

Forschungsbericht 2024-23

**Auto-Tuned Multi-Objective
Structured H_2/H_∞ Control Design
Methodology for Lidar-Based
Gust Load Alleviation**

Davide Cavaliere

Deutsches Zentrum für Luft- und Raumfahrt
Institut für Flugsystemtechnik
Braunschweig



DLR

Deutsches Zentrum
für Luft- und Raumfahrt

Forschungsbericht 2024-23

Auto-Tuned Multi-Objective Structured H_2/H_∞ Control Design Methodology for Lidar-Based Gust Load Alleviation

Davide Cavaliere

Deutsches Zentrum für Luft- und Raumfahrt
Institut für Flugsystemtechnik
Braunschweig

207 Seiten
66 Bilder
6 Tabellen
99 Literaturstellen



Herausgeber:

Deutsches Zentrum
für Luft- und Raumfahrt e. V.
Wissenschaftliche Information
Linder Höhe
D-51147 Köln

ISSN 1434-8454
ISRN DLR-FB-2024-23
Erscheinungsjahr 2024
DOI: [10.57676/0v1b-0a85](https://doi.org/10.57676/0v1b-0a85)

Erklärung des Herausgebers

Dieses Werk wird unter den Bedingungen der Creative Commons Lizenz vom Typ Namensnennung 4.0 International, abrufbar über <https://creativecommons.org/licenses/by/4.0/legalcode> , zur Nutzung überlassen.

Lizenz



Creative Commons Attribution 4.0 International

Flexible Flugzeuge, Böenlasten, Doppler-Wind-Lidar, Aktive Böenlastabminderung, Vorschausteuerung, Aeroelastik, Flugdynamik, Strukturierte H_2/H_∞ Reglerentwurf

Davide CAVALIERE

DLR, Institut für Flugsystemtechnik, Braunschweig, Deutschland

Automatisierte Multikriterielle Strukturierte H_2/H_∞ Reglerentwurf-Methodik für Lidar-basiertes Böenlastabminderung

Dissertation, Technische Universität Braunschweig

Flugzeugstrukturen werden zum Teil durch Böenlasten dimensioniert. Böenlastabminderung (GLA) zielt darauf ab, die Effizienz von Flugzeugen zu verbessern, indem diese Lasten durch eine Regelungsfunktion reduziert werden. Die dynamische und unvorhersehbare Natur von Böen sowie Systemverzögerungen und -dynamiken stellen jedoch eine Herausforderung dar. Doppler-Wind-Lidar-Sensoren ermöglichen eine frühzeitige Erkennung atmosphärischer Störungen und verbessern die erreichbare GLA-Leistung erheblich. Eine Lidar-basierte GLA-Reglerentwurfsmethodik muss die geschätzten Windfeldinformationen optimal nutzen und gleichzeitig Systembeschränkungen, sekundäre Ziele, Robustheit und Zertifizierbarkeit berücksichtigen. Strukturierte H_∞ -optimale Vorschausteuerung (preview control) hat sich als geeignet erwiesen, erfordert jedoch aufgrund der schlechten Verbindung zwischen der H_∞ -Norm und der vorübergehenden Reaktion des Systems im Zeitbereich, sowie der Komplexität der lidar-basierten Windschätzung ein schwieriges und zeitaufwändiges Entwurfsprozess. Dieser Arbeit zielt darauf ab, diese Probleme durch die Erweiterung der strukturierten optimalen Vorschausteuerungsmethodik zu beseitigen. Es wird eine neue Art Böenlastspezifikation basierend auf der H_2 -Norm und ein neuartiger Böen-Impulsfilter eingeführt, der eine systematische Abstimmung der Böenreaktion ermöglicht. Es wird eine Methode zur linearen Modellierung des Lidar- und Windschätzungssystems entwickelt, damit die Reglersynthese seine Eigenschaften berücksichtigt. Schließlich wird auf Basis der beiden vorherigen Entwicklungen ein automatisiertes iteratives Regelungsoptimierungsschema vorgeschlagen und entwickelt.

Flexible aircraft, Gust loads, Doppler wind lidar, Active load alleviation, Preview control, Aeroelasticity, Flight dynamics, Structured H_2/H_∞ control

(Published in English)

Davide CAVALIERE

German Aerospace Center (DLR), Institute of Flight Systems, Braunschweig, Germany

Auto-Tuned Multi-Objective Structured H_2/H_∞ Control Design Methodology for Lidar-Based Gust Load Alleviation

Dissertation, Technische Universität Braunschweig

Aircraft structures are sized, in part, by gust loads. Gust load alleviation (GLA) seeks to improve aircraft efficiency by reducing these loads through active control; however, the dynamic and unpredictable nature of gusts, combined with system delays and dynamics, creates a challenging control problem. Doppler wind lidar sensors allow atmospheric disturbances to be detected with ample lead time, strongly improving attainable GLA performance. A lidar-based GLA control design methodology must aim to make full use of estimated wind field information while addressing system limitations, secondary objectives, actuator usage, robustness, and certifiability. Structured H_∞ -optimal preview control has proven suitable but often requires a difficult and time-consuming design process due to the poor connection between the H_∞ norm of a transfer function and its corresponding time-domain gust response, as well as the difficulty of addressing the dynamics and uncertainties of the lidar-based wind estimation system. This thesis aims to address these obstacles by extending the structured optimal preview control methodology. A new type of discrete gust specification based on the H_2 norm and a novel discrete gust impulse filter is introduced, allowing the discrete gust response to be systematically tuned. A method for modeling the lidar and wind estimation system is developed, integrating its characteristics into control synthesis. Finally, on the basis of the previous two developments, an automated iterative control tuning scheme is proposed and developed.

TU Braunschweig – Niedersächsisches
Forschungszentrum für Luftfahrt

Berichte aus der Luft- und Raumfahrttechnik

Forschungsbericht 2024-17

**Auto-Tuned Multi-Objective Structured H_2/H_∞
Control Design Methodology for
Lidar-Based Gust Load Alleviation**

Davide Cavaliere

Deutsches Zentrum für Luft- und Raumfahrt
Institut für Flugsysteme (AS)
Braunschweig

Diese Veröffentlichung wird gleichzeitig in der Berichtsreihe „NFL
- Forschungsberichte“ geführt.

Diese Arbeit erscheint gleichzeitig als von der Fakultät für Maschinenbau der Technischen Universität Carolo-Wilhelmina zu Braunschweig zur Erlangung des akademischen Grades eines Doktor-Ingenieurs genehmigte Dissertation.

Auto-Tuned Multi-Objective Structured H_2/H_∞ Control Design Methodology for Lidar-Based Gust Load Alleviation

Von der Fakultät für Maschinenbau
der Technischen Universität Carolo-Wilhelmina zu Braunschweig

Zur Erlangung der Würde

eines Doktor-Ingenieurs (Dr.-Ing.)

genehmigte Dissertation

von: Davide Cavaliere, M.Sc.

geboren in: Rom, Italien

eingereicht am: 9. Mai 2024

mündliche Prüfung am: 15. Oktober 2024

Vorsitz: Prof. Dr.-Ing. Ulrich Römer

Gutachter: Prof. Dr.-Ing. Stefan Levedag

Prof. Dr.-Ing. Peter Hecker

Abstract

Aircraft structures are sized, in part, by gust and turbulence loads. Gust load alleviation (GLA) seeks to improve aircraft efficiency by reducing these loads through active control. However, their dynamic and unpredictable nature as well as the limitations imposed by system delays and dynamics make for a challenging control problem. Doppler wind lidar sensors allow atmospheric disturbances to be detected with ample lead time, significantly improving achievable GLA performance. Due to relatively high levels of measurement noise, such sensors require an estimator which produces a wind profile along the aircraft's direction of travel.

A lidar-based GLA control design methodology must aim to make full use of this information while respecting system limitations, addressing secondary objectives, avoiding excessive actuator use, and ensuring an adequate level of robustness, not to mention questions of certifiability. Structured H_∞ optimal preview control has proven, in many respects, to be suitable for this purpose, however it has several drawbacks. Firstly, the certification specifications partly define gust loads in terms of time-domain peak responses to discrete gusts. The link between a system's transitory time-domain response and its H_∞ norm is indirect and often poor, so tuning H_∞ specifications to meet the time-domain requirements is seldom trivial. Secondly, the lidar-based wind estimation system necessarily introduces dynamic effects and noise-induced uncertainties into the wind estimate. These have heretofore been neglected in control synthesis, forcing the designer to indirectly compensate for them by iteratively adjusting the control specifications. Together, these issues result in a difficult and unnecessarily time-consuming design process, potentially degrading its achievable performance and limiting the degree to which it can be scaled up to industrial problems.

This thesis aims to address these obstacles by developing and expanding the structured optimal preview control methodology. A new type of

discrete gust specification based on the H_2 norm and a novel discrete gust impulse filter is introduced, allowing the discrete gust response to be systematically tuned. A method for analytically computing a linear model of the combined lidar and wind estimation system is developed, and the resulting model is added to the control problem, ensuring that control synthesis takes its characteristics into account. Finally, on the basis of the previous two developments, an automated iterative control tuning scheme is proposed and developed, largely eliminating the need for manually adjusted control specifications. Though not demonstrated in the present work, this method opens the door to integrated aeroservoelastic control design in multidisciplinary aircraft design.

Zusammenfassung

Flugzeugstrukturen werden zum Teil durch Böen- und Turbulenzlasten dimensioniert. Ziel der Böenlastabminderung (Gust Load Alleviation, GLA) ist es, die Effizienz von Flugzeugen zu verbessern, indem diese Lasten durch eine aktive Regelungsfunktion reduziert werden. Die dynamische und unvorhersehbare Natur der Böen und Turbulenzen sowie die durch Zeitverzögerungen und Systemdynamik bedingten Einschränkungen stellen jedoch ein anspruchsvolles Regelungsproblem dar. Doppler-Wind-Lidar-Sensoren ermöglichen die Erkennung atmosphärischer Turbulenzen mit ausreichender Vorlaufzeit, was die erreichbare GLA-Leistung erheblich verbessert. Aufgrund des relativ hohen Messrauschens benötigen solche Sensoren einen Schätzer, der ein Windprofil entlang der Flugrichtung des Flugzeugs erstellt.

Eine Lidar-basierte GLA-Reglerentwurfsmethodik muss darauf abzielen, diese Informationen vollständig zu nutzen und gleichzeitig Systembeschränkungen zu respektieren, sekundäre Ziele anzugehen, eine übermäßige Nutzung von Aktuatoren zu vermeiden und ein angemessenes Maß an Robustheit sicherzustellen, sowie auch Fragen der Zertifizierbarkeit berücksichtigen zu können. Die strukturierte optimale H_∞ -Vorschausteu-erung (*preview control*) hat sich in vielerlei Hinsicht für diesen Zweck als geeignet erwiesen, hat jedoch einige Nachteile. Erstens definieren die Zertifizierungsspezifikationen teilweise Böenlasten in Form von Spitzenreaktionen im Zeitbereich auf diskrete Böen. Der Zusammenhang zwischen der vorübergehenden Reaktion eines Systems im Zeitbereich und seiner H_∞ -Norm ist indirekt und oft schlecht, sodass die Abstimmung von H_∞ Spezifikationen zur Erfüllung der Zeitbereichsanforderungen selten trivial ist. Zweitens bringt das Lidar-basierte Windschätzungssystem dynamische Effekte und messrauschbedingte Unsicherheiten in die Windschätzung ein. Diese wurden bisher bei der Reglersynthese vernachlässigt, was den Designer dazu zwang, sie indirekt durch iterative Anpassung der

Regelungsspezifikationen zu kompensieren. Zusammengekommen führen diese Probleme zu einem schwierigen und unnötig zeitaufwändigen Designprozess, der möglicherweise die erreichbare Leistung beeinträchtigt und die Skalierbarkeit auf industrielle Probleme einschränkt.

Dieser Arbeit zielt darauf ab, diese Probleme durch die Entwicklung und Erweiterung der strukturierten optimalen Vorschausteuerungsmethodik zu beseitigen. Es wird eine neue Art Böenlastspezifikation basierend auf der H_2 -Norm und ein neuartiger diskreter Böen-Impulsfilter (Discrete Gust Impulse Filter) eingeführt, der eine systematische Abstimmung der diskreten Böenreaktion ermöglicht. Es wird eine Methode zur analytischen Berechnung eines linearen Modells des kombinierten Lidar- und Windschätzungssystems entwickelt und das resultierende Modell zum Regelungsproblem hinzugefügt, um sicherzustellen, dass die Reglersynthese seine Eigenschaften berücksichtigt. Schließlich wird auf der Grundlage der beiden vorherigen Entwicklungen ein automatisiertes iteratives Regelungsoptimierungsschema vorgeschlagen und entwickelt, das die Notwendigkeit manuell angepasster Regelungsspezifikationen weitgehend eliminiert. Obwohl dies in dieser Arbeit nicht demonstriert wird, öffnet diese Methode auch die Tür zu einem integrierten aeroservoelastischen Reglerentwurf im multidisziplinären Flugzeugentwurf.

Acknowledgements

This thesis has been completed during my time as a research scientist at the DLR (German Aerospace Center) Institute of Flight Systems, and in particular as part of the Sustainable and Energy-Efficient Aviation (SE²A) Cluster of Excellence funded by the Deutsche Forschungsgemeinschaft (DFG).

I would like to thank Prof. Dr.-Ing. Stefan Levedag for his supervision, confidence, and support of this work. I would furthermore like to express my gratitude to Prof. Dr.-Ing. Peter Hecker, for his much-appreciated time and effort in reviewing this thesis, and to Prof. Dr.-Ing. Ulrich Römer for serving as chair of the examination board.

I must extend my deepest gratitude to my advisor, Dr.-Ing. Nicolas Fezans, without whose tireless support and invaluable insight this work would not have been possible. Special thanks to the heads of the department of Flight Dynamics and Simulation, Per Ohme and Dr.-Ing. Holger Duda, as well as Modeling and Simulation group leader Dr.-Ing. Jana Schwital, for their continuous support. I also wish to thank the Gust Load Alleviation team and especially Christian Wallace and Daniel Kiehn, whose work on control design methods and lidar modeling and simulation before and during my PhD have been crucial for much of my own work. On that note, let me also acknowledge Dr.-Ing. Patrick Vrancken from the DLR Institute of Atmospheric Physics for his work on the lidar surrogate model, and Dr.-Ing. David Quero from the DLR Institute of Aeroelasticity for graciously providing an aeroelastic model of the CRM.

I also have the pleasure of acknowledging my colleagues in the SE²A Cluster. I would particularly like to thank my fellow PhDs in project B2.3, Yannic Beyer and Julius Schultz, whose contributions to the flight dynamic model and wind estimation algorithm, respectively, play a key role in this thesis. I would also like to extend my gratitude to the members of Junior Research Group B2 and in particular Dr.-Ing. André Bauknecht

and Salvatore Asaro. Although our work together on fluidic flow actuators sadly makes only a passing appearance in this thesis, I consider it one of the defining experiences of my PhD.

On a personal note, I am deeply indebted to my friends the world over, and particularly to Hemal Bharadva, for getting me through these difficult times in (more or less) one piece. Finally, I thank my parents, Francesco and Assunta, and my sister Flavia; their support and encouragement has never wavered.

Contents

List of Figures	xvii
List of Tables	xviii
Nomenclature	xix
1 Introduction	1
1.1 Motivation	1
1.1.1 Sustainability and efficiency	1
1.1.2 Structural sizing	3
1.1.3 Active Control Technologies	4
1.2 Gust Load Alleviation Functions	7
1.2.1 Design goals	7
1.2.2 Control design considerations	8
1.2.3 Historical examples of GLA	10
1.2.4 Lidar-based GLA	13
1.3 Control Design for Lidar-Based GLA	14
1.3.1 Control design methods	15
1.3.2 Modeling lidar-based wind estimates	17
1.4 Objectives of this thesis	18
1.5 Publications	20
1.6 Structure of this thesis	22
I Background	23
2 Dynamics of Flexible Aircraft	25
2.1 Elements of aeroservoelastic flight dynamics	26
2.1.1 Rigid-body longitudinal flight dynamics	27
2.1.2 Structural dynamics	28

2.1.3	Aeroelasticity	29
2.1.4	Structural loads	30
2.1.5	Control systems	31
2.2	Model description	34
2.2.1	The SE ² A Mid-Range Aircraft	34
2.2.2	Flight dynamics model	35
2.2.3	Linearized models	36
2.3	Characteristics and Dynamics of the SE ² A MR	37
2.3.1	Reference model	37
2.3.2	Structural modes	37
2.3.3	Aeroelastic modes	38
2.3.4	Frequency responses	38
3	Atmospheric Turbulence and Gust Loads	43
3.1	Introduction to aviation turbulence	44
3.1.1	Classification of turbulence sources	44
3.1.2	Engineering models of atmospheric disturbances	45
3.2	Certification specifications and definitions	47
3.2.1	CS25.341(a): Discrete gusts	47
3.2.2	CS25.341(b): Continuous turbulence	48
3.3	Gust loads envelopes	50
4	Airborne Doppler Lidar-Based Wind Estimation	55
4.1	Wind lidar sensor system	57
4.1.1	Doppler wind lidars	57
4.1.2	Measurement geometry	58
4.1.3	Measurement noise	60
4.2	Wind field estimation	61
4.2.1	Wind field model	61
4.2.2	Measurement database	62
4.2.3	Estimation algorithm	64
4.3	Wind field resampling	68
5	Structured H_2/H_∞ Control Design for GLA	71
5.1	Optimization-based multiobjective structured control synthesis	72
5.1.1	H_2 and H_∞ optimal control	72
5.1.2	Structured controllers	75

5.1.3	Nonsmooth optimization-based control synthesis	75
5.1.4	Signal-based specifications	76
5.1.5	Robust control	77
5.2	Control design process	78
5.2.1	GLA synthesis problem	78
5.2.2	Controller structure	80
5.2.3	Model order reduction	81
5.2.4	Performance evaluation	82
5.3	GLA performance specifications	82
5.3.1	Directly-specified frequency-domain requirements	82
5.3.2	Dealing with discrete gust time-domain requirements	83
 II Contributions		 87
 6 Linear Modeling of the Lidar-based Wind Estimation System		 89
6.1	Concept definition	90
6.1.1	Legacy model	91
6.1.2	Proposed formulation of \mathbf{F}_{lidar}	92
6.2	Computing the linear filter	94
6.2.1	Frequentist interpretation of the MAP estimate	94
6.2.2	Reference measurement database	95
6.2.3	Construction of reference model \mathbf{A}	97
6.2.4	Assembling the subfilters	98
6.2.5	Simplifications and limitations	103
6.3	Validation	105
6.3.1	Evaluation of wind field estimation filter \mathbf{F}_{WFE}	107
6.3.2	Evaluation of zero-mean filter \mathbf{F}_{ZM}	110
 7 Discrete Gust Performance Specifications		 119
7.1	Discrete gusts in the frequency domain	121
7.2	Discrete gust impulse filter definition	123
7.2.1	Full-order FIR filter	123
7.2.2	Order reduction and parametrization	123
7.3	DGIF-weighted H_2 control specifications	127
7.3.1	Intuition	127

7.3.2	Controller parameter gradient	129
7.3.3	Iterative optimization	133
8	Automated Control Design	141
8.1	Concept definition	142
8.1.1	Multi-objective parameter synthesis	142
8.1.2	Definitions	143
8.1.3	Control design process	144
8.2	Control specifications	146
8.2.1	Discrete gusts	146
8.2.2	Continuous turbulence	146
8.2.3	Ride quality	147
8.2.4	Actuator use and bandwidth	148
8.3	Control design loop	148
8.3.1	Performance evaluations	149
8.3.2	Design loop iteration	150
9	Application	155
9.1	Baseline Design Case	156
9.2	Nominal Design	158
9.3	Robust Design	158
9.3.1	Design problem	158
9.3.2	Design results	161
9.3.3	Comparison to Nominal design	162
10	Conclusion and Outlook	167
10.1	Conclusion	167
10.2	Outlook	170

List of Figures

1.1	Example of wing box weight distribution and its sizing requirements.	6
1.2	Estimated reduction in structural strength requirements for the B-1 enabled by SMCS.	11
1.3	Lidar-based turbulence detection for GLA.	13
2.1	Collar's aeroelastic triangle.	27
2.2	Typical structural mode shapes.	29
2.3	Aeroelastic aircraft as a feedback system.	30
2.4	Load types and reference frame at a cut load station in the left wing.	31
2.5	Schematic representation of an aeroservoelastic aircraft with an active control system.	32
2.6	SE ² A MR flight envelope. Red lines are iso-EAS and blue lines are iso-Mach.	34
2.7	View of SE ² A MR structural model, in blue, and aerodynamic strips, in gray.	35
2.8	Reference control surface configuration	37
2.9	Mode shapes of the first four symmetrical structural modes for the MTOW mass case. The deformed airframe is colored according to vertical displacement: yellow upwards and blue downwards. The undeformed shape is light gray and semitransparent.	39
2.10	Pole map with altitude 6000 m and varying air-speed for the MTOW mass case.	40
2.11	Frequency responses of the transfer function between vertical wind input and wing root bending moment.	41

2.12	Frequency responses of the transfer function between commanded actuator deflections and wing root bending moment.	41
3.1	Discrete gust vertical wind speed profile	48
3.2	The von Kármán PSD defined in CS 25	49
3.3	Overall open-loop gust- and maneuver-load envelopes (in blue and red, respectively) as well as the gust load envelope of the reference flight point (in purple).	51
3.4	Peak upward gust-induced WRBM for the MTOW mass case across the flight envelope.	52
4.1	Schematic view of the lidar-based feedforward gust load alleviation system.	56
4.2	Doppler effect in wind lidar sensors.	58
4.3	Lidar measurement geometry.	59
4.4	Conical scan pattern measurement geometry.	59
4.5	Wind field model	62
4.6	Example of measurements contained in the measurement database. The estimation window is indicated by dashed lines and colors indicate measurement noise standard deviations. Each bin occupies the space between adjacent wind field nodes.	64
4.7	Estimates	65
4.8	Controller wind field definition	68
4.9	Controller wind field reinterpolation over time. The wind field is estimated at t_0 and t_0+100 ms (with an estimation rate of 10 Hz); in between, the controller wind field is reinterpolated from the previous estimate at 100 Hz.	69
5.1	Control synthesis problem in standard form.	73
5.2	Generic GLA control design process.	79
5.3	Control synthesis problem including lidar-based wind estimation.	80
5.4	Standard feedforward controller structure for lidar-based GLA.	81

5.5	Example of an H_∞ specification and its corresponding discrete gust performance on the WRBM of the reference model. Blue solid and dotted lines are with and without the tuned controller, respectively, and the dashed black line is the equivalent H_∞ template.	84
6.1	Schematic view of the proposed linear filter.	90
6.2	Tapped delay line $\mathbf{D}(z)$	91
6.3	Legacy formulation of the linear lidar system model \mathbf{F}_{lidar}	91
6.4	Schematic view of the proposed linear filter.	93
6.5	Geometrical relation between y_i , ξ_i , and $\boldsymbol{\theta}$	98
6.6	Surface plots of \mathbf{K}_{WFE} values before and after reinterpolation. Darker blue indicates a greater value.	101
6.7	Comparison of estimated wind field \mathbf{w}_{est} with the filtered wind field $\mathbf{w}_{filt,est}$ over 3 sequential estimation windows.	108
6.8	PSD evaluation metric computed at Node 13 (near the aircraft nose).	110
6.9	Comparison of bandwidth of PSD evaluation metric across the estimation window.	111
6.10	Standard deviation of wind estimate due to measurement noise.	113
6.11	Comparison of estimation uncertainty samples from noise-only simulations over four sequential estimation windows.	114
6.12	Power spectral density analysis of measurement noise uncertainty.	116
7.1	Power spectral densities of discrete gusts in logarithmic scale, with the von Kármán PSD plotted for comparison.	122
7.2	Characteristics of discrete gust power spectral densities.	122
7.3	Comparison of reduced-order continuous-time filters with the full-order filter.	125

7.4	Comparison of the DGIF-weighted (with $H = 70$ m) frequency response and impulse response (blue) with their respective cumulative 2-norms (solid red) and the derivatives thereof (dashed red). Based on the transfer function from vertical gust input to wing root bending moment of the reference model.	128
7.5	Distribution of gradient cosines for an $H = 70$ m discrete gust.	131
7.6	Distribution of gradient cosines for an $H = 30$ m discrete gust.	132
7.7	Iterative optimization for the wing root bending moment and an $H = 70$ m discrete gust.	135
7.8	Iterative optimization for the wing root torsional moment and an $H = 70$ m discrete gust.	136
7.9	Iterative optimization for the wing root bending moment and an $H = 30$ m discrete gust.	137
7.10	Iterative optimization for the wing root torsional moment and an $H = 30$ m discrete gust.	138
8.1	High-level schematic of the design and synthesis loops, requirements, and specifications.	144
8.2	Automated GLA control design framework	145
8.3	Characteristics of $W_{bw}(s)$ calculated with a sample set of parameter values.	149
8.4	Flowchart of the control design loop iteration rules.	150
8.5	Flowchart of the specification update rules.	151
8.6	Flowcharts of the specification update operations.	152
9.1	Requirement performance η_R at the beginning and end of the nominal control design, plotted in parallel coordinates. Requirements are met when $\eta_R \leq 1$	159

9.2	Gust load envelope of the reference flight point using from the nominal controller, plotted against the overall open-loop gust- and maneuver-load envelopes. Black chevrons indicate load requirements.	160
9.3	Requirement performance η_R over the course of the robust control design, plotted in parallel coordinates. Requirements are met when $\eta_R \leq 1$	163
9.4	Gust load envelopes of the models considered in the robust design case, plotted against the overall open-loop gust- and maneuver-load envelopes. Black chevrons indicate load requirements.	164
9.5	Comparison of $H = 70$ m discrete gust response of the reference model in open-loop, with the nominal controller, and with the robust controller.	165

List of Tables

1.1	Description of Active Control Technologies	5
2.1	SE ² A MR characteristics	34
6.1	Default lidar system parameters	106
7.1	Reduced-order DGIF parameters	126
9.1	Baseline control requirements	157
9.2	Open requirements for the robust design case.	161

Nomenclature

Abbreviations

ACT	=	Active Control Technologies
CCV	=	Control Configured Vehicle
EAS	=	Equivalent Air Speed
EASA	=	European Union Aviation Safety Agency
FAA	=	Federal Aviation Administration
FIR	=	Finite Impulse Response filter
GLA	=	Gust Load Alleviation
LOS	=	Line Of Sight
LTI	=	Linear Time Invariant
MAP	=	Maximum a posteriori
ML	=	Maximum likelihood
MLA	=	Maneuver Load Alleviation
MLW	=	Maximum Landing Weight
MTOW	=	Maximum Takeoff Weight
MZFW	=	Maximum Zero-Fuel Weight
MZFW	=	Maximum Zero-Payload Weight
OEW	=	Operating Empty Weight
PRF	=	Pulse Repetition Frequency (of the lidar)
PAP	=	Power Aperture Product
PSD	=	Power Spectral Density
RMS	=	Root mean square
TAS	=	True Air Speed

General

α	=	Angle of attack
h_{MO}	=	Maximum operating altitude.
\vec{V}_K	=	Aircraft flight path velocity vector
V_{TAS}	=	True air speed
\vec{V}_W	=	Steady wind velocity
V_{s1g}	=	1- <i>g</i> stall speed.

Gusts and turbulence

f_{gust}, ω_{gust}	=	Discrete gust characteristic frequency
H	=	Discrete gust gradient (half-length)
L	=	Continuous turbulence scale length
Ω	=	Reduced frequency ω/V_{TAS} , m^{-1}
U	=	Discrete gust velocity profile, m/s
U_{ds}	=	Discrete gust peak velocity, m/s
U_{ref}	=	Reference discrete gust peak velocity, m/s
U_σ	=	Continuous turbulence limit intensity, m/s

Lidar sensors and wind estimation

\wedge (operator)	=	Divided by individual σ_i
A	=	Forward model of the wind field estimation problem
d_i	=	Model discrepancy for measurement i
δ_{act}	=	Commanded control surface deflections
d_{lead}	=	Equivalent measurement noise (unit white noise) input signal
\mathbf{d}_{est}	=	Vector of equivalent measurement noise within the estimation window
$\mathbf{d}_{filt,est}$	=	Measurement noise uncertainty within the estimation window
$D(z)$	=	Tapped delay line
η_{apert}	=	Lidar sensor aperture angle

ϵ_i	=	Measurement noise of measurement i
f_{scan}, ϕ_{scan}	=	Lidar sensor scanning rotation rate and angle
\mathbf{F}_{WFE}	=	LTI filter approximating the lidar and wind field estimation system's response to turbulent wind fields.
\mathbf{F}_{ZM}	=	LTI filter approximating the lidar and wind field estimation system's response to measurement noise.
γ_1, γ_2	=	1st- and 2nd-derivative precision weights
\mathbf{K}_{interp}	=	Interpolation matrix which mimics controller wind field resampling
μ	=	Mean value (of a normal distribution)
n_{bins}	=	Number of measurement points along the lidar LOS
n_{meas}	=	Total number of measurements available for wind field estimation
n_{nodes}	=	Number of nodes in the estimated wind field
$\mathcal{N}(\mu, \Sigma)$	=	Normal distribution with mean μ and covariance Σ
Φ	=	Power Spectral Density (PSD)
R_i	=	LOS range of a lidar measurement
R_{min}	=	Minimum measurement range
ΔR	=	Range gate length
σ_i	=	Measurement noise standard deviation of measurement i
σ_w	=	Standard deviation of continuous turbulence field
Σ	=	Covariance (of a normal distribution)
$\boldsymbol{\theta}$	=	Vector of estimated wind field parameters (i.e. vertical wind speeds)
$\theta_{p,i}, \theta_{p+1,i}$	=	Wind field nodes adjacent to measurement i
τ_{lead}, τ_{lag}	=	Fore and aft limits (in flight time relative to the aircraft nose) of the estimation window
Δt_θ	=	Relative time between adjacent estimation wind field nodes
Δt_{ctrl}	=	Controller sampling time
\mathbf{u}_{meas}	=	Relative wind speed measurements collected by the lidar sensor
\mathbf{w}_{ctrl}	=	Controller wind field
\mathbf{w}_{est}	=	Estimated wind field
$\mathbf{w}_{filt,ctrl}$	=	Filtered vertical wind speeds at the controller wind field nodes' position

$\mathbf{W}_{filt,est}$	=	Filtered vertical wind speeds within the estimation window
\mathbf{W}_{turb}	=	Vertical wind speed caused by atmospheric turbulence
$\mathbf{W}_{turb,ctrl}$	=	Combined filtered noise and vertical wind speeds at the controller wind field nodes' positions
$\mathbf{W}_{turb,est}$	=	True vertical wind speeds within the estimation window
$w_{turb,lead}$	=	True vertical wind speed at the forwardmost point in the estimation window
$w_{turb,ctrl,lead}$	=	True vertical wind speed at the forwardmost node of the controller wind field
$\mathbf{W}_{turb,AC}$	=	Vertical wind speeds impacting the aircraft
$w_{turb,0}$	=	True vertical wind speed at the aircraft nose
\vec{x}_k	=	Flight path direction
x_{lead}, x_{lag}	=	Fore and aft limits of the wind field estimation window
Δx_θ	=	Distance between adjacent estimation wind field nodes
Δx_{ctrl}	=	Distance corresponding to one controller time step at constant airspeed
$\Delta x_{a,i}$	=	Distance between measurement i and $\theta_{p,i}$
$\Delta x_{b,i}$	=	Distance between measurement i and $\theta_{p+1,i}$
ξ_i	=	Estimated vertical wind speed at the location of measurement i
z_i	=	Measured vertical wind speed component of measurement i

Gust load alleviation control design

$\ G(s)\ _2$	=	H_2 norm of system $G(s)$
$\ G(s)\ _\infty$	=	H_∞ norm of system $G(s)$
η_{tgt}	=	Target value of output variable normalized with its open-loop value
f_s	=	Discrete sampling rate, Hz
F_{CT}	=	Filter approximating the von Kármán continuous turbulence PSD
F_{DG}	=	Discrete gust impulse filter
F_g	=	Flight profile alleviation factor
K_{weight}	=	Static gain of a weighting function
ω_{ro}	=	Roll-on frequency (for bandwidth specification), rad/s
ω_{bw}	=	Roll-off frequency (for bandwidth specification), rad/s
Φ_{turb}	=	von Kármán continuous turbulence PSD
Φ_{gust}	=	Discrete gust PSD
$\sigma(G(j\omega))$	=	Singular values of dynamical system $G(j\omega)$
t_{gust}	=	Discrete gust characteristic time, s
T_s	=	Discrete sampling time
$W_{CT}(s)$	=	Continuous turbulence specification weighting filter
$W_{DG}(s)$	=	Discrete gust specification weighting filter
$W_{bw}(s)$	=	Actuator bandwidth specification weighting filter
z_{tgt}	=	Targeted value of output variable

Chapter 1

Introduction

1.1 Motivation

1.1.1 Sustainability and efficiency

At the time of writing, civil aviation is deeply concerned with its impact on anthropogenic climate change. Although it currently only contributes a relatively modest proportion of all greenhouse gas emissions (between 3 and 6% [1]), air traffic is projected to continue growing rapidly [1] while the technologies needed for ‘decarbonization’ are far from mature [2]. To mitigate the future impact of aviation on the environment and perhaps even bring it to net zero, a wide array of new concepts and technologies have been proposed or re-proposed. These include electric and hydrogen propulsion, unconventional layouts such as blended wing bodies, high aspect ratio wings, laminar flow, and many others.

One approach to the problem is to improve aircraft efficiency, thus directly reducing emissions for conventionally-fueled aircraft and improving the feasibility of alternative energy sources. The classical constant-speed Breguet range equation below provides some basic intuition on where efforts may be focused:

[1] van der Sman *et al.*, 2021. A report on sustainable aviation in Europe.

[2] Eurocontrol, 2023. A think paper on the feasibility of zero-emission long-haul aircraft.

$$R = V \times \underbrace{\frac{1}{TSFC}}_{\text{Propulsive efficiency}} \times \underbrace{\frac{L}{D}}_{\text{Aerodynamic efficiency}} \times \underbrace{\ln\left(\frac{W_{empty} + W_{payload} + W_{fuel}}{W_{empty} + W_{payload}}\right)}_{\text{Structural efficiency}}$$

where R is the instantaneous range;¹ V is the ground speed; $TSFC$ is the thrust-specific fuel consumption; L and D are lift and drag, respectively; and $W_{payload}$, W_{fuel} , and W_{empty} are the payload weight, fuel weight, and aircraft empty weight, respectively. Broadly speaking, the terms of the equation can be separated into those related to *propulsive* efficiency, *aerodynamic* efficiency, and *structural* efficiency. R may thus be considered an indicator of aircraft efficiency: given the same fuel load, payload, and airspeed, a more efficient aircraft flies farther. Alternatively, for the same instantaneous range, a more efficient aircraft either requires less fuel or carries a heavier payload.

Historically, improvements in propulsive efficiency have played a leading role in the substantial improvement in aircraft efficiency over the last 60 years, largely driven by the transition from turbojets to turbofans with progressively higher bypass ratios [3]. Gains in aerodynamic efficiency largely depend on reductions in drag. They may be obtained by reducing parasitic drag via, e.g., natural or hybrid laminar flow, or by reducing induced drag via, e.g., higher aspect ratio wings. Structural efficiency instead refers to the fraction of the aircraft mass occupied by the airframe itself, represented here by W_{empty} . This, in turn, depends on material qualities, loads, aeroelastic stability constraints, manufacturing capabilities, and more. All three types of efficiency are, to some extent, interdependent. For example, increasing wing aspect ratio to improve aerodynamic efficiency often requires a stronger, heavier structure, reducing its structural efficiency. Finding an overall ‘sweet spot’ while taking all aspects into consideration is challenging, and changes in the aircraft’s top-level requirements or the introduction of new technologies naturally affect this sweet spot.

¹ The underlying assumption is that all values remain constant; the actual range must account for variables changing over the course of the flight, e.g., changes in altitude and flight speed affecting $TSFC$.

[3] IEA, 2009. A report on energy use in transportation, including commercial aircraft.

1.1.2 Structural sizing

An aircraft's structure must be designed to withstand the most extreme structural loads which it is certified to encounter. These primarily consist of (but are not limited to) maneuver loads, gust loads, and landing loads [4, Part C]. It must furthermore provide satisfactory aeroelastic characteristics in terms of stability (flutter, divergence, and control reversal) and vibrations (for ride quality and fatigue), as well as ensure that sufficient residual strength remains in the presence of fatigue damage. The structure at any given point in the aircraft must therefore be sized, at a minimum, to satisfy the most critical of these requirements.²

To improve the structural efficiency of a certain aircraft configuration with a given mission profile (design speeds, range, payload, etc.), either the structural design must be able to meet the sizing requirements with less weight, or the sizing requirements must be relaxed. The gain in structural efficiency can also be 'invested' in design features which improve some other aspect of the aircraft (e.g., aerodynamic efficiency) without adding more structural mass, or, more likely, some compromise between the two. This could range from modest changes, such as wing tip extensions or a small increase in takeoff weight, to more substantial improvements, such as high aspect ratio wings, or avoiding design compromises such as the original Boeing 747's 'triangular' lift distribution [5].

The former may be accomplished through advanced materials such as composites. This fact has driven the development of composite aircraft over the last several decades, such as the Boeing 787 and Airbus A350. The latter approach instead calls for an alleviation of the design requirements. This may be done passively, for instance via hinged wingtips [6], or it can be achieved via active control.

² This argument is simplified for the sake of clarity. In reality, structural design is also subject to many other considerations, including for instance manufacturing constraints, engine failure loads, pressurization loads, actuation loads, lightning protection, and more.

[4] EASA, 2023. Certification specifications for large fixed-wing aircraft in the European Union.

[5] Allison *et al.*, 1978. A Boeing report on technologies for improving aircraft efficiency.

[6] Castrichini *et al.*, 2015. Paper discussing folding wing tips for passive load alleviation.

1.1.3 Active Control Technologies

Starting principally in the late 1960s, the term Active Control Technology (ACT) and accompanying terms such as Control Configured Vehicles (CCV) and ‘fly-by-wire’ started to become commonplace in the aeronautical literature. The precise definitions and scope of these terms vary considerably, but it can be said without loss of generality that they are fundamentally concerned with improving aircraft performance through the active use of control surfaces. In most cases, these control functions serve to relax aircraft design requirements by augmenting or replacing some function of the bare airframe. Reference [7] provides one classification of the principal ACT functions, summarized (with some modification) below in Table 1.1 along with a broad description of the expected benefits. As with most aspects of aircraft design, these control functions tend to be interrelated, although the nature and extent of the interaction strongly depend on the specific aircraft configuration and the implementation of the system. For example, the aeroelastic damping provided by a flutter suppression function can also improve ride quality and reduce gust and fatigue loads.

Figure 1.1, adapted from a Boeing design study presented in [5], is an example of the potential effect of ACT on structural weight. Figure 1.1a shows the distribution of structural mass across the span of a typical wing sized primarily by maneuver loads. If it is then resized considering the use of MLA, as in Fig. 1.1b, flutter requirements or gust loads become critical for large parts of the wing. In this instance, GLA and flutter suppression could be used to obtain a further improvement in wing box weight by pushing down the gust load and flutter clearance curves. Instead, the design in Fig. 1.1a cannot be improved without first employing MLA because it is sized almost everywhere by maneuver loads. A more modern and detailed example for an Airbus-style widebody aircraft is discussed in [8].

In the decades since the advent of ACT, the adoption of these technologies in transport aircraft has had mixed success. Some form of stability

³ This definition is not universal; some sources will use GA and GLA interchangeably, and with varying meanings.

[7] Schoenman *et al.*, 1975. A Boeing report on the potential impact of ACT on aircraft design.

[8] Bucher, 2021. Final report of German federal research project LuFo V-2 Con.Move.

