. Papageorgiou and Polansky, 2009 in Chapter 4 demonstrated that the proposed equivalence between $\mathcal{H}_\infty$ Loop-Shaping and Dynamic Inversion controller is only present in the first step, where the plant is augmented with weighting filters. The methodology to tune the DI controller was successfully recreated, but it was demonstrated that the controller solution presented falls within the realm of $\mathcal{H}_2$ optimal control. It is also shown that using the observer-based $\mathcal{H}_\infty$ Loop-Shaping controller to tune the DI controller yields very different results than the ones presented in the paper. Thus, the proposed procedure does not exactly equate the $\mathcal{H}_\infty$ Loop-Shaping Design Procedure to DI control, which results in the absence of a priori robustness guarantees of the $\mathcal{H}_\infty$ LSDP.

2. **How can a Hybrid IDI structured controller be optimised using the $\mathcal{H}_\infty$ Loop-Shaping Design Procedure?**

   - **What are the robustness implications of the hybrid inversion scheme?**

     The robustness characteristics of a fully SB INDI are complementary to those of the classical model-based NDI variant. As a result, balanced robustness properties can be obtained by blending both inversion techniques in a hybrid inversion scheme. Therefore, since hybrid inversion

schemes allow navigating between model-based and sensor-based inversion schematics, they can be interpreted as an extra degree of freedom to modify the open-loop gain response and, thus, address robustness.

- **What are the robustness implications of designing a controller based on the virtual law location?**

  The analysis conducted in Chapter 5 challenges the often-claimed "automatic gain-scheduled nature" of (I)NDI control laws. While it is true that from the virtual law location, dynamic inversion schemes result in consistent closed-loop responses across the flight envelope, robustness at the plant I/O remains plant-dependent. Therefore, designing a controller based on the virtual law location fundamentally blinds the designer to robustness considerations. In fact, from a robust control perspective, the virtual control law signal v has little relevance as it concerns a signal internal to the controller itself. This discussion supports the use of linear variants of NDI/INDI to design robust controllers using linear $\mathcal{H}_\infty$-synthesis tools.

- **How can the $\mathcal{H}_\infty$ Loop-Shaping problem be framed as a non-smooth non-convex $\mathcal{H}_\infty$ problem under the structural constraints imposed by Hybrid IDI architectures?**

  Zhou and Doyle, 1998 demonstrates that the classical $\mathcal{H}_\infty$ LSDP can equivalently be interpreted as the more standard $\mathcal{H}_\infty$ problem formulation of minimising the $\mathcal{H}_\infty$ norm of the frequency-weighted gain from disturbances on the plant input and output to the controller input and output. This allows rewriting the $\mathcal{H}_\infty$ problem such that the total controller $K_s = W_1 K_\infty W_2$ appears explicitly outside of the plant $P(s)$. Consequently, it is possible to frame it as a non-smooth, non-convex $\mathcal{H}_\infty$ problem where K_s is assumed to be a structured Hybrid Incremental Dynamic Inversion controller.

- **How can the design procedure be extended to include model-following requirements in accordance with flying and handling qualities?**

  The framework is extended to 2DoF control by including a feedforward element to the Hybrid IDI controller for nominal and robust performance. Flying qualities and handling qualities specifications are derived from Anonymous, 1990 and implemented in terms of the modal parameters of LOES models via reference models T_{ref}. Nominal and robust performance is specified via minimisation of a weighted $\mathcal{H}_\infty$ norm of the error signal between the output and the reference model signal. Multi-requirement $\mathcal{H}_\infty$-synthesis in `MATLAB®` `systune` was used to employ a co-design approach where the controller and the model-following filter parameters are optimised simultaneously, making the overall procedure more streamlined.

- **To what extent can the procedure be used to develop a digital pitch-rate (SISO) controller for a highly unstable aircraft?**

  The design of a digital pitch-rate controller for the X-29 aircraft constitutes a very interesting case study, given the highly unstable nature of the plane and the hardware limitations. It has been demonstrated that the tuned Hybrid IDI using the $\mathcal{H}_\infty$ LSDP and a modified-continuous design achieves robustness properties comparable to the ones found in other studies about the design of control laws for the X-29. Furthermore, the procedure offers a systematic and transparent way to trade-off robustness and performance.

- **What parallelism can be established between the proposed procedure and the stability assessment of Dynamic Inversion-based control laws using coprime factors in Hyde and Papageorgiou, 2001?**

  The stability assessment using coprime factors in Hyde and Papageorgiou, 2001 results in an equivalent γ-value to the one obtained from the tuning procedure of Hybrid IDI control laws using $\mathcal{H}_\infty$, provided the same weighting filters W_1 and W_2 are used. Nevertheless, as the analysis tool from Hyde and Papageorgiou, 2001 searches for W_1 and W_2 weights such that the structured controller $K = K_{opt}$ for the associated shaped plant G_s, W_1 and W_2 are potentially very high order terms, but the resulting $\epsilon_\mu = \epsilon_{opt}$. This is something the synthesis procedure proposed in this study can not guarantee, as the synthesized controller K is computed from chosen W_1 and W_2 weights, and so one can only hope that the structure of K is 'rich' enough such that $K \to K_{opt}$ and, therefore, $\epsilon \to \epsilon_{opt}$. However, while the former remains a robustness analysis tool, the latter is a synthesis tool which results in a $\gamma = \epsilon^{-1}$ value, which ties directly to the robustness of the system to NCF uncertainty, and the closer the values of ϵ and ϵ_{opt}, the closer the obtained controller K is to the optimal LS controller K_{opt}. Furthermore, in the procedure proposed in this work, the filters W_1 and W_2 retain their core purpose of shaping the open-loop plant and, given a successful design, bound the broken loop gain at the plant I/O. On an additional note, the findings in Hyde and Papageorgiou, 2001 regarding the robustness of various DI controllers align with those in Pollack and Van Kampen, 2023 and Chapter 5, reinforcing the notion that sensor-based incremental dynamic inversion is not necessarily more robust than model-based dynamic inversion schemes.

- **How can the procedure be extended to a MIMO case?**

  $\mathcal{H}_\infty$ LSDP can be directly extended to a MIMO case, with the greatest challenge being the design of the weighting filters W_1 and W_2. While Hyde, 1995 proposes a systematic approach to designing diagonal weights, the fact is that diagonal weights might not work as well with plants that have strong cross-coupling between channels. The non-diagonal weights procedure by G. Papageorgiou and Glover, 1997 offers a truly systematic approach to the design of weighting filters, which avoids the trial-and-error nature of choosing diagonal weights. Using diagonal weights with the procedure proposed in this study yielded good synthesis results, while non-diagonal weights produced negative but still inconclusive results. Unfortunately, the severe lack of control authority of the lateral-directional X-29 model prevents a comprehensive assessment of the potential benefits of non-diagonal weighting filters. For MIMO applications, another question that is raised is to what extent the scaled complementary filters in the hybrid inversion scheme should be a diagonal block or if there are advantages to making K_c and H_c full-blocks. While this preliminary study raises some relevant questions, the inconclusive nature of the study dictates that more research on this area should be conducted with a more appropriate aircraft model.

3. **How can Robust Feedback linearisation be combined with $\mathcal{H}_\infty$ Loop-Shaping for flight control applications?**

   The research on optimising Hybrid IDI controllers using the $\mathcal{H}_\infty$ LSDP by framing the problem as a non-smooth, non-convex $\mathcal{H}_\infty$ problem under structural constraints opened numerous research directions. Given the limited time frame for the thesis work, it was not possible to address this last research question without compromising the research on the previous one. Nevertheless, the knowledge acquired in this study supports the immense potential of this research direction. The concept of robustness feedback linearization promises a direct connection to linear $\mathcal{H}_\infty$-synthesis methods, namely $\mathcal{H}_\infty$ Loop-Shaping. In this framework, tools such as the gap metric and the ν-gap metric can be leveraged to adequately choose the LTI point around which the robust feedback linearization is performed. The following chapter provides a more extensive assessment of the potential of the control technique.

8.1.2. Concluding Remarks

The main research contribution of the present work can be succinctly described as rearranging the $\mathcal{H}_\infty$ Loop-Shaping problem such that increased flexibility to the controller is introduced, i.e. working around the explicit presence of W_1 and W_2 as part of the final controller, in a way compatible with non-smooth non-convex $\mathcal{H}_\infty$-synthesis under controller's structural constraints in `MATLAB® systune`.

Writing the $\mathcal{H}_\infty$ Loop-Shaping as a 4-block problem leads directly to a description compatible with non-smooth, non-convex $\mathcal{H}_\infty$-synthesis (recall Figure 3.14 and Equation 3.64). However, the weighting filters W_1 and W_2 are absorbed into the plant P, and thus the structure can only be imposed on K_∞ and not on the total controller $K_s = W_1 K_\infty W_2$ (see Figure 8.1). Nevertheless, rewriting the problem such that the disturbances enter the plant at its input and output allows isolating the full controller K_s out of the plant P. As a result, controller structural constraints can now be imposed on the total controller (see Figure 8.2).

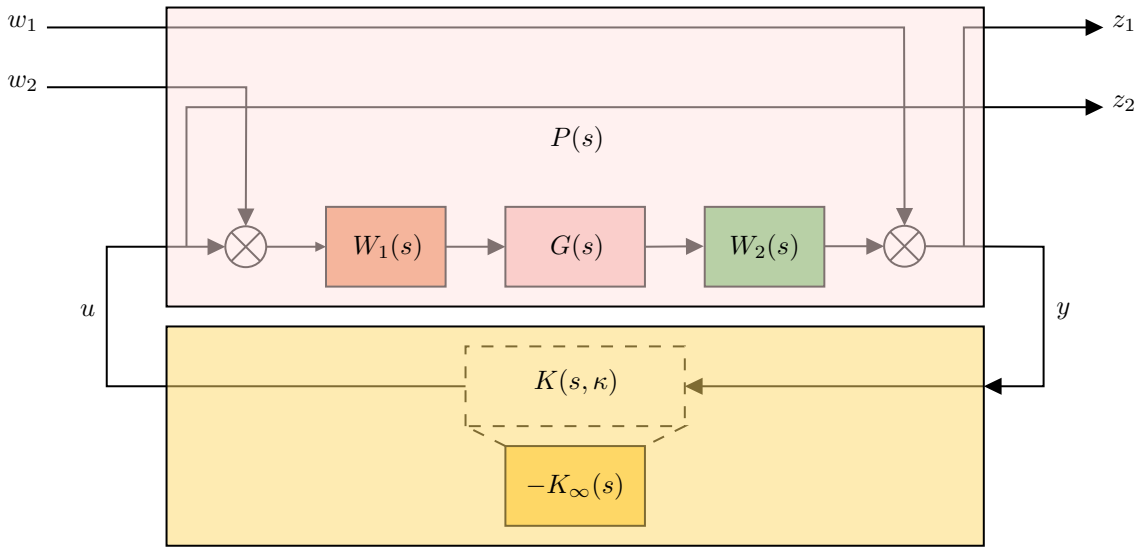

Figure 8.1: $\mathcal{H}_\infty$ Loop-Shaping 4-block problem under controller structural constraints on K_∞.

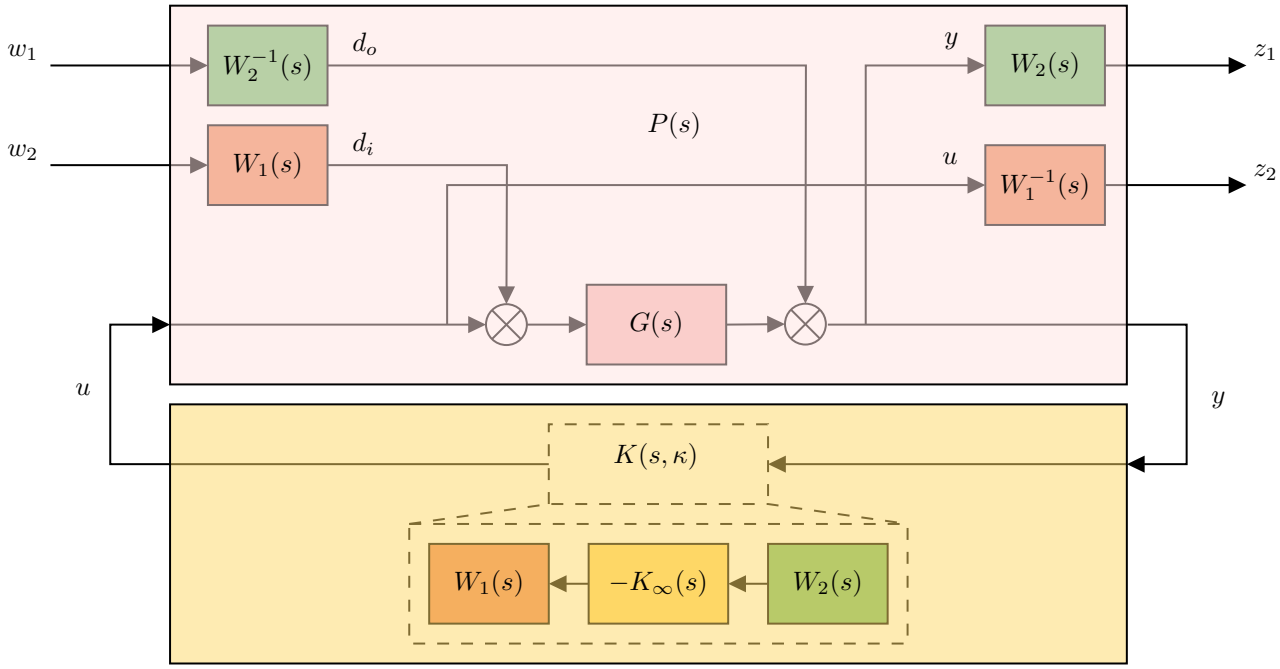

Figure 8.2: $\mathcal{H}_\infty$ Loop-Shaping equivalent problem under controller structural constraints on $K_s = W_1 K_\infty W_2$.

