



HAL
open science

Concurrent optimization of guided projectile aerodynamics and control in wind tunnel

Valentin Riss

► **To cite this version:**

Valentin Riss. Concurrent optimization of guided projectile aerodynamics and control in wind tunnel. Automatic. Université de Strasbourg, 2022. English. NNT : 2022STRAD023 . tel-04009582

HAL Id: tel-04009582

<https://theses.hal.science/tel-04009582>

Submitted on 11 Sep 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ÉCOLE DOCTORALE Mathématiques, Sciences de l'Information et de l'Ingénieur
Laboratoire des sciences de l'Ingénieur, de l'Informatique et de l'Imagerie (UMR 7357)

THÈSE présentée par :

Valentin RISS

soutenue le : **29 septembre 2022**

pour obtenir le grade de : **Docteur de l'Université de Strasbourg**

Discipline / Spécialité : Automatique

**Optimisation conjointe de l'aérodynamique et de la
commande d'un projectile guidé en soufflerie**

THÈSE dirigée par :

M. LAROCHE Edouard

Professeur, Université de Strasbourg

RAPPORTEURS :

M. ALAZARD Daniel

Professeur, ISAE-SUPAERO

M. MOSCHETTA Jean-Marc

Professeur, ISAE-SUPAERO

AUTRES MEMBRES DU JURY :

Mme. BARA Gabriela Iuliana

Maître de conférences, Télécom Physique Strasbourg

M. BASSET Michel

Professeur, Université de Haute-Alsace

M. ROUSSEL Emmanuel

Chercheur, Institut franco-allemand de recherches de Saint-Louis

M. STRUB Guillaume

Chercheur, Institut franco-allemand de recherches de Saint-Louis

“Essentially, all models are wrong, but some are useful.”

George E. P. Box, *Science and Statistics*, 1976

REMERCIEMENTS

Ces trois années et demi de thèse auront été pour moi une expérience non seulement scientifique mais aussi humaine, qui m'a permis de trouver ma place et de m'épanouir au sein de l'Institut et de la communauté scientifique.

Je tiens d'abord à remercier mes encadrants : mon directeur de thèse, Edouard Laroche, et mon collègue et mentor, Emmanuel Roussel. J'ai toujours pu compter sur vous pour répondre à mes interrogations et exercer un regard critique mais néanmoins bienveillant sur mes travaux, sans lequel ils seraient sûrement bien moins aboutis. Travailler avec vous sur ce sujet a été un plaisir et je suis enthousiaste à l'idée de poursuivre notre collaboration.

Un grand merci aux membres du jury qui m'ont fait l'honneur d'examiner mes travaux. Daniel Alazard, qui m'aura suivi depuis le début et dont les discussions que nous avons eues autour d'un café m'auront beaucoup apporté en me faisant découvrir le concept de co-conception. Jean-Marc Moschetta, dont l'expertise en aérodynamique m'aura permis de mieux cerner la complexité des écoulements autour de la maquette et les limites de mon travail de modélisation. Michel Basset, pour avoir présidé la soutenance et encouragé ces échanges constructifs, ainsi que Iuliana Bara pour avoir apporté ses compétences en automatique et alimenté la discussion. Enfin, Guillaume Strub, mon prédécesseur, qui m'a légué des bases solides aussi bien théoriques que expérimentales, sans lesquelles mes travaux n'auraient jamais pu voir le jour. A ce propos, je souhaite également remercier Philippe Bühler et Gerard Boeglen pour la réalisation des pièces de la nouvelle maquette.

Merci à mes autres collègues de l'ISL : Rémy, Paul, Nadège et le reste du groupe GNC, avec lesquelles j'ai toujours plaisir à partager une pause café et parfois mêmes quelques blagues ou anecdotes plus ou moins bien senties. Gaëtan, bien sûr, qui m'a montré que notre amitié était plus forte que la séparation de nos bureaux ou les parois en plexiglas. Mention spéciale à Sofiane, Enrico et Juliane, mes compagnons d'aventure lors des soirées JDR pizza qui m'ont permis de m'évader à des années-lumière de mes préoccupations professionnelles. Enfin, Flavien et Daniel avec qui j'ai toujours plaisir à pratiquer des activités plus ou moins sportives, alternant montées à VTT et descente de bières.

Pour finir, je dédie ce dernier paragraphe à ceux qui m'ont soutenu depuis toujours ou presque : ma famille, mes amis Vincent, Dawid et Mathias ainsi que mon chien Elyo.

RÉSUMÉ DE THÈSE

Initialement développés au début des années 1970 pour améliorer l'efficacité des systèmes d'artillerie sur des cibles mobiles et protégées telles que les blindés, les projectiles guidés ont évolué pour offrir davantage de fiabilité et de souplesse d'emploi. Ainsi, les munitions pilotées de dernière génération sont désormais capables de neutraliser avec précision des cibles fixes en environnement urbain, minimisant le nombre d'obus nécessaire et la probabilité de dommages collatéraux. Comparé à des missiles sol-sol disposant d'une précision similaire, les projectiles guidés offrent un coût réduit au prix d'une portée plus courte.




Figure 1 – Architecture des projectiles actuels (droite) et futurs (gauche)

L'Institut de Recherches Franco-Allemand de Saint-Louis, acteur majeur de la recherche dans le domaine de l'artillerie, travaille depuis de nombreuses années à l'augmentation de la portée des projectiles guidés [FB96]. Les projets actuels mettent l'accent sur l'intégration de surfaces portantes (c.f. Figure 1) lui permettant aux projectiles de planer sur de plus longues distances par rapport à une trajectoire balistique.

L'exemple de trajectoire présenté en Figure 2 combine portée accrue, flexibilité d'emploi et létalité grâce à un piqué terminal permettant une attaque par le toit. Le guidage du projectile sur cette trajectoire complexe nécessite une manœuvrabilité plus importante comparé au profil de vol des projectiles actuels. Fort heureusement, des études ont démontré que les surfaces portantes peuvent être exploitées pour générer davantage d'effort normal [VBF20] afin d'augmenter le facteur de charge et réduire le rayon de virage. Néanmoins, cela implique que l'incidence du projectile puisse être contrôlée sur une plage de




Figure 2 – Illustration d’une trajectoire avec phase de plané et attaque en piqué

variation suffisamment large : l’asservissement d’attitude doit donc maximiser le domaine de vol en incidence du projectile tout en tenant compte des limitations imposées par ses actionneurs et ses gouvernes. Ces limites se matérialisent par des phénomènes linéaires (limitation de bande passante) et non-linéaires (décrochage aérodynamique des gouvernes) qui nécessitent d’être modélisés de façon à adapter la loi de commande et l’architecture du projectile pour pouvoir suivre la trajectoire prévue.

Face à ce défi, une approche multi-disciplinaire a été retenue, visant à adapter de manière conjointe les propriétés aérodynamiques et la loi de commande du projectile guidé. Elle permet d’optimiser la marge statique pour trouver un compromis entre autorité de commande et stabilité, réduisant les risques de saturation en amplitude et en vitesse d’actionnement. Ce problème d’optimisation procédé-correcteur (également nommé co-design) [Ala+13] s’inscrit au sein d’un procédé de mise aux point des lois de commande basé sur les données ayant fait ses preuves. La méthodologie HIL pour Hardware-In-the-Loop s’articule autour d’un dispositif expérimental comprenant une maquette de projectile guidé pilotée en soufflerie et maintenue par une rotule (c.f. Fig. 3). Ce moyen d’essais est exploité pour fournir des données contribuant à la modélisation du projectile et à la validation de ses lois de commande.

Le premier chapitre de la thèse introduit le lecteur à cette méthodologie en proposant une extension des travaux existants à la commande en attitude du projectile guidé. La seconde partie présente les réponses apportées aux besoins de modélisation pour l’optimisation projectile-correcteur et la caractérisation des performances du projectile en domaine non-linéaire. Enfin, le troisième chapitre traite de la mise en place de

l'environnement de co-design dans un but de maximisation du domaine de vol, ainsi que de la validation expérimentale des performances de plusieurs architectures de projectiles.

Commande d'un projectile en attitude par l'approche HIL




Figure 3 – Le dispositif expérimental ACHILES (Automatic Control Hardware-In-the-Loop Experimental Setup)

La méthodologie de prototypage rapide de lois de commande mise au point par G. Strub [SB16] permet d'identifier un modèle linéaire de la dynamique du projectile, de synthétiser un correcteur adapté puis de le tester directement en soufflerie. Il est ainsi possible de s'affranchir des outils de mécanique des fluides numérique, potentiellement lents et coûteux en temps de calcul. De plus, cette méthode offre une validation anticipée des lois de commande, bien moins coûteuse et risquée que les essais en vol libre.

Dans le cadre de cette thèse, l'approche existante est étendue au contrôle d'attitude en modifiant la maquette afin de rendre fonctionnel le degré de liberté en roulis. Le modèle dynamique et la procédure d'identification paramétrique sont ensuite adaptés, et une nouvelle procédure de réglage des gabarits de synthèse est proposée. Enfin, les performances du correcteur d'attitude sont évaluées lors d'essais de suivi de consigne et de rejet de perturbation. La Figure 4 montre que le correcteur conçu selon l'hypothèse de découplage des axes de roulis, tangage et lacet permet de contrôler l'attitude du projectile conformément aux simulations et atteint les performances attendues. Par conséquent, et au vu du scénario de trajectoire évoqué en Figure 2, le reste de l'étude est restreinte à la dynamique en tangage du projectile pour laquelle les objectifs de performance sont les plus ambitieux.




Figure 4 – Essai de suivi de consigne en attitude (rouge : mesures, bleu : modèle, pointillés : consigne)

Modélisation aérodynamique du comportement en tangage du projectile

Dans un second temps, on cherche à étendre le domaine de vol en incidence du projectile en optimisant de manière conjointe l'ensemble projectile-correcteur. Pour cela, il est nécessaire de capturer l'influence des caractéristiques géométriques du projectile sur sa dynamique en tangage, mais aussi de quantifier l'effet des limitations aérodynamique liées à ses surfaces portantes. C'est pourquoi deux modèles distincts ont été développés dans des buts différents: Ainsi, le modèle de synthèse est de type Linéaire à Paramètres Variants (LPV) et capture le décrochage des canards sous l'angle d'une incertitude paramétrique appliquée au coefficient d'autorité des gouvernes. Le modèle de validation est lui non-linéaire, de façon à capturer plus fidèlement l'effet du décrochage des gouvernes sur la réponse en tangage et valider les performances du couple projectile-correcteur.