Name	Description	Benefits
<i>Stability augmentation</i>	Artificial improvement of inherent stability and control characteristics to ensure satisfactory safety and handling qualities.	Reduced tail surfaces, expanded aftward CG range, reduced trim loads
<i>Flutter suppression</i>	Artificial damping of flutter modes to extend the stable flight envelope.	Reduced structural mass, improved aerodynamic efficiency
<i>Maneuver load alleviation (MLA)</i>	Redistribution of forces during maneuvering flight to minimize the peak structural loads.	Reduced structural mass, improved aerodynamic efficiency
<i>Gust load alleviation (GLA)</i>	Reduction of the transient dynamic loads caused by atmospheric turbulence.	Reduced structural mass, improved aerodynamic efficiency
<i>Fatigue reduction</i>	Reduction in the fatigue damage rate by reducing the occurrence and/or amplitude of structural loading cycles.	Reduced maintenance costs, longer aircraft operating life, reduced structural mass
<i>Ride control or Gust alleviation (GA)³</i>	Improvement in crew and passenger ride quality by suppressing objectionable accelerations.	Reduced passenger discomfort and crew fatigue, reduced structural mass

Table 1.1: Description of Active Control Technologies




Figure 1.1: Example of wing box weight distribution and its sizing requirements, modified from [5].

augmentation is now nearly ubiquitous [9]. Thanks to its relative simplicity, MLA is well understood and fairly widespread. Several examples of GA functions are also flying today, such as the Airbus A330/A340 Comfort in Turbulence function [10]. Flutter control, despite much research and optimism, has not yet been developed and certified for a production aircraft [11], although large-scale research projects such as Clean Aviation CONCERTO⁴ continue to close the gap.

The challenging nature of atmospheric turbulence and gust loads has limited the impact of GLA. The constraints imposed by control system performance and aircraft dynamics are fundamentally at odds with the need to quickly produce a strong reaction to an unexpected disturbance, and they tend to strongly limit realistically achievable performance. Faster and more powerful actuators could help close the gap, but they usually come at the cost of increased weight and power consumption, as well as potentially more critical failure cases. The need to deal with aeroelastic modes further complicates control design and introduces issues with sensor and actuator placement, flutter stability (in case of feedback control), and high-order models.

⁴ <https://www.concertoproject.eu/>

[9] Balas, 2003. A survey of industrial applications of flight control law design.

[10] Hönlinger *et al.*, 1994. A paper on structural aspects of ACT, including a relatively detailed description of the A330/A340 Comfort in Turbulence function.

[11] Livne, 2018. An in-depth review of the state-of-the-art in active flutter control.

1.2 Gust Load Alleviation Functions

1.2.1 Design goals

Fundamentally, GLA consists in reducing the limit structural gust loads produced by dynamic encounters with atmospheric disturbances. The requirements for computing these limit loads are laid out in the certification specifications (e.g., EASA CS-25 [4], FAR Part 25 [12]). At present, they specify the limit gust loads in terms of both the (deterministic) peak time-domain response to a set of upward and downward discrete gusts as well as the statistical response to continuous excitation with a given turbulent energy spectrum.

GLA functions operate within the broader context of the aircraft system, so their overall design goals are not limited to reducing gust loads. These may be summarized as follows:

- *Loads*: The most extreme loads resulting from all design load cases define the structural load envelope. The principal aim of a GLA function is to diminish and restrict the parts of the load envelope defined by gust and turbulence load cases. In general, not all parts of the aircraft are sized by gust loads, not all mass cases and flight points are critical for the load envelope, and there is no structural benefit to reducing the gust loads beyond the point in which they are no longer defining the envelope. A GLA function should therefore aim to achieve a targeted reduction of the gust load envelope at selected parts of the structure as efficiently as possible.
- *Stability*: The aircraft must be stable at all times, and in particular aeroelastic stability margins must not be compromised. Aggressive GLA functions may also risk destabilizing other control loops (e.g., flutter control laws) by saturating control actuators (especially with regard to their rate limits) and driving them into a nonlinear domain.
- *Robustness*: GLA functions must be able to operate and meet performance requirements in the face of various uncertainties and degradations. These can include uncertainties in time delays, measurements, estimates, actuator dynamics, mass and mass distribution,

[12] FAA, 2020. Certification specifications for large fixed-wing aircraft in the United States.

structural damping and stiffness, unmodelled or uncertain dynamics such as fuel sloshing, and atmospheric condition, as well as measurement noise and sufficiently probable system faults.

- *Handling qualities*: Changes in the aircraft's dynamics, 'unnatural' reactions to gust encounters, and control surface actuator saturations caused by a GLA function may affect the pilot's ability to control the aircraft. Certain common-sense measures, such as limiting actuator use and ensuring, where possible, frequency separation between flight control and GLA functions, can help avoid most problems.
- *Ride quality*: Although ride qualities are not, strictly speaking, a primary goal of GLA functions, the problems are fairly closely related. Ideally, a GLA function should be able to improve some metrics of ride quality, but worsening them (as was the case with the DLR LARS system, see p. 12) may render it unacceptable.
- *Actuator use*: A GLA function which makes excessive use of control actuators may increase maintenance costs, reduce their service life, generate extra drag, and consume more energy. The controller should ideally use only as much control authority as necessary to fulfill its requirements.

1.2.2 Control design considerations

The capabilities and performance of GLA functions are closely tied to those of the control system through which they act. This system includes sensors (e.g., accelerometers, angle of attack sensors), digital computers containing the control laws, and control actuators (e.g., hydraulic actuators driving control surfaces). Sensor type, placement, and quality determine what is sensed, when it is sensed, and what type of signal processing is required. Similarly, control actuator type, placement, and dynamics determine how much control authority is available and what can be controlled with it. The 'minimum response time' of the system strongly constrains how it can react to a sudden disturbance. It is mainly defined by transmission delays, digital sampling delays, and processing delays (e.g., from sensor signal filtering), as well as the dynamics of the actuators and their rate and acceleration limits.

In the context of a pure disturbance rejection problem such as GLA, the concept of feedback and feedforward control may be framed as:

- *Feedback*: Suppress the disturbance by reacting to deviations in the sensed state of the system.
- *Feedforward*: Suppress the disturbance by acting on the expected effect of an estimated disturbance.

In a feedback approach, the key word is *react*. By definition, the disturbance can only be sensed once the gust has already impacted the airframe and begun to generate loads. The control system thus has a very short time span in which to generate the control forces needed to mitigate the gust loads. It is not difficult to see how the minimum response time can outstrip the available time. Achieving a significant reduction in peak loads on the sizing gusts (which are of very large amplitude) is difficult, and generally requires a very aggressive controller (with consequences for the system's stability and efficiency). In some cases, it may be downright impossible, due to, e.g., rate limits, limited authority, or time delays. Indeed, most historical examples of feedback GLA have relied on a damping approach which tends to have good robustness properties and acceptable performance with relatively low gains [13]. Others have relied on a strongly nonlinear approach which seeks to avoid the problem of stability entirely, for instance by using a high activation threshold on the feedback signal and holding the peak command for some time after the function is activated [14].

Feedforward controllers cannot affect the system's pole locations, so they have no effect on stability.⁵ The advantage of feedforward control lies in the ability to use direct measurements of atmospheric disturbances before they impact the aircraft. The additional lead time allows the controller to compensate for system delays and make use of slower dynamics, for instance by pitching the aircraft to control the overall angle of attack.

⁵ Strictly speaking, this is only true for a predominantly linear system. If, for example, actuators are driven to saturation [15] or a highly flexible structure deforms excessively [16], the system may be destabilized.

[13] Anderson, 1993. A paper deriving the degree of robustness needed for active aeroelastic control based on the B-1 SMCS.

[14] Rollwagen *et al.*, 1990. A paper discussing the nonlinear behavior of the L-1011, A310, and A320 load alleviation functions.

They are, however, particularly vulnerable to performance losses due to uncertainties: there is no way to ‘know’ whether the system has behaved as expected, so there is no way to compensate for any discrepancies. Their achievable performance thus depends on the quality and lead-time of the estimated disturbance, the available control authority (including actuation dynamics), as well as the precision with which the system (including its delays) can be modeled.

1.2.3 Historical examples of GLA

GLA has been extensively researched, tested, and even included in production aircraft since the introduction of ACT. Several GA systems, which share many of the same properties, enjoyed a similar success. A few of the most significant real-world examples for which information is publicly available are briefly reported below.

Boeing B-52 LAMS The Load Alleviation and Mode Stabilization program had the goal of proving that active feedback load alleviation and structural mode stabilization functions could be developed using state-of-the-art methods, and it included a flight test campaign using a Boeing B-52E [17]. The control system included feedback from accelerometers and rate gyros throughout the airframe, and used elevator, rudder, ailerons, and spoilers to control both longitudinal and lateral modes. LAMS was able to achieve significant reductions in the continuous-turbulence loads and was considered a successful program.

Lockheed C-5A ALDCS Starting with analytical investigations made during the LAMS program, a series of load alleviation functions were developed for the C-5A transport [18]. They culminated with the Active Lift Distribution Control System (ALDCS), a function which served mainly to reduce fatigue loads via active control. It used vertical accelerometers and a pitch rate gyro to drive the ailerons and inner elevators, and was able to reduce incremental wing root bending moment in continuous turbulence by around 30% [19].

[17] Burris *et al.*, 1969. Final report of the LAMS program.

[18] Disney, 1978. AGARD paper on the development of load alleviation for the C-5A.

[19] Disney, 1977. A paper on the development of load alleviation for the C-5A.

North American XB-70 and Rockwell B-1 North American (later Rockwell) developed a structural mode control system (SMCS) based on the Identical Location of Accelerometer and Force (ILAF) concept and employed it on the XB-70 [20] and B-1 [21] supersonic bombers. ILAF consists in placing accelerometers approximately in the same location as control surfaces,⁶ passing measured accelerations through ‘pseudo-integrator’ filters (essentially first-order lowpass filters) to obtain an approximate local velocity, and feeding this signal back to the control surface actuators to dampen structural modes. The XB-70 never entered service and only had an experimental version of this system [20]. The B-1 design incorporated SMCS from the start to mitigate the effect of fuselage vibrations on ride qualities at the pilot station without excessive structural weight. Figure 1.2 illustrates the reduction in fuselage sizing requirements enabled by this system; trade-off studies indicated that approximately 4.3 metric tons of structural weight (roughly 5% of the empty weight) were saved [21].




Figure 1.2: Estimated reduction in structural strength requirements for the B-1 enabled by SMCS, from [21] (© AGARD 1978).

⁶ Also sometimes denoted *co-location*.

[20] Wykes *et al.*, 1969. AGARD paper on the design and test of ILAF modal control function on the XB-70.

[21] Wykes *et al.*, 1978. AGARD paper on the B-1 SMCS.

Lockheed L1011-500 The potential for a wing span extension on the -500 variant of the L1011 prompted Lockheed to develop an Active Control System (ACS) incorporating feedback MLA and Elastic Mode Suppression (EMS) to avoid strengthening the structure [22]. The original design also included a GLA function which would have reduced gust loads by pitching the aircraft using the stabilizer; it was not sufficiently effective and was dropped [22]. The MLA system uses vertical accelerations measured in the fuselage (where elastic contributions are minimal) to deflect the ailerons, whereas the EMS system measures vertical acceleration at the wing tips and uses this as a feedback signal to the ailerons to damp wing bending motion. Altogether, Lockheed claimed a 3% reduction in fuel consumption [23].

Northrop Grumman B-2 The B-2's combination of low wing loading, large size, and requirement for sustained high speed flight at low altitudes resulted in a strong need for GLA [24]. GLA is primarily provided by a feedforward system which uses a 'gust sniffer' (leading-edge angle-of-attack sensors) to drive a dedicated control surface at the trailing edge, pitching the aircraft into gusts [24]. Using the Dryden turbulence spectrum, root-mean-square (RMS) wing bending loads are reduced by approximately 50 %.

Airbus A320 The upward wing bending moment envelope of the Airbus A320 was dominated by gust loads. In the original design, a Load Alleviation Function (LAF) was designed to reduce it to the level of the maneuver load envelope by reducing the gust loads by approximately 15% [25]. It uses a vertical accelerometer in the forward fuselage to deflect the ailerons and outer spoilers once a certain vertical load factor was exceeded.

DLR LARS The Load Alleviation and Ride-Smoothing system was developed and tested by DLR on the VFW-614 Advanced Technologies Testing Aircraft System (ATTAS) [26]. It was a GA system which aimed

[22] O'Connell, 1980. AGARD report on the design and development of the L-1011 ACS.

[23] Stauffer *et al.*, 1978. AGARD report on the fuel-saving technologies developed for the L-1011.

[24] Britt *et al.*, 2000. Paper describing the aeroservoelastic characteristics of the B-2.

[25] Payne, 1986. A paper on the design of the Airbus A320 LAF.

[26] Hahn *et al.*, 1992. A paper discussing the design and flight test results of the LARS system.

to combine feedback control of wing bending motion using vertical accelerometers and ailerons, and feedforward control of vertical load factor using angle of attack measurements driving direct lift control flaps. The feedback control proved unstable due to control system time delays, but the feedforward control successfully reduced vertical accelerations [26]. Interestingly, the aggressive dynamic lift reduction converted gust energy into longitudinal accelerations, effectively worsening ride qualities until a drag management system was developed to suppress those as well [26].

1.2.4 Lidar-based GLA

Starting with the ACLAIM [27] and AWIATOR [28] programs and continuing through to the present day [29][30] [31], wind lidar systems capable of detecting the relative wind vector well ahead of the aircraft for the purpose of GLA have been developed. Such systems make use of the Doppler effect in laser light backscattered by aerosols or air molecules to measure the relative velocity of a volume of air. In combination with an appropriate scan pattern and wind estimator, the turbulent wind profile along the flight path can be reconstructed [32].




Figure 1.3: Lidar-based turbulence detection for GLA.

[27] Soreide *et al.*, 2000. A paper describing the design and flight test results of the ACLAIM lidar system.

[28] Rabadan *et al.*, 2010. A paper about the airborne lidar system developed for AWIATOR.

[29] Vrancken *et al.*, 2022. A paper discussing modern developments in airborne direct-detection Doppler wind lidar sensors.

[30] Hamada *et al.*, 2020. A paper describing the lidar-based wind estimation and gust alleviation system investigated by JAXA.

[31] Eberle *et al.*, 2024. A paper describing the Clean Aviation Ultra-Performance Wing (UP Wing) project.

[32] Fezans *et al.*, 2017. A paper describing a method for estimating a wind field profile from lidar measurements.

The advantages in lead time of such a sensor compared to that of a more conventional angle of attack sensor may be demonstrated via a simple example. Take an aircraft in cruise at 250 m/s TAS with a distance of 15 m between the tip of the nose and the leading edge of the wing.⁷ Assuming that a gust only begins to generate aerodynamic loads as soon as it impacts the wing, a GLA controller's lead time is (at most) equal to the distance between the point of measurement and the wing leading edge. An angle of attack vane positioned at the nose would thus have a lead time of just 60 ms, while a lidar system with an effective range of 100 m has 460 ms. This allows lidar-based feedforward GLA to comfortably compensate for system delays and dynamics, and to obtain substantial levels of load reduction without requiring particularly aggressive control deflections.

The availability of the full wind profile also opens the door to capabilities such as spatial filtering. For example, the wavelet-based filtering and smooth sigmoid-based shrinkage (SSBS) presented in [33] allows the various components present in the wind profile to be separated according to frequency and amplitude. This allows small-amplitude and high-frequency terms to be filtered out without the typical phase delays involved in classical linear filtering techniques.

1.3 Control Design for Lidar-Based GLA

In selecting a control design methodology, the ideal choice is one which makes full use of the wind information provided by the lidar while taking system limitations into consideration, addressing a multiplicity of primary and secondary design goals, and ensuring an adequate level of robustness. From a practical point of view, it should also be feasible to apply it to industrial-scale problems, to implement the resulting controllers on real aircraft hardware, and to certify them.

⁷ This roughly corresponds to a mid-size airliner such as the Airbus A320.

[33] Fezans, 2018. A paper about feedforward lidar-based GLA using wavelet filtering and SSBS thresholding.

1.3.1 Control design methods

In the literature, almost all lidar-based control design methods can be sorted into one of three approaches: adaptive, model-predictive, and linear.

In an adaptive approach, controller parameters are varied in flight so as to make the controller optimal with respect to some measured or estimated performance output. For example, in both Ref. [34] and Ref. [35] it is adapted with respect to measured symmetrical wing tip vertical acceleration minus the rigid-body acceleration. In both cases, the controller itself is some variant of an FIR filter and the adaptation mechanism is based on least-squares methods.

In this context, the adaptive control has the advantage of directly addressing one of the key weaknesses of feedforward control: its vulnerability to modeling uncertainties. On the other hand, there has yet to be a realistic path to certification of adaptive control functions [36]. The difficulty of certifying adaptive systems is mainly related to the difficulty in defining comprehensive and verifiable requirements, and, depending on the algorithm, in their lack of deterministic behavior [36]. Moreover, guaranteeing a certain level of performance and stability of an adaptive controller, especially in failure cases, is challenging [37]. Several more issues concerning its practical implementation, such as choosing an appropriate adaptation policy, may be considered, however the lack of a pathway to certification is the most serious problem.

Model predictive control (MPC) generates control inputs by solving an online optimization problem with an onboard dynamic model of the aircraft. At each time step, it finds an optimal sequence of control inputs up to a predefined horizon; during the following time step, the first set of control inputs is used and the entire sequence is reoptimized using new measurements. References [38] and [39] apply MPC for GLA, and Ref. [40]

[34] Zeng *et al.*, 2010. A paper describing adaptive lidar-based feedforward control for an aeroelastic model of an F/A-18.

[35] Wang *et al.*, 2015. A paper about adaptive lidar-based feedforward control for a HALE aircraft.

[36] Bhattacharyya *et al.*, 2015. A paper about certification considerations for various types of adaptive systems.

[37] Anderson *et al.*, 2008. A paper discussing past and current issues facing adaptive control.

[38] Giessler *et al.*, 2012. A paper about lidar-based MPC for GLA.

[39] De Freitas Virgilio Pereira, 2022. A PhD thesis about MPC for control of very flexible aircraft.

[40] Sato *et al.*, 2009. A paper about lidar-based MPC for GA.

does the same for GA.

One of unique strengths of the MPC approach is that, depending on how the optimization problem is defined, it can take full advantage of the system's capabilities. This may include, for example, operating control actuators up to their rate limits with no adverse effects. The main obstacle to using MPC for GLA lies in the real-time requirement, i.e., the need to successfully solve the optimal control problem within one time step [41]. Considering the complexity of aeroelastic models; the difficulty of solving the nonlinear control problem deterministically; the link between convergence and stability proofs; the need for full state knowledge; and the tradeoff between time-step length, performance, and computational cost; it is clear that the question of implementability and certifiability of MPC controllers is not yet resolved.

Linear control approaches can cover a wide range of possibilities. In the context of lidar-based control, almost all examples in the literature adopt some form of H_2/H_∞ -optimal *preview control*. References [42] and [43] make use of LMI-based H_2 optimal control for GA; Refs. [44], [45], and [46] use structured H_∞ control for GLA; Ref. [47] uses both full-order and structured H_∞ along with μ -synthesis; Ref. [48] uses full-order H_2 and H_∞ as well as structured H_∞ synthesis; and Refs. [49] and [50] use mixed H_2/H_∞ control.

[41] Kopf *et al.*, 2018. A paper discussing the opportunities and challenges of using MPC for load alleviation.

[42] Hamada, 2013. A paper about LMI-based H_2 and impulse-to-peak preview control for GA using lidar measurements.

[43] Paku *et al.*, 2016. A paper about gain-scheduled lidar-based preview GA using LMI-based H_2 synthesis together with LPV methods.

[44] Khalil *et al.*, 2021. A paper discussing combined feedforward/feedback GLA controllers including lidar-based preview information.

[45] Khalil *et al.*, 2021. A paper about multi-channel structured H_∞ control design for lidar-based GLA.

[46] Fezans *et al.*, 2022. A paper presenting lidar-based GLA results from the CleanSky 2 NACOR project.

[47] Fournier *et al.*, 2021. A paper about robust lidar-based GLA control design using H_∞ and μ -synthesis.

[48] Ting *et al.*, 2023. A paper about H_2 and H_∞ control for a wind tunnel model simulating lidar-based GLA.

[49] Fournier *et al.*, 2022. A paper about robust lidar-based GLA using mixed H_2/H_∞ control.

[50] Cavaliere *et al.*, 2022. A paper presenting a lidar-based GLA benchmark together with a mixed H_2/H_∞ control example.

Classical and modern linear control is, by now, standard practice in commercial aircraft control design [51]. Robust control methods are mature and accepted by industry; together with gain scheduling, they can account for uncertainties and variations due to mass distribution, flight point, and configuration. From the point of view of implementation and certification, only the order and transparency of the resulting controller may be problematic. Linear methods are not able to easily deal with strong nonlinearities such as actuator rate limits, so the resulting controllers must be designed to avoid them. Moreover, as is the case for H_2 and H_∞ optimal control, their specifications are defined in the frequency domain, whereas many performance requirements (such as discrete gust loads) are defined in the time-domain. The process of fine-tuning such frequency-domain specifications to obtain specific time-domain performance is seldom trivial, and often difficult and time-consuming.

1.3.2 Modeling lidar-based wind estimates

Lidar-based control design requires some kind of model of the lidar system's output. The quality of this model can be expected to affect the performance and ease of designing the controller, however the control design methodology and tractability of the control problem limit the way in which it can be modeled.

The simplest approach, found, for example in Refs. [38], [52], [34], and [35], assumes that the lidar provides a perfectly measured vertical wind speed at a point some distance ahead of the aircraft. A slightly more sophisticated variant assumes that it also provides the vertical wind speed at a discrete number of locations along the direction of travel, essentially representing the vertical wind speed profile mentioned in Sec. 1.2.4. This can be seen, for example, in Refs. [42], [43], [44], and [48]. Some approaches also try to take measurement noise and uncertainty into account by adding an uncertain contribution to each element in the wind speed profile, e.g., in Refs. [40], [47], and [53].

In reality, airborne lidar systems are far more complex. The wind

[51] Tischler *et al.*, 2017. Textbook on flight control design focusing on the CONDUIT multi-objective parameter synthesis method.

[52] Robinson, 1996. A paper in which a simple lidar-based GA controller is designed.

[53] Hamada *et al.*, 2023. A paper about robust lidar-based GA control taking wind estimation errors into account.

estimate is affected by many factors, including measurement noise, sensor resolution, scan pattern (if any), and filtering. The approach taken for the associated wind estimation process and its configuration also tends to have a strong effect [54]. The simplified models described above are needed to make the control problem tractable, but they are poor approximations of a real system. Controllers designed with such lidar models are likely to behave differently with the real system than they do in the design system. Their resulting performance may thus be unacceptable, requiring a potentially difficult redesign to reach the required performance.

1.4 Objectives of this thesis

Section 1.1 discussed how improvements in structural efficiency can benefit aircraft efficiency as well as the role that active control (and in particular GLA) can play in improving it. Section 1.2 outlined the main goals and considerations involved in designing a GLA function, reviewed several historical examples, and introduced lidar-based GLA. Section 1.3 compared state-of-the-art control design methods for lidar-based GLA, and in particular highlighted their relative advantages and shortcomings. Altogether, it may be said that lidar-based GLA has the potential to significantly affect aircraft efficiency by strongly reducing gust loads, and that in terms of certifiability and implementability, linear robust control approaches are the most feasible.

This thesis contributes to the development of control design methods for lidar-based gust load alleviation by building on the discrete-time structured H_∞ methods developed and demonstrated by Khalil and Fezans [46][55]. These methods, though powerful, are difficult and laborious for the control designer and may easily fail to exploit the full potential of a given GLA system. Combined with the simplified lidar models discussed in Sec. 1.3.2, they sometimes require many iterations to obtain an adequate controller, each of which involves simulations and a non-trivial update of the control specifications.

When applied to industrial-scale problems, where the considered requirements and design points may number in the hundreds, these short-

[54] Kiehn *et al.*, 2022. A paper investigating the frequency response of a lidar-based wind estimation system and its dependency on several parameters thereof.

[55] Khalil, 2019. A PhD thesis on the use of structured H_∞ preview control for GLA.

comings are a significant barrier to adoption. In that context, the overarching question is: how can the control design process be improved to the point where it is feasible for complex, large-scale industrial problems? Many solutions are possible; the approach pursued by this thesis is to attempt to automate the tuning process, eliminating the need for a human designer to manually iterate over the control specifications. This leads to the first scientific question addressed here:

1. How can the lidar-based gust load alleviation control design process be automated?

To automate the tuning process, it must become possible to clearly and systematically describe the manner in which control requirements are expressed and, if necessary, iterated upon. Some requirements, such as continuous-turbulence and ride quality requirements, can be easily specified in an H_2/H_∞ framework, and pose no particular problem. Discrete gust time-domain requirements, however, are more difficult. Deciding how to define or modify, say, an H_∞ specification based on time-domain simulation results of a discrete gust encounter is not simple and prone to misspecification. Finding a way to connect discrete gust time-domain performance with frequency-domain control specifications is crucial for automating the tuning process, hence, the second scientific question:

2. How can time-domain discrete gust requirements be systematically expressed in terms of frequency-domain system norms?

A further issue is that the simplified lidar models used for control design are not representative of the actual lidar-based wind estimation system. Even if full nonlinear models of the wind estimation system are used to evaluate the tuned controller's performance in simulation, the actual control synthesis cannot take its characteristics into account. The resulting controller will often fail to perform as expected, and it is not easy to find ways to compensate for the wind estimation system's limitations. In other words, *what you tune is not what you get*. The final scientific question is then:

3. How can the lidar system's behavior and limitations be modeled in the linear control problem?

1.5 Publications

The contents of this thesis have, for the most part, been published or submitted for publication. The following list includes all publications which were produced over the course of this PhD; not all of them necessarily appear in this work. Throughout the remainder of this thesis, these publications are mainly cited in reference to the contributions of other authors. For example, the maximum-a-posteriori wind estimation algorithm developed in Ref. [56] and described in Chapter 4 is not to be regarded as a contribution of this thesis.

First author publications

- [50] Davide Cavaliere, Nicolas Fezans, Daniel Kiehn, David Quero, and Patrick Vrancken, “Gust Load Control Design Challenge Including Lidar Wind Measurements and Based on the Common Research Model,” in *AIAA SCITECH 2022 Forum*, San Diego, CA, USA & Virtual: American Institute of Aeronautics and Astronautics, Jan. 2022
- [57] Davide Cavaliere, Nicolas Fezans, and Daniel Kiehn, “Method to Account for Estimator-Induced Previewed Information Losses – Application to Synthesis of Lidar-Based Gust Load Alleviation Functions,” in *Proceedings of the 2022 CEAS EuroGNC Conference*, Berlin, Germany: CEAS-GNC-2022-063, May 2022
- [58] Davide Cavaliere and Nicolas Fezans, “Recasting Discrete 1-Cosine Gust Requirements as Frequency Domain Specifications for Load Alleviation Control Design,” in *IFASD 2022*, Madrid, Spain: Paper No. 143, Jun. 2022
- [59] Davide Cavaliere and Nicolas Fezans, “Toward Automated Gust

[56] Cavaliere *et al.*, 2024. A paper which introduces a maximum a posteriori wind estimation algorithm and uses it to compute a linear model of the lidar system.

[57] Cavaliere *et al.*, 2022. A paper deriving a linear model of the lidar system based on a regularized Gauss-Newton wind estimation algorithm.

[58] Cavaliere *et al.*, 2022. A paper introducing discrete gust impulse filters and their use for GLA control design.

[59] Cavaliere *et al.*, 2024. A paper introducing discrete gust impulse filters and discussing their properties and use for GLA control design.

Load Alleviation Control Design via Discrete Gust Impulse Filters,” *Journal of Guidance, Control, and Dynamics*, pp. 1–14, Feb. 2024

- [56] Davide Cavaliere, Nicolas Fezans, Daniel Kiehn, Julius Schulz, and Ulrich Römer, “Linear Modeling of Doppler Wind Lidar Systems for Gust Load Alleviation Design,” *Journal of Guidance, Control, and Dynamics*, Vol. 47, No. 11, pp. 2351–2368, Nov. 2024
- [60] Davide Cavaliere and Nicolas Fezans, “A Practical Approach to Automated Multiobjective Gust Load Alleviation Control Design in a Structured H_2/H_∞ Framework,” in *Proceedings of the 2024 CEAS EuroGNC Conference*, Bristol, UK: CEAS, Jun. 2024

Co-authored publications

- [61] Andre Bauknecht, Yannic Beyer, Julius Schultz, Salvatore Asaro, Khalid Khalil, Ulrich Römer, Meiko Steen, Davide Cavaliere, and Nicolas Fezans, “Novel Concepts for Active Load Alleviation,” in *AIAA SCITECH 2022 Forum*, San Diego, CA & Virtual: American Institute of Aeronautics and Astronautics, Jan. 2022
- [46] Nicolas Fezans, Christian Wallace, Daniel Kiehn, Davide Cavaliere, and Patrick Vrancken, “Lidar-Based Gust Load Alleviation - Results Obtained on the Clean Sky 2 Load Alleviation Benchmark,” in *IFASD 2022*, Madrid, Spain: Paper No. 155, Jun. 2022
- [62] Salvatore Asaro, Davide Cavaliere, Nicolas Fezans, and André Bauknecht, “System Identification and Control of a Circulation Control Airfoil for Gust Load Alleviation,” in *AIAA SCITECH 2023 Forum*, National Harbor, MD & Online: American Institute of Aeronautics and Astronautics, Jan. 2023
- [63] Yannic Beyer, Davide Cavaliere, Kjell Bramsiepe, Khalid Khalil,

[60] Cavaliere *et al.*, 2024. A paper about automated GLA control tuning using structured H_2/H_∞ methods.

[61] Bauknecht *et al.*, 2022. An article describing a set of novel technologies for active load alleviation investigated in the SE²A Cluster of Excellence.

[62] Asaro *et al.*, 2023. An article describing the system identification and feedforward GLA control of a circulation control airfoil wind tunnel model.

[63] Beyer *et al.*, 2023. A paper describing the SE²A MR flight dynamics model, its capabilities, and its gust loads.

André Bauknecht, Nicolas Fezans, Meiko Steen, and Peter Hecker, “An Aeroelastic Flight Dynamics Model for Gust Load Alleviation of Energy-Efficient Passenger Airplanes,” in *AIAA AVIATION 2023 Forum*, San Diego, CA and Online: American Institute of Aeronautics and Astronautics, Jun. 2023

1.6 Structure of this thesis

The contents of this thesis are divided into two parts: Part I includes the chapters covering background knowledge, and Part II contains the contributions of this thesis.

In Part I, Chapter 2 discusses the dynamics of flexible aircraft, and introduces the SE²A Mid-Range aircraft, together with the reference model which is used throughout this work. Chapter 3 describes the types and models of atmospheric turbulence and gusts along with the gust and turbulence load certification specifications, and presents open-loop load envelopes for the SE²A MR. Chapter 4 provides an introduction to Doppler wind lidar sensors and presents a maximum-a-posteriori wind estimation algorithm. Chapter 5 concerns structured H_2/H_∞ optimal control methods and their use for GLA control design.

In Part II, Chapter 6 describes a method by which a linear model of a lidar system may be analytically computed from the estimation algorithm presented in Chapter 4, with the aim of addressing research question 3. Chapter 7 develops a low-order parametrized filter whose impulse response is a one-minus-cosine discrete gust, and goes on to discuss its use as a weighting filter for discrete-gust H_2 performance specifications, thus providing an answer to research question 2. Chapter 8 presents a method for automatically designing GLA controllers by using heuristics to iterate over the control specification parameters, thus addressing research question 1. Finally, Chapter 9 applies the developed methods to a control design problem.

Part I

Background

Chapter 2

Dynamics of Flexible Aircraft

Contents

2.1	Elements of aeroservoelastic flight dynamics	26
2.1.1	Rigid-body longitudinal flight dynamics	27
2.1.2	Structural dynamics	28
2.1.3	Aeroelasticity	29
2.1.4	Structural loads	30
2.1.5	Control systems	31
2.2	Model description	34
2.2.1	The SE ² A Mid-Range Aircraft	34
2.2.2	Flight dynamics model	35
2.2.3	Linearized models	36
2.3	Characteristics and Dynamics of the SE ² A MR	37
2.3.1	Reference model	37
2.3.2	Structural modes	37
2.3.3	Aeroelastic modes	38
2.3.4	Frequency responses	38

Aircraft structures are made of predominantly elastic materials which deform and vibrate under loads. For a sufficiently stiff structure, the natural frequencies of the vibration modes are significantly higher than those of the rigid-body dynamics, allowing the aircraft to be treated as a rigid body. By way of the square-cube law,⁸ the larger the aircraft, the stronger the influence of flexibility, and the slower the flexible modes become; a similar trend occurs when reducing the structural strength for a given aircraft design. For modern air transports, the effects of structural dynamics and aeroelasticity on flight dynamics, loads, and control systems are impossible to ignore.

This chapter seeks, firstly, to establish a qualitative understanding of the dynamics of flexible aircraft which are relevant to the gust load alleviation problem. The subject matter is therefore treated in a primarily conceptual manner; a rigorous treatment goes beyond the necessary scope of this work. Secondly, it provides a brief description of the reference aircraft model used in this work, the SE²A Mid-Range aircraft, as well as an overview of its main characteristics and dynamics. This chapter, like the rest of this thesis, is focused on symmetrical, longitudinal dynamics due to their greater relevance to GLA.

2.1 Elements of aeroservoelastic flight dynamics

The dynamic behavior of an aircraft in flight is mainly determined by the interaction between three types of forces: elastic, inertia, and aerodynamic forces [64]. Collar's aeroelastic triangle [65], shown in Fig. 2.1, illustrates the resulting disciplines. Flight dynamics deals with 'rigid-body' interactions between aerodynamics and inertial properties;⁹ structural dynamics is concerned with *in-vacuo* vibrations; static aeroelasticity

⁸ For a proportional increase in dimensions, the surface area increases with the square of the proportion and the volume (and mass) with its cube.

⁹ This has not been true, strictly speaking, since large, flexible transonic aircraft such as the B-47 were designed in the 1950s [66], and static aeroelastic effects such as control reversal were known of even earlier. Nevertheless, whenever possible, structural dynamics and rigid-body flight mechanics are treated separately.

[64] Wright *et al.*, 2015. A textbook about aircraft loads and aeroelasticity.

[65] Collar, 1978. Paper discussing the first fifty years of aeroelasticity.

considers steadily-applied aerodynamics loads; and dynamic aeroelasticity deals with situations in which all three forces interact dynamically. If an active control system is also taken into consideration, the aircraft becomes aeroservoelastic.

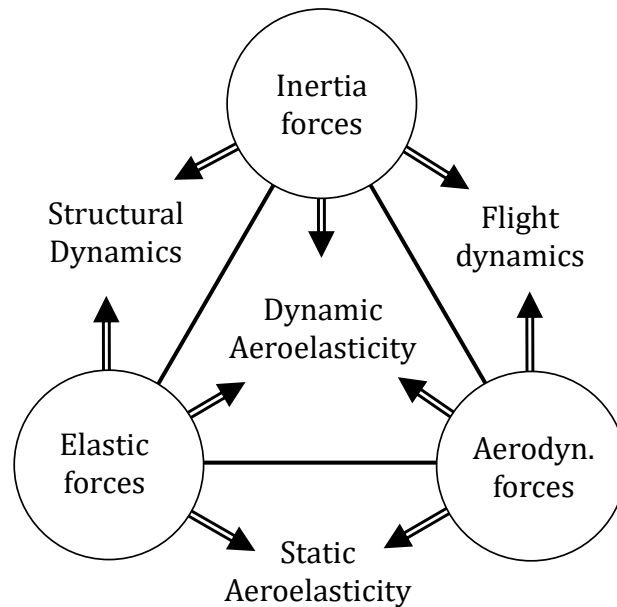


Figure 2.1: Collar's aeroelastic triangle, based on [64] and [65].

2.1.1 Rigid-body longitudinal flight dynamics

The longitudinal flight dynamics of a rigid, conventional fixed-wing aircraft at normal flight speeds are characterized by two second-order dynamic modes: the Phugoid (*PH*) and the Short Period (*SP*) modes [67]. The Phugoid mode represents the dynamic equilibrium between gravitational potential energy and kinetic energy. It is typically slow (below 0.2 Hz [67]) and poorly damped (or even mildly unstable). The Short Period mode mainly involves pitching motion and angle of attack (and therefore lift). It characterizes the aircraft's longitudinal maneuverability, as its short-term response to an elevator deflection or to a change in angle of attack is dominated by the Short Period mode. For a typical transport aircraft, it lies in the range of 0.1 to 1 Hz [67] and is well damped.

[67] Pratt, 2000. A textbook about flight control systems for fixed-wing aircraft.

2.1.2 Structural dynamics

Aircraft structures are commonly modeled using the Finite Element method, i.e., discretized into a set of ‘nodes’ connected via structural elements to neighboring nodes. Each node represents a set of degrees of freedom (displacements and rotations); they often also act as ‘lumped masses’, possessing a set of inertial properties. The elements connecting nodes define the interaction between them in the form of elastic forces, and external forces are applied directly to the nodes. The general equations of motion for the whole aircraft are thus expressed [64]:

$$\mathbf{M}\ddot{\mathbf{x}} + \mathbf{K}\mathbf{x} = \mathbf{f}_{ext} \quad (2.1)$$

where the state vector \mathbf{x} contains the set of all degrees of freedom of all structural nodes, \mathbf{M} is the mass matrix, \mathbf{K} is the stiffness matrix, and \mathbf{f}_{ext} represents the external forces acting on the structure. The inertial and elastic properties of the system are thus contained in \mathbf{M} and \mathbf{K} , respectively, and the aerodynamic forces are applied via \mathbf{f}_{ext} .¹⁰

Assuming Eq. 2.1 is linear or can be linearized (hence assuming small deformations) and setting $\mathbf{f}_{ext} = 0$, the resulting eigenvalue problem can be solved to yield the structural *normal modes*. The modes are orthogonal and are characterized by their natural frequencies (eigenvalues) and mode shapes (eigenvectors). The mode shapes describe the motion of the structural nodes involved in each mode; for example, Figure 2.2 schematically illustrates two common mode shapes. The overall motion of the structure is the sum of the contributions of all mode shapes.

The flexible structure’s equations of motion are then typically converted to a modal representation. To avoid a needlessly large system, only the modes affecting the frequency range of interest are kept, and the rest are truncated. The states of this system are the *generalized coordinates* of the flexible modes. Each coordinate is physically meaningful only in combination with its mode shape: multiplying the coordinate value by the mode shape yields the displacement of the structural nodes associated with that mode. Similarly, the external forces can be converted to generalized forces.

The structural dynamics model must then be combined with the over-

¹⁰ An additional term representing structural damping, $\mathbf{C}\dot{\mathbf{x}}$, may be included on the left-hand side of Eq. 2.1. For aircraft structures, the influence of structural damping is typically negligible [64].



Figure 2.2: Typical structural mode shapes.

all motion of the aircraft. The ‘mean axes’ methods defined, e.g., in [68], are a common approach which eliminates inertial coupling between rigid and flexible motion by defining a ‘floating’ reference system. In practice, the ‘practical’ variant of the method [68] is used, whereby the rigid-body reference frame is attached to the center of mass and the remaining (rotational) inertial coupling between the flexible and rigid-body modes is neglected. For moderately flexible aircraft (i.e., relatively small deformations), such an assumption is valid.