In order to extend the $\mathcal{H}_\infty$ LSDP to include model-following requirements in accordance with flying and handling qualities while leveraging the power of modern tools for multi-objective $\mathcal{H}_\infty$ synthesis, a co-design procedure is leveraged to simultaneously optimise the controller and the model-following filter W_{mf} parameters. Under this framework, the reference model T_{ref} is how the designer specifies flying and handling qualities requirements. The description of this $\mathcal{H}_\infty$ problem is displayed in Figure 8.3 and Equation 8.1:

$$\begin{aligned} &\min_{\kappa} \quad \max \left\{ \left\| W_{F_w} T_{w_2 \to z_2}\left(K(s,\kappa)\right) \right\|_\infty, \left\| W_{F_M} T_{w_3 \to z_3}\left(K(s,\kappa)\right) \right\|_\infty, \left\| T_{\left[\begin{smallmatrix} w_4 \\ w_5 \end{smallmatrix}\right] \to \left[\begin{smallmatrix} z_4 \\ z_5 \end{smallmatrix}\right]}\left(K(s,\kappa)\right) \right\|_\infty \right\} \\ &\text{subject to} \quad \left\| T_{w_1 \to z_1}\left(K(s,\kappa)\right) \right\|_\infty \leq 1 \end{aligned} \tag{8.1}$$

where the co-design parameters w, $M \in \kappa$, $W_{F_w} = w_{ideal}$ and $W_{F_M} = 1/M_{ideal}$ correspond to weightings to scale the co-design constraints to be competitive with the NCF robust stabilization one, and w_{ideal} and M_{ideal} correspond to parameters set by the designer.

It is important to notice that under this framework, $K(\kappa)$ can be selected to be any structured controller, such as the PID controllers widely celebrated in the industry (Åström & Hägglund, 1995) or the increasingly popular Nonlinear Dynamic Inversion and Incremental Nonlinear Dynamic Inversion controllers (Kumtepe et al., 2022). **The proposed framework offers a systematic and transparent methodology to design flight control laws with robustness guarantees while offering flexibility for the control structure to be chosen a priori.**

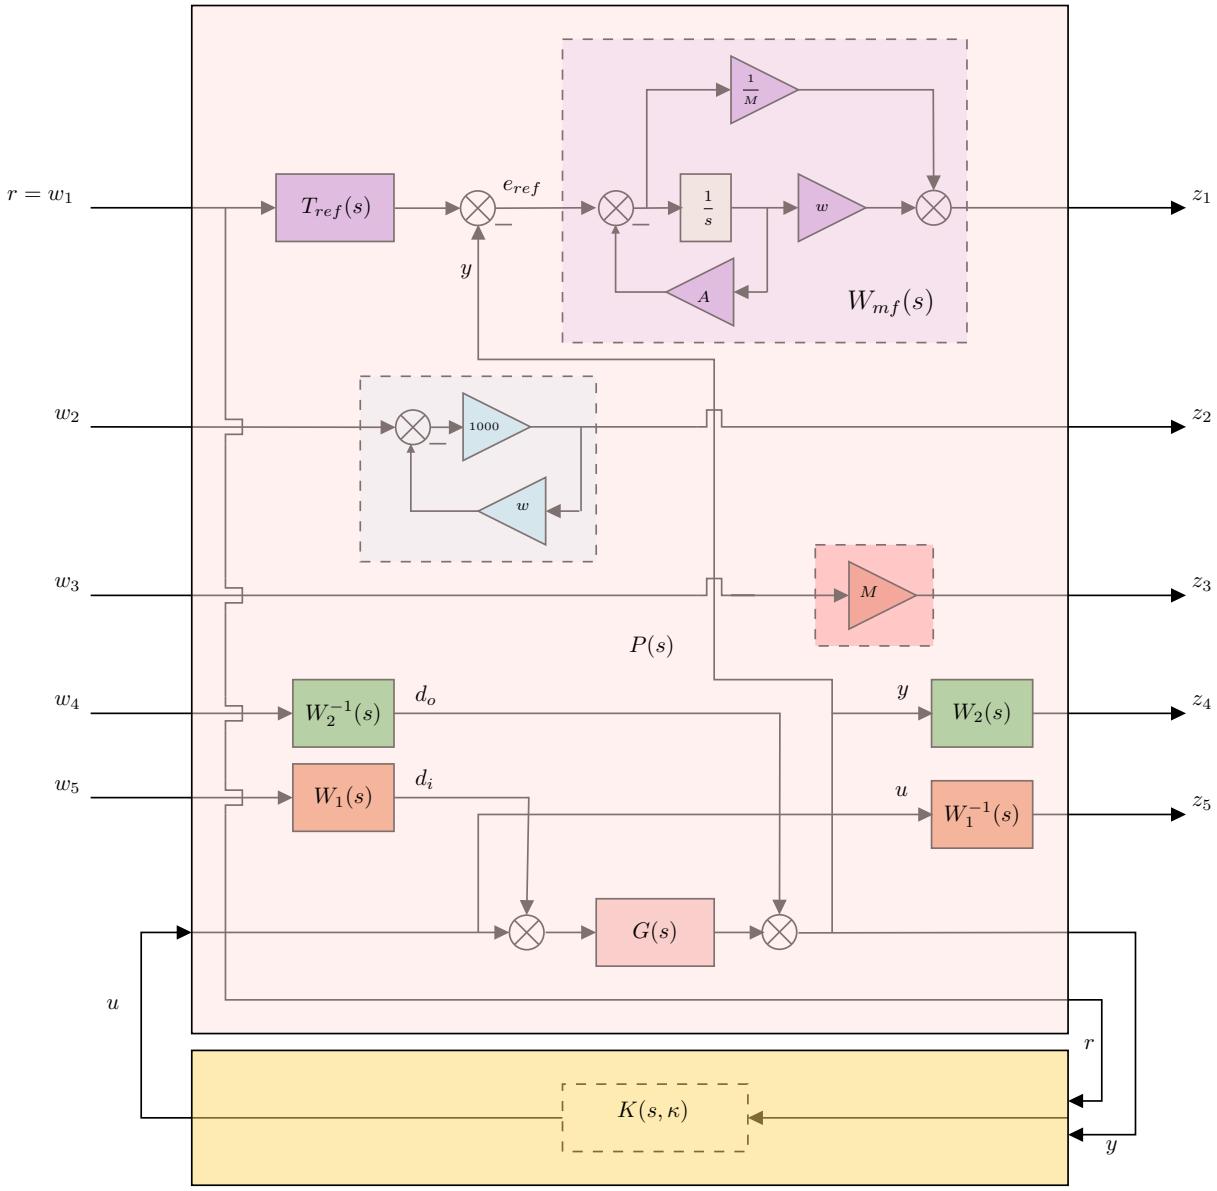

Figure 8.3: General description of non-smooth non-convex $\mathcal{H}_\infty$-synthesis under controller structural constraints against multiple requirements using the $\mathcal{H}_\infty$ LSDP and a co-design procedure for model-following requirements.

8.2. Recommendations for Future Work

8.2.1. Gain-Scheduling and LPV Control

In the present work, the robust synthesis of (I)NDI control laws was performed on a divide-and-conquer LTI basis. Consequently, the nonlinear nature of the plant is not reflected in the tuning process. One of the appealing properties of NDI is that gain scheduling is handled systematically and transparently. So, its fundamental advantage is lost by taking a step back and focusing on linear synthesis and implementation. However, as Pollack, 2024 demonstrated (and the present work corroborates), although the NDI control law is automatically scheduled, it is not scheduled in terms of robustness at the location of the plant, and so different controller design parameters may be needed as a function of the operating point from a robustness perspective.

Nevertheless, it is expected that the structure of (I)NDI leads to its design parameters exhibiting smaller variations across different flight conditions, demonstrating that reduced gain scheduling is necessary. This property was highlighted in Pollack, 2024 (even though the Boeing 747 used for the case study is not particularly nonlinear), but further exploration is recommended to verify it using the proposed methodology of integrating the hybrid control structure with the $\mathcal{H}_\infty$ LSDP approach.

Furthermore, to move away from the LTI divide-and-conquer, diving deeper into the topic of structured Linear Parameter-Varying (LPV) synthesis is recommended. The key idea of LPV control and q-LPV stemming from state transformation or other techniques is to model the nonlinear dynamics of the plant over the flight envelope as a linear parameter-varying system of the form:

$$\begin{cases} \dot{x} = A(\rho)x + B(\rho)u \\ y = C(\rho)x + D(\rho)u \end{cases} \tag{8.2}$$

where A, B, C, D are continuous (possibly nonlinear) matrix valued functions of ρ, and ρ is the vector of external scheduling variables. When the scheduling variables correspond to the system states, as is common in aerospace applications, the model is referred to as a "quasi-LPV" system. If such a quasi-LPV (q-LPV) model can be developed, and the scheduling variables are constrained within a specified range, controller synthesis procedures can be applied to design an LPV controller. This controller ensures guaranteed closed-loop stability and performance across the defined operating envelope (Leith & Leithead, 2000; Rugh & Shamma, 2000).

G. Papageorgiou and Glover, 1999 designed a robust gain-scheduled controller based on parameter-dependent Lyapunov functions and $\mathcal{H}_\infty$ Loop-Shaping, using a "quasi-LPV" model of the longitudinal dynamics of the HIRM. This LPV control system was then taken into flight in the VAAC Harrier (G. Papageorgiou et al., 2000). The research objective would be to leverage $\mathcal{H}_\infty$-synthesis non-smooth optimisation techniques to design a structured LPV Hybrid INDI (Pollack et al., 2025) while incorporating $\mathcal{H}_\infty$ Loop-Shaping as demonstrated in G. Papageorgiou and Glover, 1999 and Prempain, 2006.

8.2.2. Robust Feedback Linearization

Throughout this work, the nature of the control structure was constrained to more classical philosophies of feedback linearization (even though the idea of a hybrid approach is relatively new, its core principle is to navigate between a more classical model-based inversion scheme and a sensor-based one). It has been established that the core principle of feedback linearisation is to bring the nonlinear system to a Brunovsky canonical form. Even in a linear approach to dynamic inversion, the principle remains. However, as was shown in multiple instances, this architecture poses challenges to the combination with the $\mathcal{H}_\infty$ LSDP. Designing a 'linear' control law from the virtual location (i.e. looking to the Brunovsky canonical form of the system) "blinds" the designer to where the robustness properties of the system lie: at the actual plant input and output. Consequently, designing a robust controller for such architectures proved to be challenging.

Despite the promising results for the procedure proposed in Chapter 6 for dealing with SISO systems, observations in Chapter 7 suggested potential shortcomings of restricting one to the control structure of a Hybrid IDI when dealing with MIMO systems with coupled dynamics. Moreover, during the public defence of the doctoral thesis of Pollack, 2024, a member of the jury suggested that, in the context of robust control, it might be preferable to design the inversion so that the aircraft behaves like an average LTI point. This approach would ensure that the feedback-linearised dynamics retain a physical interpretation, as opposed to transforming the system into Brunovsky form. Such a suggestion is in agreement with the conclusions

drawn from this thesis. In fact, in the literature study conducted in Chapter 2, two promising articles (Guillard and Bourles, 2000 and Franco et al., 2006) were identified that corroborated the limitations of classic feedback linearization and proposed an alternative approach to feedback linearization, compatible with $\mathcal{H}_\infty$ Loop-Shaping, which do precisely that. It seems only natural, then, to spend some words to dive deeper into this work.

Franco et al., 2006 proposes a linearising control law that transforms a nonlinear system into its linear approximation around an operating point. The approach is anchored in the fact that it produces only a small transformation in the natural behaviour of the system, a key aspect to achieving a robust design. For the sake of clarity, consider the following description and derivation provided in Franco et al., 2006 about classical feedback linearization and robust feedback linearization:

Robust Feedback Linearization by Franco et al., 2006. The derivations and differences between Robust Feedback Linearization and Classical Feedback Linearization are quoted verbatim from Franco et al., 2006:

Consider the following nonlinear system with n states and m inputs described by:

$$\dot{x} = f(x) + g(x)u = f(x) + \sum_{i=1}^{m} g_i(x)u_i \tag{8.3}$$

where $x \in \mathbb{R}^n$ denotes the state, $u \in \mathbb{R}^m$ is the control input, and $f(x), g_1(x), \ldots, g_m(x)$ are smooth vector fields defined on an open subset of $\mathbb{R}^n$. Assume that this system satisfies the conditions for feedback linearization: there exists a vector $\lambda(x) = \begin{bmatrix} \lambda_1(x) & \cdots & \lambda_m(x) \end{bmatrix}^T$, formed by functions $\lambda_i(x)$ with relative degree r_i such that $r_1 + \cdots + r_m = n$, and the decoupling matrix of this system, given by

$$M(x) = \begin{bmatrix} L_{g_1} L_f^{r_1-1}\lambda_1(x) & \cdots & L_{g_m} L_f^{r_1-1}\lambda_1(x) \\ \vdots & \ddots & \vdots \\ L_{g_1} L_f^{r_m-1}\lambda_m(x) & \cdots & L_{g_m} L_f^{r_m-1}\lambda_m(x) \end{bmatrix} \tag{2}$$

is invertible. The output of system 8.3 is chosen as $y(x) = \lambda(x)$. The goal is to linearise this system by feedback in a neighbourhood of an operating point x_0 chosen, without loss of generality, as $x_0 = 0$. Assuming that the state is available for control purposes, the two different feedback linearization techniques can be described as follows.

Classical feedback linearization is accomplished by using a linearizing control law of the form $u_c(x, w) = \alpha_c(x) + \beta_c(x)w$, where w is a linear control, and a diffeomorphism $x_c = \phi_c(x)$, with $\alpha_c(x) = -M^{-1}(x) \begin{bmatrix} L_f^{r_1}\lambda_1(x) & \cdots & L_f^{r_m}\lambda_m(x) \end{bmatrix}^T$, $\beta_c(x) = M^{-1}(x), \phi_c^T(x) = \begin{bmatrix} \phi_{c_1}^T(x) & \cdots & \phi_{c_m}^T(x) \end{bmatrix}$, and $\phi_{c_i}^T(x) = \begin{bmatrix} \lambda_i(x) & L_f\lambda_i(x) & \cdots & L_f^{r_i-1}\lambda_i(x) \end{bmatrix}$. The resulting linear system is:

$$\dot{x}_c = A_c x_c + B_c w \tag{8.4}$$

where A_c and B_c are the matrices of the Brunovsky canonical form. The notation used here is in accordance with the one used in Franco et al., 2006, but the underlying principles are the same as the ones in Section 3.5.2.