Trois architectures ont été envisagées pour le modèle de validation. Comme son nom l'indique, la structure "data-based" fait exclusivement appel aux résultats expérimentaux extrapolés grâce à des représentation polynomiales. Au contraire, la structure "component-based" est fondée sur des formules analytiques [Nie88] issues de la théorie des écoulements potentiels appliquée aux surfaces portantes et à leurs interactions avec le fuselage. En-

fin, l'architecture "semi-local" utilise certains résultats théoriques pour capturer les non-linéarités aérodynamiques en reconstruisant la polaire de portance des canards à partir de données expérimentales. Tous les modèles implémentent dépendance continue par rapport au paramètre géométrique obtenue par l'interpolation polynomiale des coefficients aérodynamiques estimés en soufflerie. Dans le cas des configurations peu ou pas stables en boucle ouverte, une méthode d'identification en boucle fermée a dû être mise au point pour estimer les valeurs de ces coefficients.

Après comparaison avec des données expérimentales couvrant les configurations stables et instables, la solution "component based" s'est avérée limitée par l'imprécision de certains modèles analytiques. La solution "data based" offre les prédictions les plus fidèles du comportement des projectiles stables mais s'avère inadaptée aux configurations peu ou pas stables. C'est donc la structure "semi local" qui a été retenue pour évaluer l'impact du décrochage des canards sur les performances du correcteur et le domaine de vol du projectile.

Co-design du couple projectile-correcteur et validation expérimentale

La marge statique du projectile est un levier crucial pour trouver le meilleur compromis entre la rapidité du système en régime linéaire et le maintien de l'autorité de commande en régime non-linéaire. La position longitudinale de l'empennage, notée x_F , a été choisie pour adapter la stabilité statique de l'engin. L'étude de la dynamique du projectile en boucle ouverte confirme que le gain statique du procédé ainsi que la pulsation et l'amortissement des pôles dominants sont fortement influencés par ce paramètre.

Un environnement de co-design original a été mis en place pour pouvoir prendre en compte les effets aérodynamiques non-linéaires dès la phase de conception du projectile et de ses lois de commande. Un schéma de synthèse \mathcal{H}_∞ multi-objectifs permet d'implémenter le problème d'optimisation projectile-correcteur en spécifiant la fonction coût et les contraintes. L'analyse des résultats dans le domaine fréquentiel est complétée par une simulation non-linéaire de la réponse du projectile, s'appuyant sur le modèle de validation précédemment développé. Cette dernière sert à évaluer le domaine de vol du projectile asservi en tenant compte du phénomène de décrochage des gouvernes, ce qui permet de déceler ses éventuelles conséquences néfastes (saturation en braquage, emballement du correcteur) et ainsi d'adapter les gabarits de synthèse pour les atténuer. Un critère de

performance nommé "incidence max. de commande" a été défini pour comparer les performances de suivi de consigne des différentes configurations en régime non-linéaire.

De multiples architectures de projectiles à manœuvrabilité améliorée ont ainsi pu être développées. Tout d'abord, deux structures de boucles de commande ont été réglées au sein du problème d'optimisation conjointe, menant à des géométries distinctes. La meilleure des deux structures a été comparée au réglage du même correcteur pour une configuration aérodynamique correspondants aux standards de stabilité en boucle ouverte des projectiles non-guidés. Les résultats de simulation présentés en Fig. 5 montrent une nette amélioration de la réponse du projectile correspondant à une augmentation d'un facteur trois du domaine de vol en incidence de ce dernier. Enfin, la problématique de la commande du projectile en limite de stabilité a également été explorée, apportant des résultats complémentaires.




Figure 5 – Simulation du suivi de consigne du projectile optimisé (en rouge) par rapport à la configuration de référence (en bleu)

Dans un second temps, les deux architectures de co-design ainsi que la loi de commande pour projectiles à stabilité relaxée ont été implémentées sur la maquette en soufflerie afin d'être soumises à des tests de suivi de consigne et de rejet de perturbation. Ces essais ont permis d'évaluer les performances de ces projectiles asservis et de les comparer avec les prédictions fournies par le modèle non-linéaire (c.f. Fig. 6). Sur la plage de x_F correspondant aux configurations les moins stables, la variation importante de la dynamique du projectile combinée au manque d'information expérimentale disponible dégrade la précision du modèle. Ainsi, une procédure d'ajustement de certaines constantes du modèle semi-local est mise en place à partir de la réponse en boucle fermée de la configuration la moins stable. Cette nouvelle variante du modèle vient compléter le réglage original

pour prédire la réponse en régime non-linéaire du projectile sur l'ensemble de la plage de variation du paramètre géométrique. Le modèle ajusté permet de prédire le temps de réponse à 10% du projectile en limite de stabilité avec une erreur relative moyenne inférieure à 17% tandis que le modèle original estime le temps de réponse du projectile optimisé à 22% près. Ces deux modèles fournissent un encadrement de l'incidence max. de commande du projectile obtenu par co-design. Celle-ci est comprise entre $4,6^\circ$ et $7,3^\circ$ contre $2,5^\circ$ pour la configuration standard, prouvant la pertinence de cette méthodologie de conception dans le cadre de l'amélioration de la manœuvrabilité des projectiles munis de surfaces portantes.




Figure 6 – Essai de suivi de consigne d'un projectile obtenu par co-design. “*Lin mdl*” désigne le modèle de synthèse, “*NL mdl*” le modèle de validation original et “*NL retuned*” sa variante ajustée

Conclusion et perspectives

La problématique d'amélioration de la manœuvrabilité des projectiles guidés munis de surfaces portantes a été abordée par la combinaison d'un procédé d'optimisation multidisciplinaire avec une méthodologie de modélisation et de synthèse basée sur un dispositif expérimental. L'approche HIL a permis de démontrer le contrôle en attitude de ce type de projectiles innovants. La maquette a également pu fournir des données permettant la modélisation non-linéaire de la dynamique en tangage et la prise en compte du décrochage des gouvernes. Le modèle de validation a été intégré dans l'environnement de synthèse afin

de mettre au point rapidement un ensemble projectile-correcteur robuste à ce phénomène. Les résultats expérimentaux confirment les prédictions des performances en suivi de consigne et valident l'extension du domaine de vol en incidence observée en simulation par rapport à la méthodologie de conception standard. Dans le cadre de la conception de projectiles hautes performances, ces recherches fournissent une méthodologie permettant d'augmenter drastiquement la manœuvrabilité d'une configuration à portée optimisée en modifiant simultanément sa marge statique et ses lois de commande sans avoir à altérer le dimensionnement des gouvernes et des actionneurs.

Les perspectives d'amélioration et d'approfondissement de l'étude incluent une nouvelle synthèse du co-design et du projectile en limite de stabilité en utilisant le modèle non-linéaire ajusté. Ceci dans le but de réduire la sensibilité du correcteur au choix du modèle et diminuer l'incertitude sur les prédictions de l'incidence max. de commande. D'un point de vue modélisation, les structures LPV et non-linéaires pourraient être améliorées en identifiant les coefficients aérodynamiques au voisinage de plusieurs points du domaine d'incidence, sous réserve de trouver un signal d'excitation adapté. Le procédé de reconstruction de la polaire des canards pourrait être remplacé par des mesures d'effort en soufflerie qui permettraient également de vérifier la validité d'une partie des coefficients aérodynamiques. D'un point de vue commande, une approche LPV prenant en compte l'incidence de trim permettrait peut-être d'améliorer encore les performances du projectile guidé. L'intégration d'algorithmes anti-windup apporterait des propriétés de robustesse supplémentaire aux limites du domaine de vol.

CONFERENCES AND PUBLICATIONS

Articles in International Peer-Reviewed Journals

V RISS, E ROUSSEL and E LAROCHE. "Improvement of Guided Projectile Maneuverability using Airframe-Controller Optimization". In preparation.

Papers in International Peer-Reviewed Conferences with Proceedings

V RISS, E ROUSSEL and E LAROCHE. "Concurrent Airframe-Controller Optimization of a Guided Projectile fitted with Lifting Surfaces". *AIAA SCITECH 2022 Forum*. San Diego, CA, 2022.

V RISS, E ROUSSEL and E LAROCHE. "Attitude Control of a Fin-Stabilized Projectile on a Three-Axis Gimbal in Wind Tunnel". *28th Mediterranean Conference on Control and Automation*. 2020.

Papers in National Conferences with Proceedings

V RISS, E ROUSSEL and E LAROCHE. "ACHILES2 : une maquette de projectile asservi en soufflerie pour la modélisation et la commande". *6^{èmes} journées des Démonstrateurs en Automatique*. Angers, France, 2022.

NOMENCLATURE

Notation Rules

x	Scalar
\boldsymbol{x}	Vector
\mathbf{X}	Matrix
\boldsymbol{x}^T (or \mathbf{X}^T)	Transpose of \boldsymbol{x} (or \mathbf{X})

Frames and Coordinate Systems

B	Body frame
L	Local frame
W	Wind frame
$]^B$	Body coordinate system
$]^L$	Local coordinate system
$]^W$	Wind coordinate system

Geometry and Mass Properties

D	Projectile caliber (m)
S	Projectile cross section area (m ²)
S_C	Canard surface area (m ²)
x_F	Distance from fins leading edge to projectile reference datum (m)
I_{XX}, I_{YY}, I_{ZZ}	Projectile principal roll, pitch and yaw moments of inertia (kg.m ²)
I_{ZZg}	Gimbal principal yaw moment of inertia (kg.m ²)
I_{ZZt}	Total yaw moment of inertia of the projectile in its gimbal (kg.m ²)

Projectile Aerodynamics

α	Angle of attack (deg)
β	Angle of sideslip (deg)
L, M, N	Moments of the aerodynamic forces in the Body axes (N.m)
C_l, C_m, C_n	Roll, pitch and yaw aerodynamic coefficients
ϕ, θ, ψ	Euler angles (roll, pitch and yaw) (deg)
p, q, r	Body angular rates (deg/s)
$\delta_l, \delta_m, \delta_n$	Virtual roll, pitch and yaw control surfaces deflections (deg)
$\delta_1, \delta_2, \delta_3, \delta_4$	Canard 1, 2, 3 and 4 deflections (deg)
\bar{q}	Dynamic pressure (N/m ²)
\bar{v}	Freestream velocity (m/s)

Abbreviations and Acronyms

ACO	Airframe-Controller Optimization
ACHILES	Automatic Control Hardware-In-the-Loop Experimental Setup
AoA	Angle of Attack
AR	Aspect Ratio
CAD	Computer Aided Design
CAS	Control and Actuation System
CCF	Course Correction Fuze
CFD	Computational Fluid Dynamics
CG	Center of Gravity
CNC	Computer Numerical Control
CP	Center of Pressure
DoF	Degrees-of-Freedom
EoM	Equation of Motion
fs	full scale
GPS	Global Positioning System
HIL	Hardware-In-the-Loop
IMU	Inertial Measurement Unit
INS	Inertial Navigation System