2.1.3 Aeroelasticity

The external forces \mathbf{f}_{ext} may be separated into exogenous forces (e.g., control surface actuation, gusts) and ‘state-dependent’ aerodynamic forces. The latter of the two essentially forms a feedback loop with the flexible airframe. This has the effect of introducing coupling effects to the normal modes as well as modifying their dynamics. The resulting system is characterized by a new set of normal modes termed *aeroelastic* modes. These often retain some resemblance to the original structural modes, however in general their natural frequencies, damping, and mode shapes differ. Aerodynamic effects furthermore depend on airspeed and atmospheric conditions. While the structural modes are a property of the bare airframe and do not change with flight point (neglecting, e.g., changing fuel mass, variable geometry, thermal effects), the aeroelastic modes vary significantly and may even become unstable (i.e., flutter).

[68] Waszak *et al.*, 1988. A paper describing a mean-axis approach to modeling flexible aircraft flight dynamics.




Figure 2.3: Aeroelastic aircraft as a feedback system, inspired by [69].

2.1.4 Structural loads

Structural loads are the forces and moments borne by the structure, sometimes called ‘internal loads’ or ‘stress resultants’ [64]. These loads may be understood as the elastic forces which ‘balance out’ the external and inertia forces acting on the structure (cf. Eq. 2.1). Consequently, they can be computed from the elastic properties and the deformation of the structure. If the structure’s equations of motion have been converted into a modal representation, this computation can be performed using the generalized coordinates of the modes. This is known as the *mode displacement method* [70]. Loads are computed at *cut load stations*: positions where the airframe is ‘cut’, and the loads on the resulting section are computed by summing the forces and moments from the piece of the airframe which has been ‘cut free’ [70].

In a slender structure such as a wing, six load types corresponding to the six degrees of freedom may be identified: lateral shear force F_x , axial force F_y , vertical shear force F_z , vertical bending moment M_x , torsional moment M_y , and in-plane bending moment M_z .¹¹ Figure 2.4 shows a cut load station on the left wing of an aircraft together with the 6 load types.

¹¹ The load names used here are adapted to the particular case of the wing load station with the reference system shown in Fig. 2.4. These may change depending on the location of the load station and the orientation of the coordinate system used for loads computation. For example, in a wing M_y is the torsional moment, but in the fuselage it is the vertical bending moment.

[70] Reschke, 2006. A PhD thesis about methods for computing aircraft structural loads.




Figure 2.4: Load types and reference frame at a cut load station in the left wing.

Note that in this thesis, the axes of the local loads coordinate systems are always parallel to those of the body axes.

2.1.5 Control systems

Broadly speaking, an active control system may be viewed as a collection of *sensors*, *controllers*, and *actuators*. Sensors detect some property of the physical system and produce a corresponding signal, controllers process sensor signals to produce actuator commands, and actuators generate some form of control force. In combination with a flexible aircraft, this forms an *aeroservoelastic* system, schematically depicted in Fig. 2.5. The way in which this control system modifies the overall dynamics depends on its specific implementation.

In a modern commercial airliner, sensors for active control are generally limited to inertial sensors (linear accelerations, angular rates) and air data sensors (airspeed, altitude, AOA) [67]. Inertial sensors only detect airframe motion, including vibrations; the location (and orientation) in which they are installed on the airframe will affect which flexible motions




Figure 2.5: Schematic representation of an aeroservoelastic aircraft with an active control system.

they sense and how strongly they are sensed. Air data sensors detect the relative motion between the aircraft and the air around it. This motion is a combination of rigid-body inertial motion, wind speed, and flexible motion¹². Sensor dynamics can be modeled using 1st- or 2nd-order linear low-pass systems, but their natural dynamics are often fast enough to be neglected.

As with sensors, the way in which control actuators are positioned on the airframe determines the degree to which they can affect the various dynamic modes. Modern aircraft are typically equipped with hydraulically-driven control surfaces, principally including trailing-edge flaps, spoilers, and movable tailplanes [67]. This type of actuator produces both inertia forces (due to the inertia of the control surface itself) and aerodynamic forces. Their dynamics can usually be well represented by a 2nd-order low-pass linear system, although more or fewer states may be necessary, depending on the actuator characteristics and control system bandwidth [71]. Such actuators are also characterized by strong nonlinearities. Aside from the obvious asymmetry of one-sided surfaces such as spoilers, hydraulic actuators mainly suffer from deflection limits and deflection rate limits [72]. Future aircraft may be equipped with novel types of actuators, such as flu-

¹² This is true in principle, but in practice flexible motion is seldom significant; some exceptional cases, e.g., an AOA vane mounted on a flexible boom.

[71] McLean, 1990. A textbook about automatic flight control.

[72] NRC, 1997. A report about PIO/APC and the factors contributing to it.

idic flow control actuators [73][74]. Such actuators could be much faster and more linear within their operating limits than hydraulically-driven control surfaces without necessarily incurring the same penalties in weight and complexity.

Controllers generically refer to the flight control computers which contain and execute control software (control laws, filters, control logic, etc.). Almost all state-of-the-art transport aircraft are equipped with a digital flight control system. Such systems operate in discrete time: inputs and outputs are only updated at regular intervals in time called *time steps* and are held constant in between. The duration of each interval is denoted *sampling time* T_s , and the inverse of the sampling time is the system's *sampling rate* f_s . The effective upper bound of a discrete-time system's frequency response is its Nyquist frequency $f_s/2$.¹³ The most important impact of a digital system on control system dynamics is the time delay caused by the time step [75]. On average, this delay is $T_s/2$, with a corresponding decrease in phase with increasing frequency; at the Nyquist frequency this results in a phase shift of -90° .

Aircraft control functions, including flight control and load alleviation functions, are real-time applications. In other words, they must operate constantly, consistently, and in sync with the rest of the system. As a result, the required computations for the 'next step' must be completed within a fixed, known, precisely-defined period of time. In a digital system, this corresponds to an *effective time step* made up of number of system time steps, resulting in a reduced effective sampling rate (and hence Nyquist frequency).

¹³ To be precise, signals with frequency content above the Nyquist frequency are ambiguous due to aliasing [75]. Digital systems therefore usually include anti-aliasing filters to ensure that frequency content above the Nyquist frequency is negligible.

[73] Warsop *et al.*, 2018. A paper about development and testing of fluidic actuators for flight control.
 [74] Khalil *et al.*, 2024. A paper about fluidic flow control actuators for GLA.

2.2 Model description

2.2.1 The SE²A Mid-Range Aircraft

The reference aircraft used in this work is the SE²A Mid-Range (MR) Aircraft [76]. It is designed to match the role and capabilities of an Airbus A320-200 while balancing direct operating costs against CO₂-equivalent emissions to improve its overall sustainability. It aims to do so by exploiting aerodynamically efficient design features (reduced wing sweep, increased wingspan and aspect ratio, lower wing loading), advanced technologies (over-wing engines, laminar flow, load alleviation), and a lower and slower (and hence low-contrail) mission profile [76]. Its structure was sized taking only maneuver loads and flutter constraints into account; gust loads were ignored under the assumption that GLA functions would be able to reduce and contain them within the sizing envelope defined by the maneuver loads and flutter constraints [63]. The aircraft's flight envelope and design characteristics are shown in Fig. 2.6 and Table 2.1, respectively.




Figure 2.6: SE²A MR flight envelope. Red lines are iso-EAS and blue lines are iso-Mach.

Parameter	Value
V_{s1g}	76 m/s EAS
V_C	177 m/s EAS
M_C	Mach 0.77
V_D	192 m/s EAS
M_D	Mach 0.85
MTOW	64 158 kg
MZFW	55 771 kg
MLW	57 742 kg
h_{MO}	11.2 km

Table 2.1: SE²A MR characteristics

[76] Karpuk *et al.*, 2022. A paper about the aircraft design process for the SE²A Mid-Range aircraft.

2.2.2 Flight dynamics model

The SE²A MR's nonlinear aeroelastic flight dynamics model (FDM) is presented in [63]. Its aerodynamics model is based on the Nonlinear Indicial Functions Lifting Line method [63]. This is a medium-fidelity unsteady aerodynamics model which computes the aerodynamic coefficients at a series of 'strips' across the lifting surfaces using 2D nonlinear unsteady airfoil models (including dynamic stall effects) and corrects the resulting lift distribution for 3D effects by way of the lifting line method. In Fig. 2.7, the aerodynamic strips are in gray; control surfaces are in a darker shade of gray. The structure is modeled as a linear 'stick' model with 134 nodes derived from the full FEM model; this is converted into a modal representation of which the first (i.e., lowest-frequency) 30 modes are retained. The structural model is depicted in blue in Fig. 2.7; each dot is a node, and the 'sticks' connecting them are linear beam-like elements. The number and positions of the structural nodes generally do not coincide with those of the aerodynamic strips, so a pair of linear interpolation matrices serve as interface between the two [63]. A cut load station is positioned at each structural node, and loads are computed using the mode displacement method.




Figure 2.7: View of SE²A MR structural model, in blue, and aerodynamic strips, in gray.

Four mass cases are defined: Maximum Takeoff Weight (MTOW), Operating Empty Weight (OEW), Maximum Zero-Payload Weight (MZPW), and Maximum Zero-Fuel Weight (MZFW). At MTOW the wing tanks are fully fueled and the fuselage is fully loaded with passengers and cargo; at OEW the wing tanks are emptied of all usable fuel and the fuselage is emptied of passengers and cargo; at MZFW the fuselage is fully loaded and the wing tanks are empty; and at MZPW the fuselage is empty and the wings tanks are full.¹⁴ Together, these mass cases should represent most of the range of variations in longitudinal and structural dynamics modes which can be expected in regular service. The fuselage mass distribution mainly affects short period mode frequency and fuselage bending, whereas the fuel mass mainly affects wing modes. Both cause the distribution of trim loads to vary.

2.2.3 Linearized models

Linear state-space models of the aircraft are obtained by numerically linearizing the full nonlinear flight dynamics model around its trim point. Only symmetrical motion, disturbances, and control deflections are considered here, so lateral states (e.g., roll rate) are not included in the linearization process. The resulting models have between 1000 and 1500 states. These mainly include rigid-body states (pitch rate and angle, vertical velocity, and forward velocity), flexible mode generalized coordinates and velocities, unsteady aerodynamic states, downwash states, actuator states, and wind input delay states.

This final category warrants some explanation. The nonlinear flight dynamics model accepts as many independent wind inputs as there are aerodynamic strips on the aircraft (approximately 80). A linear model with just a single disturbance input is far more practical, so somehow the local wind inputs must be concentrated into a single one located at or near the nose of the aircraft. This is accomplished by way of a system of delays which convey the wind input along the length of the aircraft in accordance with its true airspeed. These delays are modeled using uneven-order Padé approximants with a relative degree of 1, which provides a good compromise between the typical all-pass Padé approximant and the

¹⁴ Note that fuel mass is modeled as a ‘fixed’ mass: it is rigidly fixed to a structural node, it does not slosh, and it does not change over time.

first-order lag systems suggested in [77].

2.3 Characteristics and Dynamics of the SE²A MR

2.3.1 Reference model

The ‘reference’ model of the SE²A MR used throughout this thesis corresponds to the flight point at 6000 m altitude and 177 m/s EAS at MTOW. This flight point is located where the flight envelope boundaries defined by V_C and M_C meet, and therefore has the highest TAS of the operational flight envelope. As will be discussed in Chapters 3 and 4, it is a critical case for gust loads and for lidar-based GLA. The reference control surface layout is shown in Fig. 2.8. This includes an elevator and four pairs of ailerons along the wing. The following sections illustrate several of its most important properties and dynamics.




Figure 2.8: Reference control surface configuration

2.3.2 Structural modes

As noted above, the structural model retains 30 modes, however for conciseness only the first (i.e., lowest-frequency) four symmetrical modes are

[77] Moulin *et al.*, 2007. A paper about modeling and GLA control of an aircraft with specialized control surfaces.

discussed here. These are referred to as Modes **1**, **3**, **5**, and **7**; their mode shapes for the MTOW mass case are plotted in Figure 2.9. Mode 1 is the first wing bending mode, Mode 3 combines wing torsion with rear fuselage bending and vertical nacelle displacement, Mode 5 is primarily in-plane wing bending together with lateral nacelle displacement, and Mode 7 combines second wing bending with fuselage bending.

2.3.3 Aeroelastic modes

The corresponding aeroelastic modes may be identified from the complete linear model. Figure 2.10 plots the phugoid, short period mode, and the first four symmetrical flexible modes for a set of flight points at an altitude of 6000 m and mass case MTOW over a range of airspeeds from V_{s1g} to V_C . For this aircraft and within this range of flight points, each of the selected structural modes dominates one aeroelastic mode, such that the aeroelastic modes can be clearly identified in terms of one of the structural modes. This is not to say that the structural modes remain orthogonal. For example, the structural Mode 1 is the dominant state in the aeroelastic Mode 1, but it also participates significantly in the short period mode and the aeroelastic Mode 3. It is also worth noting that the aeroelastic Mode 3 is the only aeroelastic mode whose damping decreases with increasing airspeed, and it is in fact the critical flutter mode for this aircraft.

2.3.4 Frequency responses

The wing root bending moment (WRBM) is one of the most important aircraft loads. To better understand the dynamics of the problem, the transfer functions from the wind input and from the commanded control surface deflections to the WRBM are examined. Figure 2.11 shows the magnitude of the frequency response from the wind input. In Fig. 2.11a it is plotted for a constant altitude and MTOW and with varying airspeed. By comparing it to Fig. 2.10 (which contains precisely the same flight points), one may identify the effect of the aeroelastic poles and their movement. In Fig. 2.11b, the mass cases are varied at the reference flight point of 6000 m and 177 m/s EAS. It is interesting to note how the payload mass affects the short period frequency, fuel weight affects Mode 7, and Mode 3 is affected by both.




Figure 2.9: Mode shapes of the first four symmetrical structural modes for the MTOW mass case. The deformed airframe is colored according to vertical displacement: yellow upwards and blue downwards. The undeformed shape is light gray and semitransparent.






Figure 2.10: Pole map with altitude 6000 m and varying airspeed for the MTOW mass case.

Figure 2.12 instead shows the transfer functions from the control surface deflection command inputs to the WRBM. Unsurprisingly, the elevator is far more effective than the wing-mounted control surfaces around the frequency of the short period mode. The main reason for this is that the wing-mounted surfaces mainly contribute by directly generating lift, whereas elevator deflections change the angle of attack of the entire airplane.



(a) Varying EAS at altitude 6000 m and at MTOW.



(b) Varying mass cases at altitude 6000 m and airspeed 177 m/s EAS

Figure 2.11: Frequency responses of the transfer function between vertical wind input and wing root bending moment.




Figure 2.12: Frequency responses of the transfer function between commanded actuator deflections and wing root bending moment.

Summary: This chapter has presented the fundamentals of flexible aircraft dynamics as well as the reference aircraft model used in this thesis. Aircraft structures are made of elastic materials, and are therefore flexible. Gust loads are a product of dynamic aeroelasticity, involving an interaction between aerodynamic, elastic, and inertia forces. As such, the ability of a control system to affect gust loads depends on its overall bandwidth and the placement of its sensors and actuators with respect to the aeroelastic mode shapes and frequencies.

The aircraft used in this thesis is the SE²A Mid-Range aircraft, a medium-sized airliner comparable in design and role to the Airbus A320, and which incorporates several features designed to minimize its environmental impact. It is modeled in a Matlab/Simulink-based medium-fidelity aeroelastic flight dynamics model which is capable of computing load cases and producing linearized models. The reference model of this aircraft is chosen at the maximum-TAS point of the operational flight envelope with the MTOW mass case, a flight condition which is critical in terms of both gust loads and preview time.

Chapter 3

Atmospheric Turbulence and Gust Loads

Contents

3.1	Introduction to aviation turbulence	44
3.1.1	Classification of turbulence sources	44
3.1.2	Engineering models of atmospheric disturbances	45
3.2	Certification specifications and definitions	47
3.2.1	CS25.341(a): Discrete gusts	47
3.2.2	CS25.341(b): Continuous turbulence	48
3.3	Gust loads envelopes	50

Gust loads and the atmospheric disturbances which cause them lie at the heart of the GLA problem. The parallel issues of understanding atmospheric turbulence, modeling it, and accounting for it in aircraft design, have been investigated since the early years of aviation.

This chapter aims to provide some background and intuition on the subject, as well as the principal definitions used to specify the GLA control problem later on. It briefly introduces the most important types of atmospheric turbulence and their causes, describes the current models, summarizes the design requirements included in the current certification specifications, and demonstrates their interaction with the flexible aircraft.

3.1 Introduction to aviation turbulence

Aviation turbulence is, simply put, the range of atmospheric disturbances scales which can meaningfully perturb the dynamics of an aircraft.

3.1.1 Classification of turbulence sources

From Ref. [78], aviation turbulence may be classified into 4 categories: *Low-level turbulence* (LLT), *Clear Air Turbulence* (CAT), *Mountain Wave Turbulence* (MWT), and *Turbulence in and near thunderstorms* (TNT). These are briefly described in the paragraphs below.

Low-Level Turbulence LLT refers to turbulence occurring within a couple thousand feet of ground level [78]. It includes mechanical forcing from winds interacting with the surface of the earth (including, e.g., buildings), dry convection from thermals, and low-level wind shear. From the point of view of commercial jet-powered transport aircraft, this type of turbulence tends to be of concern during takeoff, approach, and landing phases, and is generally not sizing in terms of gust loads.

Clear Air Turbulence CAT is, broadly defined, atmospheric turbulence which occurs away from easily detectable sources such as thunderstorms. The most common cause for this type of turbulence is the presence of shear-gravity waves, which are typically formed in strong vertical wind gradients (in other words, rapid variation in wind speed with increasing altitude). Such waves are unstable (compared to other types of gravity waves) and tend to quickly break down into turbulence. CAT is often found around the tropopause and in the vicinity of jet streams [78].

Mountain Wave Turbulence MWT refers to the effects of an air mass crossing over a large ridge. Depending on atmospheric conditions, this can cause a strong gravity wave which propagates downwind and/or vertically, as well as strong rotors and wind shear-induced turbulence. The gravity waves may eventually break, often in conditions similar to those which produce CAT, and with similar consequences [78]. MWT has been

[78] Lester, 1994. A book discussing the types and causes of aviation turbulence at a conceptual level.

known to produce strong turbulence encounters, and has been implicated in several serious accidents [78].

Turbulence in and near thunderstorms TNT is produced by strong convective systems such as thunderstorms. Aside from the strong up- and downdrafts inside the cloud, it includes gusts below the storm (e.g., downbursts), tornadoes, wind shear-induced eddies around and downwind of the system, gravity waves produced by winds crossing over the storm (as if it were a ridge), and shear-gravity waves above it [78].

3.1.2 Engineering models of atmospheric disturbances

Once larger-scale flow structures begin break down into turbulence, it becomes generally impossible to deal with the flow in terms of the ‘true’ distribution of wind velocities in space. Instead, it can be dealt with in a statistical manner by treating turbulence as a stationary Gaussian process characterized by an energy spectrum. A very well-known example of such a representation is the isotropic turbulence model, largely attributed to the work of Kolmogorov [79] and now better known as the *von Kármán* turbulence spectrum [80]. In broad terms, its spectral density curve across frequency can be broken up into three parts. At low frequencies, anisotropic large scale disturbances ‘generate’ turbulence. At high frequencies, turbulent energy is gradually dissipated by viscosity. The middle range of frequencies, called the ‘inertial subrange,’ is where most aviation turbulence occurs. In it, turbulent eddies break down into slightly smaller eddies with little or no loss of energy, creating a ‘cascade’ of turbulent energy [79]. The lower frequency bound of the inertial subrange depends on the characteristic dimensions of the ‘generating’ disturbance, and the curve with which the spectral density decays over frequency is $\Omega^{-5/3}$, i.e., $-5/3 \cdot 20$ dB/decade.

Experimental data largely supports the isotropic models of turbulence [81]. Together with the assumption of a stationary Gaussian random

[79] Vinnichenko *et al.*, 1968. A textbook about turbulence in the free atmosphere.

[80] Hoblit, 1988. A textbook about gust loads for fixed-wing aircraft.

[81] Houbolt *et al.*, 1964. A NASA report on atmospheric turbulence mainly focused on isotropic turbulence and power-spectral methods.

process, it can be treated using power-spectral methods, which are mathematically relatively simple and easy to apply [80]. In brief, they allow the root mean square (RMS) of a quantity to be found by computing the square root of the integral of its power spectral density (PSD) over frequency. For linear dynamical systems, the RMS of a given output can be found by taking as integrand the product of the PSD of the input with the squared magnitude of the system's frequency response (see, e.g., Eq. 3.8 below).

Real turbulence is, in general, not stationary and Gaussian, and it is only well-modeled by isotropic turbulence in a statistical sense. In terms of gust loads, three particular deficiencies stand out. Firstly, the distribution of wind speeds in a particular turbulence encounter tends to be 'strong-tailed,' i.e., more extreme wind speeds are more likely to be encountered than a Gaussian distribution would predict [82]. Secondly, for more extreme wind speeds, the spectral density of turbulence does not scale with frequency according to the $-5/3$ law predicted by the isotropic model, but is rather approximately a $-4/3$ law [82]. Finally, the highest-intensity turbulence encounters tend to occur in short bursts during which the statistical equilibrium is not reached [82]. The resulting RMS values for high-intensity turbulence may therefore differ from those for moderate turbulence, depending on the specific dynamics of the aircraft and the velocity profiles of the turbulence encounters.

One way to deal with these issues is to use time-domain evaluations of 'tuned' isolated discrete gust encounters. Such isolated gusts have a deterministic velocity profile, and are defined over a range of gust lengths. Their amplitudes are tuned so as to match the $-4/3$ spectral density slope over a range of frequencies. This is, of course, only a partial solution, as many velocity profiles are possible, and it is difficult to say a-priori which velocity profile is most critical for a given aircraft and output. Methodologies such as the Statistical Discrete Gust have attempted to address this issue [82], but so far have not been adopted.

[82] Jones, 2004. A report documenting the history and justification for the Statistical Discrete Gust method.

3.2 Certification specifications and definitions

Gust load requirements are defined identically in EASA CS 25.341 [4] and FAR Part 25 § 341 [12]; for brevity, only CS 25 will be referenced for the remainder of this work. They include both PSD-based continuous turbulence requirements and time-domain discrete gust requirements. The two most important specifications for sizing the structure are CS 25.341(a) (Discrete Gust Design Criteria) and CS25.341(b) (Continuous Turbulence Design Criteria). Their definitions are recalled below.¹⁵

3.2.1 CS25.341(a): Discrete gusts

Discrete gusts are modeled as one-dimensional 1–cosine vertical gusts:

$$U = \frac{U_{ds}}{2} \left(1 - \cos \left(\frac{\pi x}{H} \right) \right) = \frac{U_{ds}}{2} \left(1 - \cos \left(\frac{\pi t V_{TAS}}{H} \right) \right) \quad (3.1)$$

The gust speed profile U , shown in Fig. 3.1a, depends on the peak gust amplitude U_{ds} , gust gradient H , and gust penetration distance x , which may also be expressed as gust penetration time t multiplied by true airspeed V_{TAS} . The distance between the beginning and the end of the gust is $2H$. H is required to vary between 9 m (30 ft) and 107 m (350 ft).

The gust amplitude U_{ds} is defined as:

$$U_{ds} = U_{ref} F_g \left(\frac{H}{107} \right)^{1/6} \quad (3.2)$$

Figure 3.1b illustrate the relationship between U_{ds} and H . Reference speed U_{ref} decreases with altitude, decreasing linearly from 17.07 m/s (56 ft/s) EAS at sea level to 13.41 m/s (44 ft/s) EAS at 4,572 m (15,000 ft), and to 6.36 m/s (20.86 ft/s) EAS at 18,288 m (60,000 ft) [4]. These values are valid for airspeeds up to V_C ; at V_D , they are reduced by 50%. The flight profile alleviation factor F_g is designed to reduce the severity of the requirements

¹⁵ Note that only vertical gusts and turbulence are discussed hereafter due to their greater impact on wing loads, however the specifications also require lateral and (for the structure supporting wing-mounted engines) ‘round-the-clock’ gusts.

depending on the probability of the aircraft finding itself at a particular altitude. It increases linearly with altitude, going from a sea-level value $F_{g,SL}$ to 1:

$$F_g = F_{g,SL} + (1 - F_{g,SL}) \frac{h}{h_{MO}} \quad (3.3)$$

$$\begin{aligned} F_{g,SL} &= \frac{1}{2}(F_{gz} + F_{gm}) \quad (3.4) \\ &= \frac{1}{2} \left(\left(1 - \frac{h_{MO}}{250,000} \right) + \sqrt{\frac{MZFV}{MTOW} \tan \left(\frac{\pi}{4} \frac{MLW}{MTOW} \right)} \right) \end{aligned}$$

where h is the aircraft's altitude and h_{MO} its ceiling.

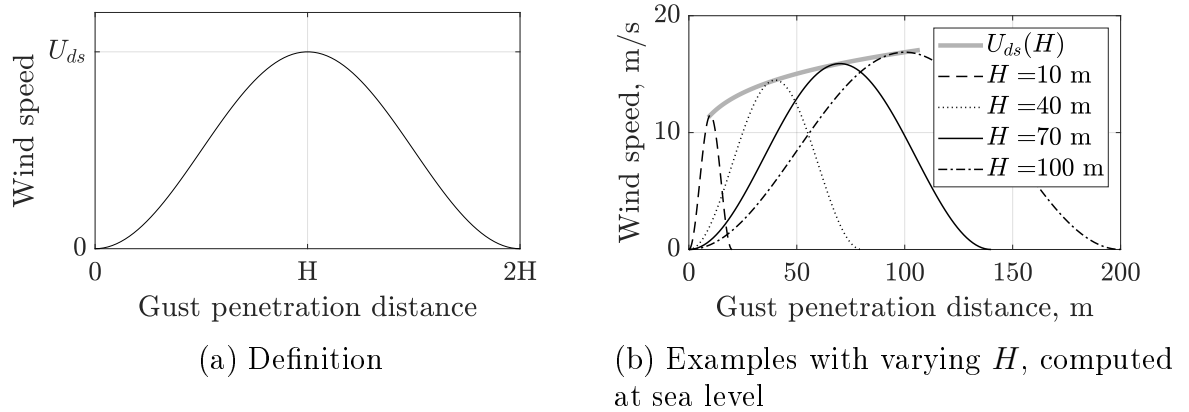


Figure 3.1: Discrete gust vertical wind speed profile

Discrete gust loads are determined in simulation. The limit values (sizing loads) are taken as the peak response across all values of H for both upward and downward gusts. If the behavior of the system is linear or nearly so, the incremental response for an upward gust is precisely equal and opposite to that of a downward gust of the same length.

3.2.2 CS25.341(b): Continuous turbulence

Continuous turbulence requirements are defined as power-spectral requirements with all assumptions noted above, i.e., that turbulence is a stationary Gaussian random process. It is characterized by the von Kármán

vertical turbulence PSD Φ_{turb} , shown in Fig. 3.2:

$$\Phi_{turb}(\Omega) = \frac{L}{\pi} \frac{1 + \frac{8}{3} (1.339 \Omega L)^2}{\left[1 + (1.339 \Omega L)^2\right]^{\frac{11}{6}}} \quad (3.5)$$

$\Omega = \frac{\omega}{V_{TAS}}$ is the reduced frequency and $L = 762$ m (2,500 ft) is the turbulence scale length. The PSD in Eq. 3.5 is defined as a function of the reduced frequency Ω , however it is sometimes useful to express it in terms of ω (rad/s) or f (Hz), for which the PSD must be rescaled [80, p. 42].

$$\Phi_{turb}(\omega) = \frac{L}{\pi V_{TAS}} \frac{1 + \frac{8}{3} (1.339 \omega L / V_{TAS})^2}{\left[1 + (1.339 \omega L / V_{TAS})^2\right]^{\frac{11}{6}}} \quad (3.6)$$

$$\Phi_{turb}(f) = \frac{2L}{V_{TAS}} \frac{1 + \frac{8}{3} (1.339 2\pi f L / V_{TAS})^2}{\left[1 + (1.339 2\pi f L / V_{TAS})^2\right]^{\frac{11}{6}}} \quad (3.7)$$




Figure 3.2: The von Kármán PSD defined in CS 25

Considering an output quantity z , e.g., the wing root bending moment, of a linear system with frequency response to vertical wind $G(\Omega)$, the limit incremental load Δz_{lim} in continuous turbulence is computed as:

$$\Delta z_{lim} = U_{\sigma} \sqrt{\int_0^{\infty} |G(\Omega)|^2 \Phi_{turb}(\Omega) d\Omega} \quad (3.8)$$

where U_{σ} is the limit turbulence intensity, defined as $U_{\sigma} = U_{\sigma,ref} F_g$. F_g is the same as in Eq. 3.4. The reference turbulence intensity $U_{\sigma,ref}$ decreases linearly with altitude. At airspeeds up to V_C , it goes from from 27.43 m/s (90 ft/s) TAS at sea level to 24.08 m/s (79 ft/s) TAS at 7,315 m (24,000 ft), and then remains constant up to 18,288 m (60,000 ft). Between V_C and V_D , it decreases linearly with airspeed such that it is reduced by 50% at V_D . The square root term in Eq. 3.8 constitutes the RMS of the output z in response to a white noise vertical wind input with a RMS of 1. It is worth noting that Φ_{turb} in Eq. 3.5 is defined for a one-sided frequency spectrum, hence the integral in Eq. 3.8 is bounded between 0 and ∞ .

3.3 Gust loads envelopes

The gust loads envelope can be obtained by computing the gust loads at all flight conditions and configurations required by CS 25. Figure 3.3 shows the gust and maneuver loads envelopes for the SE²A MR, as well as the gust loads corresponding to the reference flight point from Sec. 2.3 (p. 37). In this case, the maneuver loads envelope includes only symmetrical static and dynamic pitch maneuvers as defined in CS 25 [4].

The total loads are the sum of the trim loads and the upward and downward incremental loads. Looking at the reference gust loads, note that they are always symmetrical around the trim load, i.e., the incremental limit loads are equal in the upward and downward direction. The aircraft is in fact, almost perfectly linear with respect to wind disturbances at this flight point. In terms of the wing vertical bending moment (Fig. 3.3a), the reference model has both relatively large incremental loads and large upward trim loads due to the heavy mass case (MTOW). It is clear, then, that the reference flight point is sizing for the upward wing root bending moment, hence its criticality.

Figure 3.4 offers another perspective by plotting the peak upward gust-induced WRBM across the flight envelope. The red circle indicating the flight point with the maximum value shows, again, that the reference flight point is the most critical for this load. This plot also reveals a common

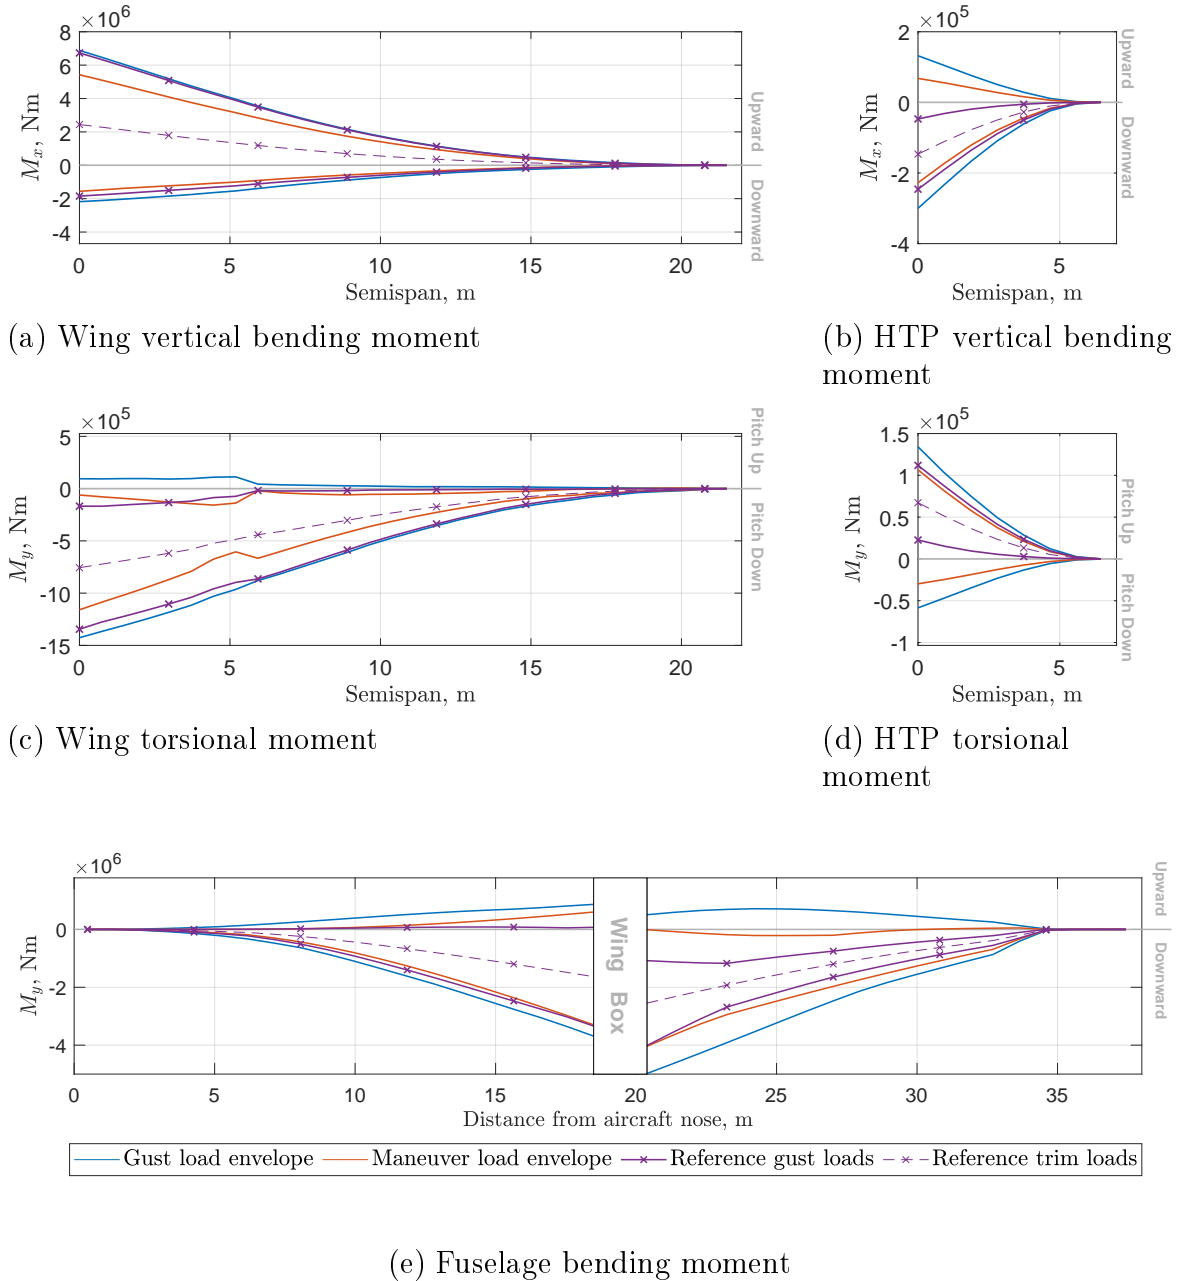


Figure 3.3: Overall open-loop gust- and maneuver-load envelopes (in blue and red, respectively) as well as the gust load envelope of the reference flight point (in purple).

(though not universal) trend in gust loads: the peak load increases with airspeed up to V_C . Between V_C and V_D , this trend is reversed as the peak gust velocities are reduced by 50%.




Figure 3.4: Peak upward gust-induced WRBM for the MTOW mass case across the flight envelope.

Summary: Atmospheric turbulence comes in many forms. For the purposes of aircraft engineering and certification, it is mainly modeled in terms of the stochastic response to continuous turbulence with a von Kármán spectral density curve, and of the peak time-domain response to a family of discrete gusts. The gust load envelopes of the SE²A MR aircraft and of the reference model have been presented.

Chapter 4

Airborne Doppler Lidar-Based Wind Estimation

Contents

4.1	Wind lidar sensor system	57
4.1.1	Doppler wind lidars	57
4.1.2	Measurement geometry	58
4.1.3	Measurement noise	60
4.2	Wind field estimation	61
4.2.1	Wind field model	61
4.2.2	Measurement database	62
4.2.3	Estimation algorithm	64
4.3	Wind field resampling	68

Airborne Doppler wind lidar sensors allow aircraft to detect relative wind speeds at a distance from the aircraft. Using such a technology to detect atmospheric disturbances well ahead of the aircraft gives a GLA function a substantial lead time, allowing it to effectively preempt gust encounters. Such a system must be capable of producing good estimates of the vertical wind speed profile and delivering them to the active control function in a timely manner.

This chapter describes one such system based on Refs. [29] [32] [56] which is used as the basis for this work. This system is designed for wind estimates with a preview distance over 100 m ahead of the aircraft and with a sampling rate sufficient for active control of the most important rigid-body and flexible modes. Only vertical turbulence and gusts are considered here due to their greater importance for gust loads, however such a system could also be used for lateral and longitudinal disturbances if necessary. Figure 4.1 shows a schematic overview of the considered lidar-based GLA system.

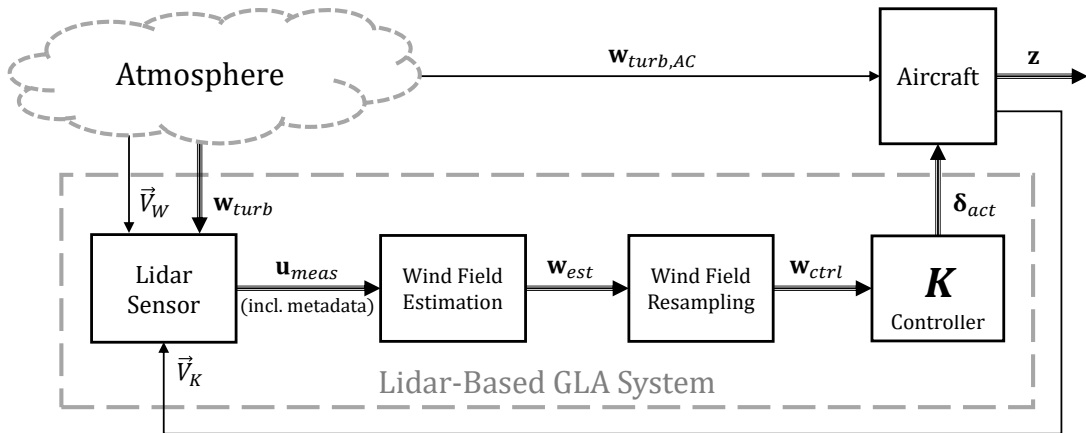


Figure 4.1: Schematic view of the lidar-based feedforward gust load alleviation system.

$\mathbf{w}_{turb,AC}$ is the part of the turbulent wind field which impacts the aircraft's aerodynamic surfaces, generating gust loads. The lidar sensor produces measurements \mathbf{u}_{meas} of the relative airspeed along its line of sight (LOS), which includes components of the aircraft's inertial velocity vector \vec{V}_K , of the steady wind speed \vec{V}_W , and of the true vertical turbulence \mathbf{w}_{turb} . These measurements, along with their associated metadata,

i.e., Earth-relative position, LOS direction, measurement range, expected measurement noise, etc. [32], are collected at a high rate (on the order of hundreds to thousands per second) and stored in a database. A wind field estimation algorithm uses the contents of this database to estimate a vertical wind field profile along the flight path \mathbf{w}_{est} at a far slower rate (around 10 Hz). The GLA controller \mathbf{K} , operating at the sampling rate of the active control system (taken to be around 100 Hz), uses the controller wind field \mathbf{w}_{ctrl} to compute actuator commands δ_{act} . A wind field resampling process serves to convert \mathbf{w}_{est} to \mathbf{w}_{ctrl} . In the following sections, each of the blocks in this diagram are described.