On the other hand, for robust feedback linearization, the linearised system has the form:

$$\dot{x}_r = A_r x_r + B_r v \tag{8.5}$$

with $A_r = \partial_x f(0)$ and $B_r = g(0)$, which corresponds to the linear approximation of the nonlinear system 8.3.
The implementation of robust feedback linearization is achieved through a linearizing control law given as $u(x, v) = \alpha(x) + \beta(x)v$, where v is a linear control, and a diffeomorphism $x_r = \phi(x)$, with $\alpha(x) = \alpha_c(x) + \beta_c(x)LT^{-1}\phi_c(x), \beta(x) = \beta_c(x)R^{-1}, \phi(x) = T^{-1}\phi_c(x), L = -M(0)\partial_x\alpha_c(0), T = \partial_x\phi_c(0)$, and $R = M^{-1}(0)$. The functions $\alpha(x), \beta(x)$, and $\phi(x)$ satisfy

$$\partial_x\alpha(0) = 0, \beta(0) = I, \text{ and } \partial_x\phi(0) = I. \tag{8.6}$$

The author then uses the concept of "local W-stability" to prove that the robustness properties of the McFarlane-Glover control (for linear systems) are maintained when used together with the robust feedback linearization to the nonlinear system. The concept of "W-stability" will not be discussed in detail, but more details can be found in the original articles Guillard and Bourles, 2000 and Franco et al., 2006.

This allows the designer to design a $\mathcal{H}_\infty$ Loop-Shaping controller based on the LTI model of the operating point, following standard guidelines. Indeed, the designer can even make use of the procedure outlined in Section 7.2 to systematically design non-diagonal filters W_1 and W_2 without worrying about the implications of compatibility with the structure of more traditional model/sensor-based dynamic inversion controllers. The question then becomes how to choose the operating point. For that, given the connections of $\mathcal{H}_\infty$ Loop-Shaping to the ν-gap metric, a central point on the flight envelope can be chosen such that the worst case ν-gap metric between this point and the other points of the flight envelope is the lowest.

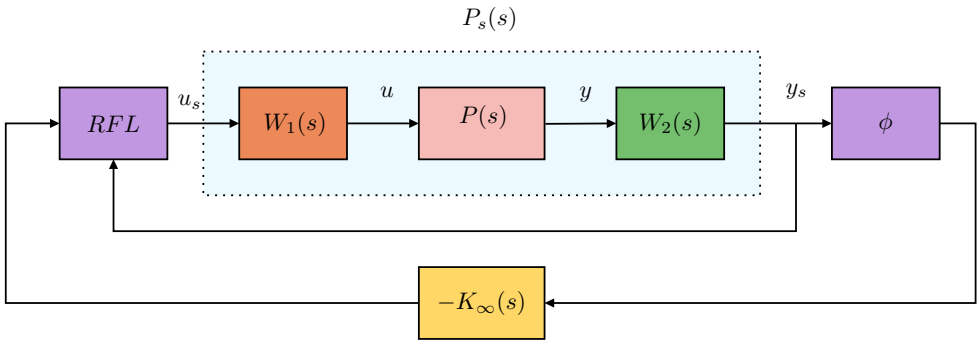

Figure 8.4: Implementation of the robust feedback linearization together with the McFarlane-Glover controller. Recreated from Franco et al., 2006.

It is expected that with this procedure, the flight envelope region, which guarantees closed-loop stability, is expanded (when compared to simply using the linear $\mathcal{H}_\infty$ Loop-Shaping controller) and that the performance of the controller is improved in the nonlinear domain.

8.2.3. Other Recommendations

- **Assess variation of gain-scheduled (linear) Hybrid IDI parameters across a flight envelope:** Since the proposed procedure is based on linear $\mathcal{H}_\infty$-synthesis methods, it still falls within the divide-and-conquer paradigm, which means that to cover the flight envelope, gain-scheduling is required. However, the modular architecture of the dynamic inversion controller is expected to present some advantages in this regard. It is of interest to explore how the gain-scheduled Hybrid IDI parameters

vary across the flight envelope and reflect on the differences from a purely nonlinear implementation (Hybrid INDI).

- **Exploit the Scaled Complementary Filter (SCF) control architecture to simplify implementation:** The SCF control architecture avoids explicit differentiation of the output feedback, which decreases the controller complexity and eases the implementation of the flight control laws.
- **Validate the procedure to tune Hybrid IDI control laws using the $\mathcal{H}_\infty$ LSDP with a flight test campaign:** Validation of this novel procedure to tune Hybrid IDI controllers requires its use in a full-cycle development of control laws, from linear synthesis, to nonlinear simulations, to finally flight test campaigns. A myriad of issues might arise in nonlinear simulations or actual hardware implementations that are not contemplated in linear control synthesis and analysis. Thus, only by doing this is it possible to attest to the potential of the technique to design modular and robust flight control laws.
- **Improve the way Handling Qualities are incorporated into the reference model T_{ref}:** In the current work, the choice of the reference model T_{ref} is based on the CAP and the Gibson criteria. However, as this was not the primary research goal, a very superficial approach was taken into consideration. There are, however, other criteria that can be considered. Furthermore, in future studies, pilot assessments during flight test manoeuvres should also be taken into consideration when assessing HQ levels.
- **Improve the structure of the model-following filter W_{mf} for objective separation of the parameters influence:** As it stands now, the DC gain of $W_{mf}^{-1}(s)$ is given by $\frac{A}{w}$, and since w is a tunable element of the filter in the co-design procedure, fixing the DC gain requires a priori knowledge of the optimized value of w. It would be advisable to restructure the model-following filter such that the designer can directly set the DC gain of $W_{mf}^{-1}(s)$ via the A parameter.
- **Use of EMF (Explicit Model Following) DI control architectures:** The Hybrid IDI feedforward term proposed in the present research deviates from what is commonly found in the literature for NDI and INDI controllers, which generally make use of EMF architectures. This was justified on the basis that the feedforward element allows the decoupling of the robustness of the system from the model following requirements, which is extremely useful in this initial study on the technique. However, having demonstrated the procedure's success in tuning a pitch-rate Hybrid IDI controller, the next step is to assess if equally successful results can be attained with an EMF architecture.
- **Further study on expanding the procedure to a MIMO case:** The unfortunate choice of a MIMO plant with severe control authority deficiencies did not allow for fully grasping the potential of using non-diagonal weighting filters. Future studies should continue research in this direction with an adequate MIMO plant to better assess the compatibility of Hybrid IDI structures and non-diagonal weighting filters. Moreover, the effects of the scaled complementary filter elements on the broken-loop response at the plant I/O still need to be better understood, and the potential of using full-block structures on these elements should be further studied.

References

Acquatella, P., Falkena, W., Van Kampen, E., & Chu, Q. (2012). Robust Nonlinear Spacecraft Attitude Control using Incremental Nonlinear Dynamic Inversion. *AIAA Guidance, Navigation, and Control Conference. Minneapolis, Minnesota, USA*, 4623. https://doi.org/10.2514/6.2012-4623

Adams, R., & Banda, S. (1993). Robust Flight Control Design Using Dynamic Inversion and Structured Singular Value Synthesis. *IEEE Transactions on Control Systems Technology*, *1*(2), 80–92. https://doi.org/10.1109/87.238401

Anonymous. (1990). *Flying Qualities of Piloted Aircraft. MIL-STD-1797A* (tech. rep.). US Department of Defense (DoD).

Anonymous. (2002). *Department of Defense Handbook: Airworthiness Certification Criteria. MIL-HDBK-516C* (tech. rep.). US Department of Defense (DoD).

Apkarian, P., & Gahinet, P. (1995). A Convex Characterization of Gain-Scheduled $\mathcal{H}_\infty$ Controllers. *IEEE Transactions on Automatic Control*, *40*(5), 853–864. https://doi.org/10.1109/9.384219

Apkarian, P., & Noll, D. (2006a). Nonsmooth H_∞ Synthesis. *IEEE Transactions on Automatic Control*, *51*(1), 71–86. https://doi.org/10.1109/tac.2005.860290

Apkarian, P., Noll, D., & Rondepierre, A. (2008). Mixed H_2/H_∞ Control via Nonsmooth Optimization. *SIAM Journal on Control and Optimization*, *47*(3), 1516–1546. https://doi.org/10.1137/070685026

Apkarian, P., Gahinet, P., & Becker, G. (1995). Self-scheduled $\mathcal{H}_\infty$ Control of Linear Parameter-varying Systems: a Design Example. *Automatica*, *31*(9), 1251–1261. https://doi.org/10.1016/0005-1098(95)00038-X

Apkarian, P., & Noll, D. (2006b). Nonsmooth Optimization for Multidisk H_∞ Synthesis. *European Journal of Control*, *12*(3), 229–244. https://doi.org/10.3166/ejc.12.229-244

Apkarian, P., & Noll, D. (2007). Nonsmooth optimization for multiband frequency domain control design. *Automatica*, *43*(4), 724–731. https://doi.org/10.1016/j.automatica.2006.08.031

Apkarian, P., Gahinet, P., & Buhr, C. (2014). Multi-Model, Multi-Objective Tuning of Fixed-Structure Controllers. *2014 European Control Conference (ECC). Strasbourg, France*, 856–861. https://doi.org/10.1109/ECC.2014.6862200

Apkarian, P., Dao, M., & Noll, D. (2015). Parametric Robust Structured Control Design. *IEEE Transactions on Automatic Control*, *60*(7), 1857–1869. https://doi.org/10.1109/TAC.2015.2396644

Apkarian, P., & Noll, D. (2017). The H_∞ Control Problem is Solved. *Aerospace Lab*, (13), pages 1–11. https://doi.org/10.12762/2017.al13-01

Apkarian, P., Noll, D., & Falcoz, A. (n.d.). How we saved the Rosetta mission [Retrieved March 31, 2025]. https://www.math.univ-toulouse.fr/~noll/PAPERS/rosetta.pdf

Åström, K., & Wittenmark, B. (1997). *Computer-Controlled Systems: Theory and Design*. Prentice Hall.

Åström, K., & Hägglund, T. (1995). *PID Controllers: Theory, Design, and Tuning*. ISA - The Instrumentation, Systems; Automation Society.

Bacon, B., Ostroff, A., & Joshi, S. (2001). Reconfigurable NDI Controller Using Inertial Sensor Failure Detection & Isolation. *IEEE Transactions on Aerospace and Electronic Systems*, *37*(4), 1373–1383. https://doi.org/10.1109/7.976972

Bacon, B., Ostroff, A., & Joshi, S. (2000). Nonlinear Dynamic Inversion Reconfigurable Controller Utilizing a Fault Tolerant Accelerometer. *19th DASC. 19th Digital Avionics Systems Conference. Proceedings (Cat. No. 00CH37126). Philadelphia, Pennsylvania, USA*, *2*, 6F5–1. https://doi.org/10.1109/dasc.2000.884920

Bacon, B., & Ostroff, A. (2000). Reconfigurable Flight Control Using Nonlinear Dynamic Inversion with a Special Accelerometer Implementation. *AIAA Guidance, Navigation, and Control Conference and Exhibit. Denver, Colorado, USA*, 4565. https://doi.org/10.2514/6.2000-4565

Balas, G., Chiang, R., Packard, A., & Safonov, M. (2022). *Robust Control Toolbox™ Reference*. The MathWorks, Inc.

Balas, G. (2003). Flight Control Law Design: An Industry Perspective. *European Journal of Control*, *9*(2-3), 207–226. https://doi.org/10.3166/ejc.9.207-226

Balas, G., Fialho, I., Packard, A., Renfrow, J., & Mullaney, C. (1997). On the Design of LPV Controllers for the F-14 Aircraft Lateral-Directional Axis During Powered Approach. *Proceedings of the 1997 American Control Conference (Cat. No. 97CH36041). Albuquerque, New Mexico, USA*, *1*, 123–127. https://doi.org/10.1109/ACC.1997.611768

Balas, G., Mueller, J., & Barker, J. (1999). Application of Gain-Scheduled, Multivariable Control Techniques to the F/A-18 System Research Aircraft. *Guidance, Navigation, and Control Conference and Exhibit. Portland, Oregon, USA*, 4206. https://doi.org/10.2514/6.1999-4206

Bates, D., & Postlethwaite, I. (2002). *Robust Multivariable Control of Aerospace Systems*. Delft University Press.

Blight, J., Dailey, R., & Gangsaas, D. (1994). Practical control law design for aircraft using multivariable techniques. *International Journal of Control*, *59*(1), 93–137. https://doi.org/10.1080/00207179408923071

Bosworth, J. (1992). *Linearized Aerodynamic and Control Law Models of the X-29A Airplane and Comparison With Flight Data*. National Aeronautics; Space Administration, Office of Management, Scientific; Technical Information Program. https://ntrs.nasa.gov/citations/19920009932

Brockett, R. (1978). Feedback Invariants for Nonlinear Systems. *IFAC Proceedings Volumes*, *11*(1), 1115–1120. https://doi.org/10.1016/s1474-6670(17)66062-2

Bugajski, D., Enns, D., & Elgersma, M. (1990). A Dynamic Inversion Based Control Law with Application to the High Angle-of-Attack Research Vehicle. *Guidance, Navigation and Control Conference. Portland, Oregon, USA*, 3407. https://doi.org/10.2514/6.1990-3407

Bugajski, D., & Enns, D. (1992). Nonlinear Control Law with Application to High Angle-of-Attack Flight. *Journal of Guidance, Control, and Dynamics*, *15*(3), 761–767. https://doi.org/10.2514/3.20902

Byrnes, C., & Isidori, A. (1984). A Frequency Domain Philosophy for Nonlinear Systems, with Applications to Stabilization and to Adaptive Control. *The 23rd IEEE Conference on Decision and Control. Las Vegas, Nevada, USA*, 1569–1573. https://doi.org/10.1109/cdc.1984.272345

Canin, D. (2019). F-35 High Angle of Attack Flight Control Development and Flight Test Results. *AIAA Aviation 2019 Forum. Dallas, Texas, USA*, 3227. https://doi.org/10.2514/6.2019-3227

Chiang, R., Safonov, M., Haiges, K., Madden, K., & Tekawy, J. (1993). A Fixed H_∞ Controller for a Supermaneuverable Fighter Performing the Herbst Maneuver. *Automatica*, *29*(1), 111–127. https://doi.org/10.1016/0005-1098(93)90176-t

Clarke, R., Burken, J., Bosworth, J., & Bauer, J. (1994). X-29 Flight Control System: Lessons Learned. *International Journal of Control*, *59*(1), 199–219. https://doi.org/10.1080/00207179408923075

Cordeiro, R., Marton, A., Azinheira, J., Carvalho, J., & Moutinho, A. (2022). Increased Robustness to Delay in Incremental Controllers Using Input Scaling Gain. *IEEE Transactions on Aerospace and Electronic Systems*, *58*(2), 1199–1210. https://doi.org/10.1109/taes.2021.3123215

Demirkiran, R., Isci, H., & Turkoglu, E. (2024). Handling Quality-Oriented Tuning Procedure of a Dynamic Inversion Control Law via a Robust Control Technique. *AIAA SCITECH 2024 Forum. Orlando, Florida, USA*, 2563. https://doi.org/10.2514/6.2024-2563

Dobos-Bubno, S., & Hartsook, L. (1977). *Validation of MIL-F-9490D - General Specification for Flight Control System for Piloted Military Aircraft. Volume II. YF-17 Lightweight Fighter Validation* (Final Report No. ADA062008). Northrop Corp Hawthorne CA Aircraft Div.