I\O	Input\Output
LMI	Linear Matrix Inequality
LPV	Linear Parameter-Varying
LQR	Linear Quadratic Regulator
LRGP	Long Range Guided Projectile
LTI	Linear Time Invariant
LUT	Look-Up Table
MoI	Moment of Inertia
MRL	Multiple Rocket Launcher
NRMSE	Normalized Root-Mean-Square Error
PCO	Plant-Controller Optimization
RHP	Right Hand Plane
RMS	Root-Mean-Square
RSS	Relaxed Static Stability
SISO	Single Input Single Output
UAV	Unmanned Air Vehicle
VFT	Virtual Flight Testing

TABLE OF CONTENTS

Introduction	25
The Evolution of Guided Projectiles	26
Modeling and Control of Aerospace Vehicles	31
Experimental Setups for Control Design	34
Motivations and Contributions	37
1 HIL Approach to Attitude Control of a Guided Projectile	39
1.1 Experimental Setup	40
1.1.1 Wind tunnel and gimbal assembly	40
1.1.2 Projectile airframe	43
1.1.3 Autopilot hardware and software	47
1.2 Linear Model of the Projectile Angular Motion	48
1.2.1 Frames of reference	49
1.2.2 Control surfaces allocation	50
1.2.3 Non-linear equations of motion	51
1.2.4 Control-oriented model	53
1.3 Identification of the Angular Dynamics	57
1.3.1 Model postulation	58
1.3.2 Experiment design and preliminary work	59
1.3.3 Estimation of the aerodynamic derivatives	61
1.4 Attitude Controller Design and Validation	64
1.4.1 Controller structure and synthesis framework	64
1.4.2 Experimental verification of closed-loop performance	68
1.5 Conclusion	71
2 Parameter-Varying Models of the Projectile Dynamics	73
2.1 Introduction to Modeling	74
2.1.1 Model purposes and requirements	74
2.1.2 Modeling strategy	75

TABLE OF CONTENTS

2.2	Experimental Data Collection	76
2.2.1	Trim map recording	76
2.2.2	System identification of the variable-stability projectile	77
2.3	Linear-Parameter-Variable Model of the Projectile Pitch Dynamics	83
2.3.1	State-space structure	83
2.3.2	Surrogate models of the aerodynamic and inertia properties	84
2.3.3	Minimum-size linear fractional representation	87
2.4	Introduction to Non-Linear Aerodynamics Modeling	89
2.4.1	Rationale and means	89
2.4.2	Fundamentals of the projectile pitch dynamics	90
2.5	Theoretical Groundwork for Analytical Aerodynamics Predictions	90
2.5.1	Simplification of the fluid mechanics equations	91
2.5.2	Potential flow theory	92
2.5.3	Slender body theory	97
2.5.4	Lifting line theory	98
2.5.5	Wing-body interference factors	99
2.5.6	Apparent-mass formula for pitch damping	101
2.6	Estimation of Aerodynamic Forces on Projectile Components	102
2.6.1	Equivalent angle of attack method for wing-body interactions	103
2.6.2	Data-based non-linear canard lift model	105
2.6.3	Analytical formula for fins lift	109
2.6.4	Semi-empirical body lift and center of pressure	110
2.7	Parameter-Dependant Models of the Aerodynamic Moment	111
2.7.1	Panorama of the non-linear modeling approaches	111
2.7.2	Component-based model	112
2.7.3	Semi-local model	113
2.7.4	Data-based model	116
2.8	Model Assessment and Selection	120
2.8.1	Open-loop fidelity assessment	120
2.8.2	Extension to statically-unstable projectiles	125
2.8.3	Model comparison and selection	128
2.9	Conclusion	129

3	Integrated Design of Projectile Aerodynamics and Control	131
3.1	Motivations and Use Case	131
3.1.1	Example of a flight scenario	131
3.1.2	Properties and limitations of gliding projectiles	132
3.1.3	Purpose of plant-controller optimization	133
3.2	State-of-the-Art of Plant-Controller Optimization	134
3.2.1	Concept	134
3.2.2	Optimization strategies	135
3.2.3	Optimization problem setup	136
3.2.4	Control design	138
3.3	Airframe-Controller Design Methodology	140
3.3.1	Selection of the tunable geometrical parameter	141
3.3.2	Models of the adjustable-stability projectile	144
3.3.3	Multi-objective synthesis and optimization	146
3.3.4	Non-linear simulation of closed-loop performance	148
3.4	Synthesis and Simulation of Agile Projectiles	150
3.4.1	Co-design with state-space controller	152
3.4.2	Co-design with dual-loop controller	160
3.4.3	Simulation benchmark against stable configurations	166
3.4.4	Control of projectiles with relaxed static stability	170
3.5	Experimental Testing and Validation of Design Candidates and Methodology	176
3.5.1	Wind-tunnel testing of the state-space co-design	177
3.5.2	Wind-tunnel testing of the RSS projectile	181
3.5.3	Wind-tunnel testing of the dual-loop co-design	187
3.5.4	Conclusions on model fidelity and agile projectile design	192
3.6	Conclusion	195
	Conclusion and Perspectives	197
	Bibliography	203
	Appendix	214
A.1	Equations of the Semi-Local Aerodynamic Model	214

INTRODUCTION

Since the beginning of ages, projectiles have been employed as weapons, complementing melee weapons by providing some stand-off range between the attacker and the target. The concept of artillery as a means to propel large projectiles over long distances, can be traced back to the ballista imagined by the Greeks and later employed by the Romans [Ale46]. Since then black powder and its derivatives have become the preferred source of energy, siege engines giving way to smoothbore guns propelling cannon balls such as the bombard pictured in Figure 7.




Figure 7 – Early developments of artillery: from ballista (left) to bombard (right)

The advances in metallurgy sparked by the Industrial Revolution [Bas92] prompted the introduction of rifled barrels allowing spin-stabilization of shells filled with explosive charges. In the 19th century, progress in interior ballistics and energetic materials led to a significant increase in initial velocities which, along with the introduction of streamlined ogive-cylinder projectiles (c.f fig. 8), cast artillery into the role of a very potent weapon during the subsequent world wars.

Towards the end of World War Two, Axis Powers designed the first guided missile systems by integrating steering mechanisms into flying weapons. The V2 named after the german *Vergeltungswaffe* (retaliation weapon) was the first guided ballistic missile, outmatching any contemporary gun in range and payload [Bec95]. The allied side saw the introduction of another type of artillery systems, with unguided Multiple Rocket Launchers (MRLs) being massively used on the Eastern front. Since then, guns, missiles




Figure 8 – Modern artillery: the Model 1897 75 mm gun and its ammunition

and rocket launchers have coexisted on the battlefield, each type of system bringing its own advantages and downsides.

Compared to mortars and howitzers, MRLs are known for the psychological effect of their devastating salvos [Pre16] and feature a simpler mechanical design at a cost of a wider dispersion. Surface-to-surface missiles, on the other hand, are able to offer unmatched range, payload and accuracy but their cost remain several orders of magnitude greater than classical artillery shells.




Figure 9 – Rocket-powered artillery: the V2 missile (left) and the BM13 MRL (right)

The Evolution of Guided Projectiles

It is not until the early 1970s that guidance, navigation and control systems have made their way into the confines of large-caliber ammunition. Developed to increase the lethality of artillery fires against mobile and armored targets [MS77], the 155 mm M712 Copperhead was the first mass-produced guided artillery projectile [Wat07] which entered service in 1982. However, the sensor technology employed by this munition to find its

target made its tactical employment cumbersome and weather-dependent.




Figure 10 – Layout of the M712 Copperhead 155 mm guided shell

Since then, the focus has shifted from moving target engagement to precision strikes with low collateral damage. The M982 Excalibur shell came more than twenty years later, bringing a new GPS/INS-based guidance scheme which allows all-weather, fire-and-forget engagement of fixed targets [MA12]. This new system enables very high accuracy independent of the firing range as well as trajectory shaping for employment in urban areas. Thanks to these new capabilities, the number of shots required to obtain the desired effect on target is greatly reduced, lightening the logistic burden, improving the survivability of the firing unit and adding to the effect of surprise [Cos95].




Figure 11 – The Excalibur projectile (left) and its range-independent dispersion (right)

As for other recent developments, two distinct trends have emerged in order to answer the contradictory requirements of cost reduction and performance improvement. The former aspect is covered by the advent of Course Correction Fuzes (CCF) which consists in add-on guidance kits that can be retrofitted to an existing stockpile of unguided ammunition in order to provide precision strike capabilities at the lowest cost. Besides, since this kits can be fitted in the field, it allows the battery commander to choose between

guided or unguided projectiles according to the mission requirements. Those designs are generally very compact and feature few moving parts to reduce technological risk and manufacturing costs, trading off some accuracy compared to more complex designs. A prime example is the Northrop-Grumman M1156 Precision Guidance Kit (PGK): it features a GPS-guided and roll-decoupled fuze actuated by a single electromagnetic brake, as depicted in fig. 12.



Figure 12 – The PGK course correction fuze

The latter trend consists in extending the range of guided ammunition in order to stretch the sphere of action of land- and surface-based artillery systems. This paradigm brings numerous advantages [Cos95]: the time spent moving into position is reduced, increasing the availability of fire support. Moreover, the firing unit is less vulnerable to counter-battery fire and has more flexibility for engaging multiple targets at various ranges.

The range increase can be achieved through various ways: the Vulcano family of ammunition claims to hit targets as far as 70 km away by using high muzzle velocities coupled with very slender sub-caliber projectiles. The Norwegian manufacturer Nammo plans to reach ranges up to 150 km with a ramjet-powered 155 mm guided shell (fig. 13). Other technical solutions include rocket-assisted projectiles and drag reduction through base bleed.



Figure 13 – The 155 mm Vulcano (top) and Nammo extreme-range (bottom) projectiles




Figure 14 – Range extension by gliding [Dec+18]

ISL, the French-German Research Institute of Saint-Louis, has a long experience in investigating new concepts for large caliber ammunition. Range extension of guided projectiles has been a topic of interest for more than 25 years, resulting in a variety of designs [TW17]. Unlike the two previously-mentioned designs, ISL's focus has been set on integrating lifting surfaces to extend the reach of artillery shells. In this case, the plan is to

perform a glide phase at a shallow flight path angle and fly at maximum lift-to-drag ratio (c.f. fig. 14).

The first concepts such as the M2PA (*Munition à Portée et Précision Accrue*) retained the spin-stabilization principle used for unguided shells, and used the angle of attack (AoA) induced by the spinning motion to generate lift and steer the projectile [FB96]. The reciprocating control canards shifted the center of pressure forwards and aft of the center of gravity to drive a complex oscillatory motion of the projectile (fig. 15) that would allow the main wings to provide additional lift and extend the gliding range.




Figure 15 – Principle of the M2PA steering [FB96]: the plot shows the orientation of the nose with respect to the body

More recent projects switched to fin-stabilized, non-spinning architectures in order to further improve aerodynamic performance by maintaining a constant AoA for best lift-to-drag ratio [Olm16]. Tri-surfaces configurations have been studied, with the rear fins providing stability during the ballistic ascent while the main wings are deployed for the glide phase. Four movable canards at the front provide steering during both glide and terminal phases.