4.1 Wind lidar sensor system

4.1.1 Doppler wind lidars

Lidar (LIght Detection and Ranging) sensors function by emitting light, usually in the form of a laser, and measuring the backscattered returns. A Doppler lidar measures the relative velocity to the target along its line-of-sight (LOS) by measuring the shift in the frequency of the backscattered signal with respect to the emitted laser caused by the Doppler effect. Doppler wind lidar sensors exploit this method to measure the relative speed between the lidar sensor and a volume of air. For a wind lidar, backscatter is caused by a small fraction of the laser energy reflecting off of many individual particles subject to random motion. The resulting backscatter returns are distributed over a range of frequencies, as in Fig. 4.2, and the Doppler effect corresponds to the shift in this distribution with respect to the emitted laser frequency. There are two main types of Doppler wind lidars: *coherent* and *direct detection* [83].

Coherent sensors rely on *aerosol* backscatter (also known as Mie scattering). The aerosol backscatter spectrum is typically narrow, so it is relatively easy to identify the Doppler shift using electronic spectral analysis (heterodyne detection). If aerosol density is sufficiently high, it also has a relatively high signal-to-noise ratio (SNR) [83]. However, aerosol density tends to be very low at high altitudes, e.g., at typical aircraft cruising altitudes, resulting in unacceptably low SNR.

[83] Vrancken, 2016. A chapter in a book about aviation turbulence dealing with airborne lidar sensors.

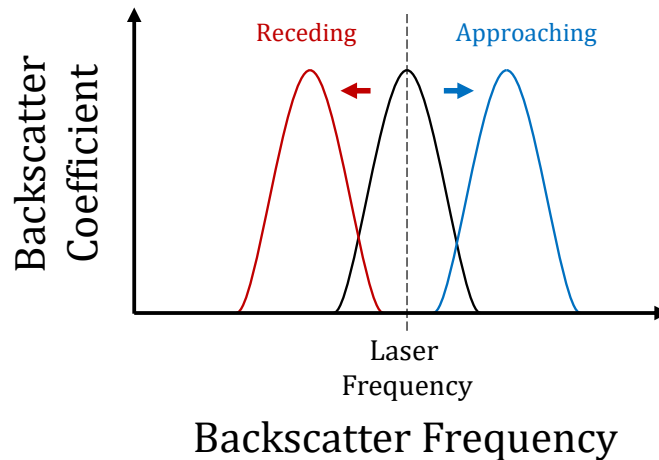


Figure 4.2: Doppler effect in wind lidar sensors.

Direct detection sensors instead use *molecular* backscatter (Rayleigh scattering) from air molecules. Molecular backscatter has a broad spectrum due to the strong influence of Brownian motion as well as relatively low backscatter coefficients, so heterodyne methods are not suitable. Instead, either *edge detection* or *fringe imaging* are used [83]. The former typically uses a pair of spectral filters whose passbands are offset above and below the laser frequency; by comparing the strength of their outputs, the direction of the Doppler shift can be determined. The latter method instead induces the backscattered signal to interfere with itself and measures the result with an interferometer; a Doppler shift creates a detectable change in the interference pattern [83]. The ability to operate in the absence of aerosols makes direct-detection sensors attractive, if not obligatory, for use in active GLA systems, however their effective range is limited by the poor ratio between backscatter intensity and emitted laser power. The lidar sensor considered hereafter is, in fact, a direct detection Doppler wind lidar with an effective range up to a couple hundred meters.

4.1.2 Measurement geometry

Figure 4.3 depicts the basic lidar measurement geometry. The lidar sensor emits laser pulses along its LOS; between pulses, backscattered radiation is measured. The range of a given return is determined from the elapsed time

between the laser pulse and the arrival of the backscattered signal. Returns within the minimum measurement range R_{min} are ignored. Following that, all returns arriving over a span of ΔR , the range gate length, are collected to produce a single measurement. Each laser ‘shot’ therefore produces a number of measurements n_{bins} along the LOS equal to the number of range gates between R_{min} and the maximum measurement range R_{max} , and the relative velocity is measured along the LOS with a resolution of ΔR .




Figure 4.3: Lidar measurement geometry.

Measurements consist of the relative wind speed along the LOS u_{meas} , which is the projection of the 3D relative velocity \vec{V}_{rel} onto the LOS axis. To allow the wind field estimation process to discriminate the vertical turbulence from other wind components, measurements must be taken in multiple different LOS directions. Here, a conical scan pattern is used (Fig. 4.4), in which the LOS rotates (approximately) around the aircraft’s longitudinal axis. This type of scan pattern is also in, e.g., Refs. [32] [84] [50]; it has the advantage of allowing the wind velocity in all three axes to be discriminated.




Figure 4.4: Conical scan pattern measurement geometry.

[84] Fezans *et al.*, 2019. A paper about combined feedback and lidar-based feedforward GLA.

From Fig. 4.4, the LOS direction relative to the axis of rotation is defined by two angles: the aperture angle η_{apert} and the scan angle ϕ_{scan} . Making the assumption that only vertical turbulence w_{turb} is present (i.e., no lateral or longitudinal turbulence), that the aircraft is in symmetrical level flight, and that the AOA is small enough that $\cos \alpha \approx 1$, the relative wind speed u_{meas} measured by the lidar at measurement point i is:

$$u_{meas,i} = V_{TAS} \cos \eta_{apert} + w_{turb,i} \sin \eta_{apert} \sin \phi_{scan,i} + \epsilon_i \quad (4.1)$$

where ϵ_i is the measurement noise, and the true airspeed V_{TAS} represents the combination of inertial speed \vec{V}_K and steady wind \vec{V}_W . The assumptions above are not strictly necessary, but they greatly simplify Eq. 4.1 by making the longitudinal axis and the air path parallel. In practice, the position and orientation of measurements are saved in absolute spatial coordinates together with the measured speeds. This allows small deviations to be compensated, however the assumptions regarding symmetrical level flight and negligible AOA are suitable for the high-speed cruise conditions critical to GLA. Situations such as steeply banked turns or high-AOA flight are more complicated but at the same time less critical in terms of gust loads.¹⁶ For the sake of simplicity, V_{TAS} shall furthermore be assumed equal to V_K for the remainder of this work.

4.1.3 Measurement noise

The key measure of sensor performance is the wind speed measurement noise. In this type of sensor, measurement noise is dominated by ‘shot noise’, a Poisson process related to the discrete number of backscattered photons which impact the detector during each measurement [85]. The number of photons involved is large enough that the resulting Poisson distribution can be approximated with a Gaussian distribution, hence it assuming the measurement noise is Gaussian. For an individual measurement i , the noise standard deviation σ_i can be computed from a simplified surrogate model [29, Eq. 8]:

$$\sigma_i = \kappa_{sens} \cdot \kappa_{atm} \cdot R_i \sqrt{\frac{f_{refresh}}{PAP \cdot \Delta R}} \quad (4.2)$$

¹⁶ Simultaneous gust and maneuver loads are not currently considered in the certification specifications. Low-speed, high-AOA flight is also usually not critical, see for instance Fig. 3.4, p. 52.

[85] Rye *et al.*, 1993. A paper on backscatter peak estimation for Doppler wind lidar sensors.

The three terms in Eq. 4.2 represent the effects of the specific hardware implementation of the sensor (κ_{sens}), of the local atmospheric conditions (κ_{atm}), and of the principal lidar system design variables, respectively. The system design variables include the measurement range R_i (referring to the middle point of that measurement's range gate, cf. fig. 4.3), the range gate length ΔR , the refresh rate $f_{refresh}$, and the power aperture product PAP .

PAP is the product of the lidar's transmitted power P_{Tx} and of the area of the receiving telescope's objective A_{Rx} ; it is a general indicator of the sensor's performance and sizing. Transmitted power P_{Tx} is, in turn, the product of pulse energy E_p and pulse repetition frequency PRF . PRF is the rate at which laser pulses are emitted, whereas $f_{refresh}$ is the rate at which measurements are produced. The difference is that the returns from a number of successive pulses may be averaged to reduce their noise. However, due to the fact that the LOS moves constantly in the scan pattern considered here, PRF is considered equal to $f_{refresh}$. ΔR determines the length of time over which photons for a given measurement are collected; the choice of ΔR is effectively a compromise between measurement noise and spatial resolution. The most important variable is the measurement range, which is linear with respect to the noise standard deviation.

4.2 Wind field estimation

4.2.1 Wind field model

The wind field profile is estimated by identifying the parameters of a wind field model. This model is shown in Fig. 4.5. It consists of a number n_{nodes} of evenly-spaced nodes aligned with the aircraft's predicted path. The space between the first and last nodes is the *estimation window* and is defined relative to the aircraft by a lead time τ_{lead} and lag time τ_{lag} . The model's parameter vector $\boldsymbol{\theta}$ has n_{nodes} elements whose values correspond to the estimated wind field \mathbf{w}_{est} , i.e., $\boldsymbol{\theta} = \mathbf{w}_{est}$. Each parameter θ_j therefore represents the value of the estimated vertical wind speed $w_{est,j}$ at the position of wind field node j . Between nodes, the estimated vertical wind speed $\xi(\boldsymbol{\theta})$ is the linear interpolation of the values at the two adjacent nodes. The nodes are numbered from aft to fore, such that θ_1 is at $-\tau_{lag}$ while $\theta_{n_{nodes}}$ is at $+\tau_{lead}$.




Figure 4.5: Wind field model

Consequently, this model implicitly assumes that the vertical wind speed is uniform in the vertical and spanwise directions, such that the vertical wind field profile varies only along the predicted flight path. Moreover, under the assumptions made above (Sec. 4.1.2), the predicted path is simply the flight path direction \vec{x}_k and the positions of the boundaries of the estimation window along the predicted path are $x_{lag} = -\tau_{lag} V_K$ and $x_{lead} = +\tau_{lead} V_K$. It is worth noting that the use of V_K in this context is tied to the ‘frozen turbulence’ assumption, i.e., that the turbulent wind field is fixed in space, such that the relative motion between the aircraft and a given gust is uniquely defined by the aircraft’s inertial velocity. This is a common assumption and one which is supported by the gust and turbulence definitions in the certification specifications, however an airmass-fixed wind field is perhaps more realistic. Insofar as the wind field model is concerned, such a change could be easily accommodated by replacing V_K with V_{TAS} and using the air path direction as the predicted path.

4.2.2 Measurement database

As the lidar’s LOS moves through its scan pattern, the sensor collects measurements at the pulse repetition frequency (PRF) and stores them inside a database along with a set of metadata including their LOS directions, Earth-relative positions, and expected standard deviation (computed from, e.g., Eq. 4.2). All measurements whose positions fall within the estimation window can then be extracted from the database and used for estimation; those which end up aft of the estimation window are discarded.

The measurement database may be considered in terms of the distribution and quality of the measurements it contains. PAP and ΔR are generally fixed for a given lidar configuration, whereas R_i increases along the LOS. Looking at Eq. 4.2 and Fig. 4.3, this implies that each of the n_{bins} measurements collected from each lidar ‘shot’ has a fixed measurement noise σ_i proportional to its LOS range,¹⁷ and that all measurements belong to one of the discrete n_{bins} measurement ranges. There are therefore only n_{bins} levels of measurement noise present in the measurement database.

Figure 4.6 shows an example of the positions and standard deviations of the measurements contained in the database. In this example, a PRF of 500 Hz is used, with $n_{bins} = 9$, $\eta_{apert} = 15^\circ$, $\Delta R = 15$ m, $R_{min} = 60$ m, and $R_{max} = 180$ m. The dashed lines indicate the bounds of the estimation window, which in this case is defined by $\tau_{lag} = 0.3$ s and $\tau_{lead} = 0.55$ s. Considering $V_{TAS} = 240$ m/s (corresponding to the critical flight point), this results in $x_{lag} = -72$ m and $x_{lead} = 132$ m. In this case, the entire database would be filled completely within just over one second. It is the time that the forwardmost measurement at $x = 180 \cos(\eta_{apert}) \approx 171$ m requires to reach the rear end of the estimation window at $x = -78$ m, i.e., 249 m / 240 m/s.

The individual bins of Fig. 4.6 are defined by the space between adjacent nodes of a wind field model with $n_{nodes} = 33$. Color denotes the value of σ_i : dark blue for measurements with the lowest noise and yellow for those with the highest. Note how the density of measurements increases from 0 at the forwardmost measurement point, increases until the aftmost LOS measurement point at $x \approx 55$ m, and then remains approximately constant (with minor variations due to binning) until x_{lag} . Similarly, the average quality of measurements in the bins ahead of the aftmost measurement point (i.e., $x > 55$ m) decreases linearly. This can be expected to produce a poorer estimate far ahead of the aircraft.

A further observation may be made regarding the effect of airspeed. Holding all other variables constant, an increase in airspeed results in an increase in the length of the estimation window as well as in the distance between successive ‘shots.’ At the same time, the measurement geometry remains unchanged, so the relative distances to the measurement points along the lidar LOS are constant. The result is that, firstly, x_{lead} extends

¹⁷ Assuming κ_{sens} , κ_{atm} remain constant.

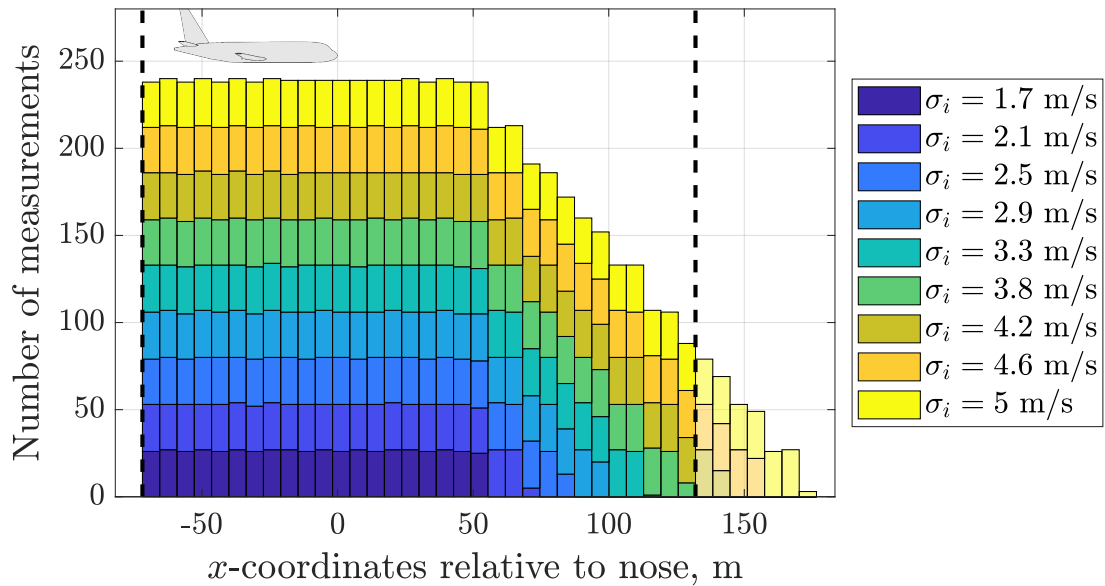


Figure 4.6: Example of measurements contained in the measurement database. The estimation window is indicated by dashed lines and colors indicate measurement noise standard deviations. Each bin occupies the space between adjacent wind field nodes.

forward, and secondly, the average density of measurements across the entire database is lower. Together, these effects degrade the quality of the estimate with increasing airspeed.

4.2.3 Estimation algorithm

The wind estimation algorithm solves an inverse problem: determine the wind model parameters θ given measurements \mathbf{y} such that the resulting estimated wind speed profile matches the true wind speed profile as closely as possible. At the same time, it should avoid excessive noise in the estimate. The estimation algorithm in Reference [84] uses a Gauss-Newton formulation of the least-squares estimation problem augmented with Tikhonov regularization to reduce the influence of measurement noise. Thanks to the fact that the wind field model is linear in parameters, this estimation algorithm converges in its first iteration. Figure 4.7 shows a few examples of estimates using this algorithm both with and without regularization

and measurement noise. Without noise, the unregularized estimate is an excellent match, but with noise, it is clearly unacceptable noisy. Regularization almost completely removes the influence of noise from the estimate, however it does so at the cost of losing part of the true gust information.




Figure 4.7: Comparison between regularized and non-regularized vertical wind field estimates (\mathbf{w}_{est}), both with and without measurement noise. Adapted from [56].

In Ref. [56], a new *maximum a posteriori* (MAP) estimation algorithm is developed under the assumption that measurement noise is Gaussian. It is shown there to be equivalent to the regularized least-squares approach of Ref. [84], however it also allows the noise-induced uncertainty in the estimate to be evaluated, a useful property which is put to use in Chapter 6. The remainder of this section recalls the essential points of the MAP algorithm from [56].

From Bayes' law, the probability density of $\boldsymbol{\theta}$ given n_{meas} measurements \mathbf{y} can be stated [56]:

$$p(\boldsymbol{\theta}|\mathbf{y}) = \frac{p(\mathbf{y}|\boldsymbol{\theta}) p(\boldsymbol{\theta})}{p(\mathbf{y})} \quad (4.3)$$

Where $p(\boldsymbol{\theta}|\mathbf{y})$ is the posterior distribution, $p(\mathbf{y}|\boldsymbol{\theta})$ the likelihood function, $p(\boldsymbol{\theta})$ the prior, and $p(\mathbf{y})$ the marginal likelihood. $p(\mathbf{y})$ is not dependent on $\boldsymbol{\theta}$ and serves only as a normalization constant, so it is sufficient to consider $p(\boldsymbol{\theta}|\mathbf{y}) \propto p(\mathbf{y}|\boldsymbol{\theta}) p(\boldsymbol{\theta})$. The MAP estimate $\boldsymbol{\theta}_{MAP}$ is then the peak of the posterior distribution $p(\boldsymbol{\theta}|\mathbf{y})$. By determining the likelihood function $p(\mathbf{y}|\boldsymbol{\theta})$ and selecting a suitable prior $p(\boldsymbol{\theta})$ which guarantees a reasonably smooth estimate, $\boldsymbol{\theta}_{MAP}$ can be computed.

From the wind field model definition in Sec. 4.2.1, the value of the estimated vertical wind ξ at a position within the estimation window is obtained from the linear interpolation of the adjacent wind field nodes. The relation between a given measurement y_i and its estimated equivalent ξ_i is [56]:

$$y_i = \xi_i(\boldsymbol{\theta}) + d_i + \epsilon_i \quad (4.4)$$

where d_i is a model discrepancy and ϵ_i is the measurement noise. For the considered wind model, and under the assumption of a 1-dimensional wind field, d_i consists of the interpolation error, and so only involves components of the turbulence with spatial frequencies above that of the wind field model; hence it is neglected. The *forward model* \mathbf{A} links $\boldsymbol{\theta}$ to \mathbf{y} . Considering the wind field model is linear in parameters, the problem is linear and so [56]:

$$\hat{\mathbf{y}} = \mathbf{A}\boldsymbol{\theta} + \boldsymbol{\epsilon} \quad (4.5)$$

where $\hat{y}_i = y_i/\sigma_i$ is a measurement scaled by its own expected noise standard deviation, and:

$$A_{i,j} = \frac{\partial \xi_i(\boldsymbol{\theta})}{\partial \theta_j} \sigma_i, \quad i = 1, \dots, n_{meas}, \quad j = 1, \dots, n_{nodes} \quad (4.6)$$

The likelihood function is then [56]:

$$p(\hat{\mathbf{y}}|\boldsymbol{\theta}) \propto \exp\left(-\frac{1}{2}\|\hat{\mathbf{y}} - \mathbf{A}\boldsymbol{\theta}\|^2\right) \quad (4.7)$$

The smoothness of the estimated wind field is regulated by employing *Gaussian smoothness priors* [56]. These penalize the first and second derivatives through finite-difference matrices $\boldsymbol{\Gamma}_1$ and $\boldsymbol{\Gamma}_2$:

$$\boldsymbol{\Gamma}_1 = \begin{bmatrix} -1 & 1 & 0 & \dots & 0 \\ 0 & -1 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & -1 & 1 \end{bmatrix}_{(n_{nodes}-1 \times n_{nodes})} \quad (4.8)$$

$$\boldsymbol{\Gamma}_2 = \begin{bmatrix} -1 & 2 & -1 & 0 & \dots & 0 \\ 0 & -1 & 2 & -1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & -1 & 2 & -1 \end{bmatrix}_{(n_{nodes}-2 \times n_{nodes})} \quad (4.9)$$

The products $\mathbf{\Gamma}_1\boldsymbol{\theta}$ and $\mathbf{\Gamma}_2\boldsymbol{\theta}$ are thus proportional to the finite-difference approximation of the first and second spatial derivatives of $\boldsymbol{\theta}$, respectively. The precision matrices (i.e., the inverse of the covariance matrices) of the Gaussian smoothness priors are then constructed as [56]:

$$\mathbf{Q}_1 = \mathbf{\Gamma}_1^T \mathbf{\Gamma}_1, \quad \mathbf{Q}_2 = \mathbf{\Gamma}_2^T \mathbf{\Gamma}_2. \quad (4.10)$$

The precision weights γ_1 and γ_2 allow the first and second derivative penalties to be weighted relative to one another and to the rest of the problem, such that the overall precision matrix of the prior is [56]:

$$\mathbf{Q}_{12} = \gamma_1 \mathbf{Q}_1 + \gamma_2 \mathbf{Q}_2 \quad (4.11)$$

The prior is then [56]:

$$p(\boldsymbol{\theta}) = \frac{|\mathbf{Q}_{12}^{-1}|^{-1/2}}{(2\pi)^{n_{nodes}/2}} \exp\left(-\frac{1}{2}\boldsymbol{\theta}^T \mathbf{Q}_{12} \boldsymbol{\theta}\right) \quad (4.12)$$

$p(\boldsymbol{\theta})$ is thus normally distributed, with $p(\boldsymbol{\theta}) \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_{12}^{-1})$.

From Bayes' law (Eq. 4.3), the posterior finally reads [56]:

$$p(\boldsymbol{\theta}|\hat{\mathbf{y}}) \propto p(\mathbf{y}|\boldsymbol{\theta}) p(\boldsymbol{\theta}) \quad (4.13)$$

$$\propto \exp\left(-\frac{1}{2}\|\hat{\mathbf{y}} - \mathbf{A}\boldsymbol{\theta}\|^2 - \frac{1}{2}\boldsymbol{\theta}^T \mathbf{Q}_{12} \boldsymbol{\theta}\right) \quad (4.14)$$

The problem is linear with Gaussian measurement noise and a Gaussian prior, so the posterior is also Gaussian and the MAP estimate $\boldsymbol{\theta}_{MAP}$ corresponds to the mean vector of the posterior distribution.

$$p(\boldsymbol{\theta}|\hat{\mathbf{y}}) \sim \mathcal{N}(\boldsymbol{\theta}_{MAP}, \boldsymbol{\Sigma}) \quad (4.15)$$

where $\boldsymbol{\Sigma}$ is the estimate covariance matrix. Putting together Eqs. 4.14 and 4.15, $\boldsymbol{\theta}_{MAP}$ and $\boldsymbol{\Sigma}$ are obtained as [56]:

$$\boldsymbol{\theta}_{MAP} = (\mathbf{A}^T \mathbf{A} + \mathbf{Q}_{12})^{-1} \mathbf{A}^T \hat{\mathbf{y}} \quad (4.16)$$

$$\boldsymbol{\Sigma} = (\mathbf{A}^T \mathbf{A} + \mathbf{Q}_{12})^{-1} \quad (4.17)$$

4.3 Wind field resampling

In general, the digital sampling rate of the control system differs from the real-time estimation rate. The estimation rate is typically far lower due to the relatively large number of computations involved in computing \mathbf{A} . Here, in fact, the control system rate and estimation rate are taken to be 100 Hz and 10 Hz, respectively. Furthermore, the estimated wind inputs to the controller, denoted as the *controller wind field* \mathbf{w}_{ctrl} , do not generally correspond to the nodes of the estimated wind field \mathbf{w}_{est} .

In this work, the controller wind field is generally defined as shown in Fig. 4.8. The space between nodes is defined by the airspeed and the control system sampling time T_s . The number of *preview* nodes h_{pre} and *postview* nodes h_{post} define the maximum extent forward and aft with respect to the reference point at the nose of the aircraft; the total number of nodes in the controller wind field is then $h_{pre} + h_{post} + 1$.




Figure 4.8: Controller wind field definition

To produce the necessary controller wind field at the correct rate, the estimated wind field must therefore be resampled. This can be done by linearly reinterpolating the estimated wind at the control system sampling rate. To account for aircraft motion between estimation steps, the interpolation coordinates are moved forward along the estimated wind field. Figure 4.9 shows an example of this for an estimation rate of 10 Hz, a control system rate of 100 Hz ($T_s = 0.01$ s); the controller wind field \mathbf{w}_{ctrl} is defined with $h_{pre} = 40$ and $h_{post} = 25$. \mathbf{w}_{est} is estimated at t_0 and at $t_0 + 100$ ms; note how the reinterpolated \mathbf{w}_{ctrl} appears to ‘slide’ forward between estimation steps.

This method of resampling introduces a couple of significant nonlinearities. Firstly, at each estimation step, the resampled output may ‘jump,’




Figure 4.9: Controller wind field reinterpolation over time. The wind field is estimated at t_0 and t_0+100 ms (with an estimation rate of 10 Hz); in between, the controller wind field is reinterpolated from the previous estimate at 100 Hz.

especially in the forwardmost nodes, as the local estimate is updated. In Fig. 4.9, this can be observed in the transition from t_0+90 ms to t_0+100 ms approximately around $x = 125$ m. The strength of these jumps depends on the parameters of the wind estimation system, and is particularly affected by the smoothing parameters and the relative strength of measurement noise in the estimate. Secondly, beyond the end of the estimation window, the resampled values exponentially drop to zero; resampled wind field nodes which ‘slide’ beyond the end will repeatedly and sharply drop to zero between estimation steps.

Summary: Airborne Doppler wind lidar sensors allow the wind profile along the direction of flight to be estimated with a significant anticipation time. This chapter has presented the working principles of Doppler wind lidar sensors as well as their measurement geometry, followed by a description of a maximum a posteriori wind estimation algorithm and of the resampling process used to convert wind estimates into controller inputs.

Chapter 5

Structured H_2/H_∞ Control Design for GLA

Contents

5.1	Optimization-based multiobjective structured control synthesis	72
5.1.1	H_2 and H_∞ optimal control	72
5.1.2	Structured controllers	75
5.1.3	Nonsmooth optimization-based control synthesis	75
5.1.4	Signal-based specifications	76
5.1.5	Robust control	77
5.2	Control design process	78
5.2.1	GLA synthesis problem	78
5.2.2	Controller structure	80
5.2.3	Model order reduction	81
5.2.4	Performance evaluation	82
5.3	GLA performance specifications	82
5.3.1	Directly-specified frequency-domain requirements	82
5.3.2	Dealing with discrete gust time-domain requirements	83

Active gust load alleviation for an industrial use case is a complex and demanding control problem. Much like flight control functions, it must take into consideration several competing design objectives, a wide range of flight conditions and configurations, realistic control law architectures, and robustness against noise and uncertainties. There is the further complication of high-order aeroelastic models, which may contain thousands of states and which introduce complex aeroelastic dynamics into the design problem. A control design framework capable of dealing with such issues is therefore needed.

This chapter serves to present the control design framework used here. Firstly, the optimization-based multi-objective, structured robust synthesis method is introduced. Secondly, the overall control design process is elaborated. Finally, the problem of selecting appropriate control specifications to achieve design objectives, especially for robustness and time-domain performance, is discussed.

5.1 Optimization-based multiobjective structured control synthesis

5.1.1 H_2 and H_∞ optimal control

A generic control problem is depicted in standard form in Figure 5.1. The *plant* P is controlled by *controller* K , which uses the *sensed outputs* \mathbf{y} to compute *control inputs* \mathbf{u} . In an optimal control problem, the performance specifications against which the controller is optimized are defined in terms of transfer functions from the *exogenous inputs* \mathbf{w} (e.g., noise disturbances, tracking commands) to the *performance outputs* \mathbf{z} . In H_2 and H_∞ optimal control, the H_2 and/or H_∞ norms of the specified transfer functions are minimized.




Figure 5.1: Control synthesis problem in standard form.

H_∞ specifications

The H_∞ norm is defined as [86, p. 158]:

$$\|G(s)\|_\infty = \max_{\omega} \bar{\sigma}(G(j\omega)) = \max_{\omega} |G(j\omega)| \quad (5.1)$$

(in SISO case)

where $\bar{\sigma}(G(j\omega))$ is the maximum singular value of a dynamical system $G(j\omega)$. It thus corresponds to the greatest singular value σ across all frequencies, which for a single-input-single-output (SISO) system is simply the maximum value of the magnitude of its transfer function. The H_∞ norm can therefore be interpreted as the maximum gain for a steady sinusoidal input. Another useful characteristic of the H_∞ norm is that it satisfies the multiplicative property, i.e., $\|A(s)B(s)\|_\infty \leq \|A(s)\|_\infty \cdot \|B(s)\|_\infty$ [86, p. 158].

H_∞ specifications are expressed by augmenting the transfer function with a weighting function in the form of a linear filter. Intuitively speaking, this can be viewed as imposing an upper-bound template, equal to the inverse of the weighting function, on the magnitude of the transfer function. If the magnitude of the resulting closed-loop transfer function is smaller than the template for all frequencies, the specification is met. The weighting function may simply consist of a scalar gain, but this is often too simplistic. Instead, using dynamic filters introduces a frequency-

[86] Skogestad *et al.*, 2010. A textbook on multivariable feedback control.

dependent weighting. The magnitudes of specified transfer functions can thus be ‘shaped’ across the frequency domain.

H_2 specifications

The H_2 norm is defined from the Frobenius norm of a transfer function matrix integrated across frequency, and can be expressed in terms of the system’s singular values [86, p. 157]:

$$\|G(s)\|_2 = \sqrt{\frac{1}{2\pi} \int_{-\infty}^{\infty} \sum_i \sigma_i^2(G(j\omega)) d\omega} = \sqrt{\frac{1}{2\pi} \int_{-\infty}^{\infty} |G(j\omega)|^2 d\omega} \quad (5.2)$$

(in SISO case)

The H_2 norm thus resembles a kind of Euclidean norm, and can be thought of as the ‘average’ response of the system. To be finite, $G(s)$ must be stable and strictly proper (i.e., no direct feedthrough).

In defining specifications, it may be interpreted in terms of the time-domain response in two ways:

- *Stochastic*: as the expected RMS of the response to a white noise input.¹⁸
- *Deterministic*: as the energy of the impulse response. In fact, by Parseval’s theorem, the SISO case in Eq. 5.2 is also [86, p. 157]:

$$\|G(s)\|_2 = \sqrt{\int_0^{\infty} |g(t)|^2 dt} \quad (5.3)$$

where $g(t)$ is the impulse response of $G(s)$.

As with H_∞ specifications, H_2 specifications can be defined using a weighting function, however in this case the multiplicative property does not apply. In fact, the H_2 norms of individual subsystems do not imply anything in terms of the H_2 norm of the combined system [86, p. 159]. Moreover, when defining specifications based on the *deterministic* interpretation of the norm, it is important to keep in mind that systems with different impulse responses can have the same frequency response magnitude.¹⁹

¹⁸ This interpretation is the basis of the LQG problem, which is itself a special case of an H_2 optimal control problem [86, pp. 355-356].

¹⁹ As a trivial example, consider all-pass filters.

5.1.2 Structured controllers

Control law architectures used in industry typically have a distinct structure [51]. Such structures are driven by design philosophies, physical interpretations of the control problem, engineering experience, realistic sensor and actuator architectures, computational limits, the difficulty of designing gain schedules,²⁰ and a general desire to limit complexity. Simple, clearly structured, and physically meaningful controllers have the advantage of being more interpretable, an indispensable quality when control design methods are limited to classical frequency-domain control methods and manual ‘tinkering,’ but it is also a useful property for certification and for dealing with some types of nonlinearities. Unfortunately, optimal control problems including structured controllers are generally non-convex.

In contrast, several ‘modern’ control approaches can only provide large, unstructured controllers. The early H_∞ control methods based on algebraic Riccati equations or LMIs, for example, produce full-order (i.e., equal to the order of the plant) controllers [87]. This has the important property of yielding a convex optimization problem and hence a globally optimal solution, but the full-order controller is often too large to be usable, especially for a high-order system. Ad-hoc approaches to get around this problem, such as using model order reduction methods to reduce the full-order controller to a more reasonable size, are inconsistent at best [87].

5.1.3 Nonsmooth optimization-based control synthesis

The control synthesis tool used in this thesis is the Matlab Control System Toolbox function *sys tune*. This tool is based on the non-smooth optimization-based structured H_2/H_∞ synthesis framework developed principally by Apkarian, Noll, and Gahinet [87].

This method of synthesis allows non-convex, non-smooth control problems to be tackled, overcoming the limitations of older H_2/H_∞ approaches based on algebraic Riccati equations or LMIs [87]. This allows the designer

²⁰ Gain scheduling refers to predetermined variations in control parameters with respect to some property of the system, e.g., altitude or airspeed. This helps maintain control performance in the face of nonlinear variations in dynamics throughout the flight envelope.

[87] Apkarian *et al.*, 2017. A paper discussing a nonsmooth optimization based approach to H_∞ (and H_2) control.

great flexibility in selecting controller structure and control specifications. Such tools have also proven themselves effective at dealing with relatively high-order or otherwise difficult problems [88], a distinctive strength when considering the size of aeroelastic aircraft models.

One of its main weaknesses is, of course, that it uses a local optimization strategy to solve a non-convex problem, implying that it is prone to falling into local optima. This is not always an issue, provided the local optimum is good enough to meet all requirements. If it isn't, the control gains must be randomly resampled (often several times) to try to find a starting point which leads to a better optimum, but there is no guarantee that the global optimum has been found. In practice, not all synthesis problems are equally difficult; a judicious choice of specifications, controller structure, and initial controller parameters may often ease the tuning procedure.

5.1.4 Signal-based specifications

As mentioned above, control specifications are mainly expressed in terms of weighted H_2 and H_∞ performance channels. Older full-order H_2 and H_∞ synthesis methods require weighting filters to be integrated into the design model, and the resulting specifications are effectively condensed into a single MIMO specification. Expressing the full set of competing design objectives as a single MIMO transfer function is often difficult and unintuitive, especially in the H_2 case. Simply defining them all as a set of SISO transfer functions is usually simpler and more practical.

In a multiobjective signal-based framework such as *systune*, multiple independent performance channels of any kind may be defined between any subset of signals in \mathbf{w} and \mathbf{z} . During control synthesis, the value of the overall objective function is then equal to the worst case (i.e., maximum) over all specifications [89, Eq. 1]. *Systune* furthermore allows for a distinction between *hard* and *soft* specifications: *hard* specifications act as constraints and *soft* specifications as objectives. The optimization problem then seeks to minimize the worst-case *soft* specification subject to satisfying the *hard* specifications. More formally, from [89]:

[88] Apkarian *et al.*, 2021. An article discussing nonsmooth optimization-based control design.

[89] Apkarian *et al.*, 2014. A paper about multi-model, multiobjective nonsmooth optimization-based control.

$$\begin{aligned}
& \underset{k \in \mathcal{K}}{\text{minimize}} && \max_{i \in \mathcal{I}} \|T_{w_i \rightarrow z_i}\| \\
& \text{subject to} && \max_{j \in \mathcal{J}} \|T_{w_j \rightarrow z_j}\| \leq 1,
\end{aligned} \tag{5.4}$$

where \mathcal{K} is the set of all tunable controller parameters, $T_{w \rightarrow z}$ denotes a weighted transfer function from w to z , $\|\cdot\|$ indicates an H_2 or H_∞ norm, and \mathcal{I} and \mathcal{J} refer to the sets of soft and hard specifications, respectively. In practice, minimization of the soft specifications is usually stopped when they pass below 1 or so to avoid overdesign and to speed up the design process.

The definitions above already imply that some sort of normalization is necessary to make the specifications comparable to one another. It is, in fact, necessary to ensure that the weighting functions used to define all control specifications are scaled such that the specification is met when $\|T_{w \rightarrow z}\| \leq 1$.

5.1.5 Robust control

This control design framework offers several options for taking robustness considerations into account. For feedback control, classical stability margins and disk margins may be directly specified as optimization goals. H_∞ specifications allow for loop-shaping design and for robust stability according to the small-gain theorem. If parametric uncertainties in the form of real uncertain parameters or complex dynamic uncertainties are included in the plant model, the synthesis process can take them into account during the optimization, ensuring a parametrically robust controller [90].

Another practical alternative for parametric robustness is to use a multi-model formulation of the control problem, whereby a single controller is tuned against multiple models of the system at the same time [87]. This approach is particularly useful when dealing with complex, high-order systems with uncertainties which cannot be easily modeled in the standard framework, or for modeling system faults. Taking the example of the large aeroelastic models needed for GLA, variations in the aeroelastic properties of the system due to uncertainty in the aircraft's mass distribution (as occurs with varying payloads and fuel load) can change the

[90] Apkarian *et al.*, 2015. A paper about including parametric robustness in nonsmooth optimization-based control.

mode shapes and natural frequencies of aeroelastic modes as well as the aircraft's trimmed flight shape. The resulting uncertainty may require hundreds of uncertain parameters to model correctly, not to mention the effort and difficulty of doing so. In contrast, it is usually fairly simple to generate several models with varying mass distributions. Along similar lines, the set of models in a multi-model problem are only required to have the same inputs and outputs but are unconstrained in terms of internal structure and model order. It is often necessary to reduce the order of large models and many model order reduction methods (e.g., balanced truncation) produce models with a 'black box' internal structure. Here, again, a parametric uncertainty approach is a good deal more challenging than a multi-model approach.

The obvious drawback to a multi-model approach is that it provides no concrete guarantees on performance 'between' or 'beyond' the selected models. It is then up to the designer to select a set of models which adequately cover the range of uncertainty, e.g., by including the 'corner cases,' and to carefully verify afterward that it is sufficiently robust.

5.2 Control design process

Figure 5.2 schematically illustrates a generic GLA control design process. Starting from a selected set of full-order linear models, first each model is processed to obtain reduced-order discrete-time models. Then the synthesis problem is assembled and the controller is tuned according to the control specifications. A performance evaluation determines whether the outcome is satisfactory. If not, the designer must identify the cause and adjust the control specifications, controller structure, or model reduction criteria accordingly. Several aspects are discussed in more detail below; in particular, the problem of selecting and adjusting the controller specifications is discussed in the next section.

5.2.1 GLA synthesis problem

Figure 5.3 illustrates the GLA synthesis problem in standard form. In this case, the plant \mathbf{P} is composed of two blocks. The aircraft system \mathbf{G}_{AC} models the aeroelastic aircraft's natural dynamics as well as sensor dynamics, actuator dynamics, and system delays. \mathbf{F}_{lidar} instead models




Figure 5.2: Generic GLA control design process, loosely based on [46]

the wind field estimation system. A detailed discussion of its contents and characteristics is reserved for Ch. 6; for now it is sufficient to note that the exogenous inputs \mathbf{w} include a vertical wind disturbance input well ahead of the aircraft and a noise input, and that it produces the estimated wind field vector \mathbf{w}_{est} as well as the model gust input w_{turb} . The sensed outputs \mathbf{y} also include the outputs of onboard sensors \mathbf{y}_{sens} . The performance outputs \mathbf{z} can include any outputs of \mathbf{G}_{AC} ; they often include cut loads, local accelerations, and control surfaces rates and positions.

The entire system is modeled in discrete time (i.e., in the z -domain) for two main reasons. The first is that \mathbf{F}_{lidar} depends on a series of discrete time delays to function accurately (see Ch. 6). The second is that digital flight control systems run in discrete time, so tuning the controller directly in discrete time can help ease its implementation.