Doyle, J. (1978). Guaranteed Margins for LQG Regulators. *IEEE Transactions on Automatic Control*, *23*(4), 756–757. https://doi.org/10.1109/tac.1978.1101812

Doyle, J. (1982). Analysis of feedback systems with structured uncertainties. *IEE Proceedings D (Control Theory and Applications)*, *129*(6), 242–250. https://doi.org/10.1049/ip-d.1982.0053

Doyle, J. (1984). Lecture Notes in Advances in Multivariable Control. https://apps.dtic.mil/sti/citations/ADA225164

Doyle, J., Glover, K., Khargonekar, P., & Francis, B. (1989). State-Space Solutions to Standard $\mathcal{H}_\in$ and $\mathcal{H}_\infty$ Control Problems. *IEEE Transactions on Automatic Control*, *34*(8), 831–847. https://doi.org/10.1109/9.29425

Droste, C., & Walker, J. (2010). *The general dynamics case study on the f-16 fly-by-wire flight control system*. American Institute of Aeronautics; Astronautics. https://doi.org/10.1007/978-1-4613-8961-3

Elgersma, M. (1988). *Control of Nonlinear Systems Using Partial Dynamic Inversion*. University of Minnesota.

El-Sakkary, A. (1985). The Gap Metric: Robustness of Stabilization of Feedback Systems. *IEEE Transactions on Automatic Control*, *30*(3), 240–247. https://doi.org/10.1109/TAC.1985.1103926

Enns, D., Bugajski, D., Hendrick, R., & Stein, G. (1994). Dynamic inversion: an evolving methodology for flight control design. *International Journal of control*, *59*(1), 71–91. https://doi.org/10.1080/00207179408923070

Enns, D., & Keviczky, T. (2006). Dynamic Inversion Based Flight Control for Autonomous RMAX Helicopter. *2006 American Control Conference. Minneapolis, Minnesota, USA*, 8–pp. https://doi.org/10.1109/ACC.2006.1657330

Fer, H., & Enns, D. (1997). Nonlinear Longitudinal Axis Regulation of F-18 HARV using Feedback Linearization and LQR. *Proceedings of the 36th IEEE Conference on Decision and Control. San Diego, California, USA*, *5*, 4173–4178. https://doi.org/10.1109/cdc.1997.649488

Francis, B. (1987). *A Course in H_∞ Control Theory*. Springer Berlin, Heidelberg. https://doi.org/10.1007/BFb0007371

Franco, A., Bourles, H., De Pieri, E., & Guillard, H. (2006). Robust Nonlinear Control Associating Robust Feedback Linearization and H_∞ Control. *IEEE Transactions on Automatic Control*, *51*(7), 1200–1207. https://doi.org/10.1109/tac.2006.878782

Friang, J.-P., Duc, G., & Bonnet, J.-P. (1998). Robust Autopilot for a Flexible Missile: Loop-Shaping H_∞ Design and Real ν-Analysis. *International Journal of Robust and Nonlinear Control*, *8*(2), 129–153. https://doi.org/10.1002/(sici)1099-1239(199802)8:2<129::aid-rnc325>3.0.co;2-x

Gadewadikar, J., Lewis, F., Xie, L., Kucera, V., & Abu-Khalaf, M. (2007). Parameterization of all stabilizing $\mathcal{H}_\infty$ static state-feedback gains: Application to output-feedback design. *Automatica*, *43*(9), 1597–1604. https://doi.org/10.1016/j.automatica.2007.02.005

Gadewadikar, J., Lewis, F., & Abu-Khalaf, M. (2006). Necessary and Sufficient Conditions for $\mathcal{H}_\infty$ static Output-Feedback Control. *Journal of Guidance, Control, and Dynamics*, *29*(4), 915–920. https://doi.org/10.2514/1.16794

Gahinet, P., & Apkarian, P. (2011). Structured H_∞ Synthesis in MATLAB [18th IFAC World Congress. Milano, Italy]. *IFAC Proceedings Volumes*, *44*(1), 1435–1440. https://doi.org/10.3182/20110828-6-IT-1002.00708

Gahinet, P., & Apkarian, P. (1994a). A Linear Matrix Inequality Approach to H_∞ Control. *International Journal of Robust and Nonlinear Control*, *4*(4), 421–448. https://doi.org/10.1002/rnc.4590040403

Gahinet, P., & Apkarian, P. (1994b). A linear matrix inequality approach to H_∞ control. *International Journal of Robust and Nonlinear Control*, *4*, 421–448. https://doi.org/10.1002/rnc.4590040403

Gahinet, P., & Apkarian, P. (2013). Automated Tuning of Gain-Scheduled Control Systems. *52nd IEEE Conference on Decision and Control. Florence, Italy*, 2740–2745. https://doi.org/10.1109/CDC.2013.6760297

Gal-Or, B. (1990). *Vectored Propulsion, Supermaneuverability and Robot Aircraft*. Springer New York, NY. https://doi.org/10.1007/978-1-4613-8961-3

Garg, S. (1993). Robust Integrated Flight/Propulsion Control Design for a STOVL Aircraft Using H-infinity Control Design Techniques. *Automatica*, *29*(1), 129–145. https://doi.org/10.1016/0005-1098(93)90177-u

Garg, S., & Ouzts, P. (1991). Integrated Flight/Propulsion Control Design for a STOVL Aircraft using H-Infinity Control Design Techniques. *1991 American Control Conference. Boston, Massachusetts, USA*, 568–576. https://doi.org/10.23919/ACC.1991.4791435

Georgiou, T., & Smith, M. (1989). Optimal Robustness in the Gap Metric. *Proceedings of the 28th IEEE Conference on Decision and Control. Tampa, Florida, USA*, 2331–2336. https://doi.org/10.1109/CDC.1989.70590

Gera, J., & Bosworth, J. (1987). Dynamic Stability and Handling Qualities Tests on a Highly Augmented, Statically Unstable Airplane. *Guidance, Navigation and Control Conference. Monterey, California, USA*, 2258. https://doi.org/10.2514/6.1987-2258

Ghaoui, L. E., Carrier, A., & Bryson Jr, A. E. (1992). Linear Quadratic Minimax Controllers. *Journal of Guidance, Control, and Dynamics*, *15*(4), 953–961. https://doi.org/10.2514/3.20929

Gibson, J. (1999). *Development of a design methodology for handling qualities excellence in fly by wire aircraft* [Thesis]. Delft University of Technology, The Netherlands. https://repository.tudelft.nl/record/uuid:6b564b35-cb74-436c-8c47-845bfbbb8b4d

Glover, K., & McFarlane, D. (1988). Robust Stabilization of Normalized Coprime Factors: An Explicit H_∞ Solution. *1988 American Control Conference. Atlanta, Georgia, USA*, 842–847. https://doi.org/10.23919/acc.1988.4789839

Glover, K., & Doyle, J. (1988). State-space formulae for all stabilizing controllers that satisfy an H_∞-norm bound and relations to relations to risk sensitivity. *Systems & Control Letters*, *11*(3), 167–172. https://doi.org/10.1016/0167-6911(88)90055-2

Glover, K., Vinnicombe, G., & Papageorgiou, G. (2000). Guaranteed Multi-Loop Stability Margins and the Gap Metric. *Proceedings of the 39th IEEE Conference on Decision and Control (Cat. No. 00CH37187). Sydney, New South Wales, Australia*, *4*, 4084–4085. https://doi.org/10.1109/cdc.2000.912354

Grauer, J. (2025). Bounds on Robustness Limitations for Unstable Aircraft Using the Bode Sensitivity Integral. *AIAA SCITECH 2025 Forum. Orlando, Florida, USA*, 2243. https://doi.org/10.2514/6.2025-2243

Gregory, I. (1998). Dynamic Inversion to Control Large Flexible Transport Aircraft. *Guidance, Navigation, and Control Conference and Exhibit. Boston, Massachusetts, USA*, 4323. https://doi.org/10.2514/6.1998-4323

Grondman, F., Looye, G., Kuchar, R., Chu, Q., & Van Kampen, E. (2018). Design and Flight Testing of Incremental Nonlinear Dynamic Inversion based Control Laws for a Passenger Aircraft. *2018 AIAA Guidance, Navigation, and Control Conference. Kissimmee, Florida, USA*, 0385. https://doi.org/10.2514/6.2018-0385

Guillard, H., & Bourles, H. (2000). Robust feedback linearization. *Proc. 14th International Symposium on Mathematical Theory of Networks and Systems. Perpignan, France*. https://doi.org/10.13140/2.1.4639.8728

Harris, J. (2018). F-35 Flight Control Law Design, Development and Verification. *2018 Aviation Technology, Integration, and Operations Conference. Atlanta, Georgia, USA*, 3516. https://doi.org/10.2514/6.2018-3516

Hauser, J., Sastry, S., & Meyer, G. (1992). Nonlinear Control Design for Slightly Non-minimum Phase Systems: Application to V/STOL Aircraft. *Automatica*, *28*(4), 665–679. https://doi.org/10.1016/0005-1098(92)90029-F

Honeywell Technology Center, Lockheed Martin Skunk Works and Lockheed Martin Tactical Aircraft Systems. (1996). *Application of Multivariable Control Theory to Aircraft Control Laws, Final Report: Multivariable Control Design Guidelines. Tech. rep. WL-TR-96-3099* (tech. rep.). Wright-Patterson AFB, OH, USA: Wright Laboratory.

Horn, R., & Johnson, C. (2012). *Matrix Analysis*. Cambridge University Press. https://doi.org/10.1017/CBO9780511810817

Huang, Y., Zhang, Y., Pool, D., Stroosma, O., & Chu, Q. (2022). Time-Delay Margin and Robustness of Incremental Nonlinear Dynamic Inversion Control. *Journal of Guidance, Control, and Dynamics*, *45*(2), 394–404. https://doi.org/10.2514/1.g006024

Hyde, R., & Glover, K. (1993). Taking H_∞ Control into Flight. *Proceedings of 32nd IEEE Conference on Decision and Control. San Antonio, Texas, USA*, 1458–1463 vol.2. https://doi.org/10.1109/CDC.1993.325429

Hyde, R. (1995). H_∞ *Aerospace Control Design: A VSTOL Flight Application*. Springer London. https://doi.org/10.1007/978-1-4471-3049-9

Hyde, R., & Papageorgiou, G. (2001). Analysing the stability of NDI-based flight controllers with LPV methods. *AIAA Guidance, Navigation, and Control Conference and Exhibit. Montreal, Canada*, 4039. https://doi.org/10.2514/6.2001-4039

Ikeda, Y., McDowell, P., & Hargis, R. (1990). Robust Controller Design for an Advanced Fighter Aircraft. *29th IEEE Conference on Decision and Control. Honolulu, Hawaii, USA*, 2600–2606. https://doi.org/10.1109/cdc.1990.203342

Iwasaki, T., & Skelton, R. (1994). All Controllers for the General $\mathcal{H}_\infty$ Control Problem: LMI Existence Conditions and State Space Formulas. *Automatica*, *30*(8), 1307–1317. https://doi.org/10.1016/0005-1098(94)90110-4

Jiali, Y., & Jihong, Z. (2016). An Angular Acceleration Estimation Method Based on the Complementary Filter Theory. *2016 IEEE International Instrumentation and Measurement Technology Conference Proceedings. Taipei, Taiwan*, 1–6. https://doi.org/10.1109/i2mtc.2016.7520548

Joos, H. (1999). A methodology for multi-objective design assessment and flight control synthesis tuning. *Aerospace Science and Technology*, *3*(3), 161–176. https://doi.org/10.1016/S1270-9638(99)80040-6

Kanade, S., & Mathew, A. (2013). 2 DOF H-Infinity Loop Shaping Robust Control for Rocket Attitude Stabilization. *International Journal of Aerospace Sciences*. https://doi.org/10.5923/j.aerospace.20130203.02

Kehoe, M., Bjarke, L., & Laurie, E. (1990). An in-flight interaction of the x-29a canard and flight control system. *Orbital Debris Conference: Technical Issues and Future Directions. Baltimore, Maryland, USA.*, 1240. https://doi.org/10.2514/6.1990-1240

Khalil, H. (2002). *Nonlinear Systems*. Prentice Hall.

Khargonekar, P., & Shim, D. (1994). Maximally robust state-feedback controllers for stabilization of plants with normalized right coprime factor uncertainty. *Systems & Control Letters*, *22*(1), 1–4. https://doi.org/10.1016/0167-6911(94)90019-1

Kim, C., Ji, C., Koh, G., & Kim, B. (2021). Stability Margin and Structural Coupling Analysis of a Hybrid INDI Control for the Fighter Aircraft. *International Journal of Aeronautical and Space Sciences*, *22*(5), 1154–1169. https://doi.org/10.1007/s42405-021-00394-8

Kim, C., Ji, C., Koh, G., & Choi, N. (2023). Review on Flight Control Law Technologies of Fighter Jets for Flying Qualities. *International Journal of Aeronautical and Space Sciences*, *24*(1), 209–236. https://doi.org/10.1007/s42405-022-00560-6

Kim, C., Ji, C., Bang, J., Koh, G., & Choi, N. (2024). A Survey on Structural Coupling Design and Testing of the Flexible Military Aircraft. *International Journal of Aeronautical and Space Sciences*, *25*(1), 122–145. https://doi.org/10.1007/s42405-023-00643-y

Kolesnikov, E. (2005). NDI-based Flight Control Law Design. *AIAA Guidance, Navigation, and Control Conference and Exhibit. San Francisco, California, USA*, 5977. https://doi.org/10.2514/6.2005-5977

Koschorke, J., Falkena, W., Van Kampen, E., & Chu, Q. (2013). Time Delayed Incremental Nonlinear Control. *AIAA Guidance, Navigation, and Control (GNC) Conference. Boston, Massachusetts, USA*, 4929. https://doi.org/10.2514/6.2013-4929

Kumtepe, Y., Pollack, T., & Van Kampen, E. (2022). Flight Control Law Design using Hybrid Incremental Nonlinear Dynamic Inversion. *AIAA SciTech 2022 Forum. San Diego, California, USA*, 1597. https://doi.org/10.2514/6.2022-1597

Lanzon, A., Freddi, A., & Longhi, S. (2014). Flight Control of a Quadrotor Vehicle Subsequent to a Rotor Failure. *Journal of Guidance, Control, and Dynamics*, *37*(2), 580–591. https://doi.org/10.2514/1.59869