Figure 16 – Tri-surface gliding projectile concept [Olm16]

Modeling and Control of Aerospace Vehicles

Aerodynamics prediction

The accurate prediction of the aerodynamic forces applied on a projectile at a given flight condition is critical to the maximization of its performance indexes such as the lift-to-drag ratio. The point of application of the resulting lift force and its derivatives w.r.t the projectile attitude and angular rates conditions the stability of the flying body and the damping of its transient motions. Also, the forces and moments generated by control surfaces deflections determine the equilibrium manifold of projectile attitude angles, contributing to the manoeuver performance of the airframe.

This is why a number of methods have been developed to predict either completely or partially the aerodynamic characteristics of a given design with the adequate level of fidelity. Starting from the simpler tools, potential flow theory gives analytical solutions for the forces and moments around simple geometries [Nie88]. These results can be combined to predict the aerodynamics of more complex body shapes, using a “component buildup” approach. Such tools were extensively developed around the 1950’s [AS53] [Bry53], at a

time when computing power was not sufficient to solve more complex forms of the fluid dynamics equations. However, they are still commonly used at the pre-design stage because of their quick run time and simplicity of use, and form the basis of successful computer programs such as the US Air Force Missile DATCOM [VJ84], the ONERA MISSILE code [Den98] and the PRODAS projectile design software [WH03].

With the advent of computational design and finite elements modeling, a family of more generic tools has emerged in the form of Computational Fluid Dynamics (CFD) software. Such a concept revolves around meshing the fluid domain around the geometry of interest and solving a set of fluid mechanics equations such as Euler or Reynolds-Averaged Navier-Stokes at the local scale [AW95]. This technique provides the three-dimensional flow structure and pressure distribution around any type of geometry (c.f. fig 17) and may resolve viscous phenomena such as vortex shedding or skin friction which are outside the scope of the potential flow theory. However, numerical effects due to the discretization of the problem must be kept in check and the computational burden is much higher than for analytical methods [Dec+18]. Also, since direct numerical simulation of the full unsteady, tri-dimensional Navier-Stokes equations is still unfeasible for most engineering applications [Aga99], one may rely on turbulence modeling to solve the Reynolds-averaged equation system, introducing a degree of uncertainty into the prediction of transition from laminar to turbulent flow regime. Thus, prediction of viscous phenomena such as boundary layer separation may be affected by the choice of the turbulence model.




Figure 17 – CFD visualization of the pressure field around a projectile [Dec+18]

An intermediate fidelity solution exists in the form of panel methods. This approach is an offspring of potential flow theory where the pressure field around the flying body is modeled by a distribution of singularities on its surface. Such formulation allows to

consider more complex geometries than the analytical design tools but remains less computationally intensive than CFD. However, its domain of application is usually limited to moderate angles of attack and viscous effects may not be included [ME81].

Control law design

Since most low spin rate projectiles such as the latest ISL long range concepts share common dynamics with missiles, the body of literature available on the topic can be leveraged to find proven controller structures and tuning methodologies [MR05; ÇLi11]. Over the years, the design practices have evolved from successive SISO loop closures employed on the early guided projectiles [MS77] to non-linear controllers based on adaptive methods such as model predictive control [OC08] or dynamic inversion [Tip+20]. However, robust control techniques like structured \mathcal{H}_∞ synthesis have proven very popular for such class of systems due to their ability to provide stability, performance and robustness guarantees and support low-order controller structures which are well suited to embedded systems [AGB95; TP21]. Modern software tools such as the MATLAB routine *systeme* [GA11] provide a convenient framework for setting up the synthesis objectives and solving the associated non-smooth optimization problem.

Plant-controller optimization

New design methods can be applied to guided projectiles in order to improve their maneuvering performance, taking advantage of the advances in aerodynamics modeling and control theory. Plant-controller optimization, also known as co-design, consists in tuning both controller gains and physical parameters influencing the open-loop plant dynamics. It has been applied to a variety of domains such as chemistry [LF94], robotics [RWH06] and powertrains [RP99][Sil+16]. In aerospace, the method has mainly been used in preliminary design for different purposes such as mass reduction of large flexible structures [HLD85], reduction of avionics requirements for satellite attitude control [Ala+13], optimization of an airborne-wind-energy system [NDV17] and aircraft control surfaces sizing under handling qualities constraints [NK96] [Den+17].

Experimental Setups for Control Design

In the literature, most of the experimental setups featuring a controlled flight vehicle mock-up in wind tunnel are employed for validation of closed-loop performance. The concept of Virtual Flight Testing (VFT) described by fig. 18 was first theorized by Ratliff in 1995 [RM95] and later demonstrated by Lawrence in 2002 [LM02] using a missile mock-up in an open-section wind tunnel (fig. 19). Such setup not only records the rotational motion of the test article but also measures the air loads applied on the airframe. This data can be fed to a model of the missile translational dynamics and simulate the free-flight trajectory of the vehicle. It may even be possible to include the seeker and guidance system dynamics by re-projecting a synthetic image of the scene on the focal plane of the seeker and updating it according to the measured attitude and the computed position with respect to the target.




Figure 18 – The Virtual Flight Testing concept [RM95]

Successful controller synthesis requires a model which captures the projectile dynamics over a wide range of frequencies with adequate accuracy. Experimental setups can provide valuable data in order to create a model structure or identify the coefficients of an existing formulation. The dynamic testing rig proposed by Gatto [GL06] and described in fig. 20 was designed to estimate the static and dynamic aerodynamic derivatives of a combat aircraft mock-up. The mounting hardware allowed measurement of the rotational motion around the 3 degrees of freedom while a six-component balance recorded the forces and moments applied on the mock-up. Steady deflections of the control surfaces were used to determine the static derivatives while sinusoidal inputs were injected to estimate the dynamic coefficients.



Figure 19 – The US Air Force VFT demonstrator [LM02]

Experimental setups which have been used for both modeling and validation purposes are much less common. One of the best examples may be the missile rig proposed by Fresconi [Fre+18] which features a canard-controlled airframe on a sting-mounted gimbal (c.f. fig. 21). Full attitude motion was obtained from embedded magnetic position sensors while a 5-axis force balance was integrated inside the sting. This setup was used to estimate the unknown parameters of the dynamic model of the mock-up and check the performance of the attitude controller. However, no insight on the autopilot tuning process is provided in the above-mentioned report.




Figure 20 – The M2370 mock-up on its dynamic testing rig [GL06]




Figure 21 – The ARL canard-controlled missile mock-up [Fre+18]

Motivations and Contributions

The design of guided projectiles dynamics and control poses numerous challenges for the sequential (aerodynamics then control) development methodology which is commonly used. The configuration aerodynamics have a strong influence on the magnitude and bandwidth of the control action required to adjust the projectile attitude and follow the guidance commands. Excessive control deflections and/or flight at high angles of attack may result in a loss of control effectiveness due to the stall of the control surfaces. Controllers operating at those off-design conditions may perform poorly and fail to track the attitude reference with sufficient accuracy to achieve mission success. The physical phenomenon of flow separation, which is the root cause of stall, remains a modeling challenge, even for state-of-the-art computational fluid dynamics simulations. Thus, the control system deficiencies may only be discovered at the flight test stage which occurs very late into the development process. In that case, the aerodynamic configuration of the projectile will have to be altered, adding considerable cost and time overruns.

The aim of this work is to propose a more effective control design methodology by accounting for the stall condition at an early stage with a control-oriented, non-linear model of the projectile aerodynamics. A plant-controller optimization scheme is implemented to increase the maximal angle of attack at which the airframe attitude can be controlled, opening the way for more agile projectiles. Additionally, control of statically unstable configurations is investigated to further enhance the maneuvering performance of guided artillery shells.

The Hardware-In-the-Loop (HIL) approach proposed by Strub [Str16] intends to speed up the control law design process by directly identifying the projectile model from wind tunnel data and providing experimental closed-loop verification and validation in the preliminary design phase. In the first chapter of this work, the existing framework is extended to full attitude control of the projectile mock-up. First, the experimental setup is redesigned to allow roll motion while the linear model is reformulated for 3 degrees-of-freedom. Then, the system identification procedure is renewed and a new experimental campaign is conducted in order to collect the required data. The fit of the model outputs and the quality of its estimated parameters are assessed before inclusion in the synthesis process. As for controller tuning, the skid-to-turn autopilot is updated with an additional roll channel and a new procedure is introduced to determine the frequency templates of the weighting filters. Finally, the performance of the attitude controller is assessed through

reference tracking and disturbance rejection experiments.

The remaining two chapters of the thesis cover the issue of operation at larger angles of attack and the mitigation of control surfaces stall. This challenge is tackled with a multi-disciplinary approach, tuning the airframe static stability along with the controller gains in order to maximize performance while accounting for the actuators and control surfaces limitations. The second chapter focuses on the development of aerodynamic models required for Airframe-Controller Optimization (ACO) and validation of the configuration candidates in non-linear domain. The focus is set on the pitch axis which is subject to the most stringent maneuverability requirements: the experimental setup is leveraged to record the equilibrium manifold of the projectile and identify linear dynamic models for a range of stable and unstable configurations: a Linear Parameter-Varying (LPV) formulation is derived from this set of dynamic models and tailored for synthesis and optimization. Additionally, three non-linear model candidates are proposed, going from an implementation of the analytical formulas from the potential flow theory to a purely data-based model structure. The models outputs are benchmarked against open-loop measurements in order to select the best candidate for assessing the closed-loop response of a family of projectile configurations.

The final chapter presents the ACO methodology and its application to three different cases, ranging from a fairly stable projectile to a more innovative configuration featuring relaxed static stability. Simulation results and experimental validation are provided for each of the designs, leading to some concluding thoughts on model fidelity and agile projectile configurations. A novel design framework has been set up to account for the effects of canard stall early at the design stage, speeding up the development process and reducing the risk of uncovering performances deficiencies at the validation stage. This is achieved by augmenting the multi-objective \mathcal{H}_∞ synthesis scheme with a non-linear simulator of the projectile response which is able to quantify the impact of canard stall on closed-loop performance.

HIL APPROACH TO ATTITUDE CONTROL OF A GUIDED PROJECTILE

The Hardware-In-the-Loop methodology aims at overcoming the challenges associated with control laws design for guided projectiles, namely the determination of an adequate dynamic model and the lack of experimental validation in the preliminary design phase. It consists in an actuated and instrumented projectile mock-up which is placed in the test section of a subsonic wind tunnel, held by a 3 degrees-of-freedom gimbal. This installation, named ACHILES for Automatic Control Hardware-In-the-Loop Experimental Setup, operates in conjunction with a modeling and controller synthesis framework implemented on MATLAB and Simulink [Str+14]. Custom scripts allow to perform parameter identification of grey-box model structures, as well as multi-objective \mathcal{H}_∞ synthesis of the projectile structured controller. Previous work carried out by Strub includes the design of a skid-to-turn pitch and yaw controller [SB16] as well as its extension to variable-air-speed operation thanks to a gain-scheduling strategy [Str+18]. However, roll control has not been attempted due to excessive friction in the legacy experimental setup, and trimming at low airspeed and high angle of attack is hampered by canard stall [Str16].