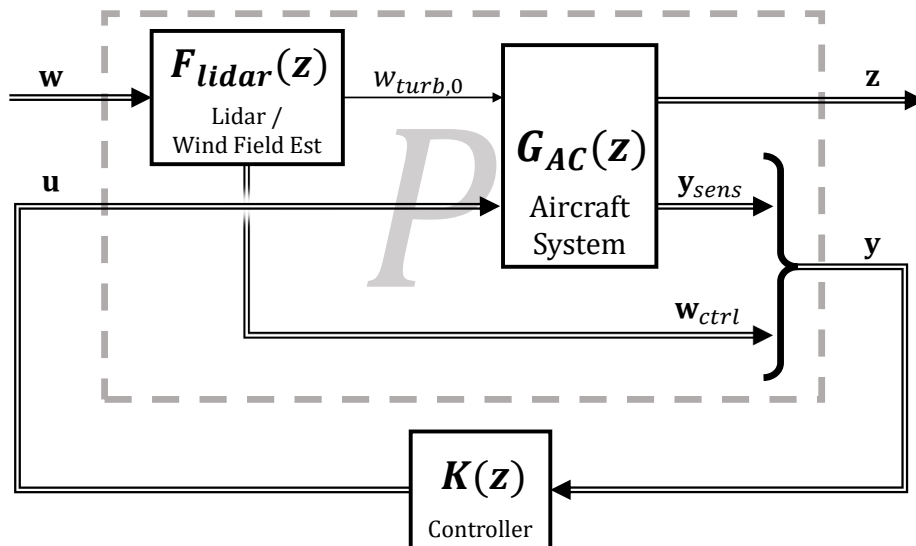


Figure 5.3: Control synthesis problem including lidar-based wind estimation.

5.2.2 Controller structure

The choice of controller structure encompasses a very broad range of possibilities. For example, a fairly simple option is to use a state space system with tunable A , B , C , D matrices, or perhaps to adopt a historical feedback GLA designs such as those of the L-1011 [22] or the A320 [25].

For the purposes of this work, a ‘standard’ feedforward controller for lidar-based control K_{FF} , depicted in Figure 5.4, is defined. It consists of a simple gain matrix with one row for each actuator command channel, and one column for each element in the controller wind field \mathbf{w}_{ctrl} provided by the wind estimation system. The commanded actuator deflections δ_{cmd} are therefore the summed elementwise product of the wind field estimate with the respective rows of K_{FF} .

This results in a controller which resembles a set of finite impulse response (FIR) filters, much like those used in [91]. Though simple, it has many degrees of freedom and is capable of very good performance if the provided wind field is sufficiently good. It furthermore benefits from a high degree of transparency: simply plotting the gains of each row immediately affords the designer insight into the relation between input wind field and actuator commands.




Figure 5.4: Standard feedforward controller structure for lidar-based GLA.

5.2.3 Model order reduction

Full-order aeroelastic models can have up to several thousand states. Aside from adding a large computational cost to the control synthesis process, such large models are usually more detailed than necessary. By reducing the order of the models, the dynamics and response characteristics of interest may be preserved while strongly reducing the complexity of the control problem.

[91] Wildschek *et al.*, 2009. A paper describing an adaptive feedforward GA controller and its flight test results on DLR ATTAS.

5.2.4 Performance evaluation

The performance evaluation at the end of the control design process serves to evaluate all aspects of the design process. This includes evaluating whether the reduced-order design model is a sufficiently representative of the full-order linear and nonlinear models, whether the choice of controller structure and models is adequate for the required performance and robustness, whether the control specifications adequately capture the requirements, and more. In consequence and depending on the progress of the design cycle, this evaluation may make use of various model types and fidelities, and involves both frequency- and time-domain (i.e., discrete gust encounter) evaluations.

5.3 GLA performance specifications

One of the key tasks of a control designer is to translate performance requirements into control specifications compatible with the selected control design tools. The designer's skill in performing this task is closely tied to the time and effort required to design a controller. In the case of the structured H_2/H_∞ optimal control framework discussed above, frequency-domain requirements are relatively easy to express directly, however time-domain requirements such as the discrete gust loads must be dealt with indirectly.

5.3.1 Directly-specified frequency-domain requirements

Several types of requirements can be directly specified in the frequency domain using H_2 and H_∞ norms.

For example, the definition of continuous turbulence requirements (Eq. 3.8, p. 50) and that of the H_2 norm (Eq. 5.2, p. 74) are essentially the same. All that is needed to define an equivalent H_2 specification is a linear weighting filter whose frequency response corresponds that of the von Kármán frequency spectrum. The passenger comfort metrics defined in ISO 2631-1 are another example of this. They are characterized by a set of frequency weightings which express human sensitivity to vibration and motion sick-

ness [92], and are equally easy to express as H_2 specifications.

H_∞ specifications can be used to impose roll-off and roll-on behavior on the controller. High-frequency dynamics tend to be more uncertain mainly due to unmodeled dynamics, greater sensitivity to time delays, and approximation error from model order reduction. In addition, atmospheric disturbances are transitory, so a GLA function should not be active at very low frequencies; this has the added advantage of making it resistant to measurement bias. A GLA controller's gain should therefore only be high in a certain frequency band, and tend to 0 at frequencies above and below this band. Roll-off specifications serve to impose a low-pass behavior (i.e., negative slope w.r.t. frequency) on the controller's gain above a certain frequency, preventing it from interacting with high-frequency modes. A roll-on specification instead imposes high-pass behavior (i.e., positive slope w.r.t. frequency) below a specified frequency. H_∞ specifications can also be used to restrict the degree to which adjacent ailerons can move in opposition to one another, as in [46, pp. 11-12, criteria 7-9].

5.3.2 Dealing with discrete gust time-domain requirements

There is no direct connection between discrete gust time-domain peak responses and the H_2/H_∞ norms, so formulating control specifications to address discrete gust requirements requires an indirect approach.

The pure H_∞ approach to discrete gust specifications

Past approaches have attempted to use H_∞ specifications, however this tends to resemble an art more than a science. It involves selecting weighting filters so as to reduce certain peaks in the frequency response of a given transfer function which are suspected of contributing to the discrete gust response. From the resulting tuned controller and simulation results, the designer must then decide how to modify the specifications to overcome the remaining shortcomings, i.e., which transfer functions should be specified, which modes/frequency-domain peaks to target, how the weighting filters should be defined or modified. There are no reliable methods for selecting any aspect of these tasks, so the designer has to proceed iter-

[92] Kubica *et al.*, 1999. An article discussing control design for improving passenger comfort.

actively, relying on very broad guidelines, intuition, and experience. An example of such a specification placed on the reference model's wing-root bending moment transfer function and its resulting discrete gust performance is shown in Fig. 5.5. In this instance, the template has been shaped to target the loads occurring in the frequencies around the short-period mode.




Figure 5.5: Example of an H_∞ specification and its corresponding discrete gust performance on the WRBM of the reference model. Blue solid and dotted lines are with and without the tuned controller, respectively, and the dashed black line is the equivalent H_∞ template.

Such an approach can produce good results, as in [46], but it is often difficult and prone to misspecification. For a given system configuration, the combination of available sensors and actuators might be ideal for controlling certain flexible modes but very poor for others. Similarly, some modes may not be very significant in terms of the gust response for the considered gust lengths and transfer paths. Certain modes may also be poorly damped, resulting in a large H_∞ norm. A poorly-defined H_∞ specification could attempt to suppress the ‘wrong’ peaks, resulting in increased computational time and cost while possibly yielding worse results, overdesigning, or even causing the synthesis process to fail. Adjusting weighting functions to deal with such issues is unintuitive and nontrivial, and may end up requiring many iterations on behalf of the designer. The problem is exacerbated with a growing number of objectives and models: the frequency response of each transfer function for each model may vary widely, so the H_∞ specifications need to be fine-tuned for each individual transfer function. It is also worth noting that the poor connection between the H_∞ norm and the nonperiodic time-domain response implies

that robust H_∞ performance does not offer any significant guarantees of robust time-domain performance.

Using the H_2 norm

The connection between the H_2 norm and the nonperiodic time-domain response is somewhat more explicit than that of the H_∞ norm thanks to its *deterministic* interpretation. The unweighted H_2 norm can be related to the time-domain peak by way of matched filter theory (MFT) [93]. MFT reveals that the maximum possible time-domain peak in response to an input with finite energy can be expressed in terms of the H_2 norm. This worst-case input is, in fact, the impulse response of the system played backwards in time. This does not often resemble a 1–cosine discrete gust, so this approach is generally too conservative to be practical.

An indirect approach, e.g., using manually tuned weighting filters, is also possible. An H_2 specification imposes no hard constraints, allowing the synthesis process to focus on the parts of the frequency response it can most easily affect. This characteristic provides some protection against the risks of misspecification associated with an H_∞ approach, but it does not necessarily simplify the process of fine-tuning the weighting filters.

Summary: This chapter has presented the structured H_2/H_∞ optimal control framework and its applications to lidar-based GLA problems. A standard feedforward GLA controller structure has been defined. Most importantly, the limitations of H_2 and H_∞ control specifications have been discussed and linked to the difficulty of connecting them to discrete gust time-domain requirements.

[93] Zeiler, 1997. A paper describing the Matched Filter Concept for gust load requirements.

Part II

Contributions

Chapter 6

Linear Modeling of the Lidar-based Wind Estimation System

Contents

6.1	Concept definition	90
6.1.1	Legacy model	91
6.1.2	Proposed formulation of \mathbf{F}_{lidar}	92
6.2	Computing the linear filter	94
6.2.1	Frequentist interpretation of the MAP estimate	94
6.2.2	Reference measurement database	95
6.2.3	Construction of reference model \mathbf{A}	97
6.2.4	Assembling the subfilters	98
6.2.5	Simplifications and limitations	103
6.3	Validation	105
6.3.1	Evaluation of wind field estimation filter \mathbf{F}_{WFE}	107
6.3.2	Evaluation of zero-mean filter \mathbf{F}_{ZM}	110

Any method for designing lidar-based controllers must somehow model the lidar system's output in a way which is compatible with the intended control framework. Given the sensitivity of feedforward control to system uncertainties, the quality of this model has a direct impact on the control design process. To make use of the methods discussed in Ch. 5, the model must furthermore be linear time-invariant (LTI).

This chapter presents one of the main contributions of this thesis: a method for analytically computing a linear model of a lidar-based wind estimation system. First, the legacy method of modeling the lidar system is recalled and the proposed approach is presented. Following that, the process for deriving the model is discussed step-by-step, together with a discussion of its inherent assumptions and limitations. Finally, a series of demonstrations are performed to evaluate how well the proposed linear filter models the full nonlinear lidar-based wind estimation system.

The material presented in this chapter has previously been published in References [57] and [56].

6.1 Concept definition

Referring to Fig. 5.3 (p. 80), the lidar system model \mathbf{F}_{lidar} is represented by the block in Fig. 6.1. Here, the input \mathbf{w} refers generically to the disturbance inputs; at a minimum, it must include a (single) vertical wind disturbance ahead of the aircraft $w_{turb,lead}$. The outputs include the aircraft model wind input $\mathbf{w}_{turb,0}$ and the controller wind field \mathbf{w}_{ctrl} .




Figure 6.1: Schematic view of the proposed linear filter.

6.1.1 Legacy model

The ‘legacy’ formulation of \mathbf{F}_{lidar} as used, e.g., in [44], relies on a *tapped delay line* $\mathbf{D}(z)$, depicted in Fig. 6.2. A tapped delay line consists of a series of discrete unit delays z^{-1} with *taps* between individual delays to extract intermediate values. Accordingly, \mathbf{F}_{lidar} is shown schematically in Fig. 6.3. It accepts as input the true vertical wind speed at the forward-most node of the controller wind field $w_{turb,ctrl,lead}$. $\mathbf{D}(z)$ then essentially propagates this signal along the length of the controller wind field, so the number of unit delays in $\mathbf{D}(z)$ corresponds to the length of the controller wind field divided by the discrete sampling time. The taps in $\mathbf{D}(z)$ are placed such that its outputs correspond to the controller wind field nodes, and the aircraft model input $\mathbf{w}_{turb,0}$ is extracted from the resulting vector. If the standard controller is used (see Sec. 5.2.2, p. 80), $\mathbf{D}(z)$ is fully tapped, however the number and location of the taps may be restricted as needed.




Figure 6.2: Tapped delay line $\mathbf{D}(z)$.




Figure 6.3: Legacy formulation of the linear lidar system model \mathbf{F}_{lidar}

The legacy \mathbf{F}_{lidar} may therefore be understood as an all-pass filter and $\mathbf{w}_{turb,ctrl}$ is effectively the true wind field. This accomplishes the essential goal of producing a wind field profile, but neglects the effects of the sensing, estimation, and resampling process described in Ch. 4. During control design, this implies that the controller is synthesized as if it had access to the full spectrum of wind information up to its Nyquist frequency, and depending on how the control problem is defined, it may try to actively control the aircraft up to that frequency. Furthermore, measurement noise uncertainty in the estimated wind field is not taken into account in any way.

It is clear that such perfect wind information is never actually available and that coupling such a controller to a realistic lidar system could result in a significant loss of performance. Such issues then require the designer to iterate several times, first checking the synthesized controller's performance in a higher-fidelity simulation environment which includes the full nonlinear lidar and wind field estimation system, and then adjusting the control specifications so as to indirectly compensate for the characteristics of the lidar system.

For example, the designed controller may end up relying on the forwardmost parts of the estimate to take advantage of their greater lead time. The quality of available measurements is, however, substantially worse well ahead of the aircraft (cf. Fig. 4.6, p. 64), so that part of the estimate is typically subject to greater measurement noise uncertainty as well as a stronger smoothing effect.²¹ The resulting controller ends up particularly sensitive to both effects, potentially degrading its performance beyond acceptable limits.

6.1.2 Proposed formulation of \mathbf{F}_{lidar}

A new form of \mathbf{F}_{lidar} which approximates the smoothing effects of the wind field estimation algorithm as well as the measurement noise uncertainty would significantly ease the control design process, reducing the necessary time and effort and potentially improving the resulting controller's robustness and performance. The proposed multiple-input multiple-output (MIMO) filter is shown in Figure 6.4.

²¹ The strength of the smoothing term \mathbf{Q}_{12} (see Sec. 4.2.3, p. 67) is equal across the estimation window, but if the available measurements are locally noisier on average, the smoothing effect is relatively stronger due to the noise scaling in Eq. 4.6.




Figure 6.4: Schematic view of the proposed linear filter.

The disturbance input vector \mathbf{w} now contains two signals: the true wind $w_{turb,lead}$ at (or beyond) the forwardmost node of the *estimated* wind field (i.e., x_{lead}) and white noise input d_{lead} . The *wind field estimation* subfilter \mathbf{F}_{WFE} applies the smoothing effect of the wind field estimation algorithm to the true wind disturbance, yielding the filtered controller wind field profile $\mathbf{w}_{filt,ctrl}$. The *zero-mean* subfilter \mathbf{F}_{ZM} instead accounts for the measurement noise-induced uncertainties in the estimate by filtering the unit-variance white noise signal d_{lead} so as to produce an estimation noise profile $\mathbf{d}_{filt,ctrl}$ which is representative in amplitude and frequency content of the actual uncertainties. The output of the filter is the combined wind field $\mathbf{w}_{cbd,ctrl}$, obtained from an elementwise sum of $\mathbf{w}_{filt,ctrl}$ and $\mathbf{d}_{filt,ctrl}$.

The tapped delay lines $\mathbf{D}_{WFE}(z)$ and $\mathbf{D}_{ZM}(z)$ serve as the dynamic part of the filters. $\mathbf{D}_{WFE}(z)$ takes as input $w_{turb,lead}$ and outputs the true vertical wind profile across the estimation window $\mathbf{w}_{turb,est}$. $\mathbf{D}_{ZM}(z)$ produces the full noise vector \mathbf{d}_{est} across the estimation window from d_{lead} . The true wind speed at the aircraft model input $\mathbf{w}_{turb,0}$ is extracted (or interpolated) from $\mathbf{w}_{turb,est}$.

In \mathbf{F}_{WFE} , the smoothing matrix \mathbf{K}_{WFE} applies the smoothing effect of the wind field estimation algorithm to the true wind profile $\mathbf{w}_{turb,est}$, producing the filtered wind field $\mathbf{w}_{filt,est}$. In \mathbf{F}_{ZM} , the uncertainty matrix \mathbf{K}_{ZM} approximates the measurement noise-related uncertainty in the estimated wind field by transforming the white noise vector \mathbf{d}_{est} into the

estimation noise profile $\mathbf{d}_{filt,est}$. Finally, the interpolation matrix \mathbf{K}_{interp} approximates the wind field resampling process, transforming $\mathbf{w}_{filt,est}$ and $\mathbf{d}_{filt,est}$ into their equivalents in the controller wind field coordinates, $\mathbf{w}_{filt,ctrl}$ and $\mathbf{d}_{filt,ctrl}$. Seeing as the estimated and controller wind field coordinates are identical for both subfilters, both instances of \mathbf{K}_{interp} are also identical.

The following section discusses how the components of the filter can be analytically computed and assembled from the parameters of the lidar sensor and the wind estimation algorithm. The most complex components, \mathbf{K}_{WFE} and \mathbf{K}_{ZM} , are directly derived from a modified interpretation of the MAP formulation of the wind estimation problem presented in Sec. 4.2.3, along with a ‘reference’ measurement database.

6.2 Computing the linear filter \mathbf{F}_{lidar}

6.2.1 Frequentist interpretation of the MAP estimate

In Section 4.2.3 (p. 64), Equations 4.16 and 4.17 provided the expressions for computing the MAP estimate and its uncertainty in the form of a covariance matrix Σ . Put differently, given a set of noisy measurements, these yield the most likely value of the model parameters $\boldsymbol{\theta}_{MAP}$ and the uncertainty as to its deviation from the (unknown) true wind profile, also known as the *epistemic* uncertainty [56].²² To construct \mathbf{F}_{lidar} according to the definition above, a different type of uncertainty is required, i.e., given the estimate from a set of ‘true’ (i.e., noiseless) measurements, how do noisy estimates vary around this ‘true’ estimate?

This corresponds to the frequentist properties of the Bayesian MAP estimate, also developed in [56], and which may be analyzed as follows. Consider r sets of measurements of the same wind field: $\hat{\mathbf{y}}^{(1)}$, $\hat{\mathbf{y}}^{(2)}$, \dots , $\hat{\mathbf{y}}^{(r)}$. Each set of measurements contains the same ‘true’ information $\hat{\mathbf{y}}_{true}$ and differs from the others only in the (Gaussian) noise present in the individual measurements, so the sets of measurements are themselves characterized by a normal distribution around the ‘true’ measurements,

²² To be precise, it also takes into consideration the prior. This is important when γ_1 or γ_2 are large, as this results in a very smooth estimate with low uncertainty, but which may deviate strongly from the true wind field profile.

i.e., $\hat{\mathbf{y}} \sim \mathcal{N}(\hat{\mathbf{y}}_{true}, \mathbb{I})$ [56].²³ r MAP estimates $\boldsymbol{\theta}_{MAP}$ can be computed from these sets of measurements, resulting in a distribution of the MAP estimates Θ_{MAP} . The transformation from $\hat{\mathbf{y}}$ to $\boldsymbol{\theta}_{MAP}$ is linear (cf. Eq. 4.16), so Θ_{MAP} is also normally distributed [56]:

$$\Theta_{MAP} \sim \mathcal{N}(\boldsymbol{\theta}_{MAP,true}, \boldsymbol{\Sigma}_{MAP}) \quad (6.1)$$

in which the ‘true’ estimate $\boldsymbol{\theta}_{MAP,true}$ and frequentist covariance matrix $\boldsymbol{\Sigma}_{MAP}$ are defined:

$$\boldsymbol{\theta}_{MAP,true} = \mathbf{B}\hat{\mathbf{y}}_{true} \quad (6.2)$$

$$\boldsymbol{\Sigma}_{MAP} = \mathbf{B}\mathbf{B}^T \quad (6.3)$$

$$\text{with } \mathbf{B} = (\mathbf{A}^T \mathbf{A} + \mathbf{Q}_{12})^{-1} \mathbf{A}^T$$

Samples can be drawn from this distribution with [56]:

$$\boldsymbol{\theta}_{MAP} = \boldsymbol{\theta}_{MAP,true} + \tilde{\mathbf{L}}\boldsymbol{\zeta}_{noise} \quad (6.4)$$

where $\tilde{\mathbf{L}}$ is the square root of $\boldsymbol{\Sigma}_{MAP}$, i.e., $\boldsymbol{\Sigma}_{MAP} = \tilde{\mathbf{L}}\tilde{\mathbf{L}}^T$, and $\boldsymbol{\zeta}_{noise}$ is a vector with n_{nodes} elements of unit variance Gaussian white noise, i.e., $\boldsymbol{\zeta}_{noise} \sim \mathcal{N}(0, \mathbb{I})$. Several methods exist to compute the square root of a matrix; here $\tilde{\mathbf{L}}$ is obtained via the Schur decomposition method from Ref. [94].²⁴

Equation 6.4 is finally the desired result. It shows that a noisy MAP estimate can be decomposed into its ‘true’ estimate and a uncertain contribution due to the noise, and that the measurement noise can be represented by an equivalent vector with the same number of elements as $\boldsymbol{\theta}$.

6.2.2 Reference measurement database

In Sec. 4.2.3, the wind field estimation is performed using the contents of the measurement database, which contains a set of measurements along

²³ Note that the individual measurements in $\hat{\mathbf{y}}$ are divided by their noise standard deviation, so the remaining Gaussian noise has unit variance.

²⁴ In the notation of [94], the square root of a matrix A can be computed via its Schur decomposition $A = QSQ^H$ as $A^{1/2} = QS^{1/2}Q^H$. $S^{1/2}$ is then simply computed via Cholesky decomposition.

[94] Björck *et al.*, 1983. A paper describing a Schur method for computing the square root of a matrix.

with their metadata: measurement positions, LOS directions, and expected noise standard deviations. For the computation of \mathbf{F}_{lidar} , the measurements themselves are not needed, as the wind data and measurement noise are provided by the input signals $w_{turb,lead}$ and d_{lead} . As will be shown in the following sections, the measurement database is only required to compute \mathbf{A} , for which only the position (i.e., x -coordinate relative to the aircraft), scan angle $\phi_{scan,i}$, aperture angle η_{apert} , and expected standard deviation σ_i of each measurement are needed.

Accordingly, a *reference* measurement database containing only the necessary metadata can be built up. Considering a constant airspeed and no steady wind (see Sec. 4.1.2, p. 58), sets of LOS measurements are evenly spaced along the predicted path with V_K/PRF , so given number of measurement points and along the LOS and their positions, the x -coordinates can be easily computed. As before, the wind field is assumed one-dimensional, varying only along the x direction, so y - and z -coordinates are not needed. Measurements which fall outside of the estimation window do not participate in the estimation process, so these are ignored. The aperture angle η_{apert} is the same for all measurements. The expected standard deviation σ_i can be calculated from Eq. 4.2 (p. 60), in which all variables are constant for all measurements except for measurement LOS range R_i , which varies along the LOS.

The value of scan angle ϕ_{scan} assigned to each measurement is less trivial. In the general case, ϕ_{scan} varies throughout the scan pattern, and the distribution of its values among the measurements depends on the scan rate, PRF , and V_K . From Eq. 4.1 (p. 60), the vertical wind component of a measurement i is proportional to $\sin \phi_{scan,i}$, so the distribution of ϕ_{scan} can locally affect the quality of the estimate.

For example, measurements taken near $\phi_{scan} = 0^\circ$ or 180° have a $\sin \phi_{scan}$ near 0, and therefore a very small vertical wind component. If some part of the estimation window contains a local concentration of such measurements, the estimate can be expected to be locally poorer. Indeed, the forwardmost parts of the estimation window tend to have a very low number of available measurements (cf. Fig. 4.6), a relatively large proportion of which come from the most recently collected measurements, so that part of the estimation window is vulnerable to concentrations of similar ϕ_{scan} values. The ‘leading’ scan angle of each estimation (i.e., the one corresponding to the current LOS) changes with each estimation step unless the scan rate and PRF are both integer multiples of the estimation

rate If the scan rate is too slow or the estimation window too small, the ‘average’ value of $\sin \phi_{scan}$ may vary significantly between windows.

The reference measurement database is necessarily constant (as a consequence of the filter being LTI), so the value of ϕ_{scan} attributed to each measurement point must remain constant as well. The goal is for \mathbf{F}_{lidar} to be representative of the average estimation quality, so arbitrarily choosing a ‘leading’ scan angle and extrapolating ϕ_{scan} for all measurement points based on the scan rate is not ideal. Instead, a single ‘effective’ value of $\sin \phi_{scan}$ which can be assigned to all measurements is sought. Equation 4.14 (p. 67) reveals that the wind field estimation problem essentially amounts to a least-squares optimization, so an excellent approximation can be found by taking the square root of the expected value of $\sin^2 \phi_{scan}$, i.e., the root mean square of $\sin \phi_{scan}$. For the conical scan pattern discussed in Sec. 4.1.2, this is:

$$\sin \phi_{scan,i} \equiv (\sin \phi_{scan})_{RMS} = \frac{1}{\sqrt{2}}, \quad i = 1, \dots, n_{meas} \quad (6.5)$$

6.2.3 Construction of reference model \mathbf{A}

Using the reference measurement database, the reference model \mathbf{A} can be constructed. Recalling Equation 4.6, each element of \mathbf{A} is defined as:

$$A_{i,j} = \frac{\partial \xi_i(\boldsymbol{\theta})}{\partial \theta_j} \frac{1}{\sigma_i} \quad (4.6)$$

From Eq. 4.1 (p. 60), one may note that V_{TAS} and η_{apert} are the same for all LOS measurements. η_{apert} is generally known, and V_{TAS} is measured separately (e.g., via the aircraft’s air data and inertial reference system or via a separate estimation process using lidar measurements). Measurement y_i can then be defined as the projection of the vertical wind speed on the LOS by removing the V_{TAS} component from $u_{meas,i}$:

$$\begin{aligned} y_i &= u_{meas,i} - V_{TAS} \cos \eta_{apert} \\ &= w_{turb,i} \sin \eta_{apert} \sin \phi_{scan,i} + \epsilon_i, \quad i = 1, \dots, n_{meas} \end{aligned} \quad (6.6)$$

$\xi_i(\boldsymbol{\theta})$ is the linear interpolation of the $\boldsymbol{\theta}$ values of the two adjacent wind field nodes, although it must also include the trigonometric term $\sin \eta_{apert} \sin \phi_{scan,i}$ to remain consistent with y_i . The geometric relationship is illustrated in Fig. 6.5. Nodes p and $p + 1$ are the nodes adjacent

(along the predicted path) to measurement i , and $\Delta x_{a,i}$, $\Delta x_{b,i}$ are the distances along the predicted path between measurement i and nodes p and $p+1$, respectively. This gives:

$$\xi_i(\boldsymbol{\theta}) = \left[\frac{\Delta x_{b,i}}{\Delta x_{a,i} + \Delta x_{b,i}} \theta_{p,i} + \frac{\Delta x_{a,i}}{\Delta x_{a,i} + \Delta x_{b,i}} \theta_{p+1,i} \right] \cdot \sin \eta_{apert} \sin \phi_{scan}, \quad i = 1, \dots, n_{meas} \quad (6.7)$$




Figure 6.5: Geometrical relation between y_i , ξ_i , and $\boldsymbol{\theta}$.

Each row of \mathbf{A} corresponds to one measurement i ($i \in [1, n_{meas}]$) and each column to one parameter θ_j ($j \in [1, n_{nodes}]$). From Eqs. 4.6 and 6.7, each row \mathbf{A}_i is then:

$$\mathbf{A}_i = \left[0, \dots, 0, \frac{\Delta x_{b,i}}{\Delta x_{a,i} + \Delta x_{b,i}}, \frac{\Delta x_{a,i}}{\Delta x_{a,i} + \Delta x_{b,i}}, 0, \dots, 0 \right] \cdot \frac{\sin \eta_{apert} \sin \phi_{scan}}{\sigma_i}, \quad i = 1, \dots, n_{meas} \quad (6.8)$$

6.2.4 Assembling the subfilters

The components of the *wind field estimation* \mathbf{F}_{WFE} and *zero-mean* \mathbf{F}_{ZM} subfilters from Fig. 6.4 can now be computed and assembled. These include the filter matrices \mathbf{K}_{WFE} and \mathbf{K}_{ZM} , the tapped delay lines $\mathbf{D}_{WFE}(z)$ and $\mathbf{D}_{ZM}(z)$, and the interpolation matrix \mathbf{K}_{interp} .

Smoothing matrix \mathbf{K}_{WFE}

$\boldsymbol{\theta}_{MAP,true}$ from Eq. 6.2 is the MAP estimate computed using only perfect measurements $\hat{\mathbf{y}}_{true}$. With the added assumption that the true vertical

wind varies linearly between estimation nodes, $\hat{\mathbf{y}}_{true}$ can be expressed as $\mathbf{A}\boldsymbol{\theta}_{true}$. Consequently:

$$\boldsymbol{\theta}_{MAP,true} = \mathbf{B}\mathbf{A}\boldsymbol{\theta}_{true} \quad (6.9)$$

Considering that $\mathbf{w}_{turb,est}$ from Fig. 6.4 represents the true values of the wind field profile within the estimation window, it can be substituted for $\boldsymbol{\theta}_{true}$ in Eq. 6.9. The filtered wind field $\mathbf{w}_{filt,est}$ corresponds to $\boldsymbol{\theta}_{MAP,true}$. Equation 6.9 now reads $\mathbf{w}_{filt,est} = \mathbf{B}\mathbf{A}\mathbf{w}_{turb,est}$, leading to:

$$\mathbf{K}_{WFE} = \mathbf{B}\mathbf{A} \quad (6.10)$$

Uncertainty matrix \mathbf{K}_{ZM}

Variations in the MAP estimate due to measurement noise are represented by the estimate uncertainty $\boldsymbol{\theta}_{noise}$. From Eq. 6.4, samples of $\boldsymbol{\theta}_{noise}$ are obtained by:

$$\boldsymbol{\theta}_{noise} = \tilde{\mathbf{L}}\boldsymbol{\zeta}_{noise} \quad (6.11)$$

recalling that $\tilde{\mathbf{L}}$ is the square root of the frequentist MAP covariance matrix Σ_{MAP} , and $\boldsymbol{\zeta}_{noise}$ is a vector of n_{nodes} samples of Gaussian white noise with unit variance. $\boldsymbol{\zeta}_{noise}$ may be thought of as a kind of ‘average’ noise contribution from all the measurements available at each wind field node. The noise vector \mathbf{d}_{est} from Fig. 6.4 can thus replace $\boldsymbol{\zeta}_{noise}$, and \mathbf{K}_{ZM} is:

$$\mathbf{K}_{ZM} = \tilde{\mathbf{L}} \quad (6.12)$$

Discrete-time implementation

The filter matrices \mathbf{K}_{WFE} and \mathbf{K}_{ZM} are coupled with the tapped delay lines $\mathbf{D}_{WFE}(z)$ and $\mathbf{D}_{ZM}(z)$, respectively. Both of them must use the same discrete sampling time, and the ‘equivalent noise’ elements of \mathbf{d}_{est} are logically connected to the respective true wind elements in $\mathbf{w}_{turb,est}$, so $\mathbf{D}_{WFE}(z)$ and $\mathbf{D}_{ZM}(z)$ are necessarily identical.

The number of delays in $\mathbf{D}_{WFE}(z)$ and $\mathbf{D}_{ZM}(z)$ depends mainly on the required discrete sampling time. Both \mathbf{K}_{WFE} and \mathbf{K}_{ZM} , as derived above, are based on the original parameter vector $\boldsymbol{\theta}$, which has n_{nodes} elements and corresponds to the original wind field model. Their inputs

$\mathbf{w}_{turb,est}$ and \mathbf{d}_{est} would, therefore, also need to have n_{nodes} elements. The time delay between individual elements in the vectors would need to be equal to $(\tau_{lag} + \tau_{lead})/(n_{nodes} - 1)$.

The control system sampling time T_s (which is used for \mathbf{F}_{lidar} as well) is generally a property of the control system and bears no relation to the wind field mode. Consequently, the sampling time of $\mathbf{D}_{WFE}(z)$ and $\mathbf{D}_{ZM}(z)$ is also T_s , and \mathbf{K}_{WFE} and \mathbf{K}_{ZM} must be modified so as to accept input vectors spaced with T_s . This basically involves replacing the original wind field nodes with new nodes spaced with T_s . The number of new nodes is then:

$$n_{nodes,c} = \frac{\tau_{lag} + \tau_{lead}}{T_s} + 1 \quad (6.13)$$


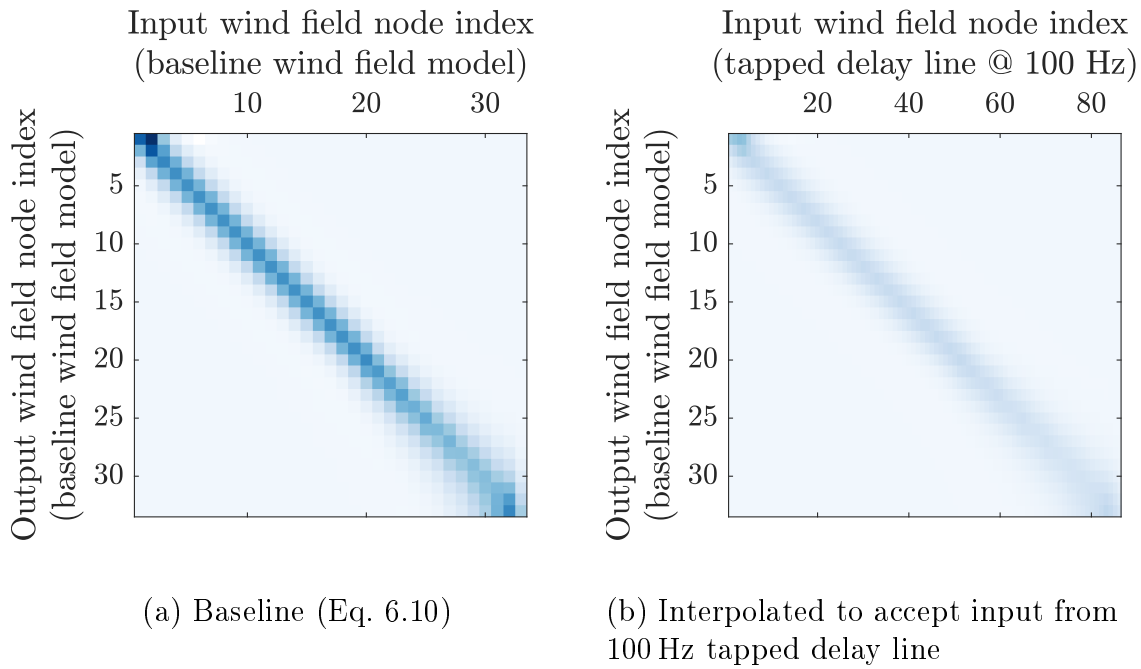
$\mathbf{D}_{WFE}(z)$ and $\mathbf{D}_{ZM}(z)$ thus have $(n_{nodes,c} - 1)$ delays, and \mathbf{w}_{turb} and \mathbf{d}_{est} have $n_{nodes,c}$ elements. The new nodal coordinates are used to linearly reinterpolate the rows of \mathbf{K}_{WFE} and \mathbf{K}_{ZM} , each of which must then be normalized such that the sum of each row remains unchanged. An example of the effect of this reinterpolation and normalization process on the values of \mathbf{K}_{WFE} is shown in Figs. 6.6a and 6.6b for a $T_s = 10$ ms (100 Hz).

Furthermore, \mathbf{K}_{ZM} must account for the change in PSD of the white noise signal. A discrete-time white noise signal has a flat power spectrum which is band-limited to the Nyquist frequency. Signal variance is equal to the integral of the power spectrum over frequency, so for a constant signal variance, an increase in sampling rate proportionally reduces the PSD. d_{lead} is unit-variance white noise regardless of sampling time, so to obtain the same frequency response magnitude at the output, \mathbf{K}_{ZM} must be multiplied by $\sqrt{n_{nodes,c}/n_{nodes}}$.

Controller wind field interpolation \mathbf{K}_{interp}

Section 4.3 (p. 68) introduced the need for wind field resampling. In brief, the estimated wind field and the controller wind field differ in sampling rate and wind field node coordinates. The estimated wind field is converted to the controller wind field by moving the controller wind field coordinates relative to the estimated wind field coordinates between estimation steps and linearly interpolating the estimated wind at the controller sampling rate T_s .

In the filter, the effect of this step is produced by introducing the



(c) Reinterpolated and truncated to match the controller wind field node coordinates.

Figure 6.6: Surface plots of \mathbf{K}_{WFE} values before and after reinterpolation. Darker blue indicates a greater value.

interpolation matrix \mathbf{K}_{interp} . In Fig. 6.4, the two instances of \mathbf{K}_{interp} convert the filtered wind field $\mathbf{w}_{filt,est}$ and estimate uncertainty $\mathbf{d}_{filt,est}$ from estimated wind field model coordinates into their controller wind field equivalents $\mathbf{w}_{filt,ctrl}$ and $\mathbf{d}_{filt,ctrl}$. \mathbf{K}_{interp} linearly interpolates the input vector at the controller wind field coordinates. However, in the filter, there are no differences in sampling time, and the output coordinates of the controller wind field cannot change over time (as it is LTI), so the effect of the relative movement of the coordinates cannot be reproduced entirely.

As an approximation, the interpolation coordinates of the controller wind field nodes can be shifted half an estimation step forward, e.g., 50 ms if the estimation rate is 10 Hz. The logic is that, over the course of an estimation step, a controller wind field node slides forward along the estimated wind field with a constant velocity equal to V_K before being reset at the beginning of the next estimation step. On average, then, the node coordinate will find itself halfway. The resulting outputs should therefore be approximately representative of the average characteristics of the controller wind field. The node coordinates cannot go beyond the forward end of the estimation window, so nodes whose ‘shifted’ coordinates exceed the limit are simply set equal to the last node.

The combination of \mathbf{K}_{WFE} (or \mathbf{K}_{ZM}) with \mathbf{K}_{interp} effectively amounts to a row-wise linear reinterpolation of \mathbf{K}_{WFE} , shown for example in the transition from Fig. 6.6b to Fig. 6.6c. In practice, the full transformation is performed in a single 2-D interpolation of the baseline \mathbf{K}_{WFE} , i.e., passing directly from Fig. 6.6a to Fig. 6.6c.

6.2.5 Simplifications and limitations

Several assumptions have been made above to obtain the desired results. The most important of them are recalled and discussed here.

Measurement noise

The measurement noise ϵ_i is assumed to have a Gaussian distribution with standard deviation σ_i . As discussed in Sec. 4.1.3, this assumption is based on the fact that the predominant source of measurement noise, ‘shot noise’, is Gaussian in essence. In reality, the measurement noise distribution may deviate from the Gaussian due to various secondary effects, e.g., Brillouin scattering [29]. A departure from the Gaussian distribution will generally violate the assumptions used in the Bayesian MAP. In practice, \mathbf{F}_{WFE} only relies on the true part of the measurements, and remains unaffected, however the representativeness of \mathbf{F}_{ZM} is more or less degraded, depending on the precise distribution.

The noise distribution is also assumed to be centered, i.e., $E(\epsilon_i) = 0$, such that it does not introduce a bias into the measurements. The impact of a non-centered noise $E(\epsilon_i) \neq 0$ on the real estimate could be mitigated relatively easily using, e.g., an online autocalibration procedure or some form of high-pass filtering.

The noise vector \mathbf{d}_{est} may adequately represent an ‘equivalent’ magnitude of the noise-related uncertainties, but it cannot entirely account for the noise in thousands of independent measurements. One particular effect it cannot recreate is that of measurements being added within the measurement range of the sensor. This causes the average noise content of the measurement database to vary considerably at locations ahead of the minimum detection range ($x > 50$ m in Fig. 4.6, p. 64), which in turn affects the estimate.

Scan angle approximation

As anticipated in Sec. 6.2.2, some lidar system configurations may be susceptible to local concentrations of measurements with similar ϕ_{scan} values. The scan angle approximation from Eq. 6.5 entirely ignores such effects. Consequently, situations in which there are local concentrations of measurements with $|\sin \phi_{scan,i}| < 1/\sqrt{2}$ will produce estimates which are locally poorer than those predicted by the lidar, and vice-versa.