Lavretsky, E., & Wise, K. (2024). *Robust and Adaptive Control With Aerospace Applications*. Springer Cham. https://doi.org/10.1007/978-3-031-38314-4

Lawrence, D., & Rugh, W. (1995). Gain Scheduling Dynamic Linear Controllers for a Nonlinear Plant. *Automatica*, *31*(3), 381–390. https://doi.org/10.1016/0005-1098(94)00113-W

Leith, D., & Leithead, W. (2000). Survey of gain-scheduling analysis and design. *International Journal of Control*, *73*(11), 1001–1025. https://doi.org/10.1080/002071700411304

Lesprier, J., Biannic, J.-M., & Roos, C. (2014). Nonlinear structured $\mathcal{H}_\infty$ controllers for parameter-dependent uncertain systems with application to aircraft landing. *2014 IEEE Conference on Control Applications (CCA). Juan Les Antibes, France*, 433–438. https://doi.org/10.1109/cca.2014.6981384

Limebeer, D., Kasenally, E., & Perkins, J. (1993). On the Design of Robust Two Degree of Freedom Controllers. *Automatica*, *29*(1), 157–168. https://doi.org/10.1016/0005-1098(93)90179-w

Lombaerts, T., & Looye, G. (2011). Design and flight testing of manual nonlinear flight control laws. *AIAA Guidance, Navigation, and Control Conference. Portland, Oregon, USA*, 6469. https://doi.org/10.2514/6.2011-6469

Lombaerts, T., Kaneshige, J., Schuet, S., Aponso, B., Shish, K., & Hardy, G. (2020). Dynamic Inversion based Full Envelope Flight Control for an eVTOL Vehicle using a Unified Framework. *AIAA Scitech 2020 Forum. Orlando, Florida, USA*, 1619. https://doi.org/10.2514/6.2020-1619

Lu, B., Wu, F., & Kim, S. (2006). Switching LPV Control of An F-16 Aircraft via Controller State Reset. *IEEE Transactions on Control Systems Technology*, *14*(2), 267–277. https://doi.org/10.1109/TCST.2005.863656

Magni, J.-F., Bennani, S., & Terlouw, J. (1997). *Robust Flight Control: a Design Challenge* (Vol. 110). Springer. https://doi.org/10.1007/BFb0113842

Marques, J. (2024). *Design and simulator evaluation of a structured 2dof H_∞ loop-shaping control law for a business jet* [Master's thesis, TU Delft]. https://repository.tudelft.nl/record/uuid:9cd19d98-f1e9-4615-94c3-b812309d55db

McFarlane, D., Glover, K., & Noton, M. (1988). Robust Stabilization of a Flexible Space Platform: An H_∞ Coprime Factor Approach. *1988 International Conference on Control - CONTROL 88. Oxford, UK*, 677–682.

McFarlane, D., & Glover, K. (1992). A Loop Shaping Design Procedure Using H_∞ Synthesis. *IEEE Transactions on Automatic Control*, *37*(6), 759–769. https://doi.org/10.1109/9.256330

McFarlane, D., & Glover, K. (1990). *Robust Controller Design Using Normalized Coprime Factor Plant Descriptions*. Springer Berlin, Heidelberg. https://doi.org/10.1007/BFb0043199

Meyer, G., & Cicolani, L. (1975). *A Formal Structure for Advanced Automatic Flight-Control Systems* (Report). NASA. https://ntrs.nasa.gov/citations/19750015534

Michailidis, M. G., Kanistras, K., Agha, M., Rutherford, M., & Valavanis, K. (2018). Nonlinear Control of Fixed-Wing UAVs with Time-Varying Aerodynamic Uncertainties Via μ-Synthesis. *2018 IEEE Conference on Decision and Control (CDC). Miami Beach, Florida, USA*, 6314–6321. https://doi.org/10.1109/CDC.2018.8619011

Miller, C. (2011). Nonlinear Dynamic Inversion Baseline Control Law: Flight-Test Results for the Full-scale Advanced Systems Testbed F/A-18 Airplane. *AIAA Guidance, Navigation, and Control Conference. Portland, Oregon, USA*, 6468. https://doi.org/10.2514/6.2011-6468

Nichols, R., Reichert, R., & Rugh, W. (1993). Gain Scheduling for H-Infinity Controllers: A Flight Control Example. *IEEE Transactions on Control systems technology*, *1*(2), 69–79. https://doi.org/10.1109/87.238400

Panza, S., Lovera, M., Sato, M., & Muraoka, K. (2020). Structured μ-Synthesis of Robust Attitude Control Laws for Quad-Tilt-Wing Unmanned Aerial Vehicle. *Journal of Guidance, Control, and Dynamics*, *43*(12), 2258–2274. https://doi.org/10.2514/1.G005080

Papageorgiou, C., & Glover, K. (2005). Robustness Analysis of Nonlinear Flight Controllers. *Journal of Guidance, Control, And Dynamics*, *28*. https://doi.org/10.2514/1.9389

Papageorgiou, G. (1998). *Robust Control System Design: H_∞ Loop Shaping and Aerospace Applications* [Doctoral dissertation, University of Cambridge, Department of Engineering]. https://doi.org/10.17863/CAM.19503

Papageorgiou, G., & Glover, K. (1997). A Systematic Procedure for Designing Non-Diagonal Weights to Facilitate H_∞ Loop Shaping. *Proceedings of the 36th IEEE Conference on Decision and Control. San Diego, California, USA*, *3*, 2127–2132. https://doi.org/10.1109/cdc.1997.657081

Papageorgiou, G., & Glover, K. (1999). Design of a Robust Gain Scheduled Controller for the High Incidence Research Model. *Guidance, Navigation, and Control Conference and Exhibit. Portland, Oregon, USA*, 4276. https://doi.org/10.2514/6.1999-4276

Papageorgiou, G., Glover, K., D'Mello, G., & Patel, Y. (2000). Taking robust LPV control into flight on the VAAC Harrier. *Proceedings of the 39th IEEE conference on decision and control (Cat. No. 00CH37187). Sydney, Australia*, *5*, 4558–4564. https://doi.org/10.1109/CDC.2001.914633

Papageorgiou, G., & Polansky, M. (2009). Tuning a dynamic inversion pitch axis autopilot using McFarlane–Glover loop shaping. *Optimal Control Applications and Methods*, *30*(3), 287–308. https://doi.org/10.1002/oca.864

Paul, R., Garrard, W., Paul, R., & Garrard, W. (1997). Dynamics and Control of Tailless Aircraft. *22nd Atmospheric Flight Mechanics Conference. New Orleans, Louisiana, USA*, 3776. https://doi.org/10.2514/6.1997-3776

Pavel, M., Shanthakumaran, P., Chu, Q., Stroosma, O., Wolfe, M., & Cazemier, H. (2020). Incremental Nonlinear Dynamic Inversion for the Apache AH-64 Helicopter Control. *Journal of the American Helicopter Society*, *65*(2), 1–16. https://doi.org/10.4050/jahs.65.022006

Pérez, C. A., Theodoulis, S., Sève, F., & Goerig, L. (2022). Automatic weighting filter tuning for robust flight control law design [18th IFAC Workshop on Control Applications of Optimization CAO 2022]. *IFAC-PapersOnLine*, *55*(16), 400–405. https://doi.org/10.1016/j.ifacol.2022.09.057

Pollack, T., & Van Kampen, E. (2023). Robust Stability and Performance Analysis of Incremental Dynamic-Inversion-Based Flight Control Laws. *Journal of Guidance, Control, and Dynamics*, *46*(9), 1785–1798. https://doi.org/10.2514/1.g006576

Pollack, T., Theodoulis, S., & Wang, X. (2025). *Quasi-LPV Transformations for Robust Gain Scheduling of Incremental Nonlinear Dynamic Inversion-based Controllers*.
Submitted to LPVS2025.

Pollack, T. (2024). *Advances in Dynamic Inversion-Based Flight Control Law Design Methods* [Doctoral dissertation]. Delft University of Technology, The Netherlands. https://repository.tudelft.nl/record/uuid:28617ba0-461d-48ef-8437-de2aa41034ea

Pollack, T., Theodoulis, S., & Van Kampen, E. (2024). Commonalities between robust hybrid incremental nonlinear dynamic inversion and proportional-integral-derivative flight control law design. *Aerospace Science and Technology*, *152*, 109377. https://doi.org/10.1016/j.ast.2024.109377

Postlethwaite, I., Prempain, E., Turkoglu, E., Turner, M., Ellis, K., & Gubbels, A. (2005). Design and flight testing of various $\mathcal{H}_\infty$ controllers for the Bell 205 helicopter [Aerospace IFAC 2002]. *Control Engineering Practice*, *13*(3), 383–398. https://doi.org/10.1016/j.conengprac.2003.12.008

Prempain, E. (2006). Gain-Scheduling $\mathcal{H}_\infty$ Loop Shaping Control of Linear Parameter-Varying Systems [5th IFAC Symposium on Robust Control Design]. *IFAC Proceedings Volumes*, *39*(9), 215–219. https://doi.org/10.3182/20060705-3-FR-2907.00038

Raab, S., Zhang, J., Bhardwaj, P., & Holzapfel, F. (2018). Proposal of a Unified Control Strategy for Vertical Take-off and Landing Transition Aircraft Configurations. *2018 Applied Aerodynamics Conference. Atlanta, Georgia, USA*, 3478. https://doi.org/10.2514/6.2018-3478

Reigelsperger, W., & Banda, S. (1998). Nonlinear simulation of a modified F-16 with full-envelope control laws. *Control Engineering Practice*, *6*(3), 309–320. https://doi.org/10.1016/s0967-0661(98)00024-0

Roskam, J. (1998). *Airplane Flight Dynamics and Automatic Flight Controls*. Design Analysis; Research Corporation.

Rugh, W., & Shamma, J. (2000). Research on gain scheduling. *Automatica*, *36*(10), 1401–1425. https://doi.org/10.1016/S0005-1098(00)00058-3

Russell, R. (2003). *Non-Linear F-16 Simulation Using Simulink and Matlab* (Technical Report No. Ver. 1.0). University of Minnesota. Minneapolis, MN.

Sefton, J., & Glover, K. (1990). Pole/zero cancellations in the general $\mathcal{H}_\infty$ problem with reference to a two block design. *Systems & Control Letters*, *14*(4), 295–306. https://doi.org/10.1016/0167-6911(90)90050-5

Seiler, P., Packard, A., & Gahinet, P. (2020). An Introduction to Disk Margins [Lecture Notes]. *IEEE Control Systems Magazine*, *40*(5), 78–95. https://doi.org/10.1109/MCS.2020.3005277

Sève, F., Theodoulis, S., Wernert, P., Zasadzinski, M., & Boutayeb, M. (2017). Gain-Scheduled $\mathcal{H}\_\backslash\{\sqcup\dagger$ Loop-Shaping Autopilot Design for Spin-Stabilized Canard-Guided Projectiles. *Aerospace Lab*, (13), 1–24. https://doi.org/10.12762/2017.AL13-03

Shamma, J., & Athans, M. (1991). Guaranteed Properties of Gain Scheduled Control for Linear Parameter-varying Plants. *Automatica*, *27*(3), 559–564. https://doi.org/10.1016/0005-1098(91)90116-J

Sieberling, S., Chu, Q., & Mulder, J. (2010). Robust Flight Control Using Incremental Nonlinear Dynamic Inversion and Angular Acceleration Prediction. *Journal of Guidance, Control, and Dynamics*, *33*(6), 1732–1742. https://doi.org/10.2514/1.49978

Simplício, P., Pavel, M., Van Kampen, E., & Chu, Q. (2013). An acceleration measurements-based approach for helicopter nonlinear flight control using Incremental Nonlinear Dynamic Inversion. *Control Engineering Practice*, *21*(8), 1065–1077. https://doi.org/10.1016/j.conengprac.2013.03.009

Skogestad, S., & Postlethwaite, I. (2005). *Multivariable Feedback Control: Analysis and Design*. John Wiley & Sons.

Smeur, E., Chu, Q., & De Croon, G. (2016). Adaptive Incremental Nonlinear Dynamic Inversion for Attitude Control of Micro Air Vehicles. *Journal of Guidance, Control, and Dynamics*, *39*(3), 450–461. https://doi.org/10.2514/1.g001490

Smeur, E., Bronz, M., & De Croon, G. (2020). Incremental Control and Guidance of Hybrid Aircraft Applied to a Tailsitter Unmanned Air Vehicle. *Journal of Guidance, Control, and Dynamics*, *43*(2), 274–287. https://doi.org/10.2514/1.g004520

Smit, M., & Craig, I. (1998). Robust Flight Controller Design Using H_∞ Loop-Shaping and Dynamic Inversion Techniques. *Guidance, Navigation, and Control Conference and Exhibit. Boston, Massachusetts, USA*, 4132. https://doi.org/10.2514/6.1998-4132

Smith, P. (1998). A Simplified Approach to Nonlinear Dynamic Inversion Based Flight Control. *23rd Atmospheric Flight Mechanics Conference. Boston, Massachusetts, USA*, 4461. https://doi.org/10.2514/6.1998-4461

Snell, A., Enns, D., & Garrard, W. (1992). Nonlinear Inversion Flight Control for a Supermaneuverable Aircraft. *Journal of Guidance, Control, and Dynamics*, *15*(4), 976–984. https://doi.org/10.2514/3.20932

Stein, G. (2003). Respect the Unstable. *IEEE Control Systems Magazine*, *23*(4), 12–25. https://doi.org/10.1109/MCS.2003.1213600

Stevens, B., Lewis, F., & Johnson, E. (2015). *Aircraft Control and Simulation: Dynamics, Controls Design, and Autonomous Systems*. John Wiley & Sons. https://doi.org/10.1002/9781119174882

Stilwell, D., & Rugh, W. (1999). Interpolation of Observer State Feedback Controllers for Gain Scheduling. *IEEE transactions on automatic control*, *44*(6), 1225–1229. https://doi.org/10.1109/9.769379

Stilwell, D., & Rugh, W. (2000). Stability preserving interpolation methods for the synthesis of gain scheduled controllers. *Automatica*, *36*(5), 665–671. https://doi.org/10.1016/S0005-1098(99)00193-4

Su, R. (1982). On the linear equivalents of nonlinear systems. *Systems & Control Letters*, *2*(1), 48–52. https://doi.org/10.1016/s0167-6911(82)80042-x

Theodoulis, S. (2008). *Robust control in a nonlinear context for large operating domains* [Doctoral dissertation, Université Paris Sud-Paris XI]. https://theses.hal.science/tel-00352237v1