This chapter intends to introduce the HIL design process to the reader by proposing an extension of the existing framework to full 3 DoF attitude control of the projectile mock-up. The first section presents the redesign of the experimental setup in order to reduce roll bearing friction and enable tuning of the projectile geometrical properties, while the second part details the derivation of the linear model structure in its new 3 degrees-of-freedom form. Section three describes the system identification procedure, including the estimation of the model parameters and the experimental assessment of the model fidelity. The final section reports the new developments on the skid-to-turn autopilot, including the addition of the roll channel and a new procedure to adjust the frequency templates of the weighting filters. Reference tracking and disturbance rejection experiments are performed to evaluate the performance of the attitude controller.

1.1 Experimental Setup

The ACHILES is a complex and multi-disciplinary system, with the aerodynamics, mechanical design and embedded electronics playing a critical role in its ability to emulate guided projectile in free flight. The hardware design is the result of a trade-off between performance requirements specific to each subsystem, each of them competing for the limited internal volume of the mock-up. Therefore, the projectile can be declined in multiple configurations adapted to a specific type of experiment.

1.1.1 Wind tunnel and gimbal assembly

Subsonic wind tunnel

The wind tunnel and most of the gimbal structure have been left unchanged with reference to the setup used by Strub [Str16]. The tunnel pictured in fig. 1.1 is of closed-return design, featuring an open test section measuring 70 cm by 90 cm and a 40 kW DC motor capable of accelerating the airflow up to 50 m/s. In the following study, the airspeed is set to 30 m/s, leading to a Mach number of 0.1 and a Reynolds number based on fin chord length of approximately 100,000. These experimental conditions differ from long-range guided projectiles flying in the transonic or supersonic regimes, implying that the structure of the flow field and the forces and moments applied on the mock-up may be greatly dissimilar to those encountered during live-firing experiments. However, valuable insight can still be gained from benchmarking aerodynamic configurations or controller designs, assessing the relevance of new control and design methodologies before eventually upgrading to more representative flight conditions. Moreover, new concepts of shorter-ranged projectiles such as guided mortar rounds or loitering ammunition may encounter similar velocities as what can be simulated in the tunnel, opening the door for higher-fidelity experimental setups based on this installation.

The wind tunnel low operating cost makes it particularly adapted to conceptual studies, where the outcome is uncertain and the number of tests required may be large. Also, the moderate forces applied on the mock-up allow the use of additive manufacturing to speed up the prototyping phase and broaden the scope of the tested configurations. The dimensions of the test section enable to reproduce the large angles of attack and sideslip that may be encountered by agile projectiles. The mock-up does not need to be miniaturized, so that development costs can be reduced by using modified or off-the-shelf hardware components such as the servo-actuators or the IMU. This is why this facility has been

selected to evaluate a novel design methodology based on plant-controller optimization and expand the scope of projectile control towards larger angles of attack, as described in chapter three.




Figure 1.1 – ISL subsonic wind tunnel layout [Jaeggy, 1982]

Three Degrees-of-Freedom (3 DoF) gimbal

The gimbal is composed of three sub-assemblies (excluding the projectile) which are connected to each other by ball bearings allowing only one degree of angular relative motion. The supporting frame rests on the floor thanks to four adjustable stands, its width and height providing ample clearance with respect to the airflow. The yaw frame is composed of two carbon fiber rods closed by a pair of CNC-machined aluminum rails and is supported by a pair of ball bearings at its top and bottom. The pitch rods are stainless steel bars connecting the roll sleeve to the yaw frame and fitted with bearings at their outer ends. Finally, the roll sleeve links the projectile mock-up to the pitch rods. Two variants of this assembly have been used: the legacy variant allows to adjust the longitudinal position of the body in the gimbal, but its excessive friction prevents the roll axis to be used. To overcome this issue, a smaller diameter bearing pictured in fig. 1.3

has been directly integrated into the mock-up, reducing friction to acceptable levels at the expense of the longitudinal adjustment capability.




Figure 1.2 – 3 degrees-of-freedom gimbal assembly [Strub, 2016]

Compared to a more conventional setup based on a sting and an inboard ball joint [Nel16], the ACHILES gimbal offers significantly greater amplitudes of motion with 360° in roll and pitch and $\pm 45^\circ$ in yaw, versus the respective 360° , $\pm 9^\circ$ and $\pm 9^\circ$ allowed by Nelson’s design. Note that both support structures are intrusive and may induce discrepancies with flight test results due to the sting interfering with the projectile base or, in the ACHILES case, the pitch rods shedding vorticity and interacting with the downstream fins. However, the ISL setup is bulkier, leading to an additional drawback: as the yaw frame extends outboard, its moment of inertia relative to the rotation axis outweighs the one of the projectile and dramatically decreases the natural frequency of the yawing motion. Thus, in order to correctly estimate the lateral aerodynamic coefficients, the yaw frame inertia must be accounted for in the model of the projectile rotational motion.




Figure 1.3 – CAD of the redesigned roll sleeve

1.1.2 Projectile airframe




Figure 1.4 – General layout of the ACHILES mock-up configured for ACO

Most features of the mock-up design illustrated by fig. 1.4 were carried over from earlier iterations of the experimental setup [Str16]. The projectile body is cylindrical with a caliber of 80 mm and a total length between 400 mm and 435 mm depending on the

type of roll sleeve fitted. An hemispherical nose contributes to the mock-up predictable aerodynamic behavior in subsonic flight and provides additional internal volume for ballasting the projectile. The model accommodates two sets of four lifting surfaces each, one dedicated to attitude control and the other to lift generation and airframe static stability. The canards located at the front provide pitch, roll and yaw moments in order to control the projectile attitude, while the fins situated aft of the center of mass generate a restoring moment driving the AoA and sideslip to zero. To simplify the manufacturing and modeling process, both fins and canards were shaped like uncambered flat planes with respective thicknesses of one and two millimeters. As the degree of maturity of the autopilot design has increased over the course of the project, the layout of the mock-up evolved to promote similarities with other innovative projectile concepts and tackle control issues associated with complex aerodynamic phenomena. Two mock-up configurations have been tested, one focused on attitude control and the other on airframe-controller optimization.

"3 DoF" layout

The first configuration is the closest to the projectile geometry introduced by Strub [Str16]. However, its redesigned roll sleeve allows modeling and control of the attitude dynamics, hence this setup has been retained to perform the experiments presented the first chapter of this thesis. This mock-up features oversized canards as well as large fins protruding aft of the projectile base for maximum stability (c.f. fig. 1.5). Its geometrical characteristics are summarized in Table 1.1 while its mass and inertia properties are given by Table 1.2. More details on the determination of the mock-up CG and MoI can be found in section 1.3.2.

Fins base chord: 115 mm	Fins tip chord 55 mm	Canards chord: 60 mm
Caliber: 80 mm	Fins span: 110 mm	Canards span: 60 mm
Tot. length: 435 mm	Fins TE pos: 470 mm	Canards TE pos: 132 mm

Table 1.1 – Projectile mock-up geometry (3 DoF layout)

Mass: 3.1252 kg	CG pos: 187 mm
I_{XX} : $4.206 \cdot 10^{-3}$ kg.m ²	I_{ZZ} : $6.143 \cdot 10^{-2}$ kg.m ²

Table 1.2 – Measured mass and inertia properties (3 DoF layout)




Figure 1.5 – "3 DoF" projectile mock-up layout

"ACO" layout




Figure 1.6 – "Dragonfly" long-range guided projectile concept [Decroq, 2019]

In order to pave the way for future applications, the ACHILES mock-up has been modified to emulate some of the most significant design features of innovative Long-Range Guided Projectile (LRGP) configurations currently investigated at ISL. The configuration of interest, named Dragonfly and pictured in fig. 1.6 is a fin-stabilized subsonic projectile featuring moderate aspect ratio wings designed to perform gliding flight over long distances. In an effort to promote similarities between the ACHILES and this new airframe, the aspect ratio of both fins and canards have been increased. However, due to the limited

size of the test section, ratios in excess of two were found to be impractical. Also, concerns about the lateral stability of this preliminary design have led to the vertical fins being retained on the HIL mock-up, unlike the Dragonfly concept. To allow the static stability of the mock-up to be readily adjusted over a wide range of variation, the aft part of the projectile has been redesigned. The legacy roll bearing was retrofitted to enable the balancing of the mock-up without an external counterweight and the original tail assembly was replaced with a much lighter 3D printed version which longitudinal position could be adjusted over a travel range of 175 mm.

During early testing of this configuration, it has been observed that the higher canards aspect ratios increase the non-linearity of the control response and promote earlier onset of canards stall. As this issue may be critical for projectile maneuverability, this layout has been chosen to model such complex phenomena and attempt to mitigate them using plant-controller optimization, respectively in chapters two and three. The geometry of this new mock-up variant pictured on fig. 1.7 is presented in Table 1.3 while its mass and inertia characteristics are described in Table 1.4.




Figure 1.7 – "ACO" projectile mock-up layout

	Fins chord 45 mm	Canards chord: 30 mm
Caliber: 80 mm	Fins span: 90 mm	Canards span: 60 mm
Tot. length: 400 mm	Fins TE pos: 225 to 400 mm	Canards TE pos: 118 mm

Table 1.3 – Projectile mock-up geometry (ACO layout)

Mass: 2.0415 kg	CG pos (aft fins): 123 mm
I_{XX} : Not measured	I_{ZZ} : $2.4789 \cdot 10^{-2}$ kg.m ²

Table 1.4 – Measured mass and inertia properties (ACO layout)

1.1.3 Autopilot hardware and software

Embedded hardware

The actuators, sensors and embedded computer are directly carried over from Strub’s experimental setup [Str+14; SB16]. The main attitude reference is provided by a Microstrain™ 3DM-GX3-25 Inertial Measurement Unit (IMU) whose specifications are detailed in Table 1.5. In order to cope with disturbed magnetometer data due to the proximity of the metallic chassis of the gimbal, the IMU output is merged with an external heading measurement through a complementary filter. This external heading reference is obtained from an incremental coder mounted on top of the yaw frame with a resolution of 0.045° .