Model discrepancy

Section 6.2.4 introduces an assumption that the true wind field only varies linearly between nodes. The true wind field has no such restrictions, so this assumption results in a sort of low-pass effect. The cutoff frequency of this effect is inversely proportional to the distance between wind field nodes.

Angle approximations

The assumption that the rotational axis of the lidar sensor is aligned with the aircraft's flight path (first introduced in Sec. 4.1.2) may cause the linear filter to produce different results than the actual estimator in some situations. As noted already in Sec. 4.1.2, the real wind estimation system saves the position and orientation of measurements in absolute spatial coordinates, so it can partially adapt to dynamic changes in aircraft motion and orientation, whereas \mathbf{F}_{lidar} cannot.

Spatial resolution

ΔR is the span over which the measured LOS wind speed is averaged (see Sec. 4.1). This span also defines the distance between adjacent measurements along the LOS, so it also defines the spatial resolution of the measurements. However, the multitude of overlapping measurements in the database creates a super-resolution effect, such that the effective spatial resolution of the system is not uniquely defined by ΔR . The actual resolution and accuracy of the estimation process are not obvious and rely on several interrelated variables.

\mathbf{F}_{lidar} is built up under the assumption that measurements are point-like, so neither the limiting effect of ΔR nor the effect of super-resolution are modeled. This primarily affects the filter's characteristics at relatively high spatial frequencies. For reasonably high values of γ_1 and γ_2 , as is the case for the evaluations in the following section, the effects of smoothing attenuate the response at the relevant frequencies.

Resampling nonlinearities

The wind field resampling process, discussed in Sec. 4.3, introduces several effects which cannot be adequately modeled in \mathbf{F}_{lidar} . Firstly, as noted already in Sec. 6.2.4, the relative motion between the controller wind field coordinates and the estimated wind field can only be approximated by a static offset which reproduces the average effect. Secondly, Sec. 4.3 mentions two important nonlinearities: jumps in the estimate at the end of each estimation step, and interpolation coordinates sliding beyond the forward edge of the wind field model. Both of these can cause unexpected deteriorations in the controller wind field and cannot be modeled in the filter. Generally speaking, the controller wind field coordinates should be chosen so as to avoid parts of the estimation window which are sensitive to these issues in order to avoid unexpected behavior and performance losses.

6.3 Validation

This section presents several evaluations of the time- and frequency-domain characteristics of \mathbf{F}_{lidar} with the aim of demonstrating the degree to which it is representative of a realistic lidar-based wind estimation system. The objective of these evaluations is to demonstrate that \mathbf{F}_{lidar} , when included in the control problem, allows the control synthesis software to properly take the lidar system's properties into account. In consequence, four types of results are presented:

- *Full Est*: Wind field estimate produced at 10 Hz by the full nonlinear lidar simulation and wind estimation algorithm.
- *Full Reint*: Resampled wind field produced at 100 Hz by the full nonlinear lidar simulation and wind estimation algorithm. This is representative of the input to a GLA controller.
- *Lin Analysis*: Analytical results derived from the linear filter. This is representative of the properties 'seen' by control synthesis tools.
- *Lin Sim*: Simulated results using the linear filter.

Generally speaking, a good match between *Full Reint* and *Lin Analysis* is desired. Not all types are used in all cases; only the relevant information is selected for each demonstration. When both *Full Est* and *Full Reint* are present, this is used to highlight the effect of wind field resampling. *Full Est* is usually excluded from frequency-domain evaluations as it has a Nyquist frequency of only 5 Hz. *Lin Analysis* is only valid for frequency-domain evaluations; time domain evaluations of \mathbf{F}_{lidar} are necessarily limited to *Lin Sim*. In frequency-domain evaluations, *Lin Sim* is computed using the same signal processing techniques as the *Full* results so as to help demonstrate the validity of the analytical techniques.

A single set of lidar and wind estimation parameter values, listed in Table 6.1, is used for the following evaluations. V_K corresponds to the true airspeed of the reference flight point (see Sec. 2.3, p. 37), and σ_i/R_i is computed using Eq. 4.2 (p. 60) for the reference flight point altitude (which affects κ_{atm}). Unless stated otherwise, the controller wind field coordinates are chosen equal to those of the wind field model. Moreover, unless stated otherwise, \mathbf{F}_{lidar} makes use of the output coordinate shift proposed in Sec. 6.2.4.

Table 6.1: Default lidar system parameters

Parameter	PRF	PAP	R_{min}	ΔR	n_{bins}
Value	500 Hz	0.05 Wm ²	60 m	15 m	9
Parameter	n_{nodes}	τ_{lead}	τ_{lag}	γ_1	γ_2
Value	33	0.55 s	0.3 s	0.8	1.37
Parameter	η_{apert}	f_{scan}	V_K	σ_i/R_i	
Value	15°	13 Hz	240 m/s	0.0242 s ⁻¹	




In several of the results below, continuous turbulence simulations are used to provide a realistic, stochastic, arbitrarily long wind field. This turbulent wind field is a Gaussian, pseudo-random signal conforming to the von Kármán spectrum specified in CS-25.341.b [4] (see Sec. 3.2.2, p. 48). It is generated before the simulation using an inverse Fourier transform, and is parametrized with a standard deviation σ_w . The pseudo-random property allows the exact same wind profile to be generated each time, improving the comparability of the evaluations.

On the basis of Eq. 6.4 (p. 95), \mathbf{F}_{WFE} and \mathbf{F}_{ZM} are evaluated separately, in Sections 6.3.1 and 6.3.2, respectively. This means that in Section 6.3.1, simulations are performed without measurement noise, and measurement noise is generally ignored. Similarly, in Section 6.3.2, no turbulent wind is present, and all simulations contain only measurement noise.

6.3.1 Evaluation of wind field estimation filter \mathbf{F}_{WFE}

Time/space-domain

Figure 6.7 compares the actual estimated wind field (with no measurement error) with the output of the linear filter for both a discrete gust with $H = 30$ m (Fig. 6.7a), a discrete gust with $H = 100$ m (Fig. 6.7b), and a section of continuous turbulence (Fig. 6.7c). The gusts and turbulence are plotted in spatial coordinates, and each window shows the estimated wind field at a given point in time. As the aircraft moves forward, the estimate also proceeds forward along the true wind field. All windows are chosen at the start of an estimation step, such that no resampling effects are present. In Fig. 6.7b, the match is nearly perfect. Figures 6.7a and 6.7c contain more significant discrepancies, likely due to their greater high-frequency content, however the match remains quite good.

(a) Discrete gust with $H = 30$ m(b) Discrete gust with $H = 100$ m

(c) Continuous turbulence (von Kármán spectrum)

Figure 6.7: Comparison of estimated wind field \mathbf{w}_{est} with the filtered wind field $\mathbf{w}_{filt,est}$ over 3 sequential estimation windows.

Power spectral analysis

In the frequency domain, the filter may be compared to the full estimation by performing a power-spectral analysis similar to the evaluation metric used in [54]. In essence, a continuous turbulence wind field long enough to be statistically significant at all relevant frequencies (≈ 1000 s) is simulated with no measurement noise. The power spectral density (PSD) of the error between estimated wind speed and true wind at the position of a selected wind field node $\Phi_{\Delta w}$ is calculated and then normalized with the PSD of the true wind Φ_{turb} . The resulting curve $\Phi_{\Delta w}/\Phi_{turb}$ indicates the relative precision of the wind field estimate across the frequency spectrum. A value near 0 indicates that the estimated wind follows the true wind almost perfectly at those frequencies, whereas a value near 1 indicates that it is not tracked at all.

Figure 6.8 shows an analysis of this type evaluated at a wind field node just ahead of the aircraft nose. The *Full Reint* and *Lin Sim* curves are computed by simulation as described above. The *Lin Analysis* curve is instead derived directly from the transfer function of the filter. This is done by augmenting the filter with an output Δw_{filt} which subtracts the filtered wind speed at the evaluated node from the true wind speed at the same location, taking the magnitude of this transfer function $w_{turb,lead} \rightarrow \Delta w_{filt}$, and squaring it to find its PSD. For this wind field node, the filter matches *Full Reint* quite well.

The match between filter and estimate can be evaluated at all points in the estimation window by computing a ‘bandwidth’ frequency, i.e., the frequency at which the evaluation metric passes above 0.5 (see Fig. 6.8). This is plotted in Fig. 6.9; to demonstrate the effect of the controller wind field coordinate shift proposed in Sec. 6.2.4, the *Lin Analysis* and *Lin Sim* bandwidths in Fig. 6.9a are computed without the coordinate shift, whereas those in Fig. 6.9b include the shift.

A comparison of Figs. 6.9a and 6.9b shows that the linear results appear to shift to the left and their bandwidths are slightly reduced. This results in a marked improvement in the match between the linear evaluations and *Full Reint*, essentially confirming that the coordinate shift performs as expected. The ‘crooked’ appearance of the *Lin Analysis* and *Lin Sim* values for nodes 5 to 16 is caused by the fact that the distribution of measurements in the reference database is invariant and slightly uneven (cf. Fig. 4.6, p. 64); in the full system, this distribution varies slightly over

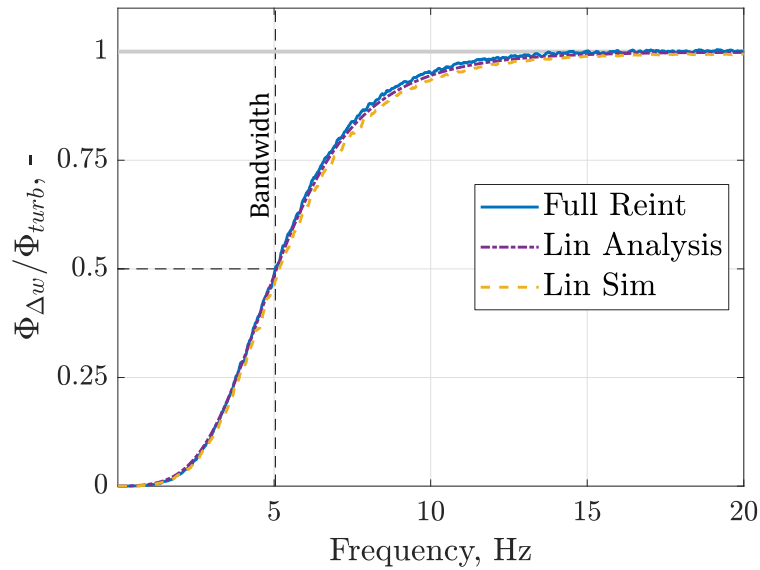


Figure 6.8: PSD evaluation metric computed at Node 13 (near the aircraft nose).


time, leading to the smoother appearance of the *Full Reint* results. Note also how the *Lin Analysis* results in Fig. 6.9b for nodes 32 and 33 have the same value; this is because the coordinate shift moves node 32 slightly beyond the forward edge of the estimation window.

6.3.2 Evaluation of zero-mean filter F_{ZM}


Demonstrating the validity of F_{ZM} is not trivial. Unlike F_{WFE} , for which the comparison can be made by evaluating how the estimate relates to the true wind, there is no ‘true’ signal against which the noise-induced uncertainty can be compared. The demonstrations proposed in this section aim to evaluate their similarity in terms of three key properties. Firstly, the amplitude of the uncertainties at all points in the estimation window should be similar. Secondly, the uncertainty appearing in a certain position in space should maintain some ‘continuity’ across estimates. Finally, the frequency content of the uncertainty should be similar.

Uncertainty bounds

As established in Sec. 6.2.1 (p. 94), the uncertainty in the wind estimate is normally distributed around 0. This distribution is illustrated in Fig. 6.10, in which the standard deviation of the measurement noise uncertainty



(a) Without coordinate shift.



(b) With coordinate shift.

Figure 6.9: Comparison of bandwidth of PSD evaluation metric across the estimation window.


across the estimation window is plotted. *Full Est*, *Full Reint*, and *Lin Sim* results are computed from long (≈ 1000 s) simulations with measurement noise only, whereas the *Lin Analysis* curve is simply the H_2 -norm of the set of transfer functions from d_{lead} to each output of \mathbf{F}_{lidar} . Here, again, Fig. 6.10a does not have the coordinate shift and Fig. 6.10b does.

In Fig. 6.10a, with no coordinate shift, the linear results show good agreement with *Full Est*. As before, *Full Reint* appears shifted to the left due to the resampling effect. Figure 6.10b shows that the coordinate shift partly mitigates this difference, better aligning the linear results with *Full Reint*. Forward of node 29, there is a wide discrepancy where *Full Reint* suddenly decreases. This is the point at which the resampled nodes begin to reach the end of the estimation window and return to zero before the end of each estimation step. This effectively violates the assumption of normally distributed values, reduces their average absolute value, and thus apparently reduces the uncertainty.


Time/space domain continuity

Estimation uncertainty originates in the noise present in individual measurements. Measurements remain in the database until they fall behind the aft end of the estimation window (τ_{lag}), so their individual contributions continue to appear in the estimate over several estimation steps. This presumably implies that certain ‘features’ of the uncertainty profile which appear in one estimate will persist in subsequent estimates, remaining fixed in space and apparently sliding aftwards over time. The fact that noise vector \mathbf{d}_{est} is built up using a tapped delay line means that a given element of \mathbf{d}_{est} will also move down the entire length of the estimation window, creating a similar effect of ‘continuity’.

Figure 6.11 attempts to show this by plotting a sequence of four sequential noise-only wind estimates for both *Full Reint* and the *Lin Sim*. Each estimate is separated from the next by 200 ms (two estimation steps). In both cases, certain features do persist, although most features which appear at the forward end of the estimation window disappear in the following estimation steps as more measurements are added.



(a) Without coordinate shift.



(b) With coordinate shift.

Figure 6.10: Standard deviation of wind estimate due to measurement noise.



(a) *Full Reint*(b) *Lin Sim*


Figure 6.11: Comparison of estimation uncertainty samples from noise-only simulations over four sequential estimation windows.

Nodal power spectral analysis


The frequency content of the noise uncertainty can be evaluated by computing the PSD of the estimated wind $\Phi_{\Delta w}$ at individual points in the estimation window, not unlike the evaluation metric used for \mathbf{F}_{WFE} above. For *Full Est*, *Full Reint*, and *Lin Sim*, this is computed by using a long (≈ 1000 s) noise-only simulation. For *Lin Analysis*, the magnitude of the transfer function from d_{lead} to the output is squared. In all cases, the PSDs are then normalized by dividing them by their integrals across frequency. Figure 6.12a shows the results for node 25.

A new ‘bandwidth’ frequency is once again defined to allow the PSDs to be evaluated across the estimation window. It is defined as the frequency at which the cumulative integral over frequency of the PSD passes above 50% of its total value. This is plotted in Fig. 6.12b.

The bandwidth frequencies agree well up to node 15, which is a short distance ahead of the aircraft. Then they diverge until a nearly constant gap of approximately 0.3 Hz is established from nodes 20 to 26, and then a narrowing up to node 31. Starting from node 31, the same nonlinear resampling effects discussed above dominate the *Full Reint* results, and the coordinate shift results in constant values for *Lin Analysis*. The gap which is observed between nodes 16 and 31 clearly correlates with the parts of the estimation window in which measurements are added (cf. Fig. 4.6, p. 64). This suggests that the noise contributions from the new measurements cause ‘jumps’ in the estimate which shift the frequency distribution upwards, and which cannot be reproduced by the linear filter.



(a) Normalized PSD at Node 25.



(b) PSD bandwidth frequency (cumulative integral 50% of total value) across the estimation window

Figure 6.12: Power spectral density analysis of measurement noise uncertainty.

Summary: This chapter has described a method for analytically computing a linear filter which models the time- and frequency-domain properties of a lidar-based wind estimation system. This filter is composed of two subfilters: a wind field estimation filter \mathbf{F}_{WFE} which models its effect on the true wind information, and a zero-mean filter \mathbf{F}_{ZM} which models the effect of measurement noise. The structure and contents of the filter are derived from the expression of the maximum a posteriori wind estimation algorithm described in Ch. 4. This expression, once reinterpreted from a frequentist perspective and in combination with a reference measurement database, allows the filter parameters to be computed directly. The resulting filter is shown to match the full nonlinear model of the lidar system very well, with the exception of a few strong nonlinearities which cannot be addressed by a linear model.

Chapter 7

Discrete Gust Performance Specifications

Contents

7.1	Discrete gusts in the frequency domain	121
7.2	Discrete gust impulse filter definition	123
7.2.1	Full-order FIR filter	123
7.2.2	Order reduction and parametrization	123
7.3	DGIF-weighted H_2 control specifications	127
7.3.1	Intuition	127
7.3.2	Controller parameter gradient	129
7.3.3	Iterative optimization	133

As discussed in Chapter 5, the difficulties inherent in expressing discrete gust time-domain requirements as frequency-domain control specifications are a significant obstacle to effective GLA control design. The ‘artisanal’ character of iteratively tuning weighting functions on the basis of time-domain simulations can lead to a difficult and time-consuming design process, not to mention the lack of any guarantees of robustness or optimality with respect to the achievable discrete-gust performance of the system. Direct specification remains impossible, however finding a systematic method to express and fine-tune weighting functions stands to significantly improve the situation and opens the door to an automated design process.

This chapter proposes such a method, based on a *discrete gust impulse filter* (DGIF). The main property of a DGIF is that it has an impulse response which closely resembles a discrete gust profile of a given length. It thus replicates both the time-domain and frequency-domain properties of a one-minus-cosine gust. A similar approach is pursued in [95] and [96], where a weighting filter for a single discrete gust length is numerically identified. The following sections aim to derive such a filter, to define a parametrized expression for computing it, to demonstrate that it matches discrete gusts in the time- and frequency-domains, and to make a compelling case that such filters do, in fact, make for better discrete gust specifications.

Part of the results presented in this chapter have previously been published in [58] and [59].

[95] Fonte *et al.*, 2015. An article discussing feedback GLA design for a regional aircraft.

[96] Ripipi *et al.*, 2013. A paper about improved rational matrix approximations for aeroelastic modeling.

7.1 Discrete gusts in the frequency domain

First, the frequency-domain properties of one-minus-cosine discrete gust must be established.

To that end, a set of discrete gusts with varying gust gradient H are generated (Sec. 3.2.1, p. 47) and converted to the time domain by assuming a constant $V_{TAS} = 240$ m/s, equal to that of the reference flight point (see Sec. 2.3, p. 37). For each gust, the PSD Φ_{gust} is calculated via Fast Fourier Transform (FFT), including zero-padding to ensure sufficient frequency resolution. Each PSD is then normalized by dividing it by its integral over frequency as well as its squared amplitude scaling factor $(H/107)^{2/6}$ (cf. Eq. 3.2). This serves to preserve the $-4/3$ slope. The resulting PSDs are thus plotted in Fig. 7.1 in logarithmic coordinates together with the (normalized) von Kármán PSD.

The characteristic gust frequency ω_{gust}/f_{gust} (in rad/s and Hz, respectively) and time t_{gust} are a useful starting point from which to evaluate the relative characteristics of Φ_{gust} . They are defined from the contents of the cosine term in Eq. 3.1:

$$\omega_{gust} = \frac{\pi V_{TAS}}{H}, \quad f_{gust} = \frac{\omega_{gust}}{2\pi} = \frac{V_{TAS}}{2H}, \quad t_{gust} = \frac{1}{f_{gust}} \quad (7.1)$$

In Fig. 7.2, the frequencies are thus normalized with f_{gust} , and the PSD itself is normalized by dividing it by $\lim_{f \rightarrow 0} \Phi_{gust}(f)$. The left plot uses a logarithmic scaling and the right plot uses a linear scaling. At f_{gust} , the PSD has a magnitude of -12 dB ≈ 0.25 , and above f_{gust} , it has a roll-off of approximately -125 dB/decade. In the linear plot, the cumulative integral across frequency of the PSD is plotted as a percentage. This curve shows that over 90% of the frequency content is contributed by the frequencies below f_{gust} .

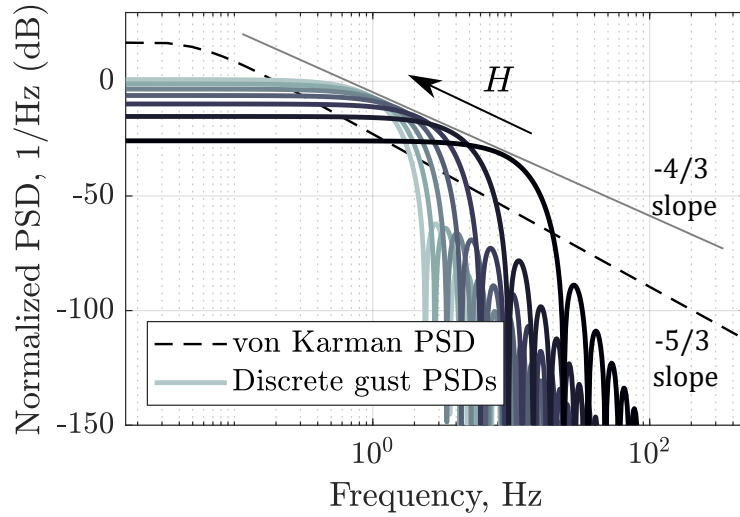


Figure 7.1: Power spectral densities of discrete gusts in logarithmic scale, with the von Kármán PSD plotted for comparison.




Figure 7.2: Characteristics of discrete gust power spectral densities.

7.2 Discrete gust impulse filter definition

7.2.1 Full-order FIR filter

The main requirement in designing the DGIF is its impulse response; by definition, for an LTI system, this automatically guarantees the correct frequency response as well. A simple approach is to construct a finite impulse response (FIR) filter. An FIR filter is a type of digital filter composed of a *tapped delay line* (see, e.g., Fig. 6.2, p. 91) and a set of constant gains, one for each tap. A desired impulse response can then be easily created by assigning the (time-discretized) values of the desired response to the gains. In this instance, the order of the filter (in other words, the number of delays in the tapped delay line) is determined by the gust gradient H , true airspeed V_{TAS} , and filter sampling time. The sampling time, by way of the Nyquist frequency, also limits the filter at higher frequencies. Judging by Fig. 7.1, the standard control system sampling rate $f_s = 100$ Hz (with corresponding Nyquist frequency of 50 Hz) should be adequate, although the shortest gusts may be slightly distorted by discrete-time effects. The frequency response of the resulting filter is consistent with the discrete gust PSD, in that $|F(j\omega)|^2 \propto \Phi_{gust}(\omega)$.²⁵

7.2.2 Order reduction and parametrization

Assuming $f_s = 100$ Hz, $V_{TAS} = 240$ m/s, and $H = 100$ m, the corresponding FIR filter requires at least 84 states to cover the entire gust length. Relative to the overall control problem, this is an enormous number of states. Consider, for instance, that the number of states in the reduced-order aircraft model used for control design is around 100, and that there may be tens or perhaps hundreds of individual specifications containing the DGIF. Changing gust length and flight point (and therefore V_{TAS}) also requires a change in filter order, unless it is always sized for the longest possible gust at the slowest possible flight speed. Moreover, for increasing flight speeds, the frequency margin between the shortest gusts and the Nyquist frequency may not always be adequate.

It is therefore considerably preferable to have a low-order, continuous-time filter. To that end, the FIR filter is first computed with a large


²⁵ Here it is only ‘proportional to’ because the scaling factor of the PSD may vary. If both terms are normalized in the same way, the expression becomes ‘equal.’

sampling rate (e.g. $f_s = 1000$ Hz) to ensure a generously high Nyquist frequency. It is then converted to continuous time using the Tustin method. The model order is then reduced by balanced truncation (via the MATLAB Control System Toolbox function *balred*).


Figure 7.3 compares the results of this order reduction for several filter orders. Figure 7.3a plots the reduced filters against the full-order filter in both frequency and time domain, whereas Fig. 7.3b plots the time-domain error and the cumulative integral of its absolute value. Time and frequency are both normalized with respect to the characteristic gust time t_{gust} and frequency f_{gust} from Eq. 7.1, respectively. The magnitudes are normalized to give a gain of 0 dB at zero frequency. The time responses are normalized with respect to the peak gust speed U_{ds} .

In the frequency domain, the reduced-order filters appear to match the full-order curve up to a certain frequency, and then transition to a -20 dB/decade roll-off. The frequency where this transition occurs may be considered the limit of validity for the reduced filter. Judging from the time-domain plot, the 3rd-order and 4th-order systems are visibly different from the full-order response, and indeed the error plots show that the former reaches error well in excess of 5% and the latter reaches 2%. Starting with the 5th-order filter, they are difficult to distinguish, and remain for the most part below 1% error. In the frequency domain, the 3rd- and 4th-order filters cannot claim to be valid beyond a normalized frequency of 4, whereas the 5th- and 6th-order filters approximate the full-order response fairly well up to a gust-normalized frequency of 8, by which point the magnitude is below -60 dB (10^{-3}). In terms of the cumulative error, each additional state approximately halves it; beyond the 6th-order filter, there is little real improvement. Consequently, the 5th- and 6th-order filters are chosen as the best compromise between order and accuracy.

The process of generating the reduced-order filters described above is relatively laborious and computationally expensive due to the model order reduction step. An expression for the reduced-order transfer function which can be parametrized in terms of the main variables would allow the filters to be computed directly. In fact, by finding the poles and zeros of the reduced filters and factorizing them with ω_{gust} , the expressions for the 5th- and 6th- order filters in Eqs. 7.2 and 7.3 are easily derived. The numerical values of their parameters are listed in Table 7.1.



(a) Frequency response and impulse response.



(b) Impulse response error and its cumulative absolute integral.

Figure 7.3: Comparison of reduced-order continuous-time filters with the full-order filter.

$$F_{DG}^{[5]} = \frac{U_{ds} \pi}{\omega_{gust}} \frac{\left(\frac{s^2}{(a_1 \cdot \omega_{gust})^2} + \frac{2 \cdot b_1 \cdot s}{a_1 \cdot \omega_{gust}} + 1 \right) \left(\frac{s^2}{(a_2 \cdot \omega_{gust})^2} + \frac{2 \cdot b_2 \cdot s}{a_2 \cdot \omega_{gust}} + 1 \right)}{\left(\frac{s}{c_1 \cdot \omega_{gust}} + 1 \right) \left(\frac{s^2}{(c_2 \cdot \omega_{gust})^2} + \frac{2 \cdot d_2 \cdot s}{c_2 \cdot \omega_{gust}} + 1 \right) \left(\frac{s^2}{(c_3 \cdot \omega_{gust})^2} + \frac{2 \cdot d_3 \cdot s}{c_3 \cdot \omega_{gust}} + 1 \right)} \quad (7.2)$$

$$F_{DG}^{[6]} = \frac{U_{ds} \pi}{\omega_{gust}} \frac{\left(\frac{s^2}{(a_1 \cdot \omega_{gust})^2} + \frac{2 \cdot b_1 \cdot s}{a_1 \cdot \omega_{gust}} + 1 \right) \left(\frac{s^2}{(a_2 \cdot \omega_{gust})^2} + \frac{2 \cdot b_2 \cdot s}{a_2 \cdot \omega_{gust}} + 1 \right) \left(\frac{s}{a_3 \cdot \omega_{gust}} + 1 \right)}{\left(\frac{s^2}{(c_1 \cdot \omega_{gust})^2} + \frac{2 \cdot d_1 \cdot s}{c_1 \cdot \omega_{gust}} + 1 \right) \left(\frac{s^2}{(c_2 \cdot \omega_{gust})^2} + \frac{2 \cdot d_2 \cdot s}{c_2 \cdot \omega_{gust}} + 1 \right) \left(\frac{s^2}{(c_3 \cdot \omega_{gust})^2} + \frac{2 \cdot d_3 \cdot s}{c_3 \cdot \omega_{gust}} + 1 \right)} \quad (7.3)$$

Table 7.1: Reduced-order DGIF parameters

	a_1	a_2	a_3	b_1	b_2	c_1	c_2	c_3	d_1	d_2	d_3
$F_{DG}^{[5]}$	1.978	7.75	—	0.0608	-0.5439	0.554	0.947	1.61	—	0.53	0.185
$F_{DG}^{[6]}$	1.9912	4.714	-12.144	-0.0123	0.2112	0.711	1.2755	2.029	0.8356	0.3928	0.1175

7.3 DGIF-weighted H_2 control specifications


7.3.1 Intuition

The most important property of a frequency-domain discrete gust control specification is the link between the resulting system norms (i.e., H_2 and H_∞ norms) and the time-domain peak. In other words: to what extent does a change in the norm, for instance due to a change in controller parameters, correlate to the corresponding change in time-domain peak and vice-versa? In the ideal case, the relationship would be perfectly proportional, such that a 50% reduction in norm implies a 50% reduction in the time-domain peak; in practice, a strong and positive correlation is sufficient. Conversely, a poor correlation between the system norms and the time-domain peaks, i.e., if one is insensitive to the other or if there is no consistent trend, would make such specifications difficult to use.


In the definition of the H_2 norm in Section 5.1.1 (p. 74), it is noted that the norm may be interpreted as the energy of the impulse (i.e., Eq. 5.3). By weighting a given transfer function with a DGIF, the resulting impulse response is equal to the discrete gust response (for the corresponding gust length H) of the unweighted transfer function. A DGIF-weighted H_2 specification can, therefore, directly specify the energy of the discrete gust response, and indirectly affect its peak.

To provide some intuition as to the link between an H_2 -DGIF specification and the time-domain peak, Figure 7.4 plots the frequency response (Fig. 7.4a) and impulse response (Fig. 7.4b) of a DGIF-weighted transfer function, together with their respective cumulative 2-norm (i.e., with a progressively increasing upper integration limit) and the derivative thereof. In Fig. 7.4a, the norm is dominated by the first peak, and it reaches 80% of its total value below 1 Hz. In Fig. 7.4b, the time-domain response is also clearly dominated by the first, positive peak, which alone contributes approximately 80% of the total norm.

The derivative of the 2-norm indicates how strongly a given portion of the response contributes to the overall norm. Accordingly, an optimization process whose cost function is the H_2 norm of the DGIF-weighted transfer function might be expected to target the parts of the response where the 2-norm derivative is high. Note how the absolute value of the peak disproportionately affects the derivative of the 2-norm: the amplitude of the second (negative) peak is approximately half of the first, but the 2-



(a) Frequency response



(b) Impulse response

Figure 7.4: Comparison of the DGIF-weighted (with $H = 70$ m) frequency response and impulse response (blue) with their respective cumulative 2-norms (solid red) and the derivatives thereof (dashed red). Based on the transfer function from vertical gust input to wing root bending moment of the reference model.

norm derivative's peak value during the first peak is over 5 times greater than during the second. Whether a control synthesis optimization process actually targets this part of the response depends on many other factors, including available actuators, system dynamics, controller structure, and other specifications in the control problem.

The remainder of this section investigates the link between DGIF-weighted H_2 norms and discrete gust time-domain peaks. The following investigations are not to be regarded as exhaustive or conclusive, but rather as indicative of general trends.

7.3.2 Controller parameter gradient

Taking a simple feedforward controller such as the standard controller from Fig. 5.4 (p. 81), one can find the gradient for the controller parameters with respect to some function by perturbing the parameters and computing the Jacobian via finite differences. The gradient computed with respect to a given system norm can then be compared to that of the time-domain peak of a given discrete gust response. For a norm which is strongly linked to the time-domain peak, the two gradients are expected to be closely aligned.

One method of comparison is to find the cosine of the angle between the two gradient vectors. Denoting the gradient with respect to the norm \mathbf{d}_N and that with respect to the time-domain peak \mathbf{d}_{TD} , the gradient cosine c is then:

$$c = \frac{\mathbf{d}_N \cdot \mathbf{d}_{TD}}{\|\mathbf{d}_N\|_2 \cdot \|\mathbf{d}_{TD}\|_2} \quad (7.4)$$





Gradient cosine c thus takes on values ranging from -1 to 1. For values near 1, the angle is very small and the gradients point nearly in the same direction; for values near -1, the angle is also small, but they point in opposite directions; and around 0, they are nearly orthogonal.


Here it is evaluated by taking the reference model (see Sec. 2.3, p. 37), a standard controller with a controller wind field of $h_{pre} = 40$ and $h_{post} = 25$ (essentially identical to the example from Sec. 4.3) and using all five actuator inputs. The initial controller gains are randomly sampled 50 times, the gradients are computed at each sampled point, and the resulting gradient cosine is determined. Considering a given SISO transfer function $Y(s)$, four system norms are taken into account: the unweighted H_2 norm $|Y(s)|_2$, the DGIF-weighted H_2 norm $|F_{DG}(s)Y(s)|_2$, the unweighted H_∞ norm $|Y(s)|_\infty$, and the DGIF-weighted H_∞ norm $|F_{DG}(s)Y(s)|_\infty$.

Figures 7.5 and 7.6 show the resulting distribution of gradient cosines for discrete gusts with $H = 70$ m and 30 m, respectively. Four load outputs are shown: the wing root bending moment, the wing root torsional moment, the HTP-root bending moment, and the HTP-root torsional moment. Each column of the figures represents a system norm; the actual datapoints are plotted to the left of the vertical axis, and their distribution is plotted to the right.


Clearly, the results vary significantly across outputs and gust lengths. Recalling that a concentration of points near $+1$ is favorable, the DGIF-weighted H_2 norm generally fares well. In most cases, the unweighted H_2 norm does equally or perhaps slightly worse; the notable exception is the wing root torsional moment for a 70 m gust (Fig. 7.5b), where the DGIF-weighted H_2 is clearly better. The H_∞ norms both show a larger spread in values compared to the H_2 norms, and in particular tend to form clusters of points. The unweighted H_∞ norm does relatively well for the 70 m gust, though it also performs poorly for the 30 m gust.

The tendency of the H_∞ norms to form scattered clusters may be interpreted as situations where the frequency-domain peak (which defines the H_∞ norm) is significantly different from the dynamic modes which drive the time-domain peak. This phenomenon naturally becomes more pronounced for shorter gust lengths, as the frequency range affected by the gust becomes wider. The H_2 norms are not immune to scattered points, however these appear more often as outliers than as clear clusters of points.


Figure 7.5: Distribution of gradient cosines for an $H = 70$ m discrete gust.




(a) Wing root bending moment



(b) Wing root torsional moment



(c) HTP root bending moment



(d) HTP root torsional moment

Figure 7.6: Distribution of gradient cosines for an $H = 30$ m discrete gust.

7.3.3 Iterative optimization

The suitability of the various system norms may be more directly assessed by way of *systune* itself. A simplified control design problem is built up using the reference model, a standard lidar-based controller defined the same way as above (Sec. 7.3.2), and a single performance output. From an initial set of controller parameters, *systune* then optimizes the controller one iteration at a time; at each iteration, the time-domain peak for a selected discrete gust length is also evaluated. In this manner, the correlation between the system norm and the time-domain peak can be evaluated over the course of the optimization. For a norm which is strongly linked to the time-domain peak, a tightly-correlated progression is expected.


25 sets of randomly sampled controller gains (including one with all gains set to zero) are generated, and 100 *systune* iterations are performed on each one for each performance output. This is done once again for four system norms: the unweighted H_2 norm, the DGIF-weighted H_2 norm, the unweighted H_∞ norm, and the DGIF-weighted H_∞ norm.

The results for the wing root bending moment and wing root torsional moment considering an $H = 70$ m gust are shown in Figs. 7.7 and 7.8, respectively; Figs. 7.9 and 7.10 show the same for an $H = 30$ m gust. For each norm and output, four plots are shown, from left to right: the progression of the (relative) value of the norm; the progression of the time-domain peak value; the norm value plotted against the corresponding time-domain peak; and a density plot of the individual points. In the first three, the optimization run corresponding to each initial controller is plotted as a line. Both norms and time-domain peaks are normalized with respect to their open-loop (i.e., no controller) values.

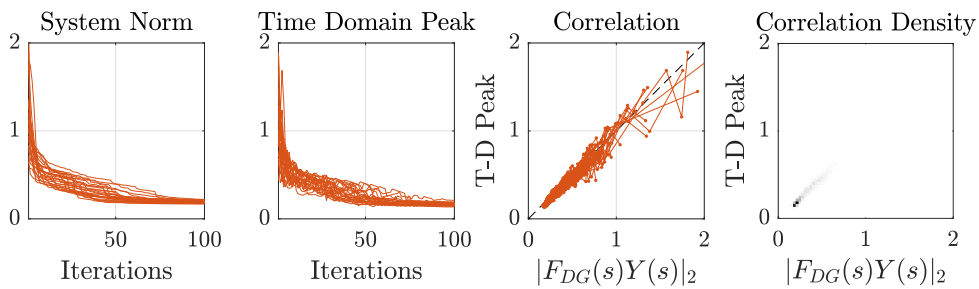
For this analysis, a given system norm is more suitable as a discrete-gust control specification if its progression remains close to the ideal 1:1 slope, plotted here as a dashed black line in the correlation plots. For the wing root bending cases (Figs. 7.7 and 7.9), all norms appear to do fairly well. In these cases, the H_∞ runs tend to reach their final values in relatively few iterations, with fairly large jumps in these first few iterations. This results in a strong concentration of data points around these final positions, and little information on the intermediate correlation between time-domain peak and norm. For both H_∞ norms, the first couple iterations appear to be poorly correlated: large changes in norm value do not correspond to large changes in the time-domain peaks. In

the unweighted H_∞ cases, the final value of the norm is nearly identical across runs, but the corresponding time domain peaks vary significantly, again indicating a poor correlation. The H_2 runs tend to decrease more gradually, making the descent slope more visible. Individual iterations are occasionally negatively correlated, however the overall trend is positive. The DGIF-weighted H_2 correlates reasonably well with the ideal slope; the unweighted H_2 shows a stronger but still positive slope.


In the torsional moment cases (Figs. 7.8 and 7.10), the differences are exacerbated. The H_∞ runs are evidently poorly correlated, and the unweighted H_2 tends to an even stronger slope. The DGIF-weighted H_2 remains remarkably consistent, clustering around the ideal slope.




(a) Unweighted H_2 norm $|Y(s)|_2$



(b) DGIF-weighted H_2 norm $|F_{DG}(s)Y(s)|_2$



(c) Unweighted H_∞ norm $|Y(s)|_\infty$



(d) DGIF-weighted H_∞ norm $|F_{DG}(s)Y(s)|_\infty$

Figure 7.7: Iterative optimization for the wing root bending moment and an $H = 70$ m discrete gust.

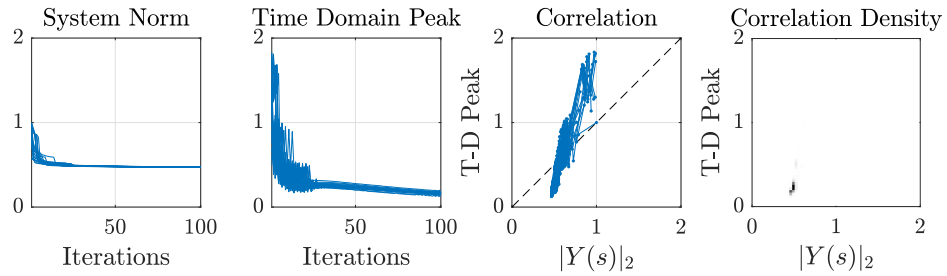
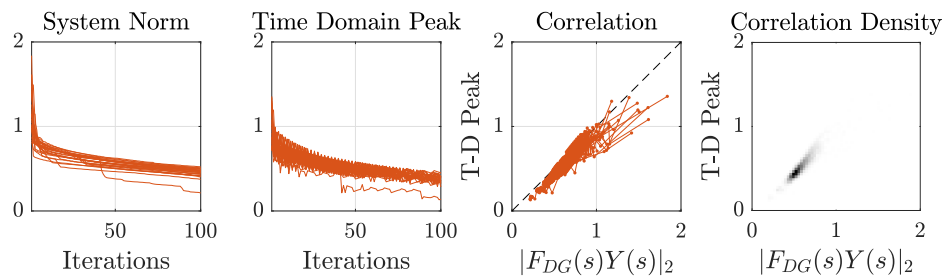




(a) Unweighted H_2 norm $|Y(s)|_2$ (b) DGIF-weighted H_2 norm $|F_{DG}(s)Y(s)|_2$ (c) Unweighted H_∞ norm $|Y(s)|_\infty$ (d) DGIF-weighted H_∞ norm $|F_{DG}(s)Y(s)|_\infty$


Figure 7.8: Iterative optimization for the wing root torsional moment and an $H = 70$ m discrete gust.