Theodoulis, S., & Duc, G. (2009). Missile Autopilot Design: Gain-Scheduling and the Gap Metric. *Journal of Guidance, Control, and Dynamics*, *32*(3), 986–996. https://doi.org/10.2514/1.34756

Tierno, J., & Glavaski, S. (1999). Phase-Lead Compensation of Pitch Axis Control Laws via McFarlane-Glover H_∞ Loop Shaping. *Proceedings of the 38th IEEE Conference on Decision and Control (Cat. No. 99CH36304). Phoenix, Arizona, USA*, *2*, 1538–1543. https://doi.org/10.1109/cdc.1999.830216

Tu, K., Sideris, A., Mease, K., Nathan, J., & Carter, J. (1999). Robust Lateral-Directional Control Design for the F/A-18. *Guidance, Navigation, and Control Conference and Exhibit. Portland, Oregon, USA*, 4204. https://doi.org/10.2514/6.1999-4204

Valasek, J., Ito, D., & Ward, D. (2001). Robust dynamic inversion Controller Design and Analysis for the X-38. *AIAA guidance, navigation, and control conference and exhibit. Montreal, Canada*, 4380. https://doi.org/10.2514/6.2001-4380

Van't Veld, R., Van Kampen, E., & Chu, Q. (2018). Stability and Robustness Analysis and Improvements for Incremental Nonlinear Dynamic Inversion Control. *2018 AIAA Guidance, Navigation, and Control Conference. Kissimmee, Florida, USA*, 1127. https://doi.org/10.2514/6.2018-1127

Vidyasagar, M., Schneider, H., & Francis, B. (1982). Algebraic and Topological Aspects of Feedback Stabilization. *IEEE Transactions on Automatic Control*, *27*(4), 880–894. https://doi.org/10.1109/tac.1982.1103015

Vidyasagar, M., & Kimura, H. (1986). Robust Controllers for Uncertain Linear Multivariable Systems. *Automatica*, *22*(1), 85–94. https://doi.org/10.1016/0005-1098(86)90107-x

Vinnicombe, G. (2000). *Uncertainty and Feedback, H_∞ Loop-Shaping and the ν-Gap Metric*. Imperial College Press. https://doi.org/10.1142/p140

Voulgaris, P. (1988). *High Performance Multivariable Control of the "Supermaneuverable" F18/HARV Fighter Aircraft*. http://hdl.handle.net/1721.1/34996

Voulgaris, P., & Valavani, L. (1991). High Performance Linear-Quadratic and H-Infinity Designs for a “Supermaneuverable” Aircraft. *Journal of Guidance, Control, and Dynamics*, *14*(1), 157–165. https://doi.org/10.2514/3.20617

Walker, G., & Allen, D. (2002). X-35B STOVL flight control law design and flying qualities. *2002 Biennial International Powered Lift Conference and Exhibit. Williamsburg, Virginia, USA*, 6018. https://doi.org/10.2514/6.2002-6018

Wang, X., Van Kampen, E., Chu, Q., & Lu, P. (2019). Stability Analysis for Incremental Nonlinear Dynamic Inversion Control. *Journal of Guidance, Control, and Dynamics*, *42*(5), 1116–1129. https://doi.org/10.2514/1.G003791

Webster, F., & Purifoy, D. (1991, July). *X-29 High Angle-of-Attack Flying Qualities* (Final Report No. AFFTC-TR-91-15). Air Force Flight Test Center, Edwards Air Force Base, California, Air Force Systems Command, United States Air Force. https://apps.dtic.mil/sti/tr/pdf/ADB157529.pdf

Weiser, C., Ossmann, D., & Looye, G. (2020). Design and flight test of a linear parameter varying flight controller. *CEAS Aeronautical Journal*, *11*(4), 955–969. https://doi.org/10.1007/s13272-020-00461-y

Wu, F., Yang, X., Packard, A., & Becker, G. (1996). Induced $\mathcal{L}_{\in}$-norm Control for LPV Systems with Bounded Parameter Variation Rates. *International Journal of Robust and Nonlinear Control*, *6*(9-10), 983–998. https://doi.org/10.1002/(SICI)1099-1239(199611)6:9/10<983::AID-RNC263>3.0.CO;2-C

Yang, Y., Wang, X., Zhu, J., Yuan, X., & Zhang, X. (2020). Robust proportional incremental nonlinear dynamic inversion control of a flying-wing tailsitter. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, *234*(16), 2274–2295. https://doi.org/10.1177/0954410020926657

Zames, G. (1966). On the Input-Output Stability of Time-Varying Nonlinear Feedback Systems-Part II: Conditions Involving Circles in the Frequency Plane and Sector Nonlinearities. *IEEE Transactions on Automatic Control*, *11*(3), 465–476. https://doi.org/10.1109/tac.1966.1098356

Zames, G. (1981). Feedback and Optimal Sensitivity: Model Reference Transformations, Multiplicative Seminorms, and Approximate Inverses. *IEEE Transactions on Automatic Control*, *26*(2), 301–320. https://doi.org/10.1109/tac.1981.1102603

Zhou, K., Doyle, J., & Glover, K. (1996). *Robust and Optimal Control*. Prentice Hall.

Zhou, K., & Doyle, J. (1998). *Essentials of Robust Control*. Prentice Hall.

Ziegler, J., & Nichols, N. (1942). Optimum Settings for Automatic Controllers. *Transactions of the American Society of Mechanical Engineers*, *64*(8), 759–765. https://davidr.no/iiav3017/papers/Ziegler_Nichols_%201942.pdf

A

IFAC ROCOND 2025 Conference Paper

Part of the research conducted in this master's thesis contributed to the development of a conference paper submitted to the IFAC Symposium on Robust Control Design (ROCOND 2025). The paper has been approved for publication.

Tuning Hybrid Incremental Dynamic Inversion Control Laws using H_∞ Loop-Shaping

Leonardo Encarnação * **Tijmen Pollack** * **Spilios Theodoulis** *

* *Delft University of Technology, 2629 HS Delft, The Netherlands*

Abstract: Nonlinear Dynamic Inversion (NDI) control techniques provide a conceptually simple and modular control framework, making it an attractive technique for designing flight control laws with shorter design cycles. However, its lack of inherent robustness guarantees shifts the burden of the design from the synthesis to the analysis part. Conversely, H_∞ Loop-Shaping provides controllers with robust stability guarantees. This work proposes a novel framework leveraging the H_∞ Loop-Shaping Design Procedure to optimize a structured linear variant of Incremental Nonlinear Dynamic Inversion (INDI) control, a Hybrid IDI controller. The Hybrid IDI controller consists of a blend between classical model-based DI and sensor-based IDI. The proposed methodology is validated through the design of a pitch-rate controller for NASA's X-29 experimental aircraft. Results demonstrate that the approach achieves robustness guarantees comparable to standard full-order H_∞ controllers while maintaining the simplicity and modular architecture of NDI-like structures, thereby combining the advantages of both techniques.

Keywords: Robust control, H-infinity Loop-Shaping, Generalized stability margin, Incremental Nonlinear Dynamic Inversion, Hybrid Incremental Nonlinear Dynamic Inversion, Flight control

1. INTRODUCTION

Nonlinear Dynamic Inversion (NDI) and H_∞ Loop-Shaping are two popular control techniques that found application in the field of flight controls. However, they are the result of different philosophies. While NDI control is the product of time-domain thought, the H_∞ Loop-Shaping is the product of frequency-domain thought, where the plant model uncertainty can be naturally and systematically treated (Papageorgiou and Polansky, 2009). What is attractive about NDI control is that it is conceptually simple and enables a modular approach to control. Motivated by this, NDI control has found application in Lockheed Martin F-35 Lightning II strike fighter jet (Harris, 2018). On the other hand, H_∞ Loop-Shaping (LS) allows the design of control laws with robustness guarantees (Hyde and Glover, 1993). With the advent of the `systune` tool in `MATLAB`® the H_∞ machinery can now be used to optimize highly structured controllers (Apkarian and Noll, 2006). Given the absence of robustness guarantees of NDI controllers (and its Incremental INDI variant), it is of interest to explore the possibility of leveraging the H_∞ Loop-Shaping Design Procedure (LSDP) to tune NDI/INDI controllers with robustness guarantees (Papageorgiou and Polansky, 2009).

The structure of this paper is as follows: Section 2 provides a comprehensive overview of the theoretical foundations of NDI and H_∞ LS. Section 3 introduces the novel framework that combines these two methodologies. In Section 4, the proposed approach is applied to design a pitch-rate controller for the X-29 aircraft. Finally, Section 5 discusses the wider significance of the procedure, reinforcing its potential to design IDI controllers with robustness guarantees.

2. THEORETICAL BACKGROUND

2.1 Nonlinear Dynamic Inversion Control

Nonlinear dynamic inversion control laws make use of the principles of feedback linearization to transform selected input-output channels into a chain of integrators of relative degree ρ. Consider the nonlinear system with n states and m inputs described by:

$$\Sigma: \quad \begin{cases} \dot{x} = f(x) + G(x)u, \\ y = h(x), \end{cases} \tag{1}$$

where the state vector $x \in R^n$, the input vector $u \in R^m$, the observation vector $y \in R^m$, and smooth mappings f, G, and h.

Writing the system relative degree as $\rho = [\rho_1, \ldots, \rho_m]^T$, the output dynamics can be described as (Wang et al., 2019):

$$\begin{aligned} y^{(\rho)} &= \begin{bmatrix} L_f^{\rho_1} h_1(x) \\ \vdots \\ L_f^{\rho_m} h_m(x) \end{bmatrix} + \begin{bmatrix} L_{g_1} L_f^{\rho_1 - 1} h_1(x) & \cdots & L_{g_m} L_f^{\rho_1 - 1} h_1(x) \\ \vdots & \ddots & \vdots \\ L_{g_1} L_f^{\rho_m - 1} h_m(x) & \cdots & L_{g_m} L_f^{\rho_m - 1} h_m(x) \end{bmatrix} u \\ &\triangleq \alpha(x) + B(x)u \end{aligned} \tag{2}$$

where $L_f^k h_i(x)$ and $L_{g_i} L_f^{k-1} h_i(x)$ represent repeated Lie derivatives of the function h_i along the vector fields f and g_i, with g_i being a column vector of the matrix G (Khalil, 2002).

For traditional feedback linearization, this expression can be used directly to construct a control law that linearizes the input-output dynamics to a set of $\sum_{i=1}^{m} \rho_i$ parallel integrators. Assuming that the control effectiveness matrix $B(x)$ is invertible, the following control law is obtained:

$$u = \hat{B}^{-1}(x)\big[y^{(\rho)} - \hat{\alpha}(x)\big]. \tag{3}$$

where the onboard model estimates of $\alpha(x)$ and $B(x)$ are represented as $\hat{\alpha}(x)$ and $\hat{B}(x)$, respectively. Here, $\nu \in R^m$ is the pseudo-control vector generated by an auxiliary virtual control law designed to meet control objectives.

To obtain an analogous control law in incremental form, a Taylor expansion of the output dynamics around the system state at time $t-\Delta t$ can be performed (Wang et al., 2019), where Δt represents the sampling interval. Denoting this condition by the subscript 0 for ease of notation yields the expression:

$$y^{(\rho)} = y_0^{(\rho)} + \frac{\partial[\alpha(x)+B(x)u]}{\partial x}\bigg|_0 \underbrace{(x-x_0)}_{\Delta x} + B(x_0)\underbrace{(u-u_0)}_{\Delta u} + R_1, \tag{4}$$

where R_1 represents the expansion remainder. Consequently, the time-scale separation assumption can be leveraged to design the incremental control input Δu, which assumes that all state-dependent and residual terms can be neglected (Grondman et al., 2018). This is often argued as justified in the case of high sampling rates and high-bandwidth actuators.

The control law is completed by adding the control vector u_0 to the resulting incremental term:

$$u = u_0 + \hat{B}^{-1}(x_0)\big[\nu - y_0^{(\rho)}\big]. \tag{5}$$

Note that compared to its nonincremental counterpart from Equation 3, the resulting control law does not require any model information on $\alpha(x)$ but uses only sensor feedback of the previous control vector and the derivative of the control variable instead.

Consider, for simplicity, a system with relative degree $\rho = 1$, such that $y^{(\rho)} = \dot{y}$. Defining $\xi^{MB} = \hat{\alpha}(x)$ and $\xi^{SB} = \dot{y} - \hat{B}u_0$ as the signals in the inversion path for model-based (MB) nonlinear dynamic inversion and sensor-based (SB) incremental nonlinear dynamic inversion, we have in the absence of On-Board model (OBM) uncertainty and disturbances (Pollack, 2024a):

$$\underbrace{\hat{\alpha}(x)}_{\xi^{MB}} = \underbrace{\dot{y} - \hat{B}(x)u_0}_{\xi^{SB}} \tag{6}$$

However, the fact that the above equality only holds in the absence of uncertainty and disturbances is a direct indication that the different inversion schemes possess different robustness properties. Pollack and Van Kampen (2023) highlight that these different robustness properties can clearly be understood by analysing the open-loop response of a MB and a SB linear dynamic inversion schematic with the same virtual control law controller. In order to stay closer to this framework, the rest of the present work will also adopt a linear version of nonlinear dynamic inversion, which will be referred to as simply Dynamic Inversion (DI). The authors demonstrate that SB inversion schemes result in an elevated open-loop frequency response compared to MB ones, resulting in higher gain at low frequencies, but also higher and often prohibitive crossover frequencies. The higher gain at low-frequencies helps to understand the reason for which Incremental Dynamic Inversion (IDI) is often claimed as more robust against aerodynamic model uncertainties. However, robustness at higher frequencies is penalized due to the insufficient roll-off. This helps to explain why unfiltered implementations of IDI control are often unsuccessful. Practical implementations of SB IDI control require filtering via synchronized Low-Pass Filters (LPF) H_c (Grondman et al., 2018):

$$\xi^{SB}_{filtered}(s) = H_c(s)[sy(s) - \hat{B}u(s)] \tag{7}$$

Pollack and Van Kampen (2023) suggests that reintroducing model information of the bare airframe dynamics in the form of a complementary augmentation element can lead to increased robustness of the control law. This concept was originally proposed in Jiali and Jihong (2016) and formulated in Kim et al. (2021) and Kumtepe et al. (2022) as a Hybrid INDI approach. However, the specific blend of model-based ξ^{MB} and sensor-based ξ^{SB} inversion used in the present work originates from Pollack (2024a); Pollack et al. (2024b) and can be defined as:

$$\begin{aligned}\xi^{HB}(s) &= (1 - K_c H_c(s))\xi^{MB}(s) + K_c H_c(s)\xi^{SB}(s) \\ &= \xi^{MB}(s) + K_c H_c(s)(\xi^{MB}(s) - \xi^{SB}(s)) \\ &= \xi^{MB}(s) + K_c H_c(s) e_\xi(s)\end{aligned} \tag{8}$$

where K_c is a scaling gain $\in$ [0,1]. Given the relationships in Equation 8, Hybrid IDI can be interpreted as a standard model-based DI with additional error compensation. Correspondingly, Hybrid IDI collapses to purely model-based design when $K_c = 0$, and $\xi^{HB}(s) \to \xi^{SB}(s)$ if $K_c = 1$ and the bandwidth of $H_c(s)$ is made sufficiently large.