Figure 1.8 – Actuators and avionics of the HIL mock-up [Str16]

The actuation mechanism has been designed so that the actuator dynamics could be accurately represented by a second-order LTI system. Each canard is actuated by a custom servomotor through a toothed belt: the hobby-grade unit is augmented with an absolute magnetic encoder feeding a P-PI cascaded controller scheme. Thanks to a mixed-sensitivity \mathcal{H}_∞ synthesis, the response of the servo-actuator matches with the reference model and is unaffected by supply voltage saturation for deflection commands up to $\pm 60^\circ$.

	Accelerometers	Rate gyros	Magnetometers
Measurement range	$\pm 5 g$	$\pm 300^\circ/\text{s}$	$\pm 2.5 \text{ Gauss}$
Non-linearity	$\pm 0.1 \% \text{ fs}$	$\pm 0.03 \% \text{ fs}$	$\pm 0.4 \% \text{ fs}$
Bias stability	$\pm 0.04 \text{ mg}$	$18^\circ/\text{h}$	–
Initial bias error	$\pm 2 \text{ mg}$	$\pm 0.25^\circ/\text{s}$	$\pm 3 \text{ mGauss}$
Scale factor stability	$\pm 0.05 \%$	$\pm 0.05 \%$	$\pm 0.1 \%$
Noise density	$80 \mu\text{g}/\sqrt{\text{Hz}}$	$0.03^\circ/\text{s}/\sqrt{\text{Hz}}$	$100 \mu\text{Gauss}/\sqrt{\text{Hz}}$
Alignment error	$\pm 0.05^\circ$	$\pm 0.05^\circ$	$\pm 0.05^\circ$
Bandwidth	225 Hz	440 Hz	230 Hz
Sampling rate	30 kHz	30 kHz	7.5 kHz

Table 1.5 – IMU Specifications [Str16]

Software environment

The embedded operating system consists of a firm-real-time distribution of Linux installed on a Gumstix Overo FireSTORM single-board computer. This solution is based on a co-kernel approach where the Xenomai kernel running the critical tasks has the highest priority when accessing the processor and peripherals. These tasks include reading the IMU data, writing the servo setpoints and communicating with the remote yaw coder board.

The development environment is based on MATLAB and Simulink software running on a dedicated Linux laptop connected to the Gumstix via Wi-Fi. The automatic C code generation functionality is leveraged to allow quick and simple implementation of the control laws on the target hardware after cross-compilation with the ARM-gcc compiler. The remote execution and real-time monitoring is performed over the wireless link before the full-resolution measured data is retrieved at the end of the test and downloaded on the development computer for further post-processing and analysis.

1.2 Linear Model of the Projectile Angular Motion

The first step of the Hardware-In-the-Loop controller design methodology is to propose a model structure and estimate the values of its key parameters from experimental data. This is known as the grey-box modeling approach [Boh06] and has been used to good effect in Strub’s previous work [Str+15] [SB16]. In order to explicitly estimate the aerodynamic derivatives of the projectile, a new model structure was derived from the non-linear equations of rotational motion and then simplified to minimize the number of

parameters to be estimated.

1.2.1 Frames of reference

The projectile motion is described using two main frames of reference represented in fig. 1.9 which share the same origin at the gimbal center of rotation. The Local frame follows a modified North-East-Down convention: X^L is chosen parallel to the wind direction, Z^L points downwards and Y^L completes the frame according to the right-hand rule. The Body frame is attached to the projectile, with X^B exiting through the nose along the fuselage revolution axis, Y^B pointing to the right and Z^B pointing downwards. The orientation of the Body frame relative to the Local frame is described by three consecutive rotations around the roll, pitch and yaw axes corresponding to the respective Euler angles (ϕ, θ, ψ) . Note that, to define the intermediate coordinate systems ¹ and ², the rotations are performed according to the “yaw, pitch then roll” sequence.




Figure 1.9 – Kinematic diagram of the gimbal showing both Body and Local frames

Additionally, the velocity vector of the projectile relative to the air mass v^{AB} is used to define the Wind coordinate system. The angle of attack α and the sideslip angle β defined in Figure 1.10 are very often used to express the aerodynamic coefficients of the airframe. In our case, by construction, the orientation of the velocity vector is coincident

with the X axis of the Local frame. Thus, the sideslip angle is the opposite of the yaw angle ψ and the angle of attack is equivalent to the pitch attitude θ .




Figure 1.10 – Wind coordinate system [Str16]

1.2.2 Control surfaces allocation

Since the ACHILES mock-up is fitted with four independent control surfaces but only has three degrees-of-freedom, the transformation between the virtual control actions $[\delta_l, \delta_m, \delta_n]^T$ on the respective roll, pitch and yaw axes, and the actual deflections $[\delta_1, \delta_2, \delta_3, \delta_4]^T$ of each actual canard is non-unique.

$$\begin{bmatrix} \delta_l \\ \delta_m \\ \delta_n \end{bmatrix} = \frac{1}{4} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\phi) & \sin(\phi) \\ 0 & -\sin(\phi) & \cos(\phi) \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 & 1 \\ -1 & 1 & 1 & -1 \\ -1 & -1 & 1 & 1 \end{bmatrix} \begin{bmatrix} \delta_1 \\ \delta_2 \\ \delta_3 \\ \delta_4 \end{bmatrix} \quad (1.1)$$

The control scheme proposed in Equation (1.1) corresponds to an X canards layout at zero bank angle. This layout gives the best control authority on the pitch and yaw axes since all of the four canards can contribute to pitch-only or yaw-only commands [Nie88, p. 226]. The projectile bank angle is accounted for during control surfaces allocation as it changes the orientation of the control surfaces and their effect on the pitch and yaw angles (i.e, if the projectile is rolled by 90°, a set of controls deflection causing pure pitch

at $\phi = 0^\circ$ now produces exclusively yaw). The actual canards deflections can be computed from the virtual control actions by taking the pseudo-inverse of the product of the two matrices.

1.2.3 Non-linear equations of motion

The projectile rotational motion is obtained by applying Euler's second law to multiple subsets of the gimbal kinematic chain. The equations of motion are derived under the following assumptions:

- The Local frame is assumed to be inertial.
- The projectile and the gimbal components are rigid bodies.
- All pivot joints are frictionless and play-free.
- The CG of the projectile is coincident with the gimbal center of rotation.
- The inertia tensor of the projectile is diagonal: $\mathbf{I}_O^B = \text{diag}(I_{XX}, I_{YY}, I_{ZZ})$.
- The inertia terms of all the gimbal parts are negligible, except for the yaw frame that has a significant moment of inertia I_{ZZg} around its rotation axis.

The non-linear equations of motion of the projectile inside the 3 DoF gimbal are derived in [Str16, p. 51]. Equations (1.2) to (1.4) express the differentials of the Body angular rates \dot{p} , \dot{q} and \dot{r} as functions of the projections L , M and N of the aerodynamic moment in the Body frame.

$$\dot{p} = \frac{1}{I_{XX}} [L + qr (I_{YY} - I_{ZZ})] \quad (1.2)$$

$$\dot{q} = \frac{1}{A_Y} \left[k_Y \left(1 + \frac{I_{ZZg} \cos^2 \phi}{I_{ZZ} \cos^2 \theta} \right) - k_Z \frac{I_{ZZg} \sin \phi \cos \phi}{I_{ZZ} \cos^2 \theta} - k_\psi I_{ZZg} \frac{\sin \phi}{\cos^2 \theta} \right] \quad (1.3)$$

$$\dot{r} = \frac{1}{A_Z} \left[-k_Y \frac{I_{ZZg} \sin \phi \cos \phi}{I_{YY} \cos^2 \theta} + k_Z \left(1 + \frac{I_{ZZg} \sin^2 \phi}{I_{YY} \cos^2 \theta} \right) - k_\psi I_{ZZg} \frac{\cos \phi}{\cos^2 \theta} \right] \quad (1.4)$$

where symbols A_Y , A_Z , k_Y , k_Z and k_ψ are defined as follows:

$$\begin{aligned}
 A_Y &= I_{YY} + I_{ZZg} \frac{1}{\cos^2 \theta} \left[\sin^2 \phi + \frac{I_{YY}}{I_{ZZ}} \cos^2 \phi \right] \\
 A_Z &= I_{ZZ} + I_{ZZg} \frac{1}{\cos^2 \theta} \left[\cos^2 \phi + \frac{I_{ZZ}}{I_{YY}} \sin^2 \phi \right] \\
 k_Y &= M - pr(I_{XX} - I_{ZZ}) \\
 k_Z &= N - pq(I_{YY} - I_{XX}) \\
 k_\psi &= \frac{1}{\cos \theta} \left[pq \cos \phi - pr \sin \phi + 2(q^2 - r^2) \tan \theta \cos \phi \sin \phi + 2qr \tan \theta (1 - 2 \sin^2 \phi) \right]
 \end{aligned}$$

Additionally, the attitude derivatives are given by the Euler angles differential equations:

$$\dot{\phi} = p + q \sin \phi \tan \theta + r \cos \phi \tan \theta \quad (1.5)$$

$$\dot{\theta} = q \cos \phi - r \sin \phi \quad (1.6)$$

$$\dot{\psi} = q \frac{\sin \phi}{\cos \theta} + r \frac{\cos \phi}{\cos \theta} \quad (1.7)$$

Change of reference frame

The expressions recalled above are functions of the projectile bank angle ϕ . However, a fundamental property of cruciform airframes can be leveraged to get rid of this dependency. As explained by Nielsen [Nie88, pp. 122-124], if the projectile is first pitched at an angle θ before being banked at an angle ϕ , then the magnitude of the normal force applied on the projectile is independent of the bank angle. As a result, the pitching and yawing moments can be expressed in the non-rolling frame $[O, X^2, Y^2, Z^2]$ of Figure 1.9 for modeling and identification purposes before being projected in the Body frame during simulations. Besides, the transformation from the frame 2 to the Body frame is included in the control allocation (Equation (1.1)) for the implementation on the HIL setup.