(a) Unweighted H_2 norm $|Y(s)|_2$



(b) DGIF-weighted H_2 norm $|F_{DG}(s)Y(s)|_2$



(c) Unweighted H_∞ norm $|Y(s)|_\infty$



(d) DGIF-weighted H_∞ norm $|F_{DG}(s)Y(s)|_\infty$

Figure 7.9: Iterative optimization for the wing root bending moment and an $H = 30$ m discrete gust.

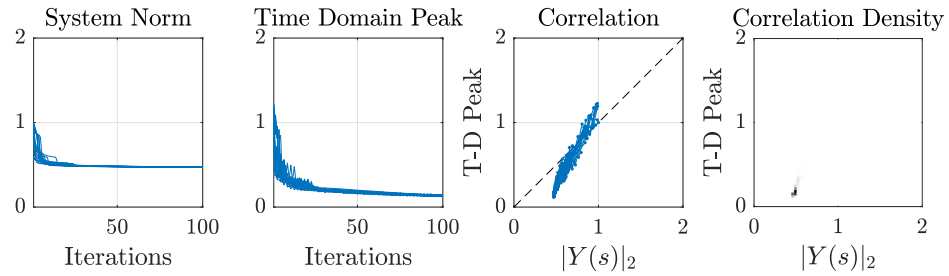
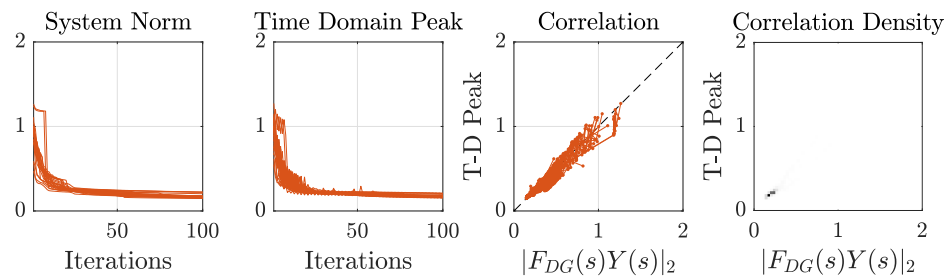
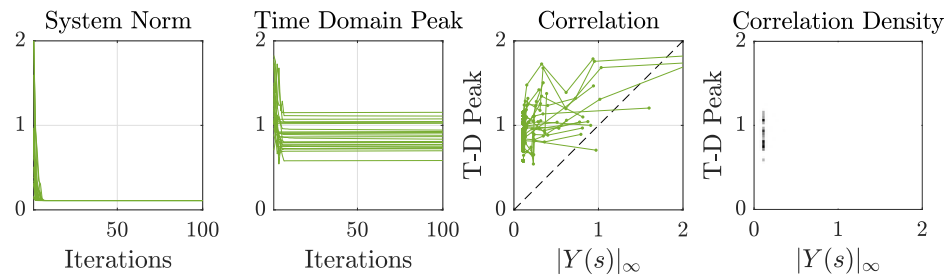

(a) Unweighted H_2 norm $|Y(s)|_2$ (b) DGIF-weighted H_2 norm $|F_{DG}(s)Y(s)|_2$ (c) Unweighted H_∞ norm $|Y(s)|_\infty$ (d) DGIF-weighted H_∞ norm $|F_{DG}(s)Y(s)|_\infty$

Figure 7.10: Iterative optimization for the wing root torsional moment and an $H = 30$ m discrete gust.

Summary: This chapter has introduced the *discrete gust impulse filter* (DGIF) and its potential for use as a control specification. The DGIF's impulse response matches a one-minus-cosine profile, thus replicating the time- and frequency-domain properties of a discrete gust. When used as a weighting function for in an H_2 control specification, it allows the energy of a gust response to be specified. Such H_2 -DGIF specifications are shown, through demonstrative examples, to consistently correlate well with the time-domain peak response.

Chapter 8

Automated Control Design

Contents

8.1	Concept definition	142
8.1.1	Multi-objective parameter synthesis	142
8.1.2	Definitions	143
8.1.3	Control design process	144
8.2	Control specifications	146
8.2.1	Discrete gusts	146
8.2.2	Continuous turbulence	146
8.2.3	Ride quality	147
8.2.4	Actuator use and bandwidth	148
8.3	Control design loop	148
8.3.1	Performance evaluations	149
8.3.2	Design loop iteration	150

With the results of the previous two chapters, two important obstacles to lidar-based GLA control design have been overcome. Firstly, now that the lidar system is well-modeled in the linear control problem, the properties and performance of the tuned controller should match those of the actual system. Secondly, DGIF-weighted performance specifications allow the designer to readily specify the time-domain peak response to a discrete gust. These specifications still require some iteration to reach the targeted time-domain performance, but the good correlation between the DGIF-weighted H_2 norm and the time-domain peaks suggests that simple iteration rules should be sufficient to automate the process.

This chapter proposes a method to automate the design process. Its main purpose is to allow the designer to express discrete gust control requirements in their natural form; translating these requirements into control specifications and iterating over them is done automatically. The following sections discuss the concept in more detail, list the most important control requirements and how they are expressed, and finally explain the workings of the automated design loop.

Part of the contents of this chapter are drawn from [60].

8.1 Concept definition

8.1.1 Multi-objective parameter synthesis

At a basic level, the automated control design process described here is a form of *multi-objective parameter synthesis*, similar to CONDUIT [51] and MOPS [97]. It is similar to them in that it aims to deal with industrial-size non-convex control problems which can include multiple competing requirements expressed in their natural form. The main difference is that it seeks to take advantage of the efficient and powerful nonsmooth-optimization-based synthesis methods implemented in *syntune* (see Ch. 5). This choice has important consequences for its implementation.

CONDUIT, for instance, is built around a Feasible Sequential Quadratic Programming (FSQP) optimizer. Such an optimizer is, in principle, “agnostic to design method and architecture” [51, p. 565]. This means it optimizes a set of parameters with respect to the design requirements, regardless of whether these parameters are the control gains themselves,

[97] Joos, 1999. A paper describing the MOPS methodology for tuning multiobjective control design.

or whether they are the tuning criteria of an inner control design method. In [51], for example, it is also used to tune Q and R matrices for LQR design, or in another instance the weighting filter parameters for full-order H_∞ synthesis. These ‘inner-loop’ control design methods are ideally relatively fast and globally optimal; otherwise, computing gradients/search directions quickly becomes problematic.

Systune, in some sense, is itself a multi-objective parameter synthesis method. As discussed in Ch. 5, it is neither globally optimal nor (compared to, e.g., LQR) particularly fast. This makes the use of generalized optimization methods (such as FSQP) in the ‘outer-loop’ difficult, if not impractical. Instead, heuristics are likely to be faster and more effective. The resulting automated design loop has distinct inner and outer optimization loops. The *systune*-based inner optimization loop synthesizes controllers according to a set of frequency-domain specifications. The outer optimization loop translates control requirements into frequency-domain specifications for the inner-loop and heuristically iterates over specification parameters to satisfy the requirements.

8.1.2 Definitions

The inner and outer loops are hence designated the *synthesis* and *design* loops, respectively, and their corresponding goals denoted *specifications* and *requirements*. A *requirement* (also denoted *Req*) belongs to the design loop and expresses goals in their ‘natural’ form, whereas a *specification* (*Spec*) refers to the frequency-domain control specifications used by the inner-level synthesis loop. Requirements have a multiplicity of specifications: each specification belongs to only one requirement, however each requirement may have a number of specifications.

Requirements are defined in three levels of *criticality*: *hard* requirements, *soft* requirements, and *open* requirements. *Hard* requirements act as *constraints*; they represent the acceptable limits of the system, and thus the *feasible* design space. *Soft* requirements act as *objectives*, i.e., the desired performance of the system, and thus define the *satisfactory* design space. The *hard* and *soft* requirements are thus akin to the hard and soft specifications defined in *systune* (see Sec. 5.1.4, p. 76), and their corresponding specifications are expressed as such. *Open* requirements define the criteria which are optimized once all soft and hard requirements are met. They serve to reduce overdesign, reallocating excess control authority

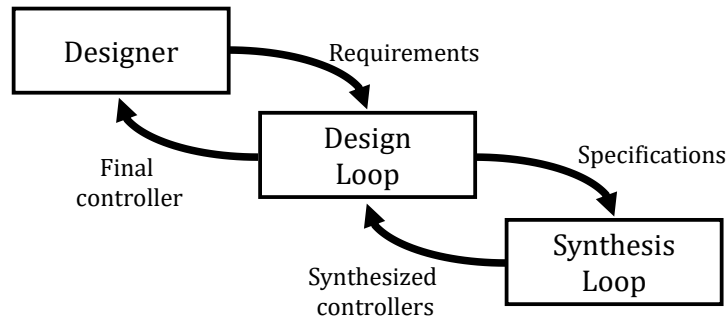


Figure 8.1: High-level schematic of the design and synthesis loops, requirements, and specifications.

to secondary objectives, such as minimizing actuator usage or improving ride comfort. Another parallel may be drawn here to CONDUIT, whose three design phases are defined similarly.²⁶ Due to the indirect relationship between the time-domain requirements and the frequency-domain specifications, it is generally impossible for the present control design process to proceed in a sequence of phases; instead, as explained below in Sec. 8.3, they are given priority in each iteration according to their criticality.

A final distinction must be made between *iterable* (*Iter*) and *non-iterable* (*Noniter*) requirements. *Noniterable* requirements are directly expressed in terms of H_2/H_∞ specifications, and so require no iterations. *Iterable* requirements are those which are specified, e.g., in the time domain, and therefore require some form of iteration. Here, only discrete-gust requirements are actually iterable, and they may only be defined for *hard* and *soft* requirements.

8.1.3 Control design process

Figure 8.2 provides an overview of the GLA control design process with the automated design tool.

Load envelopes and a database of linear models covering the flight envelope and all relevant mass cases are derived from the full nonlinear flight dynamics model (FDM). The control designer provides a set of top-level inputs to the control design tool. These include the desired flight points

²⁶ In the first, hard requirements are met; in the second, soft requirements are met subject to the hard requirements; and in the third an objective function is minimized to reduce overdesign and obtain a Pareto-optimal outcome [51].

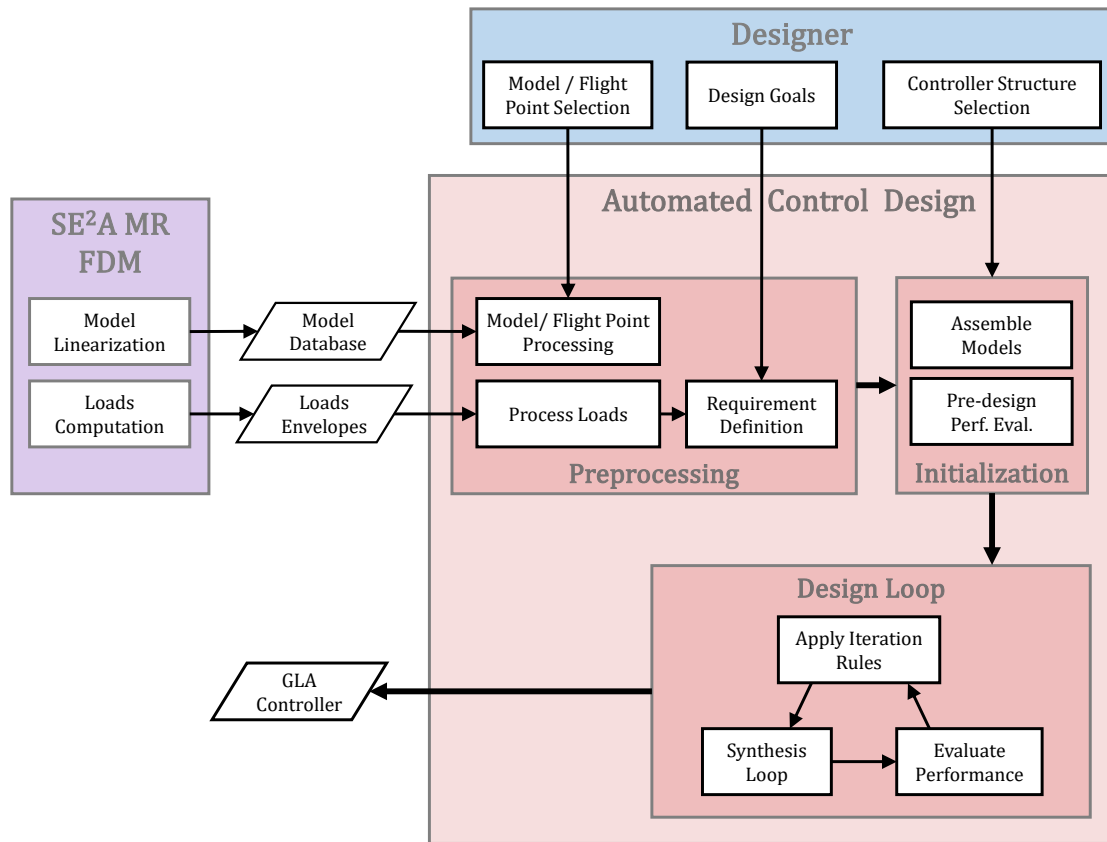


Figure 8.2: Automated GLA control design framework

and model parameters, including mass cases and system configurations (e.g., lidar system settings and time delays); the design goals, i.e., hard, soft, and open requirements; and the untuned controller, which defines its inputs (\mathbf{y}), outputs (\mathbf{u}), internal structure, tunable parameters and their initial values, and bounds on the tunable parameters, if necessary.

In the preprocessing step, the external inputs are checked against available models, the necessary model data and load envelopes are loaded, and the control requirements are assembled. The initialization step begins by preparing the linear design and evaluation models, including model order reduction, conversion to discrete time, integration of the lidar model, and additional augmentations such as time delays. Next, the initial specifications are created, using the design model to normalize them as needed. Finally, the (untuned) controller is connected to the design models and an initial performance evaluation is conducted. The control design loop then

begins, alternating between applying the iteration rules, synthesizing controllers, and evaluating the resulting performance. The loop ends when the stopping conditions defined in the iteration rules are met, producing the final tuned controller and a final post-design performance evaluation.

8.2 Control specifications

8.2.1 Discrete gusts

Discrete gust specifications are based on the H_2 -DGIF specifications explored in Ch. 7. For a generic SISO transfer function $T_{w \rightarrow z}(s)$, the H_2 weighting function $W_{DG}(s)$ is thus expressed:

$$W_{DG}(s) = K_{norm} K_{weight} F_{DG}(s) \quad (8.1)$$

with F_{DG} the DGIF corresponding to the desired airspeed and gust length, K_{weight} a scalar weighting gain, and K_{norm} the normalization factor.

The normalization factor K_{norm} is computed, if possible, such that $\|K_{norm} F_{DG} T_{w \rightarrow z}(s)\|_2 = 1$.²⁷ K_{weight} is then chosen such that for a targeted relative change η_{tgt} in the L_2 norm of the discrete gust response, $K_{weight} = 1/\eta_{tgt}$. If $\|T_{w \rightarrow z}(s)\|_2 \approx 0$, as is the case, e.g., for actuator deflections of the open-loop aircraft, the specification cannot be normalized and $K_{norm} = 1$; K_{weight} is then defined in absolute values.

Recalling that such specifications only define the L_2 norm of the discrete gust response and not its time-domain peak, the value of K_{weight} must be iteratively adjusted to reach the targeted time-domain peak. A corresponding set of iteration rules is proposed below, in Sec. 8.3.

8.2.2 Continuous turbulence

Comparing the definition of continuous turbulence requirements (Eq. 3.8, p. 50) and that of the H_2 norm (Eq. 5.2, p. 74), it is clear that continuous turbulence requirements can be expressed as H_2 specifications. All that is needed is a linear weighting filter F_{CT} whose frequency response matches that of the von Kármán PSD [80, p. 46], i.e., $|F_{CT}(s)|^2 = \Psi_{turb}$.

²⁷ Recall that the H_2 norm does not benefit from the multiplicative property (Sec. 5.1.1), so F_{DG} must be included to obtain the desired normalization.

The ‘ $-5/3$ ’ slope of the von Kármán PSD precludes a perfectly representative F_{CT} , but various approximations can be found. Here, the filter developed in [98, p. 46] is used:

$$F_{CT}(s) = \frac{1}{\sqrt{2\pi}\gamma} \frac{(1 + 2.618\gamma^{-1}s)(1 + 0.12981\gamma^{-1}s)}{(1 + 2.083\gamma^{-1}s)(1 + 0.823\gamma^{-1}s)(1 + 0.08977\gamma^{-1}s)} \quad (8.2)$$

with $\gamma = \frac{V_{TAS}}{L}$

The expression in Eq. 8.2 approximates the two-sided von Kármán PSD expressed as a function of ω , i.e. $\Phi_{turb}(\omega)/2$ (cf. Eq. 3.6, p. 49). This corresponds exactly to the definition of the H_2 norm in Eq. 5.2; the only factor left unaccounted for is $\sqrt{1/2\pi}$, which is easily included in the weighting function W_{CT} . This is then:

$$W_{CT}(s) = K_{norm}K_{weight}F_{CT}(s) = \frac{\sqrt{2\pi}U_\sigma}{z_{tgt}}F_{CT}(s) \quad (8.3)$$

where z_{tgt} is the targeted value of the output variable. Here K_{norm} may be simply be chosen as 1, seeing as z_{tgt} can be specified directly. Alternatively, once again computing it as $\|K_{norm}F_{CT}(s)T_{w \rightarrow z}(s)\|_2 = 1$ allows a relative change to be specified via $K_{weight} = 1/\eta_{tgt}$.

8.2.3 Ride quality

Following the example of [92], ride quality/passenger comfort metrics can be addressed via ISO 2631-1. This standard defines a set frequency weightings which cover vertical and horizontal vibration and motion sickness. The criteria themselves are defined as weighted 2-norms, in much the same way as the continuous turbulence requirements (see Eq. 3.8, p. 50). Consequently, H_2 specifications may be defined following precisely the same logic as the continuous turbulence specifications above.

Reference [99] provides low-order approximations of the frequency weightings which may be used as weighting functions. For reference, one example of each is reported below [99]:

[98] Barr *et al.*, 1974. FAA report on wind models for flight simulators.

[99] Zuo *et al.*, 2003. A paper in which low-order linear filters approximating ISO2631-1 weighting filters are identified.

Vertical acceleration:

$$W_k(s) = \frac{81.89s^3 + 796.6s^2 + 1937s + 0.1446}{s^4 + 80s^3 + 2264s^2 + 7172s + 21196} \quad (8.4)$$

Horizontal acceleration:

$$W_d(s) = \frac{12.66s^3 + 163.7s^2 + 60.64s + 12.79}{s^4 + 23.77s^3 + 236.1s^2 + 692.8s + 983.4} \quad (8.5)$$

Motion sickness:

$$W_f(s) = \frac{0.1457s^4 + 0.2331s^3 + 13.75s^2 + 1.705s + 0.3596}{s^5 + 7.757s^4 + 19.06s^3 + 28.37s^2 + 18.52s + 7.23} \quad (8.6)$$

8.2.4 Actuator use and bandwidth

As discussed in Sec. 5.3.1 (p. 82), it is advisable to impose upper and lower limits on the available bandwidth for actuators. This can take the form of an H_∞ template with roll-on and roll-off bounds. The resulting weighting filter W_{bw} is then expressed as:

$$W_{bw}(s) = K_{weight} \underbrace{\left(\frac{s + \omega_{ro}}{s} \right)^{n_{ro}}}_{\text{Roll-on}} \underbrace{\left(\frac{s + \omega_{bw}}{\omega_{bw}} \right)^{n_{bw}}}_{\text{Roll-off}} \quad (8.7)$$

where K_{weight} is the scalar weighting gain, ω_{ro} and n_{ro} the roll-on frequency (in rad/s) and order, and ω_{bw} and n_{bw} the roll-off frequency and order. The roll-on/roll-off orders n_{ro} and n_{bw} , respectively, are positive integers (or 0); the resulting slopes of the template are equal to $+20 \cdot n_{ro}$ dB/dec and $-20 \cdot n_{bw}$ dB/dec. A visual representation of the weighting function and its corresponding template is plotted in Fig. 8.3.

8.3 Control design loop

The design loop is shown in the bottom right of Fig. 8.2. Each iteration goes through three steps: a controller is produced by the synthesis loop, its performance is evaluated, and a set of iteration rules determines whether to exit the loop or to update the control specifications for the

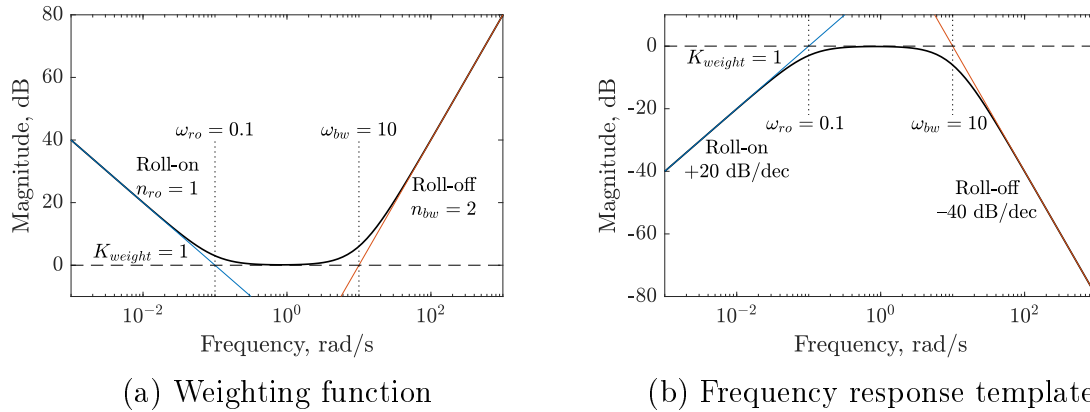


Figure 8.3: Characteristics of $W_{bw}(s)$ calculated with a sample set of parameter values.

next synthesis step. As mentioned above, the synthesis loop is *systune*; it is discussed in some detail in Ch. 5. The other two steps are described below.

8.3.1 Performance evaluations

Performance evaluations serve to evaluate the current performance of both specifications and requirements. *Specification performance* η_S refers to the values of the frequency-domain norms used to define specifications. *Requirement performance* η_R instead depends on requirement type; for hard and soft requirements, it is defined such that a requirement is met when $\eta_R \leq 1$. For requirements which are directly expressed as frequency-domain norms (i.e., noniterable requirements), $\eta_R = \eta_S$.

For iterable (i.e., discrete gust) requirements, η_R refers to the ratio between the actual and the targeted *incremental* time-domain peaks, respectively z_{peak} and z_{tgt} . Time-domain peaks are computed from a series of discrete gust simulations, for which the range of gust lengths and corresponding peak gust velocities are computed according to the certification specifications [4]. If the targeted value is specified as an absolute value, the incremental target can be found by subtracting the trim value. The performance is then computed as $\eta_R = z_{peak}/z_{tgt}$. When a specification or requirement is applied to more than one model in a multimodel control problem, η_S and η_R are computed individually for each model.

The models used for the evaluation may vary. Here, the reduced-order

design model is used for frequency-domain evaluations and the a full-order linear model for time-domain simulations.

8.3.2 Design loop iteration

Figure 8.4 is a flowchart depicting the proposed design loop iteration rules. As mentioned in Sec. 8.1.2, it is not possible to proceed linearly, first satisfying *hard*, then *soft*, and finally *open* requirements. Instead, these rules use a priority hierarchy; in other words, at each iteration, only the highest-priority category of unmet requirements is addressed. If any *hard* requirements are unmet, only *hard* requirements are addressed; if all *hard* requirements are met but some *soft* requirements are not, then only *soft* requirements are addressed; and if all *hard* and *soft* requirements are met, then *open* requirements are taken into consideration.




Figure 8.4: Flowchart of the control design loop iteration rules.

Three final outcomes are possible: *hard fail* if the *hard* requirements cannot be met, *soft fail* if the *soft* requirements cannot be met, and *success* if all *hard* and *soft* requirements are met. If *noniterable* requirements are unmet, this immediately results in failure of the design loop. If only

iterable requirements are unmet, they are updated according to a set of specification update rules, although if certain stopping conditions are met, this also results in a design loop failure.

The hard and soft specification update rules are shown in Figure 8.5. These include two types of operations: *tightening* and *relaxing*, both shown schematically in Fig. 8.6.

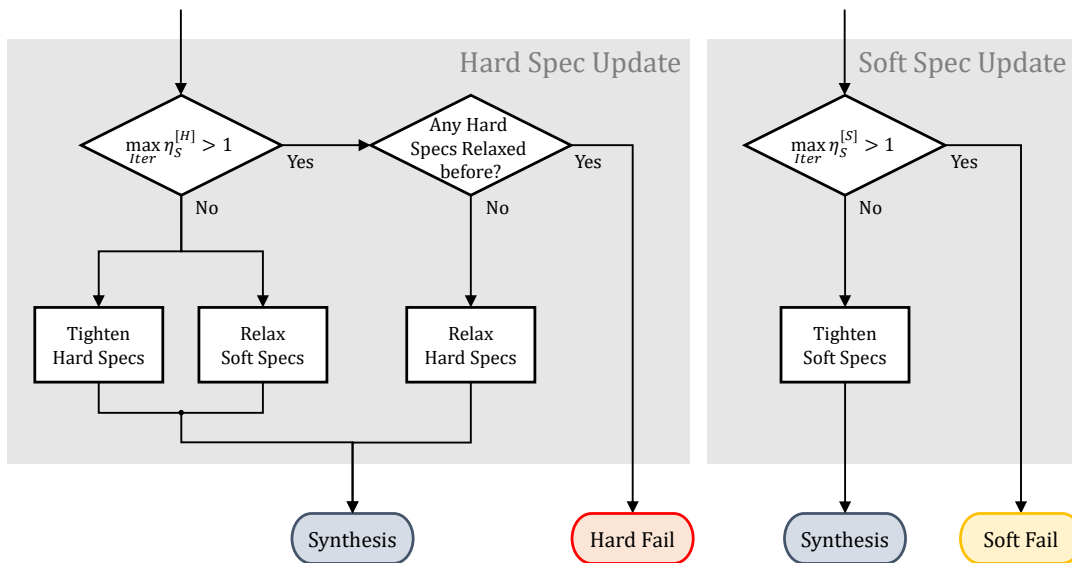
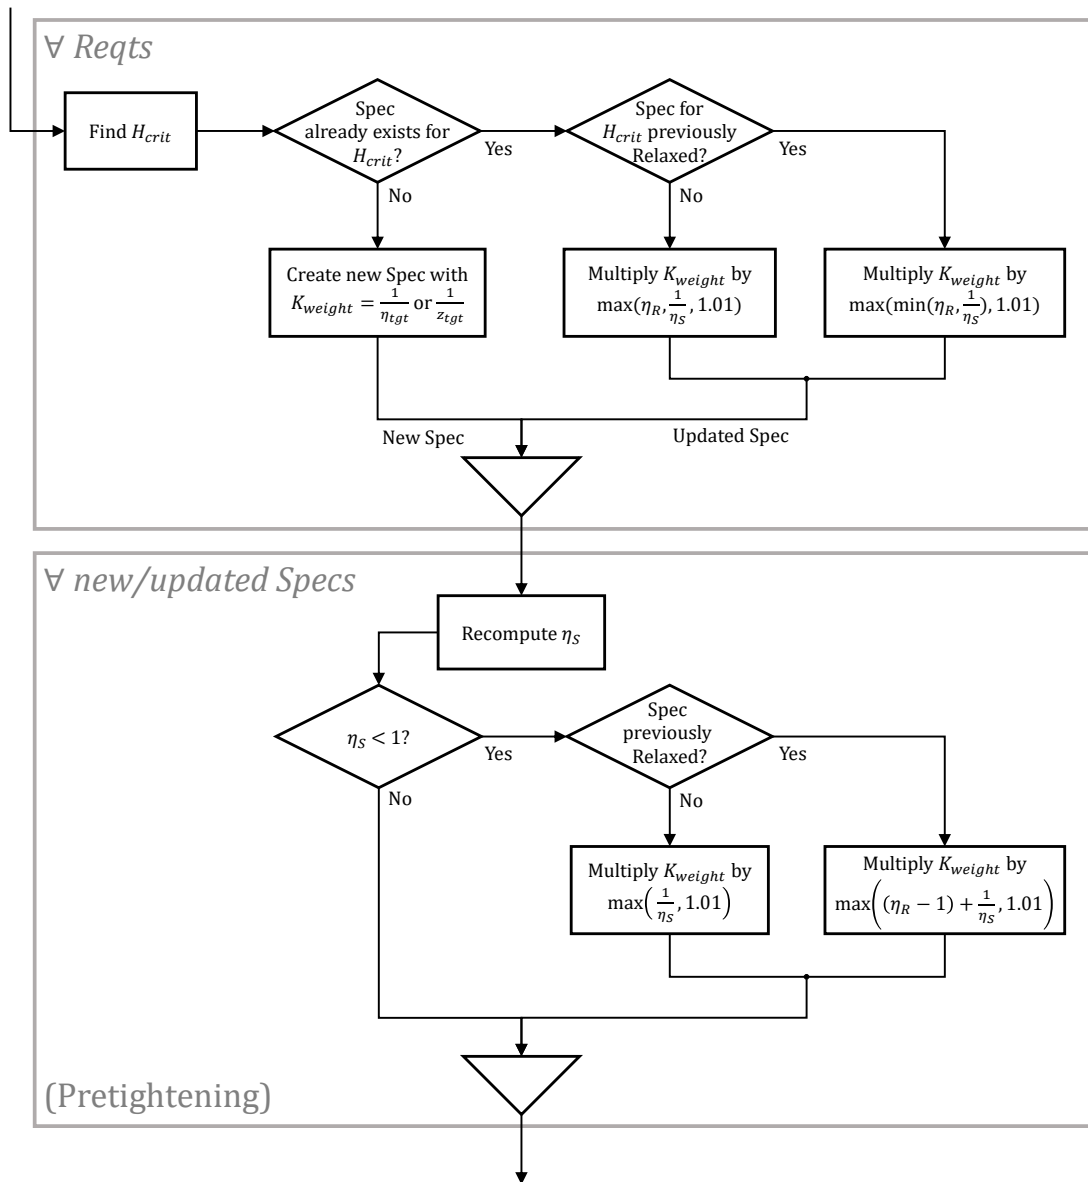


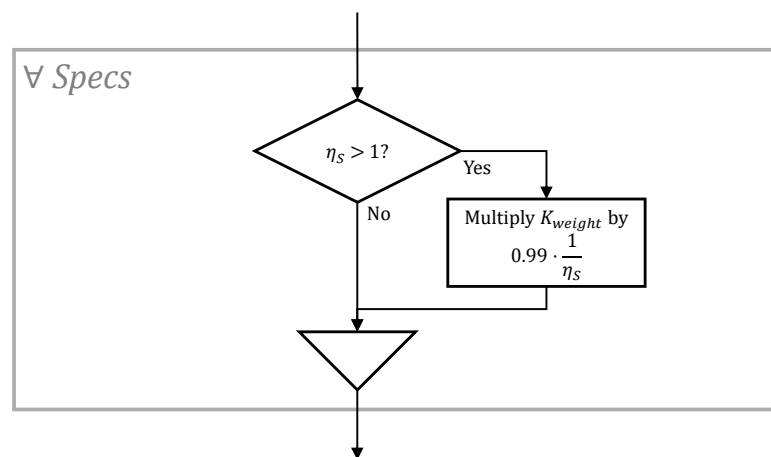
Figure 8.5: Flowchart of the specification update rules.

Tightening serves to drive down the time-domain peak, either by increasing K_{weight} on existing specifications, or by adding new ones. For each requirement, worst-case gust length H_{crit} is computed; if a specification already exists for that H , K_{weight} is increased; if not, a new one is added, with an initial guess for the value of K_{weight} . An additional *pretightening* step is performed on updated Specs. This serves to ensure that the tightened specification has an effect on the next synthesis; if the initial specification performance $\eta_S \leq 1$, it is already satisfied. *Pretightening* checks whether $\eta_S \geq 1$, and if not, increases K_{weight} to make it so. *Relaxation* instead does the opposite: it reduces K_{weight} to ensure that all concerned specifications end up with an $\eta_S \approx 0.99$ at the beginning of the next synthesis step. This serves to either ‘reset’ the specifications if they have been overtightened, or to remove them from the next synthesis step.

In both the initial *tightening* and the *pretightening* steps, the factor by which K_{weight} is multiplied depends on whether that specification has been



(a) Tighten



(b) Relax

Figure 8.6: Flowcharts of the specification update operations.

previously *relaxed*. If not, it is tightened more aggressively to attempt to converge as quickly as possible. Once a specification has been relaxed, it is assumed that it had previously overshoot the achievable limit K_{weight} value, so it is retightened more gradually to avoid overshooting again.

Looking at Figure 8.5, for both *hard* and *soft* updates, specifications are only *tightened* if $\eta_S \leq 1$. For a *hard* specification update, when hard specifications are *tightened*, soft specifications are *relaxed*. This serves to get the soft requirements ‘out of the way’, allowing the hard requirement performance to be brought into an acceptable range as quickly and with as little computational effort as possible. If there are specifications with $\eta_S > 1$, they may be relaxed only once; if it occurs again in a subsequent iteration, a *hard fail* is triggered. For a *soft* specification update, relaxation is not allowed; specifications with $\eta_S > 1$ trigger a *soft fail*.

Once all *hard* and *soft* requirements are met at least once, the design is, in essence, already a *success*. If any *open* requirements have been defined, an *open synthesis* is performed (Fig. 8.4). *Open synthesis* consists in temporarily redefining all *soft* specifications as *hard* ones (i.e., both *soft* and *hard* specifications are treated as *hard* by *systune*), and *systune* attempts to reduce the *open* specifications to 0 with no random restarts.²⁸ This could cause the *hard* or *soft* requirement performance η_R to exceed 1, so several more iterations may be needed. Once an open synthesis is completed without violating any other requirements, the design loop finally terminates with a *success*.

Summary: This chapter has presented an automated GLA control design framework. It deals with discrete gust requirements by iteratively adding and tuning DGIF-weighted H_2 specifications. A priority system of hard and soft requirements allow constraints and objectives to be distinguished, and open requirements reduce overdesign by reallocating any excess control authority. Other, non-iterated requirements can be specified in parallel, and robustness is taken into account via multi-model synthesis.

²⁸ In a normal synthesis step, if it is unable to meet the stopping conditions, the synthesis may be restarted a number of times with randomized initial controller gains to help avoid getting stuck in unfavorable local minima; the process stops when $\max \eta_S \leq 1$ or it runs out of restarts. See Ch. 5 for more details.

Chapter 9

Application

Contents

9.1	Baseline Design Case	156
9.2	Nominal Design	158
9.3	Robust Design	158
9.3.1	Design problem	158
9.3.2	Design results	161
9.3.3	Comparison to Nominal design	162

With the above tools and methods in hand, this chapter serves to demonstrate their use for a sample GLA design problem. First, a baseline design case for the SE²A Mid-Range aircraft defining the primary objectives and constraints is presented. Then, a *nominal* design accounting only for the reference flight point is performed and analyzed. Following that, in Section 9.3, the design problem is augmented with a series of robustness requirements, and the resulting controller is compared to the nominal controller.

9.1 Baseline Design Case

The baseline gust load alleviation design goals are chosen as:

1. Reduce inner-wing vertical bending gust loads to those of the symmetrical maneuver loads.
2. Do not allow incremental inner-wing torsional loads to exceed those of the open-loop gust load envelope.
3. Limit aft fuselage bending loads to those of the open-loop gust load envelope.
4. Limit HTP bending and torsional loads to those of the open-loop gust load envelope.
5. Ensure ride qualities in terms of the peak vertical acceleration and ISO 2631-1 motion sickness metric are not degraded with respect to the open loop.
6. Limit the control surface rates to 40°/s and deflections to 10°.
7. Ensure the frequency response of control surface commands tends to 0 at very low and high frequencies.

Table 9.1 lists the corresponding requirements. Inner wing loads are specified via two load stations: the wing root (with subscript $WL, 76$) and an inner-mid-wing station slightly outboard of the engine nacelles ($WL, 85$). Fuselage loads are specified at a load station just ahead of the HTP ($FU, 35$), and HTP loads are specified only at the root ($HL, 60$). Ride quality requirements are evaluated from local vertical acceleration at the pilot station ($a_{z,fus, fwd}$) and at the aftmost end of the cabin ($a_{z,fus, aft}$).

Table 9.1: Baseline control requirements

Goal	Req. ID	Output	Description	Requirement Type	Performance Target
Soft Requirements					
1	S1	$M_{x,WL,76}$	Wing root bending moment	CS 25.341(a) Discrete Gusts	\leq Maneuver load envelope
	S2	$M_{x,WL,85}$	Inner mid-wing bending moment		
2	S3	$M_{y,WL,76}$	Wing root torsional moment	CS 25.341(a) Discrete Gusts	\leq Open-loop gust load envelope
	S4	$M_{y,WL,85}$	Inner mid-wing torsional moment		
3	S5	$M_{y,FU,35}$	Aft fuselage bending moment	CS 25.341(a) Discrete Gusts	\leq Open-loop gust load envelope
4	S6	$M_{x,HL,60}$	HTP root bending moment	CS 25.341(a) Discrete Gusts	\leq Open-loop gust load envelope
	S7	$M_{y,HL,60}$	HTP root torsional moment		
5	S8	$a_{z,fus,fwd}$	Vertical accel. at pilot station.	CS 25.341(a) Discrete Gusts	\leq Open-loop peak
	S9	$a_{z,fus,aft}$	Vertical accel. at aft cabin.		
	S10	$a_{z,fus,fwd}$	Vertical accel. at pilot station.		
	S11	$a_{z,fus,aft}$	Vertical accel. at aft cabin.		
Hard Requirements					
6	H1	δa_{sym}^1	Symmetrical aileron deflection rates	CS 25.341(a) Discrete Gusts	$\leq 40^\circ/s$
	H2	δa_{sym}^2			
	H3	δa_{sym}^3			
	H4	δa_{sym}^4			
	H5	$\delta \dot{e}$	Elevator deflection rate		
	H6	δa_{sym}^1	Symmetrical aileron deflections	CS 25.341(a) Discrete Gusts	$\leq 10^\circ$
	H7	δa_{sym}^2			
	H8	δa_{sym}^3			
	H9	δa_{sym}^4			
	H10	δe	Elevator deflection		
7	H11	δa_{sym}^1	Symmetrical aileron deflection rates	H_∞ bandwidth	$K_{weight} = 20,$ $n_{ro} = 1, \omega_{ro} = 0.05 \times 2\pi \text{ rad/s},$ $n_{bw} = 2, \omega_{bw} = 10 \times 2\pi \text{ rad/s}$
	H12	δa_{sym}^2			
	H13	δa_{sym}^3			
	H14	δa_{sym}^4			
	H15	$\delta \dot{e}$			

9.2 Nominal Design

The *nominal* design case considers only the reference model (see Sec. 2.3, p. 37), i.e., the SE²A MR aircraft at 6000 m altitude, Mach 0.76, and MTOW mass case. The digital flight control system is assumed to run with a discrete sampling rate of 100 Hz ($T_s = 10$ ms) and with an overall time delay of 50 ms. The lidar parameters are identical to those shown in Table 6.1 (p. 106). The controller wind field is defined with a preview horizon $h_{pre} = 45$ and postview $h_{post} = 25$; the maximum preview distance is thus ≈ 108 m.