Since the hybrid design allows navigating between model-based and sensor-based inversion schematics, it can be interpreted as an extra degree of freedom to modify the open-loop gain response, and thus help address robustness.

2.2 H_∞ Loop-Shaping Control

The H_∞ LSDP, originally proposed by McFarlane and Glover (1992) leverages the concepts of classical Loop-Shaping and the H_∞ machinery to design controllers which robustly stabilize a plant with respect to Normalized Coprime Factor (NCF) uncertainty. It is a two-step procedure, where in the first step the plant is shaped with weighting filters W_1 and W_2, and in the second step, a robust controller K_∞ is computed which robustly stabilizes the shaped plant $G_s = W_2 G W_1$ with respect to NCF uncertainty. The first step is essential since robust stabilization alone is of limited practical value because it does not allow the designer to specify any performance requirements (Skogestad and Postlethwaite, 2005).

The robust controller K_∞ is a stabilising controller which minimizes the following H_∞ norm:

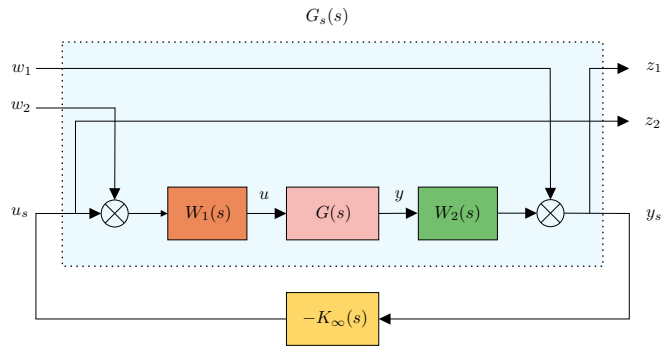

Fig. 1. NCF robust stabilization problem written as mixed sensitivity 4-block problem.

$$\begin{aligned} &\min_{\text{stab } K_\infty} \left\| \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} \rightarrow \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} \right\|_\infty = \\ &\min_{\text{stab } K_\infty} \left\| \begin{bmatrix} I_p \\ -K_\infty \end{bmatrix} (I_p + G_s K_\infty)^{-1} [I_p \ G_s] \right\|_\infty \leq \gamma \triangleq \epsilon^{-1} \end{aligned} \tag{9}$$

The lowest achievable value γ_{min} is given by (Glover and McFarlane, 1988):

$$\gamma_{\min} = \epsilon_{\max}^{-1} = \left\{ 1 - \| [N \ M] \|_H^2 \right\}^{-\frac{1}{2}} = (1 + \rho(XZ))^{\frac{1}{2}}, \tag{10}$$

where $\| \cdot \|_H$ denotes the Hankel norm, ρ denotes the spectral radius, and X and Z are the solutions to the associated Algebraic Riccati Equations (ARE's). The suboptimal full-order K_∞ can be computed from γ, where $\gamma < \gamma_{min}$, and the 2 ARE's.

3. TUNING A DI-STRUCTURED CONTROLLER USING THE LSDP

In Zhou and Doyle (1998), the authors demonstrate that the classical H_∞ LSDP can equivalently be interpreted as the more standard H_∞ problem formulation of minimising the H_∞ norm of the frequency-weighted gain from disturbances on the plant input and output to the controller input and output (see Figure 2).

In order to demonstrate this equivalence, recall that the NCF robust stabilization problem can be written as a mixed sensitivity 4-block problem, as presented in Figure 1 and Equation 9.

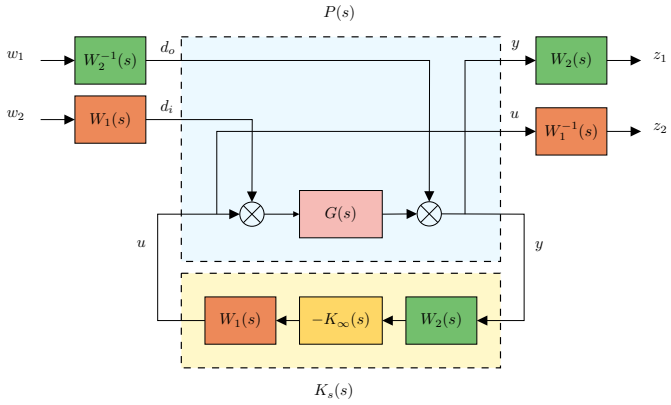

Fig. 2. Alternative structure for the NCF LSDP robust stabilization.

Using the fact that $G_s = W_2 G W_1$ and defining as the total controller $K_s = W_1 K_\infty W_2$, and assuming that the weighting filters are square and invertible, it is possible to rewrite Equation 9 as:

$$\begin{aligned} &\left\| \begin{bmatrix} I_p \\ -K_\infty \end{bmatrix} (I_p + G_s K_\infty)^{-1} [I_p \ G_s] \right\|_\infty = \\ &\left\| \begin{bmatrix} W_2 \\ -W_1^{-1} K_s \end{bmatrix} W_2^{-1} (I_p + G_s K_\infty)^{-1} W_2 [W_2^{-1} \ G W_1] \right\|_\infty \end{aligned} \tag{11}$$

Using the fact that the inner part $W_2^{-1} (I_p + G_s K_\infty)^{-1} W_2$ can be written as:

$$\left[W_2^{-1} (I_p + W_2 G W_1 K_\infty) W_2 \right]^{-1} = (I_p + G K_s)^{-1}, \tag{12}$$

and the following equivalence holds:

$$\begin{aligned} &\left\| \begin{bmatrix} I_p \\ -K_\infty \end{bmatrix} (I_p + G_s K_\infty)^{-1} [I_p \ G_s] \right\|_\infty \leq \epsilon^{-1} \Leftrightarrow \\ &\left\| \begin{bmatrix} W_2 \\ -W_1^{-1} K_s \end{bmatrix} (I_p + G K_s)^{-1} [W_2^{-1} \ G W_1] \right\|_\infty \leq \epsilon^{-1} \end{aligned} \tag{13}$$

If the multiplications in Equation 13 are expanded and defining $L_s = (I_p + G K_s)^{-1}$, the following particular weighted mixed sensitivity problem is achieved:

$$\left\| \begin{bmatrix} W_2 L_s W_2^{-1} & W_2 L_s G W_1 \\ -W_1^{-1} K_s L_s W_2^{-1} & -W_1^{-1} K_s L_s G W_1 \end{bmatrix} \right\|_\infty \leq \epsilon^{-1} \tag{14}$$

The previous result can be represented as a 4-block mixed sensitivity problem with weights on the disturbances and plant input and output, as shown in Figure 2.

$$\begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} W_2 L_s W_2^{-1} & W_2 L_s G W_1 \\ -W_1^{-1} K_s L_s W_2^{-1} & -W_1^{-1} K_s L_s G W_1 \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} \tag{15}$$

The advantage of this equivalent formulation is that it is possible to consider the inputs and outputs in the actual plant I/O, rather than the shaped plant I/O's.

One of the challenges of combining H_∞ LSDP with Dynamic Inversion is that the fundamental robustness properties are tied to the open-loop gain at the plant I/O, and not to the virtual control law location. Since DI techniques close the loop at this virtual location, combining it with H_∞ LSDP raises concerns, as shaping the plant's open-loop response with the filters W_1 and W_2 may distort the inversion path.

However, since this equivalent formulation allows to consider the disturbances and the outputs at the actual plant I/O, that difficulty is overcome. Nevertheless, the fact that the feedback controller K_s is constructed as $K_s = W_1 K_\infty W_2$ might be undesirable when highly structured control architectures are considered, like the one of Hybrid IDI. Nevertheless, it is hypothesised that if a given structured controller K_{struct} has compatible tunable elements to those of K_s, then K_{struct} can achieve similar results to K_s, by using non-smooth optimization techniques to solve the H_∞ synthesis problem under structural constraints

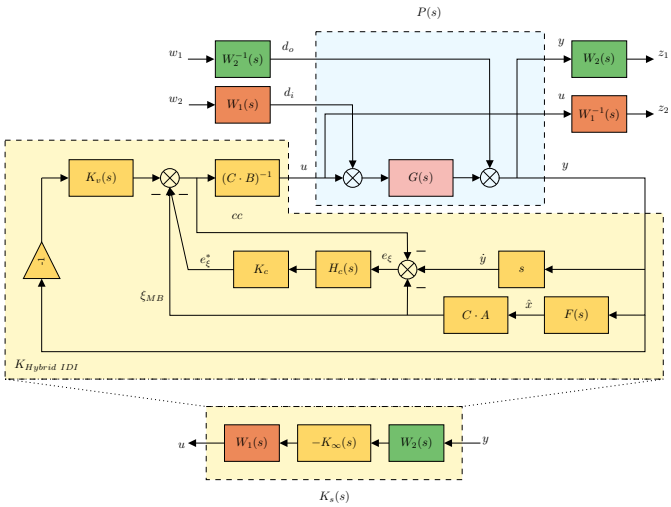

Fig. 3. Leveraging alternative structure to tune Hybrid IDI controller using the H_∞ LSDP.

(Apkarian and Noll, 2006). This effectively makes the nature of the optimization non-convex.

It has been established that there exists a full-order solution for K_∞ which minimizes the H_∞ norm from w to z. Alternatively, the H_∞ minimization solution can be reimagined as: there is an optimal stabilizable controller K_s, subject to $W_1 K_\infty W_2$, which minimizes the H_∞ norm from w to z. Reconsidering K_s as an appropriate structured controller K_{struct}, it is hypothesized that the minimisation of the H_∞ norm from w to z results in the tunable elements of K_{struct} being reconfigured such that $K_{struct} \to K_s$ and $\gamma_{K_{strcut}} \to \gamma_{K_s}$. In order to leverage this concept to tune the Hybrid IDI controller, the structure in Figure 3 was used in `systune`. The various components of the Hybrid IDI K_{H-IDI} are defined in Section 4.

4. DESIGN OF PITCH-RATE CONTROLLER FOR THE X-29

The X-29 is a forward swept-wing former NASA research airplane, which was built to demonstrate basic aerodynamic performance improvements that might be gained from new composite materials. It constitutes a particularly interesting case study due to its fundamental control limitations, which awarded it as one of the examples provided by Gunter Stein in his very famous "Respect the Unstable" Bode Prize lecture (Stein, 2003).

The flight condition (M = 0.9 and altitude = 8000 ft) chosen for the controller's design is very challenging due to the plane's violent instability. The aircraft has 3 longitudinal inputs: a canard, a symmetric flap and a strake flap. However, for simplicity, the short-period dynamics model used is retrieved from Lavretsky and Wise (2024), where a linear combination of the 3 inputs is used to achieve a single virtual input δ_e, which makes the nature of the problem SISO:

$$\begin{bmatrix} \dot{\alpha} \\ \dot{q} \end{bmatrix} = \begin{bmatrix} Z_\alpha & 1 + Z_q/V_o \\ M_\alpha & M_q \end{bmatrix} \begin{bmatrix} \alpha \\ q \end{bmatrix} + \begin{bmatrix} Z_\delta \\ M_\delta \end{bmatrix} \delta_e \tag{16}$$

where $Z_\alpha = -2.241\ s^{-1}$, $Z_q = 0.9897$, $M_\alpha = 44.74\ s^{-2}$, M_q=$-0.9024\ s^{-1}$, Z_δ=$-0.2331\ s^{-1}$, and M_δ=$-45.93\ s^{-2}$.

In order to model the actuator dynamics, a simple second-order transfer function was used (Lavretsky and Wise, 2024), with $\omega_n = 70$ rad/s and $\xi_n = 0.7$.

The open-loop system has an unstable pole at p = 5.12 rad/s. This unstable RHP pole imposes a lower bound on the allowed complementary sensitivity T bandwidth w_{BT}, with an approximate bound of $w_c > 2\text{p}$ (Skogestad and Postlethwaite, 2005). Therefore, the crossover frequency of the shaped plant G_s should be made to be at least $w_c > 2\text{p} = 10.24$ rad/s. In order to boost the low-frequency gain of the open-loop response, obtain a -20 dB slope around crossover, and attain the desired crossover frequency value of 10.2 rad/s, W_1 was set as a PI, while, for simplicity, W_2 was set to 1.

$$W_1 = K_p + \frac{K_i}{s} = \frac{0.24183 \cdot (s + 8.166)}{s} \tag{17}$$

The obtained shaped plant, given by $G_s = G_{ol} W_1$ fulfils the desired shaping, as can be observed in Figure 4.

In order to assess whether the following shaping of the plant is compatible with a robust design, a suboptimal $\gamma_{th} > \gamma_{min}$ is computed with a tolerance equal to 1e-3. The resulting $\gamma_{th} = 2.3181$ adheres to the common guideline that sets an acceptable range for γ between 1 and 4 (Hyde and Glover, 1993).

Subsequently, the non-convex H_∞-minimisation problem in Figure 3 is implemented in `systune` to optimize the Hybrid IDI parameters: $K_v(s)$, $H_c(s)$ and K_c. The virtual law controller $K_v(s)$ is defined as a PI controller in series with a first-order LPF, the synchronizing filter $H_c(s)$ is also a first-order LPF and K_c is a scaling gain $\in$ [0,1].

$$K_v(s) = \frac{K_p \cdot (s + K_i/K_p)}{s} \cdot \frac{w_{lp}}{s + w_{lp}} \tag{18}$$

$$H_c(s) = \frac{w_{H_c}}{s + w_{H_c}} \tag{19}$$

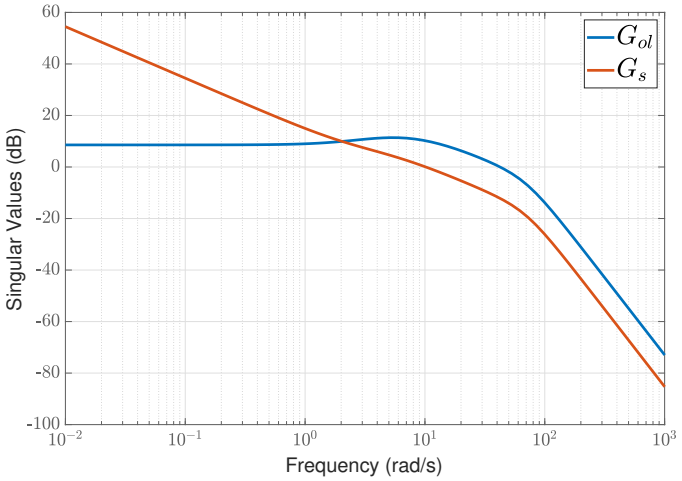

Fig. 4. Frequency response of open-loop plant G_{ol} and shaped plant G_s.