Thus, in the non-rolling frame, ϕ is can be set to zero in equations (1.3), (1.4), (1.6) and (1.7), leading to the following set of equations:

$$\dot{p} = \frac{1}{I_{XX}} [L + qr (I_{YY} - I_{ZZ})] \quad (1.8)$$

$$\dot{q} = \frac{1}{I_{YY}} [M - pr (I_{XX} - I_{ZZ})] \quad (1.9)$$

$$\dot{r} = \frac{1}{\left(I_{ZZ} + \frac{I_{ZZg}}{\cos^2 \theta}\right)} \left[N - pq (I_{YY} - I_{XX}) - \frac{I_{ZZg}}{\cos^3 \theta} (pq + 2qr \tan \theta) \right] \quad (1.10)$$

$$\dot{\phi} = p + q \sin \phi \tan \theta + r \cos \phi \tan \theta \quad (1.11)$$

$$\dot{\theta} = q \quad (1.12)$$

$$\dot{\psi} = \frac{r}{\cos \theta} \quad (1.13)$$

The state and input vectors are respectively chosen as $\mathbf{x} = [\phi, p, \theta, q, \psi, r]^T$ and $\mathbf{u} = [\delta_l, \delta_m, \delta_n]^T$. With that definition, the equations (1.8) to (1.13) can be summarized into the following form:

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t)) \quad (1.14)$$

With the state functions:

$$f_\phi = p + q \sin \phi \tan \theta + r \cos \phi \tan \theta \quad (1.15)$$

$$f_p = \frac{1}{I_{XX}} [L + qr (I_{YY} - I_{ZZ})] \quad (1.16)$$

$$f_\theta = q \quad (1.17)$$

$$f_q = \frac{1}{I_{YY}} [M - pr (I_{XX} - I_{ZZ})] \quad (1.18)$$

$$f_\psi = \frac{r}{\cos \theta} \quad (1.19)$$

$$f_r = \frac{1}{\left(I_{ZZ} + \frac{I_{ZZg}}{\cos^2 \theta}\right)} \left[N - pq (I_{YY} - I_{XX}) - \frac{I_{ZZg}}{\cos^3 \theta} (pq + 2qr \tan \theta) \right] \quad (1.20)$$

1.2.4 Control-oriented model

Linearization

A Jacobian linearization of the projectile pitch and yaw dynamics is performed around the trim point \bar{p} associated to the equilibrium conditions $(\bar{\phi}, \bar{\theta}, \bar{\psi}) = (0, 0, 0)$. In the vicinity of \bar{p} , the behavior of the non-linear system described by Equation (1.14) can be approxi-

mated by the following linear system:

$$\dot{\mathbf{x}}_\delta(t) \simeq \left. \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \right|_{\bar{\mathbf{p}}} \mathbf{x}_\delta(t) + \left. \frac{\partial \mathbf{f}}{\partial \mathbf{u}} \right|_{\bar{\mathbf{p}}} \mathbf{u}_\delta(t) \quad (1.21)$$

where $\mathbf{x}_\delta = \mathbf{x} - \bar{\mathbf{x}}$ and $\mathbf{u}_\delta = \mathbf{u} - \bar{\mathbf{u}}$ represent the deviations from the equilibrium.

The state-space representation of the projectile attitude dynamics is deduced from Equation (1.21), bearing in mind that, at the equilibrium condition, the equations (1.11) to (1.13) give $\bar{p} = \bar{q} = \bar{r} = 0$. In addition to the previously-defined state and input vectors, the outputs are chosen as the attitude angles $\mathbf{y} = [\phi, \theta, \psi]^T$. The state equation of the projectile model is then $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}$ while the output equation is simply $\mathbf{y} = \mathbf{C}\mathbf{x}$. The state, input and output matrices are described by the equations (1.22), (1.23) and (1.24).

$$\mathbf{A} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & A_{lp} & A_{l\theta} & A_{lq} & A_{l\psi} & A_{lr} \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & A_{mp} & A_{m\theta} & A_{mq} & A_{m\psi} & A_{mr} \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & A_{np} & A_{n\theta} & A_{nq} & A_{n\psi} & A_{nr} \end{bmatrix} \quad (1.22)$$

$$\mathbf{B} = \begin{bmatrix} 0 & 0 & 0 \\ B_{l\delta l} & B_{l\delta m} & B_{l\delta n} \\ 0 & 0 & 0 \\ B_{m\delta l} & B_{m\delta m} & B_{m\delta n} \\ 0 & 0 & 0 \\ B_{n\delta l} & B_{n\delta m} & B_{n\delta n} \end{bmatrix} \quad (1.23)$$

$$\mathbf{C} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \quad (1.24)$$

The coefficients of the \mathbf{A} and \mathbf{B} matrices are given by:

$$A_{lk} = I_{XX}^{-1} \frac{\partial L}{\partial k} \quad A_{mk} = I_{YY}^{-1} \frac{\partial M}{\partial k} \quad A_{nk} = I_{ZZt}^{-1} \frac{\partial N}{\partial k}$$

$$B_{lk} = I_{XX}^{-1} \frac{\partial L}{\partial k} \quad B_{mk} = I_{YY}^{-1} \frac{\partial M}{\partial k} \quad B_{nk} = I_{ZZt}^{-1} \frac{\partial N}{\partial k}$$

where $I_{ZZt} = I_{ZZ} + I_{ZZg}$ is the total yaw MoI of the projectile in its gimbal

Model simplification based on geometrical considerations

In order to keep only the major contributing terms and reduce the risk of colinearity between the parameters in the forthcoming estimation step, the number of non-zero coefficients has been reduced based on considerations on the mock-up geometry and its control allocation scheme:

- Roll control surfaces deflections are symmetric w.r.t X^B , thus $B_{m\delta l} = 0$ and $B_{n\delta l} = 0$.
- Pitch control surfaces deflections are symmetric w.r.t $[O, X^B, Z^B]$, thus $B_{l\delta m} = 0$ and $B_{n\delta m} = 0$.
- Yaw control surfaces deflections are symmetric w.r.t $[O, X^B, Y^B]$, thus $B_{l\delta l} = 0$ and $B_{m\delta n} = 0$.

Also, previous work [SB16] reports that the decoupling assumption between the pitch and yaw axes is relevant in the light of the limitations of the model identification procedure. This implies that the coefficients $A_{m\psi}$, A_{mr} , $A_{n\theta}$ and A_{nq} are set to zero. Moreover, the peculiarities of the projectile geometry indicate that this assumption may be extended to the roll axis since most of the cross-coupling effects found on aircraft configurations may be irrelevant:

- Since the wing span of the projectile is relatively small, the velocity differential induced at each wing tip by yawing motion is much smaller. Thus, the rolling moment resulting from the lift imbalance induced by this local airspeed difference can be neglected.
- As the fins are placed in the middle of the fuselage (i.e not on top nor at the bottom) and do not feature any dihedral or anhedral, the flow field above and below the fin during a sideslip condition are similar and no significant “dihedral effect” should be expected.
- Since the airfoil used on the fins is symmetrical, there is no difference in drag between the rising wing and the dropping wing during a rolling motion. Thus, no yawing moment develops as the result of a roll rate.

The above propositions respectively imply that the coefficients A_{lr} , $A_{l\psi}$ and A_{np} can be neglected. By symmetry, these considerations also apply to the terms A_{lq} , $A_{l\theta}$ and A_{mp} , meaning that the roll axis is now completely decoupled from pitch and yaw. Consequently, the state and input matrices of the attitude dynamics model reduce to:

$$\mathbf{A} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & A_{lp} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & A_{m\theta} & A_{mq} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & A_{n\psi} & A_{nr} \end{bmatrix} \quad (1.25)$$

$$\mathbf{B} = \begin{bmatrix} 0 & 0 & 0 \\ B_{l\delta l} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & B_{m\delta m} & 0 \\ 0 & 0 & 0 \\ 0 & 0 & B_{n\delta n} \end{bmatrix} \quad (1.26)$$

Decomposition into aerodynamic derivatives

The moments applied on the projectile can be expressed as functions of non-dimensional parameters C_l , C_m and C_n named aerodynamic coefficients, which are introduced in the equations (1.27) to (1.29). Their purpose is that experimental results obtained with scale models could be directly applied to full-scale airframes. Hence, they are intended to be insensitive to variations of the experimental conditions such as the freestream dynamic pressure \bar{q} and velocity \bar{q} , or to the model geometry (caliber D and cross-section S). The three moment coefficients are decomposed into aerodynamic derivatives (Equations (1.30) to (1.32)) through a first-order Taylor series expansion so that they can be substituted into the dimensional coefficients of the \mathbf{A} and \mathbf{B} matrices.

$$L = \bar{q}SD C_l \quad (1.27)$$

$$M = \bar{q}SD C_m \quad (1.28)$$

$$N = \bar{q}SD C_n \quad (1.29)$$

$$C_l = \frac{D}{\bar{v}} C_{l_p} p + C_{l_{\delta l}} \delta_l \quad (1.30)$$

$$C_m = C_{m_\alpha} \theta + \frac{D}{\bar{v}} C_{m_q} q + C_{m_{\delta m}} \delta_m \quad (1.31)$$

$$C_n = -C_{n_\beta} \psi + \frac{D}{\bar{v}} C_{n_r} r + C_{n_{\delta n}} \delta_n \quad (1.32)$$

$$\begin{aligned} A_{l_p} &= \frac{\bar{q}SD^2}{\bar{v}} \frac{C_{l_p}}{I_{XX}} & B_{l_{\delta l}} &= \bar{q}SD \frac{C_{l_{\delta l}}}{I_{XX}} \\ A_{m_\theta} &= \bar{q}SD \frac{C_{m_\alpha}}{I_{YY}} & A_{m_q} &= \frac{\bar{q}SD^2}{\bar{v}} \frac{C_{m_q}}{I_{YY}} & B_{m_{\delta m}} &= \bar{q}SD \frac{C_{m_{\delta m}}}{I_{YY}} \\ A_{n_\psi} &= \bar{q}SD \frac{-C_{n_\beta}}{I_{ZZt}} & A_{n_r} &= \frac{\bar{q}SD^2}{\bar{v}} \frac{C_{n_r}}{I_{ZZt}} & B_{n_{\delta n}} &= \bar{q}SD \frac{C_{n_{\delta n}}}{I_{ZZt}} \end{aligned}$$

Thanks to the successive simplifications justified by a-priori knowledge of the projectile aerodynamics, the number of aerodynamic parameters to be estimated reduces to eight, namely C_{l_p} , $C_{l_{\delta l}}$, C_{m_α} , C_{m_q} , $C_{m_{\delta m}}$, C_{n_β} , C_{n_r} and $C_{n_{\delta n}}$. This low number of parameters will aid the estimation process described in the next section.

1.3 Identification of the Angular Dynamics

According to the legacy system identification procedure presented in Figure 1.11, one should first select a model structure and choose a subset of its parameters to be estimated from experimental data. *A priori* and *a posteriori* identifiability analyses should then be performed to ensure that the estimation problem is feasible and well-conditioned given a sample of input-output data. However, in our case, prior knowledge of successful model structures and parameter vectors combined with strong decoupling assumption between the roll, pitch and yaw dynamics made this step unnecessary. Also, since no sensitivity analysis was performed, the optimal input design procedure could not be followed. Instead, the excitation signal was adapted from previous work [SB16] and adjusted to provide good model fit on validation data and low uncertainty on the parameter estimates.