The design process concluded successfully after 6 iterations. Figure 9.1 plots the requirement performance η_R for all requirements at the beginning (Fig. 9.1a) and end (Fig. 9.1b) of the process. Each column in the figures represents a requirement from Table 9.1. The final controller meets most requirements with some margin.

The resulting gust load envelope for the design flight point is shown in Fig. 9.2, along with the overall open-loop load envelopes and the load requirements. It is clear that the wing bending loads (Fig. 9.2a) have been reduced across nearly the entire span; in fact, only the HTP loads (Figs. 9.2b and 9.2d) and the mid-aft fuselage bending moment (Fig. 9.2e) are slightly increased, evidently due to the use of the elevator.

9.3 Robust Design

9.3.1 Design problem

Though the *nominal* design shown above appears to perform well, it has been designed and evaluated on a single flight point and configuration. A GLA controller should be robust against the various uncertainties to which it will be exposed and which cannot be reliably compensated via, e.g., gain scheduling. Here, a robust controller is designed mainly by using the multi-model method. The objective is to meet the requirements of the baseline design case subject to uncertainties in mass distribution and time delay, while minimizing actuator use and susceptibility to lidar measurement noise.

Consequently, the final result of the nominal design is taken as starting point, and 4 further models are added to the design problem: one each for








Figure 9.1: Requirement performance η_R at the beginning and end of the nominal control design, plotted in parallel coordinates. Requirements are met when $\eta_R \leq 1$.




(a) Wing vertical bending moment




(b) HTP vertical bending moment



(c) Wing torsional moment



(d) HTP torsional moment



(e) Fuselage bending moment

Figure 9.2: Gust load envelope of the reference flight point using from the nominal controller, plotted against the overall open-loop gust- and maneuver-load envelopes. Black chevrons indicate load requirements.

the remaining mass cases (OEW, MZPW, and MZFW), and another copy of the reference model with an additional 50 ms of delay (i.e., a total of 100 ms). Furthermore, Table 9.2 defines two *open* requirements. These are defined as single-input multiple-output (SIMO) H_2 requirements, which is essentially equivalent to the sum of the H_2 norms of the individual input-output pairs. The first (O1) seeks to minimize actuator usage in response to gusts and turbulence. The requirement therefore is defined with true wind input w_{turb} and includes a von Kármán continuous turbulence weighting filter F_{CT} (see Sec. 8.2.2, p. 146). The second should minimize the control response to measurement noise uncertainty in the wind estimate. It takes as input lidar white noise input d_{lead} , and is defined as an unweighted SIMO H_2 requirement. Both requirements are normalized such that their $\eta_R = 1$ at the start of the design process.

Table 9.2: Open requirements for the robust design case.

Req. ID	Input	Outputs	Description	Requirement Type
Open Requirements				
O1	w_{turb}	$\delta \dot{e}$ $\delta \dot{a}_{sym}^1$ $\delta \dot{a}_{sym}^2$ $\delta \dot{a}_{sym}^3$ $\delta \dot{a}_{sym}^4$	Actuator rates in turbulence	SIMO H_2 weighted with F_{CT}
O2	d_{lead}	$\delta \dot{e}$ $\delta \dot{a}_{sym}^1$ $\delta \dot{a}_{sym}^2$ $\delta \dot{a}_{sym}^3$ $\delta \dot{a}_{sym}^4$	Actuator rates from measurement noise	Unweighted SIMO H_2

9.3.2 Design results

The design process terminated successfully after 32 iterations. The requirement performance plots in Fig. 9.3 show its performance over the course of the design process.


The initial performance (Fig. 9.3a) corresponds to that of the *nominal* controller. Clearly, it is not sufficiently robust: the delayed model exceeds the required wing bending and HTP loads, and the MZPW case exceeds

the limit peak acceleration at the pilot station. Figure 9.3b shows the first iteration in which all *soft* and *hard* requirements are met, before the first *open* synthesis step, and Figure 9.3c shows the final results.


The open requirements are significantly improved by the end, and there are no significant degradations in the other requirements. Judging by the requirement performance, it appears that the improvement in the open requirements is correlated with an improvement in the H_∞ bandwidth requirements' performance. The load envelopes in Fig. 9.4 show that although most loads are at the limits, the original design goals are all met.

9.3.3 Comparison to Nominal design


Figure 9.5 plots the open-loop reference model's response against that of the *nominal* and *robust* controllers for a discrete gust with $H = 70$ m. Although both controllers have a similar effect on the wing loads and accelerations, they have different control strategies. The *nominal* controller clearly pursues a pitching strategy, whereby the elevator is deflected ahead of gust impact, and the aileron deflections appear mainly to support and compensate this elevator action. In the *robust* controller, the elevator oscillates at a higher frequency, evidently acting on Mode 3 (see Sec. 2.3) around 3 Hz; for comparison, look at the wing root torsional moment ($M_{y,WL,76}$). The main load alleviation effort seems to come from the ailerons, which have generally larger peak deflections than the *nominal* response. Inner aileron pairs 3 and 4 produce a direct lift effect, with negative (trailing edge up) deflections in phase with the peak wing bending response. Outer aileron pairs 1 and 2 (with some participation by 3 and 4 as well) instead appear to do the opposite, deflecting trailing-edge-down ahead of the peak gust response. This may be a combination of an attempt to use their weak effect on pitching moment together with twisting the outer wing into the gust.



(a) Initial performance with *nominal* controller




(b) Performance when *soft* and *hard* requirements are first met and before *open* optimization




(c) Final performance


Figure 9.3: Requirement performance η_R over the course of the robust control design, plotted in parallel coordinates. Requirements are met when $\eta_R \leq 1$.




(a) Wing vertical bending moment




(b) HTP vertical bending moment



(c) Wing torsional moment



(d) HTP torsional moment



(e) Fuselage bending moment

Overall Load Envelopes:	— Open-loop Gust Loads	— Maneuver Loads			
Line Types:	- - - Trim Open-loop	— Closed-loop		
Aircraft Models:	× MTOW	○ OEW	□ MZFW	◇ MZPW	△ MTOW + Delay

Figure 9.4: Gust load envelopes of the models considered in the robust design case, plotted against the overall open-loop gust- and maneuver-load envelopes. Black chevrons indicate load requirements.




Figure 9.5: Comparison of $H = 70$ m discrete gust response of the reference model in open-loop, with the nominal controller, and with the robust controller.

Summary: This chapter presented a lidar-based GLA control design case using the methods presented in this thesis. The baseline design requirements aim at a reduction in inner-wing bending moment while including constraints on actuator use, ride qualities, and other loads. An initial nominal design is performed on a single design point using only the baseline requirements. A robust design, starting from the results of the nominal design, then takes several mass cases and an additional time delay into account, as well as two open requirements aimed at minimizing actuator use. The designs are successful; the robust case is shown to meet all requirements, and the open requirements succeed in minimizing actuator excitation in response to turbulence and measurement noise.

Chapter 10

Conclusion and Outlook

10.1 Conclusion

The objectives of this thesis were expressed as three scientific questions in the introduction. These questions are repeated below, and the answers developed throughout this work are summarized for each one. Seeing as the first question depends on the results of the second two, they are discussed here in reverse order.

3. How can the lidar system's behavior and limitations be modeled in the linear control problem?

The type of lidar system considered here deals with large numbers of noisy measurements by employing a wind estimation algorithm. This, among other effects, results in low-pass dynamics with respect to both the true turbulent wind field and the measurement noise. The legacy lidar models consisting of tapped delay lines are unable to account for such behavior, essentially resulting in a control design which assumes that the lidar system provides perfect measurements of the true turbulent wind field.

Starting from the expression of the wind estimation algorithm, it is possible to analytically derive a linear filter which closely models the dynamic behavior of the full lidar system. First, a 'reference' measurement database containing only measurement metadata (position, direction, noise standard deviation), and which is representative of the contents of an actual measurement database, is built up based on the selected lidar sensor pa-

rameters. From this, the parameters of the wind estimation algorithm can be computed and used to assemble the filter.

The resulting linear filter is able to closely reproduce the time- and frequency-domain characteristics of the lidar system's response to both the true turbulent wind field and to measurement noise. If included in the control problem, this filter allows the synthesized controller to 'see' and account for the associated loss of information. Depending on how the control specifications are defined, it may also be made robust to measurement noise uncertainties. The filter is, however, still a linear system, and thus cannot adequately account for any nonlinearities.

2. How can time-domain discrete gust requirements be systematically expressed in terms of frequency-domain system norms?

Discrete gust requirements are expressed as the time-domain peak response over a family of one-minus-cosine discrete gusts. H_2 and H_∞ norms cannot be easily linked to this time-domain peak, however by way of Parseval's theorem, the H_2 norm is equivalent to the signal energy of the impulse response of a system. A discrete gust impulse filter (DGIF), i.e., a filter whose impulse response is equivalent to a one-minus-cosine discrete gust, allows an aircraft model's impulse response to be converted into a discrete gust response. A DGIF-weighted H_2 control specification can thus specify the signal energy of the response to a discrete gust of a certain length.

Although this is still not a direct connection to the time-domain peak, there is a much clearer link. To systematically express control specifications, this is enough; if such specifications are consistently well-correlated with the time-domain peak, they can be easily iterated to reach the required performance. Although it is probably not possible to prove that such a correlation can be found for all possible transfer functions, a couple of demonstrations were proposed to evaluate this property. It is shown that, compared to unweighted H_2 and H_∞ norms as well as the DGIF-weighted H_∞ norm, the DGIF-weighted H_2 norm appears to have a consistently good correlation with the time-domain peak response.

1. How can the lidar-based gust load alleviation control design process be automated?

The favorable properties of the H_2 -DGIF specifications allow a control designer to easily iterate over a control specification to reach a desired level of performance for a discrete gust of a given length. For full-scale design problems, which may involve tens to hundreds of requirements, manually iterating over the specifications is impractical; some form of automated design process is needed. The structured control synthesis tool considered here is too costly (in terms of time and computation) for generic optimization methods, such as those used in multi-objective parameter synthesis (MOPS) tools, to be feasible.

Instead, a set of relatively simple iteration rules are defined to automate this process. These iteration rules add and tune H_2 -DGIF specifications as necessary, based on time-domain simulations, to drive the discrete gust performance to its required levels. A system of hard, soft, and open requirements allow various constraints and objectives to be taken into account according to their criticality.

The resulting automated control design system is shown to be capable of handling a relatively complex design case, including robustness to mass case variations, time delays, and measurement noise uncertainties.

10.2 Outlook

The methods developed in this thesis are directed at improving control design within the context of the current certification specifications. The methods themselves may still be improved, and future developments may need to address an evolving certification landscape and more detailed aspects of a GLA function's implementation. Several possible avenues of future research are described below.

More efficient auto-tuning The set of iteration rules in the automated tuning process proposed here are only one possibility. They appear to work well, but they were chosen for transparency and reliability rather than efficiency. Improvements in the iteration rules or even a completely different approach to optimization could significantly accelerate the process. For example, under the current rules, in each discrete gust requirement, only a single control specification corresponding to a single gust length can be added or modified during each iteration. Some transfer functions are dominated by their response to one or two gust lengths, and so they may converge quickly under this system, but others are sensitive to the full range of possible gust lengths, requiring many iterations before they are fully specified. At the same time, simply specifying all gust lengths for all requirements risks creating a prohibitive number of control specifications. Finding a way to determine how many gust lengths need to be specified for each requirement during each iteration could thus minimize the needed number of iterations without an excessive computational burden.

Integrated multidisciplinary aeroservoelastic optimization To allow aircraft designs to fully exploit the advantages of a GLA function, the aircraft should be optimized taking its capabilities into account. Variations in the aircraft's structure and characteristics affect control design and performance; for instance, a lighter structure may have more torsionally-flexible wings, reducing the effectiveness of wing-mounted control surfaces and thus GLA performance. GLA controllers must therefore be constantly designed and redesigned as the aircraft evolves. The automated design process presented in this thesis could be integrated into a multidisciplinary aircraft design tool, allowing designs to account for realistic GLA performance.

Preview controller parametrization The standard preview controller used in this thesis is relatively simple, but in the examples above, it has hundreds of tunable parameters. Due to the relatively limited bandwidth of the lidar-derived wind estimates, and the greater sensitivity to phase uncertainties at higher frequencies, most of these degrees of freedom are essentially superfluous. A controller which is parametrized, for example, with polynomial functions or with wavelets, could achieve the same performance with a fraction of the tunable parameters. This would greatly simplify the control problem, improve the transparency and implementability of the resulting controller, and possibly even improve the quality of the result.

Improved random restarts The non-convex nature of structured multiobjective control synthesis implies that the initial values of the controller parameters strongly influence which local minimum the optimization process will end up in. The built-in restart function in *syntune* automatically resamples the control parameters with random values and restarts the synthesis process until the stopping conditions are met or the maximum number of restarts is reached. While this should eventually allow the synthesis procedure to get out of a poor local optimum and increase the likelihood of finding a favorable one, using randomly sampled parameters also often leads to poor or unrealistic outcomes, leading, in turn, to an unreliable, time-consuming and computationally expensive synthesis process. Finding a method to compute more meaningful initial values, perhaps on the basis of prior knowledge of the system and control problem, could yield a significant improvement.

Noise thresholding The measurement noise uncertainties in the lidar-based wind estimates are always present, even in the absence of turbulence. On a real aircraft, the control surface deflections resulting from such uncertainties in calm air are unacceptable, so they would presumably need to be filtered out or thresholded in some way. This cannot be accomplished with linear filtering, so a nonlinear component will be introduced into the GLA function. The methods developed in this thesis may be used to help design and size such a threshold function, e.g., by selecting the threshold as a function of the standard deviation of the uncertainties at each wind field node. Future developments will need to find a way to evaluate the

effect of such nonlinearities on control performance, especially when flying in light or medium turbulence, and account for them in control design.

Gust and turbulence models The current certification specifications include only one-dimensional turbulence and the only type of discrete gust profile is the one-minus-cosine gust, and the methods discussed in this thesis have been developed accordingly. True turbulent wind fields are multidimensional, and are certainly not limited to one-minus-cosine eddies. It is reasonable to expect that as aircraft become more sensitive to gust loads due to their configuration and structural properties, the certification specifications will also adopt more sophisticated models of turbulence. Vortex-shaped discrete gusts, spanwise non-uniform continuous turbulence and gusts, and Statistical Discrete Gusts, for example, are potential candidates. For wind lidar systems, three-dimensional wind fields may even be necessary due to the measurement geometry. The lidar measurement geometry, its wind estimation system, and the control design requirements will evolve accordingly, so control design methods will need to be adapted as well.

Bibliography

- [1] Elisabeth van der Sman, Bram Peerlings, Johan Kos, Rogier Lieshout, and Thijs Boonekamp, “Destination 2050: A Route to Net Zero European Aviation,” NLR-CR-2020-510, Netherlands, Feb. 2021.
- [2] Eurocontrol, “The Challenge of Long-Haul Flight Decarbonisation: When Can Cutting-Edge Energies and Technologies Make a Difference?” Think Paper 21, Eurocontrol, Aug. 2023.
- [3] IEA, “Transport, Energy, and CO₂,” International Energy Agency, 2009.
- [4] EASA, *CS-25: Certification Specifications and Acceptable Means of Compliance for Large Aeroplanes*. Cologne, Germany: European Union Aviation Safety Agency, Jan. 2023, Vol. Amendment 27.
- [5] Robert L. Allison, Brian R. Perkin, and Richard L. Schoenman, “Applications of Winglets and/or Wing Tip Extensions with Active Load Control on the Boeing 747,” in *NASA Conference Publication 2036: CTOL Transport Technology*, Langley Research Center, Hampton, Virginia: NASA, Mar. 1978.
- [6] Andrea Castrichini, Vijaya Hodigere Siddaramaiah, Dario Calderon, Jonathan E. Cooper, Thomas Wilson, and Yves Lemmens, “Nonlinear Folding Wing-Tips for Gust Loads Alleviation,” in *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Kissimmee, Florida: American Institute of Aeronautics and Astronautics, Jan. 2015.
- [7] R. L. Schoenman and H. A. Shomber, “Impact of Active Controls on Future Transport Design, Performance, and Operation,” in *Aerospace Engineering and Manufacturing Meeting*, Feb. 1975, p. 751 051.

- [8] Jochen Bucher, “Schlussbericht für das Projekt Con-Move: Konzepte, Entwürfe und Verifikation von integrierten Movable-Funktionen : Projektdauer: 01.01.2016-3.03.2020,” Airbus Operations GmbH, 2021, 1 Online-Ressource (76 Seiten, 4,64 MB).
- [9] Gary J. Balas, “Flight Control Law Design: An Industry Perspective,” *European Journal of Control*, Vol. 9, No. 2-3, pp. 207–226, Jan. 2003.
- [10] H. Hönlinger, H. Zimmermann, O. Sensburg, and J. Becker, “Structural Aspects of Active Control Technology,” in *Active Control Technology: Applications and Lessons Learned*, Turin, Italy: NATO AGARD, May 1994.
- [11] Eli Livne, “Aircraft Active Flutter Suppression: State of the Art and Technology Maturation Needs,” *Journal of Aircraft*, Vol. 55, No. 1, pp. 410–452, Jan. 2018.
- [12] FAA, CFR Title 14 Vol 1 Part 25 - Airworthiness Standards: Transport Category Airplanes, 2020.
- [13] Mark R. Anderson, “Robustness Evaluation of a Flexible Aircraft Control System,” *Journal of Guidance, Control, and Dynamics*, Vol. 16, No. 3, pp. 564–571, May 1993.
- [14] G. Rollwagen, H. Ellgoth, and G. Beuck, “Identification of Dynamic Response, Simulation and Design of a Highly Nonlinear Digital Load Alleviation System for a Modern Transport Aircraft,” in *ICAS 1990*, 1990.
- [15] Steven Ravenscroft, “Actuation systems,” in *Flight Control Systems: Practical Issues in Design and Implementation*, ser. IEE Control Engineering Series 57, Roger W. Pratt, Ed., Stevenage: IEE publishing, 2000.
- [16] Thomas E Noll, John M Brown, Marla E Perez-Davis, Stephen D Ishmael, Geary C Tiffany, and Matthew Gaier, “Investigation of the Helios Prototype Aircraft Mishap,” NASA, 2004.
- [17] P. M. Burris and M. A. Bender, “Aircraft Load Alleviation and Mode Stabilization,” AFFDL-TR-68-161, Air Force Flight Dynamics Laboratory, Wright-Patterson AFB, Ohio, USA, Nov. 1969.
- [18] T. E. Disney, “C-5A Load Alleviation,” in *Active Controls in Aircraft Design*, ser. AGARDograph 234, Neuilly-sur-Seine, France: NATO AGARD, Nov. 1978.

- [19] T. E. Disney, "C-5A Active Load Alleviation System," *Journal of Spacecraft and Rockets*, Vol. 14, No. 2, pp. 81–86, Feb. 1977.
- [20] John H. Wykes and Eldon E. Kordes, "Analytical Design and Flight Tests of a Modal Suppression System on the XB-70 Airplane," in *AGARD Conference Proceedings No. 46: Aeroelastic Effects from a Flight Mechanics Standpoint*, Marseille, France: NATO AGARD, Apr. 1969.
- [21] John H. Wykes and Christopher J. Borland, "B-1 Ride Control," in *Active Controls in Aircraft Design*, ser. AGARDograph 234, P. R. Kurzhals, Ed., Neuilly-sur-Seine, France: NATO AGARD, Nov. 1978.
- [22] Robert F. O'Connell, *Design, Development and Implementation of an Active Control System for Load Alleviation for Commercial Transport Airplane*, ser. AGARD Report 683. Neuilly-sur-Seine: AGARD, 1980.
- [23] W. A. Stauffer, R. L. Foss, and J. G. Lewolt, "Fuel Conservative Subsonic Transport," in *Active Controls in Aircraft Design*, ser. AGARDograph 234, NATO AGARD, Nov. 1978.
- [24] Robert T. Britt, Steven B. Jacobson, and Thomas D. Arthurs, "Aeroseroelastic Analysis of the B-2 Bomber," *Journal of Aircraft*, Vol. 37, No. 5, pp. 745–752, Sep. 2000.
- [25] B. W. Payne, "Designing a Load Alleviation System for a Modern Civil Aircraft," in *ICAS 1986*, 1986.
- [26] K.-U. Hahn and R. Koenig, "ATTAS Flight Test and Simulation Results of the Advanced Gust Management System LARS," in *Guidance, Navigation and Control Conference*, Hilton Head Island, SC, U.S.A.: American Institute of Aeronautics and Astronautics, Aug. 1992.
- [27] David C. Soreide, Rodney K. Bogue, L. J. Ehernberger, Stephen M. Hannon, and David A. Bowdle, "Airborne Coherent Lidar for Advanced In-Flight Measurements (ACLAIM) Flight Testing of the Lidar Sensor," in *American Meteorological Society 9th Conf. on Aviation, Range, and Aerospace Meteorology*, Orlando, Florida, USA, Sep. 2000.
- [28] Guillermo Jenaro Rabadan, Nikolaus P. Schmitt, Thomas Pistner, and Wolfgang Rehm, "Airborne Lidar for Automatic Feedforward Control of Turbulent In-Flight Phenomena," *Journal of Aircraft*, Vol. 47, No. 2, pp. 392–403, Mar. 2010.

- [29] Patrick Vrancken and Jonas Herbst, “Aeronautics Application of Direct-Detection Doppler Wind Lidar: An Adapted Design Based on a Fringe-Imaging Michelson Interferometer as Spectral Analyzer,” *Remote Sensing*, Vol. 14, No. 14, p. 3356, Jul. 2022.
- [30] Yoshiro Hamada, Ryota Kikuchi, and Hamaki Inokuchi, “LIDAR-based Gust Alleviation Control System: Obtained Results and Flight Demonstration Plan,” *IFAC-PapersOnLine*, Vol. 53, No. 2, pp. 14 839–14 844, 2020.
- [31] Adrian Eberle, Bruno Stefes, and Daniel Reckzeh, “Clean Aviation Ultra-Performance Wing (UP-Wing),” in *AIAA SCITECH 2024 Forum*, Orlando, FL: American Institute of Aeronautics and Astronautics, Jan. 2024.
- [32] N. Fezans, J. Schwithal, and D. Fischenberg, “In-Flight Remote Sensing and Identification of Gusts, Turbulence, and Wake Vortices Using a Doppler LIDAR,” *CEAS Aeronautical Journal*, Vol. 8, No. 2, pp. 313–333, Jun. 2017.
- [33] Nicolas Fezans, “An Unusual Structure for a Feedforward Gust Load Alleviation Controller,” in *Advances in Aerospace Guidance, Navigation and Control*, Bogusław Dołęga, Robert Głębocki, Damian Kordos, and Marcin Żugaj, Eds., Cham: Springer International Publishing, 2018, pp. 47–68.
- [34] Jie Zeng, Boris Moulin, Raymond de Callafon, and Martin J. Brenner, “Adaptive Feedforward Control for Gust Load Alleviation,” *Journal of Guidance, Control, and Dynamics*, Vol. 33, No. 3, pp. 862–872, May 2010.
- [35] Yongzhi Wang, Feng Li, and Andrea Da Ronch, “Adaptive Feedforward Control Design for Gust Loads Alleviation of Highly Flexible Aircraft,” in *AIAA Atmospheric Flight Mechanics Conference*, Dallas, TX: American Institute of Aeronautics and Astronautics, Jun. 2015.
- [36] S. Bhattacharyya, D. Cofer, D. Musliner, J. Mueller, and E. Engstrom, “Certification Considerations for Adaptive Systems,” in *2015 International Conference on Unmanned Aircraft Systems (ICUAS)*, Denver, CO, USA: IEEE, Jun. 2015, pp. 270–279.

- [37] Brian D.O. Anderson and Arvin Dehghani, “Challenges of Adaptive Control—Past, Permanent and Future,” *Annual Reviews in Control*, Vol. 32, No. 2, pp. 123–135, Dec. 2008.
- [38] H.-G. Giessler, M. Kopf, P. Varutti, T. Faulwasser, and R. Findeisen, “Model Predictive Control for Gust Load Alleviation,” in *IFAC Proceedings Volumes*, Vol. 45, 2012, pp. 27–32.
- [39] Mateus De Freitas Virgilio Pereira, “Constrained Control for Load Alleviation in Very Flexible Aircraft,” Ph.D. Dissertation, University of Michigan, Ann Arbor, MI, USA, 2022.
- [40] Masayuki Sato, Nobuhiro Yokoyama, and Atsushi Satoh, “Disturbance Suppression via Robust MPC Using Prior Disturbance Data Application to Flight Controller Design for Gust Alleviation,” in *Proceedings of the 48th IEEE Conference on Decision and Control (CDC) Held Jointly with 2009 28th Chinese Control Conference*, Shanghai, China: IEEE, Dec. 2009, pp. 8026–8033.
- [41] Michael Kopf, Eric Bullinger, Hans-Gerd Giessler, Stephan Adden, and Rolf Findeisen, “Model Predictive Control for Aircraft Load Alleviation: Opportunities and Challenges,” in *2018 Annual American Control Conference (ACC)*, Milwaukee, WI: IEEE, Jun. 2018, pp. 2417–2424.
- [42] Yoshiro Hamada, “Aircraft Gust Alleviation Using Discrete-Time Preview Controller with Prior Gust Information,” in *SICE Annual Conference 2013*, Nagoya, Japan: IEEE, Sep. 2013.
- [43] Yuji Paku, Takashi Shimomura, and Yoshiro Hamada, “Gain-Scheduled Preview Control for Aircraft Gust Alleviation,” in *2016 55th Annual Conference of the Society of Instrument and Control Engineers of Japan (SICE)*, Tsukuba, Japan: IEEE, Sep. 2016, pp. 633–638.
- [44] Ahmed Khalil and Nicolas Fezans, “Gust Load Alleviation for Flexible Aircraft Using Discrete-Time Preview Control,” *The Aeronautical Journal*, Vol. 125, No. 1284, pp. 341–364, Feb. 2021.
- [45] Ahmed Khalil and Nicolas Fezans, “A Multi-Channel H_∞ Preview Control Approach to Load Alleviation Design for Flexible Aircraft,” *CEAS Aeronautical Journal*, Vol. 12, No. 2, pp. 401–412, Apr. 2021.

- [46] Nicolas Fezans, Christian Wallace, Daniel Kiehn, Davide Cavaliere, and Patrick Vrancken, “Lidar-Based Gust Load Alleviation - Results Obtained on the Clean Sky 2 Load Alleviation Benchmark,” in *IFASD 2022*, Madrid, Spain: Paper No. 155, Jun. 2022.
- [47] Hugo Fournier, Paolo Massioni, Minh Tu Pham, Laurent Bako, Robin Vernay, and Michele Colombo, “Robust Gust Load Alleviation of Flexible Aircraft Equipped with Lidar,” *Journal of Guidance, Control, and Dynamics*, pp. 1–15, Sep. 2021.
- [48] Kuang-Ying Ting, Mehran Mesbahi, and Eli Livne, “Aeroservoelastic Wind Tunnel Evaluation of Preview H_2 and H_∞ Gust Load Alleviation,” *Journal of Guidance, Control, and Dynamics*, pp. 1–19, Aug. 2023.
- [49] Hugo Fournier, Paolo Massioni, Minh Tu Pham, Laurent Bako, Robin Vernay, and Michele Colombo, “Robust Gust Load Alleviation at Different Flight Points and Mass Configurations,” in *AIAA SCITECH 2022 Forum*, San Diego, CA, USA & Virtual: American Institute of Aeronautics and Astronautics, Jan. 2022.
- [50] Davide Cavaliere, Nicolas Fezans, Daniel Kiehn, David Quero, and Patrick Vrancken, “Gust Load Control Design Challenge Including Lidar Wind Measurements and Based on the Common Research Model,” in *AIAA SCITECH 2022 Forum*, San Diego, CA, USA & Virtual: American Institute of Aeronautics and Astronautics, Jan. 2022.
- [51] Mark B. Tischler, Tom Berger, Christina M. Ivler, Mohammadreza H. Mansur, Kenny K. Cheung, and Jonathan Y. Soong, *Practical Methods for Aircraft and Rotorcraft Flight Control Design: An Optimization-Based Approach*. Reston, Virginia: American Institute of Aeronautics and Astronautics, Inc., Apr. 2017.
- [52] Paul A. Robinson, “Use of Predictive Lidar Measurements in Alleviating Turbulence-Induced Disturbances of Aircraft in Flight,” in *Aerospace/Defense Sensing and Controls*, Robert G. Otto, James Lenz, and Russell Targ, Eds., Orlando, FL, May 1996, pp. 86–97.
- [53] Yoshiro Hamada, Ryota Kikuchi, and Hamaki Inokuchi, “Robust Preview Feedforward Compensation for LIDAR-based Gust Alleviation,” *IFAC-PapersOnLine*, Vol. 56, No. 2, pp. 1180–1185, 2023.

- [54] Daniel Kiehn, Nicolas Fezans, Patrick S Vrancken, and Christoph Deiler, “Parameter Analysis of a Doppler Lidar Sensor for Gust Detection and Load Alleviation,” in *IFASD 2022*, Madrid, Spain, Jun. 2022.
- [55] Ahmed Khalil, “An H_∞ Preview Control Approach to Active Load Alleviation Design for Flexible Aircraft,” Ph.D. Dissertation, TU Braunschweig, Braunschweig, Germany, 2019.
- [56] Davide Cavaliere, Nicolas Fezans, Daniel Kiehn, Julius Schulz, and Ulrich Römer, “Linear Modeling of Doppler Wind Lidar Systems for Gust Load Alleviation Design,” *Journal of Guidance, Control, and Dynamics*, Vol. 47, No. 11, pp. 2351–2368, Nov. 2024.
- [57] Davide Cavaliere, Nicolas Fezans, and Daniel Kiehn, “Method to Account for Estimator-Induced Previewed Information Losses – Application to Synthesis of Lidar-Based Gust Load Alleviation Functions,” in *Proceedings of the 2022 CEAS EuroGNC Conference*, Berlin, Germany: CEAS-GNC-2022-063, May 2022.
- [58] Davide Cavaliere and Nicolas Fezans, “Recasting Discrete 1-Cosine Gust Requirements as Frequency Domain Specifications for Load Alleviation Control Design,” in *IFASD 2022*, Madrid, Spain: Paper No. 143, Jun. 2022.
- [59] Davide Cavaliere and Nicolas Fezans, “Toward Automated Gust Load Alleviation Control Design via Discrete Gust Impulse Filters,” *Journal of Guidance, Control, and Dynamics*, pp. 1–14, Feb. 2024.
- [60] Davide Cavaliere and Nicolas Fezans, “A Practical Approach to Automated Multiobjective Gust Load Alleviation Control Design in a Structured H_2/H_∞ Framework,” in *Proceedings of the 2024 CEAS EuroGNC Conference*, Bristol, UK: CEAS, Jun. 2024.
- [61] Andre Bauknecht, Yannic Beyer, Julius Schultz, Salvatore Asaro, Khalid Khalil, Ulrich Römer, Meiko Steen, Davide Cavaliere, and Nicolas Fezans, “Novel Concepts for Active Load Alleviation,” in *AIAA SCITECH 2022 Forum*, San Diego, CA & Virtual: American Institute of Aeronautics and Astronautics, Jan. 2022.

- [62] Salvatore Asaro, Davide Cavaliere, Nicolas Fezans, and André Bauknecht, “System Identification and Control of a Circulation Control Airfoil for Gust Load Alleviation,” in *AIAA SCITECH 2023 Forum*, National Harbor, MD & Online: American Institute of Aeronautics and Astronautics, Jan. 2023.
- [63] Yannic Beyer, Davide Cavaliere, Kjell Bramsiepe, Khalid Khalil, André Bauknecht, Nicolas Fezans, Meiko Steen, and Peter Hecker, “An Aeroelastic Flight Dynamics Model for Gust Load Alleviation of Energy-Efficient Passenger Airplanes,” in *AIAA AVIATION 2023 Forum*, San Diego, CA and Online: American Institute of Aeronautics and Astronautics, Jun. 2023.
- [64] Jan R. Wright and Jonathan E. Cooper, *Introduction to Aircraft Aeroelasticity and Loads*, 2nd ed. John Wiley & Sons, Incorporated, 2015.
- [65] A. R. Collar, “The First Fifty Years of Aeroelasticity,” *Aerospace*, Feb. 1978.
- [66] Malcolm J. Abzug and E. Eugene Larrabee, *Airplane Stability and Control, Second Edition: A History of the Technologies That Made Aviation Possible*, 2nd ed. Cambridge: Cambridge University Press, 2002.
- [67] Roger W. Pratt, Ed., *Flight Control Systems: Practical Issues in Design and Implementation*, ser. IEE Control Engineering Series 57. Stevenage: IEE publishing, 2000.
- [68] Martin R. Waszak and David K. Schmidt, “Flight Dynamics of Aeroelastic Vehicles,” *Journal of Aircraft*, Vol. 25, No. 6, pp. 563–571, Jun. 1988.
- [69] Eli Livne, “Active Flutter Suppression – Plans for Assessing the Technology’s State of the Art,” in *Joint Advanced Materials and Structures Center of Excellence (JAMS) 2014 Technical Review Meeting*, Seattle, WA, USA: University of Washington, Mar. 2014.
- [70] Christian Reschke, “Integrated Flight Loads Modelling and Analysis for Flexible Transport Aircraft,” Ph.D. Dissertation, Universität Stuttgart, Stuttgart, Germany, 2006.

- [71] Donald McLean, *Automatic Flight Control Systems*, ser. Prentice Hall International Series in Systems and Control Engineering. Englewood Cliffs, N.J: Prentice Hall, 1990.
- [72] NRC, “Aviation Safety and Pilot Control: Understanding and Preventing Unfavorable Pilot-Vehicle Interactions,” NRC, National Research Council, Washington, D.C., 1997.
- [73] Clyde Warsop and William J. Crowther, “Fluidic Flow Control Effectors for Flight Control,” *AIAA Journal*, Vol. 56, No. 10, pp. 3808–3824, Oct. 2018.
- [74] Khalid Khalil and André Bauknecht, “Fluidic Flow Control Devices for Gust Load Alleviation,” *Journal of Aircraft*, pp. 1–14, Feb. 2024.
- [75] Gene F. Franklin, J. David Powell, and Michael L. Workman, *Digital Control of Dynamic Systems*, 3rd ed. ; reprinted in 2006 with corrections. Half Moon Bay, CA: Ellis-Kagle Press, 2006.
- [76] Stanislav Karpuk, Rolf Radespiel, and Ali Elham, “Assessment of Future Airframe and Propulsion Technologies on Sustainability of Next-Generation Mid-Range Aircraft,” *Aerospace*, Vol. 9, No. 5, p. 279, May 2022.
- [77] B. Moulin and M. Karpel, “Gust Loads Alleviation Using Special Control Surfaces,” *Journal of Aircraft*, Vol. 44, No. 1, pp. 17–25, Jan. 2007.
- [78] Peter F. Lester, *Turbulence: A New Perspective for Pilots*. Englewood, CO: Jeppesen Sanderson Training Systems, 1994.
- [79] N. K. Vinnichenko, N. Z. Pinus, S. M. Shmeter, and G. N. Shur, *Turbulence in the Free Atmosphere*. 1968.
- [80] Frederic M. Hoblit, *Gust Loads on Aircraft: Concepts and Applications*. Washington DC: American Institute of Aeronautics and Astronautics, Jan. 1988.
- [81] John C Houbolt, Roy Steiner, and Kermit G Pratt, “Dynamic Response of Airplanes to Atmospheric Turbulence Including Flight Data on Input and Response,” TR R-199, NASA, 1964.
- [82] J. G. Jones, “Documentation of the Linear Statistical Discrete Gust Method,” Federal Aviation Administration, Jul. 2004.

- [83] Patrick Vrancken, “Airborne Remote Detection of Turbulence with Forward-Looking LIDAR,” in *Aviation Turbulence*, Robert Sharman and Todd Lane, Eds., Cham: Springer International Publishing, 2016.
- [84] Nicolas Fezans, Hans-Dieter Joos, and Christoph Deiler, “Gust Load Alleviation for a Long-Range Aircraft with and without Anticipation,” *CEAS Aeronautical Journal*, Vol. 10, No. 4, pp. 1033–1057, Dec. 2019.
- [85] B.J. Rye and R.M. Hardesty, “Discrete Spectral Peak Estimation in Incoherent Backscatter Heterodyne Lidar. I. Spectral Accumulation and the Cramer-Rao Lower Bound,” *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 31, No. 1, pp. 16–27, Jan. 1993.
- [86] Sigurd Skogestad and Ian Postlethwaite, *Multivariable Feedback Control: Analysis and Design*, 2. ed., repr. Chichester: Wiley, 2010.
- [87] Pierre Apkarian and Dominikus Noll, “The H_∞ Control Problem Is Solved,” *AerospaceLab Journal*, Vol. Issue 13, 11 pages, 2017.
- [88] Pierre Apkarian and Dominikus Noll, “Optimization-Based Control Design Techniques and Tools,” in *Encyclopedia of Systems and Control*, John Baillieul and Tariq Samad, Eds., London: Springer London, 2021.
- [89] Pierre Apkarian, Pascal Gahinet, and Craig Buhr, “Multi-Model, Multi-Objective Tuning of Fixed-Structure Controllers,” in *2014 European Control Conference (ECC)*, Strasbourg, France: IEEE, Jun. 2014, pp. 856–861.
- [90] Pierre Apkarian, Minh Ngoc Dao, and Dominikus Noll, “Parametric Robust Structured Control Design,” *IEEE Transactions on Automatic Control*, Vol. 60, No. 7, pp. 1857–1869, Jul. 2015.
- [91] Andreas Wildschek, Rudolf Maier, Klaus-Uwe Hahn, Dirk Leissling, Michael Press, and Adolf Zach, “Flight Test with an Adaptive Feed-Forward Controller for Alleviation of Turbulence Excited Wing Bending Vibrations,” in *AIAA Guidance, Navigation, and Control Conference*, Chicago, Illinois, USA: American Institute of Aeronautics and Astronautics, Aug. 2009.
- [92] François Kubica and Béatrice Madelaine, “Passenger Comfort Improvement by Integrated Control Law Design,” in *RTO AVT Specialists’ Meeting on Structural Aspects of Flexible Aircraft Control*, Ottawa, Canada, Oct. 1999.

- [93] Thomas A. Zeiler, “Matched Filter Concept and Maximum Gust Loads,” *Journal of Aircraft*, Vol. 34, No. 1, pp. 101–108, Jan. 1997.
- [94] Åke Björck and Sven Hammarling, “A Schur Method for the Square Root of a Matrix,” *Linear Algebra and its Applications*, Vol. 52–53, pp. 127–140, 1983.
- [95] Federico Fonte, Sergio Ricci, and Paolo Mantegazza, “Gust Load Alleviation for a Regional Aircraft Through a Static Output Feedback,” *Journal of Aircraft*, Vol. 52, No. 5, pp. 1559–1574, Sep. 2015.
- [96] M. Ripepi and P. Mantegazza, “Improved Matrix Fraction Approximation of Aerodynamic Transfer Matrices,” *AIAA Journal*, Vol. 51, No. 5, pp. 1156–1173, May 2013.
- [97] H.-D. Joos, “A Methodology for Multi-Objective Design Assessment and Flight Control Synthesis Tuning,” *Aerospace Science and Technology*, Vol. 3, No. 3, pp. 161–176, Apr. 1999.
- [98] Neal M. Barr, Dagfinn Gangsaas, and Dwight R Schaeffer, “Wind Models for Flight Simulator Certification of Landing and Approach Guidance and Control Systems,” FAA FAA-RD-74-206, Federal Aviation Administration, Washington, D.C., Dec. 1974.
- [99] L. Zuo and S.A. Nayfeh, “Low Order Continuous-Time Filters for Approximation of the ISO 2631-1 Human Vibration Sensitivity Weightings,” *Journal of Sound and Vibration*, Vol. 265, No. 2, pp. 459–465, Aug. 2003.