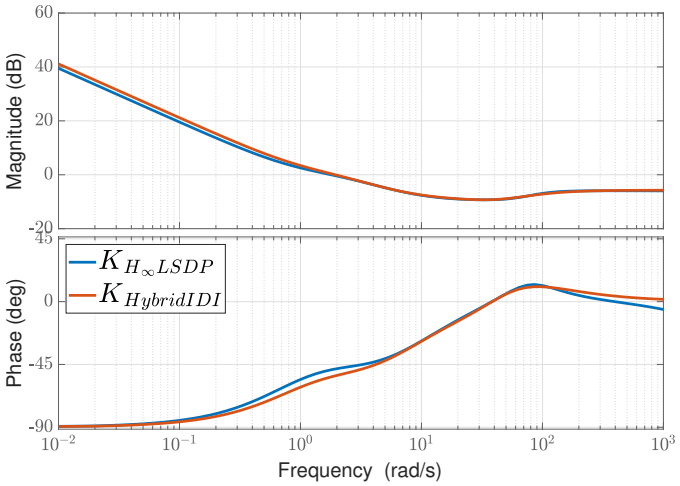

Fig. 5. Gain and Phase comparison of standard full-order H_∞ controller and tuned Hybrid IDI.

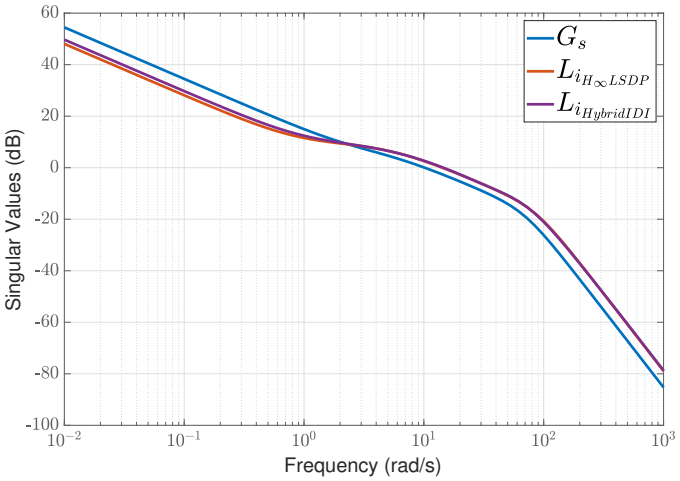

Fig. 6. Open-loop gain at plant input of standard full-order H_∞ controller and of the tuned Hybrid IDI.

Furthermore, given that DI control requires full-state feedback for the inversion, a simple angle-of-attack estimator based on the short-period model information is introduced.

$$\hat{x}(s) = [\hat{\alpha}(s)\ q(s)]^T = F(s)\ q(s) = \left[\frac{1 + Z_q/V_o}{s - Z_\alpha}\ \ 1\right]^T q(s) \tag{20}$$

Moreover, notice that instead of using feedback on the actuator position $u = \delta_{e_{meas}}$, internal Control Command (CC) feedback is used (refer to Figure 3). In IDI control schemes, feedback on the input signal CBu is necessary. Assuming ideal actuators, the input signal is equal to the commanded input $u = u_{cmd}$, which implies:

$$u = u_{cmd} = (CB)^{-1}CC \Rightarrow CC = CBu \tag{21}$$

In reality, actuators have internal dynamics, and therefore CC is merely an approximation of CBu. However, the fact that CC feedback is a location inside the control law itself makes it conceptually simpler to implement and does not require additional sensor measurements. Furthermore, the work of Pollack (2024a) suggests that no limitations in terms of robustness of the control law result from this choice of input feedback signal.

Table 1. Optimization of the various Hybrid IDI parameters using the H_∞ LSDP.

	Optimization		Virtual Control			Inversion Loop	
Parameter	γ_{th}	γ	K_p	K_i	w_{lp}	K_c	w_{H_c}
Value	2.3181	2.3622	12.7882	36.2545	52.12	0.3060	80.62

The results of the optimization are presented in Table 1. As can be observed, the obtained γ is similar to γ_{th} which attests the success of the optimization. If we first inspect the Bode plots of the full-order H_∞ LS controller and the tuned Hybrid IDI it becomes evident that they are virtually identical as demonstrated in Figure 5. Furthermore, examining the obtained open-loop gain at the plant input in Figure 6 confirms that their open-loop responses are also nearly identical. Therefore, it is demonstrated that under the hypothesis that the K_{H-IDI} has a compatible structure to that of K_s, then the non-convex H_∞ synthesis displayed in Figure 3 results in $\gamma \to \gamma_{th}$ and $K_{H-IDI} \to K_s$.

5. WIDER SIGNIFICANCE OF THE PROCEDURE

To further elaborate on the proposed procedure to tune Hybrid IDI controllers using the H_∞ LSDP, the block diagram presented in Figure 7 is used (recall that throughout the procedure W_1 and W_2 were assumed invertible).

The closed-loop implementation reduces to simply having the controller K_{H-IDI} and the plant G, as the introduced filters perfectly cancel with their respective inverses. However, by considering the disturbances and outputs entering the specified locations, the standard H_∞ LSDP framework is recovered, wherein the controller K^*_∞ is computed to robustify the shaped plant given by W_2GW_1. Referring back to the proposed procedure, it was assumed that $K_{H-IDI} \to W_1K_\infty W_2$, implying that in the ideal scenario where $K_{H-IDI} = W_1K_\infty W_2$, then $K^*_\infty = W_1^{-1}W_1K_\infty W_2W_2^{-1} = K_\infty$. However, since K_{H-IDI} is highly structured, it serves only as an approximation to $K_s = W_1K_\infty W_2$, and thus the only guarantee is that $K^*_\infty \to K_\infty$. Nevertheless, when the H_∞ norm from $[d_o\, d_i]^T \to [y_s\, u_s]^T$ is computed using this equivalent closed-loop representation, the resulting γ-value matches exactly with that obtained during the synthesis procedure displayed in Table 1. This consistency underscores the validity of the proposed method for tuning structured controllers, specifically the Hybrid IDI controller while maintaining the robustness guarantees of H_∞ LS.

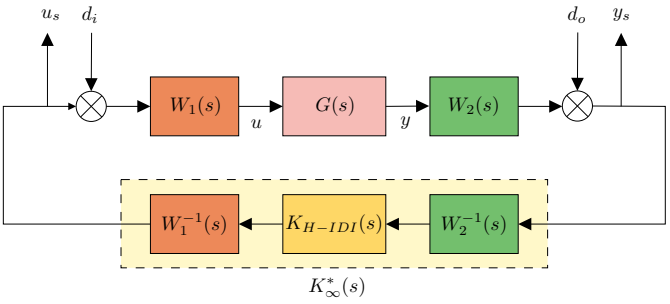

Fig. 7. Equivalent closed-loop representation with addition of "dummy" filters and their respective inverses.

Papageorgiou and Hyde (2001) make use of a similar concept to analyse the robustness of different DI controllers. The idea of the article is to generalise the analysis to any linear controller K, regardless of the design method using the NCF stability test, which is inherently integrated into the synthesis process of H_∞ LSDP.

The approach involves transforming the controller and the plant into an equivalent NCF controller and its corresponding weighted plant, as illustrated in Figure 7. The goal is to determine diagonal weights W_1, W_2 such that $W_1^{-1}KW_2^{-1}$ becomes the optimal controller for the weighted plant W_2GW_1. This is achieved by solving (Papageorgiou and Hyde, 2001):

$$\begin{aligned}\min_{W_1,W_2}\|T_{d\to e}\|_\infty = \min_{W_1,W_2}\left\|\begin{bmatrix} W_2 & \\ & W_1^{-1}\end{bmatrix}\begin{bmatrix} I \\ K\end{bmatrix}\times\right.\\ \left.\times(I-GK)^{-1}\begin{bmatrix} I & G\end{bmatrix}\begin{bmatrix} W_2^{-1} & \\ & W_1\end{bmatrix}\right\|_\infty \triangleq \frac{1}{\epsilon_\mu},\end{aligned} \tag{22}$$

where d consists of $[d_o\, d_i]^T$, e consists of $[y_s\, u_s]^T$ and W_1, W_2 are restricted to being diagonal, stable and minimum phase. The authors then make use of this tool to analyse the robustness of various DI controllers. Thus, its value remains primarily from an analysis standpoint.

If the same values for W_1 and W_2 are used, then ϵ_μ yields the same ϵ value obtained from the structured $\mathcal{H}_\infty$ minimisation. Nevertheless, the analysis tool searches for potentially very high order W_1 and W_2 weights such that the structured controller $K = K_s$ for the associated shaped plant G_s, and hence $\epsilon_\mu = \epsilon_{th}$. This study's synthesis procedure cannot guarantee this, as the synthesized controller K is computed from chosen W_1 and W_2 weights; thus one can only hope that the structure of K is 'rich' enough such that $K \to K_s$ and, therefore, $\epsilon \to \epsilon_{th}$.

On the other hand, this study's procedure retains the core purpose of W_1 and W_2 in shaping the open-loop plant and, given a successful design, bound the broken-loop gain at the plant I/O. Moreover, the synthesis results in a $\gamma = \epsilon^{-1}$ value, which ties directly to the system's robustness to NCF uncertainty, and the closer the values of ϵ and ϵ_{th}, the closer the obtained controller K is to the optimal LS controller K_s. Therefore, the procedure effectively allows designers to tune DI controllers with a-priori robustness guarantees.

6. CONCLUSION AND FUTURE WORK

This study presented the design of a longitudinal control law for the X-29 aircraft, combining the H_∞ LSDP with a Hybrid IDI structure. The framework leverages H_∞ optimization to tune a highly structured controller with the robustness guarantees of H_∞ LS. The procedure is not constrained to the specified control architecture of a Hybrid IDI and applies to other structured controllers with a compatible structure to that of $K_s = W_1K_\infty W_2$. Future studies should address the extension to 2DoF controllers, incorporating model-following requirements according to handling qualities, as well as considerations for a digital implementation. Furthermore, they should also analyse the variation of the gain-scheduled Hybrid IDI parameters across a flight envelope and reflect on the differences from a purely non-linear implementation (Hybrid INDI).

REFERENCES

Apkarian, P. and Noll, D. (2006). Nonsmooth H_∞ Synthesis. *IEEE Transactions on Automatic Control*, 51(1), 71–86.

Glover, K. and McFarlane, D. (1988). Robust stabilization of normalized coprime factors: An explicit H_∞ solution. In *1988 American Control Conference*, 842–847. IEEE.

Grondman, F., Looye, G., Kuchar, R.O., Chu, Q.P., and Van Kampen, E. (2018). Design and Flight Testing of Incremental Nonlinear Dynamic Inversion based Control Laws for a Passenger Aircraft. In *2018 AIAA Guidance, Navigation, and Control Conference*, 0385.

Harris, J.J. (2018). F-35 Flight Control Law Design, Development and Verification. In *2018 Aviation Technology, Integration, and Operations Conference*, 3516.

Hyde, R. and Glover, K. (1993). The Application of Scheduled H_∞ Controllers to a VSTOL Aircraft. *IEEE Transactions on Automatic Control*, 38(7), 1021–1039.

Jiali, Y. and Jihong, Z. (2016). An angular acceleration estimation method based on the complementary filter theory. In *2016 IEEE International Instrumentation and Measurement Technology Conference Proceedings*, 1–6. IEEE.

Khalil, H. (2002). *Nonlinear Systems*. Pearson Education. Prentice Hall, Upper Saddle River, NJ.

Kim, C.S., Ji, C.H., Koh, G.O., and Kim, B.S. (2021). Stability margin and structural coupling analysis of a hybrid INDI control for the fighter aircraft. *International Journal of Aeronautical and Space Sciences*, 22(5), 1154–1169.

Kumtepe, Y., Pollack, T., and Van Kampen, E. (2022). Flight Control Law Design using Hybrid Incremental Nonlinear Dynamic Inversion. In *AIAA SciTech 2022 Forum*, 1597.

Lavretsky, E. and Wise, K.A. (2024). *Robust and Adaptive Control With Aerospace Applications*. Springer.

McFarlane, D. and Glover, K. (1992). A Loop Shaping Design Procedure using H_∞ Synthesis. *IEEE Transactions on Automatic Control*, 37(6), 759–769.

Papageorgiou, G. and Hyde, R.A. (2001). Analysing the stability of NDI-based flight controllers with LPV methods. In *Proceedings of the AIAA Guidance, Navigation and Control Conference, Montreal*.

Papageorgiou, G. and Polansky, M. (2009). Tuning a dynamic inversion pitch axis autopilot using McFarlane-Glover loop shaping. *Optimal Control Applications and Methods*, 30(3), 287–308.

Pollack, T. and Van Kampen, E. (2023). Robust Stability and Performance Analysis of Incremental Dynamic-Inversion-Based Flight Control Laws. *Journal of Guidance, Control, and Dynamics*, 46(9), 1785–1798.

Pollack, T. (2024a). *Advances in Dynamic Inversion-based Flight Control Law Design: Multivariable Analysis and Synthesis of Robust and Multi-Objective Design Solutions*. Ph.D. thesis, Delft University of Technology.

Pollack, T., Theodoulis, S., and van Kampen, E. (2024b). Commonalities between robust hybrid incremental nonlinear dynamic inversion and proportional-integral-derivative flight control law design. *Aerospace Science and Technology*, 152.

Skogestad, S. and Postlethwaite, I. (2005). *Multivariable Feedback Control: analysis and design*. John Wiley & Sons.

Stein, G. (2003). Respect the Unstable. *IEEE Control systems magazine*, 23(4), 12–25.

Wang, X., Van Kampen, E., Chu, Q., and Lu, P. (2019). Stability Analysis for Incremental Nonlinear Dynamic Inversion Control. *Journal of Guidance, Control, and Dynamics*, 42(5), 1116–1129.

Zhou, K. and Doyle, J.C. (1998). *Essentials of Robust Control*. Prentice hall Upper Saddle River, NJ.