Figure 1.11 – System identification procedure [Str16]

1.3.1 Model postulation

The structure of the attitude model is based on the state-space representation derived in the previous section. Since the bank angle ϕ and the angular rates q and r are either integrals or derivatives of other state variables, they can be discarded to simplify the estimation problem. Therefore, the output vector used for identification is $\mathbf{y} = [p, \theta, \psi]^T$. In order to account for measurement and process noise, an error model composed of a disturbance matrix \mathbf{K} and a residuals signal $\mathbf{e}(t)$ corresponding to white Gaussian noise is added to the state and output equations. Since the outputs are a subset of the state vector (sensor dynamics are neglected), the state and measurement noise is assumed to

be fully correlated. The unknown parameter vector to be estimated is composed of the aerodynamic derivatives augmented with the coefficients of the disturbance matrix :

$$\boldsymbol{\theta} = \left[C_{lp}, C_{l\delta l}, C_{m\alpha}, C_{mq}, C_{m\delta m}, C_{n\beta}, C_{nr}, C_{n\delta n}, K_{\dot{p}p}, K_{\dot{\theta}\theta}, K_{\dot{q}\theta}, K_{\dot{\psi}\psi}, K_{\dot{r}\psi} \right]^T$$

Finally, the model structure is defined by Equations (1.33) and (1.34). The \mathbf{A} and \mathbf{B} matrices are taken from (1.25) and (1.26) while \mathbf{C} , \mathbf{D} and \mathbf{K} are detailed below.

$$\dot{\mathbf{x}}(t) = \mathbf{A}(\boldsymbol{\theta})\mathbf{x}(t) + \mathbf{B}(\boldsymbol{\theta})\mathbf{u}(t) + \mathbf{K}(\boldsymbol{\theta})\mathbf{e}(t) \quad (1.33)$$

$$\mathbf{y}(t) = \mathbf{C}\mathbf{x}(t) + \mathbf{D}\mathbf{u}(t) + \mathbf{e}(t) \quad (1.34)$$

$$\mathbf{C} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \quad (1.35)$$

$$\mathbf{D} = \mathbb{O}_3 \quad (1.36)$$

$$\mathbf{K} = \begin{bmatrix} 0 & 0 & 0 \\ K_{\dot{p}p} & 0 & 0 \\ 0 & K_{\dot{\theta}\theta} & 0 \\ 0 & K_{\dot{q}\theta} & 0 \\ 0 & 0 & K_{\dot{\psi}\psi} \\ 0 & 0 & K_{\dot{r}\psi} \end{bmatrix} \quad (1.37)$$

1.3.2 Experiment design and preliminary work

The difference between the new structure and the legacy model is that the aerodynamic derivatives are estimated in place of the dimensional coefficients of the state and input matrices. This requires the knowledge of the mock-up inertial properties, as well as the moment of inertia of the yaw gimbal. In exchange, the estimated aerodynamic properties can be directly compared to the output of CFD simulations, semi-empirical aerodynamics predictions or force balance measurements in wind-tunnel [Dec+18].

Determination of the mock-up inertial properties

The principal moments of inertia of the projectile are evaluated with a commercial measurement system operated by ISL ballistics group. The first step of the process is to weight the mock-up and estimate the position of its center of gravity. The projectile mass is given by a laboratory scale while a Space Electronics™ SE8913 device provides the CG position. Then, additional measurements gathered with a Raptor Scientific™ XR10 instrument are used to compute the axial and transverse moments of inertia of the mock-up. The measured mass and inertial properties are summarized in Tables 1.2 and 1.4 of the previous section.




Figure 1.12 – CG position measurement setup for transverse (left) and base (right) planes

Due to the design of the mechanical interface between the projectile and the measurement equipment, only the transverse MoI corresponding to I_{ZZ} could be determined. Since Strub [Str16, p. 15] measured very similar values for I_{YY} and I_{ZZ} , these two moments of inertia can be assumed equal. Additionally, the MoI of the yaw gimbal frame I_{ZZg} was estimated from the CAD model of this assembly: the result is one order of magnitude greater than the projectile transverse inertia with $I_{ZZg} = 0.69981 \text{ kg.m}^2$

Input signal building

As the model structure assumes that the rolling, pitching and yawing motions are decoupled, the parameter estimation experiments can be performed sequentially. Since the roll dynamics are first-order, the roll excitation sequence consists of a single step on the roll control channel. Conversely, as the pitch and yaw dynamics are more complex, a pseudo-random binary sequence based on Strub's previous work [SB16] is used. The

amplitude of the signal is chosen to maximize the signal-to-noise ratio while avoiding non-linearities caused by actuator saturation.

1.3.3 Estimation of the aerodynamic derivatives

Data collection

The data collection protocol includes two sets of excitation sequences to be consistent with the decoupling assumption made at the model building stage: first, the roll axis is excited while the pitch and yaw axes are locked, then the pitch and yaw axes are sequentially excited while the roll axis is locked. Each excitation sequence is then split into an estimation sequence and a validation sequence: the former is used for parameter estimation while the latter allows to assess the fit quality of the identified model.




Figure 1.13 – Identified model (color) fit on validation data (black)

On Figure 1.13, the model responses in roll rate, pitch attitude and yaw angle are plotted in color whereas the dash-dotted line represents the measurements. The identified model shows a satisfactory fit with the projectile measured states as the normalized root-mean-square fit percentages for roll rate, pitch attitude and yaw angle come out respectively at 90.1%, 76.9% and 77.9%. The roll rate dynamics are well reproduced, as well as the decaying oscillatory responses of pitch and yaw angles. The pitch response presents some moderate hysteresis that could be attributed to bearing friction, as well as process noise due to the compliance of the support structure. The additional frequency content may also come from unmodeled dynamics.

Parameter estimation

Starting from an initial guess provided by semi-empirical aerodynamic tools, the estimates of the model coefficients are obtained from a recursive algorithm based on prediction error minimization [Lju86; Lju02; WP97]. The estimation process is carried out in two steps: first, the aerodynamic derivatives are estimated without a noise model, then the coefficients of the disturbance matrix are determined by assuming that the parameter values found in the first step are fixed. The MATLAB function *greyest* of the System Identification Toolbox is used to compute the estimates of the model parameters displayed in Table 1.6.

Coef.	C_{lp}	$C_{l\delta l}$	$C_{m\alpha}$	C_{mq}	$C_{m\delta_m}$	$C_{n\beta}$	C_{nr}	$C_{n\delta n}$
Value	-24.55	1.20	-18.02	-209.5	10.78	22.65	-467.3	-13.20
Unc. (%)	5.31	4.57	2.41	11.5	2.56	1.19	14.8	1.72

Table 1.6 – Aerodynamic derivatives estimates

The sign and magnitude of the parameters are mostly consistent with the a-priori knowledge of the projectile flight dynamics: the rate damping coefficients C_{lp} , C_{mq} and C_{nr} are all negative but the yaw damping is estimated to be twice as high as for the pitch axis, which does not make sense from an aerodynamic point of view since the mock-up has tetragonal symmetry. This difference in damping could only be attributed to excessive friction in the bearings of the yaw gimbal. Apart from that, the sign of the pitch and yaw stability coefficients $C_{m\alpha}$ and $C_{n\beta}$ confirms that the projectile is statically stable on both axes (a positive AoA disturbance produces a nose-down moment while a left sideslip produces a clockwise moment). Also, the sign of the control coefficients $C_{l\delta l}$, $C_{m\delta_m}$ and $C_{n\delta n}$ agrees with the chosen conventions (i.e. a positive δ_m means that the equivalent pitch

control surface is oriented trailing edge down, so that it creates lift and thus a positive pitching moment for a canard configuration). Finally, the magnitudes of the corresponding pitch and yaw coefficients are fairly close, as expected because of the symmetry properties of the projectile configuration.

Uncertainty analysis

The *bootstrap* technique [ET94] is employed to quantify the uncertainty on the parameter estimates: the parameter estimation procedure is performed on a large number of fictive data series generated from the measured data. The bootstrapping procedure is described by Strub [Str16, pp. 68-69] and reproduced below for completeness:

- Estimate the parameter vector $\hat{\boldsymbol{\theta}}$ from the measured input and output vectors \boldsymbol{u} and \boldsymbol{y} .
- Compute the residual vector \boldsymbol{e} using these values and the noise model.
- Create fictive residuals vectors \boldsymbol{e}_f by applying random permutations to \boldsymbol{e}
- Using the original input \boldsymbol{u} , the fictive residuals \boldsymbol{e}_f and the estimated parameters vector $\hat{\boldsymbol{\theta}}$, generate fictive output vectors \boldsymbol{y}_f .

The uncertainty percentages given in Table 1.6 have been computed from 30 fictive datasets by comparing the worst-case, most different estimate of each parameter to its actual value. The pitch and yaw damping coefficients show the greatest amount of uncertainty, however such results are on a par with other measurement techniques [FLM21]. The accuracy of the other parameter estimates appears to be satisfactory, being less than 5 percent for most of them.

However, the estimated aerodynamic derivatives may deviate from the values that could be measured in free flight because of some process noise in the experimental setup. The closed-return wind tunnel is subject to periodic fluctuations of the freestream airspeed that may excite the flexible modes of the suspension structure. Moreover, the cylindrical rods holding the mock-up in place may be the cause of an unsteady vortex shedding phenomenon which can disturb the forces and moments applied on the projectile. Finally, friction in the gimbal bearings may artificially increase the damping of the projectile motion and the magnitude of the associated aerodynamic derivatives.

1.4 Attitude Controller Design and Validation

Once the dynamic model of the projectile has been successfully identified, the next step of the HIL methodology is to select a controller structure and adjust the synthesis framework to meet the performance objectives while respecting the actuator limitations. The resulting controller is then implemented on the experimental setup so that its reference tracking and disturbance rejection performance can be assessed. The closed-loop model predictions are compared to the attitude measurements in order to verify its accuracy and confirm the validity of the decoupling assumptions.

1.4.1 Controller structure and synthesis framework

The controller synthesis framework should be designed to meet reference tracking and disturbance rejection objectives. However, the controller structure has to cope with the dynamics of the open-loop projectile and the hardware limitations such as actuators saturation or computational delays: the closed-loop system must present sufficient stability margins and limit the control effort to avoid reaching the canards lift limit. Therefore, the synthesis method should attempt to find the best trade-off between the performance and stability objectives given the available actuator authority.

Controller structure postulation

The three-axis autopilot pictured in fig. 1.14 implements a dual-loop architecture with feedforward gains, which is derived from structures commonly found in the literature [Str16], [MR05].




Figure 1.14 – Controller structure

On the pitch and yaw axes, an inner regulation loop provides a proportional-derivative action using the feedback gains \mathbf{K}_R on both Euler angles and rates. This structure augments the stability of the projectile while limiting sensitivity to measurement noise by avoiding differentiation of the sensors outputs. The outer servo loop includes pure integral action in $\mathbf{K}_S(s)$ to cancel out angular static error. A feedforward controller $\mathbf{K}_F(s)$ composed of lead-lag filters on both pitch and yaw channels improves the reference tracking of the closed-loop plant. Finally, a second-order roll-off filter $\mathbf{F}_{ro}(s)$ mitigates the impact of high frequency process noise on the yaw channel to avoid the resonance of the support structure.

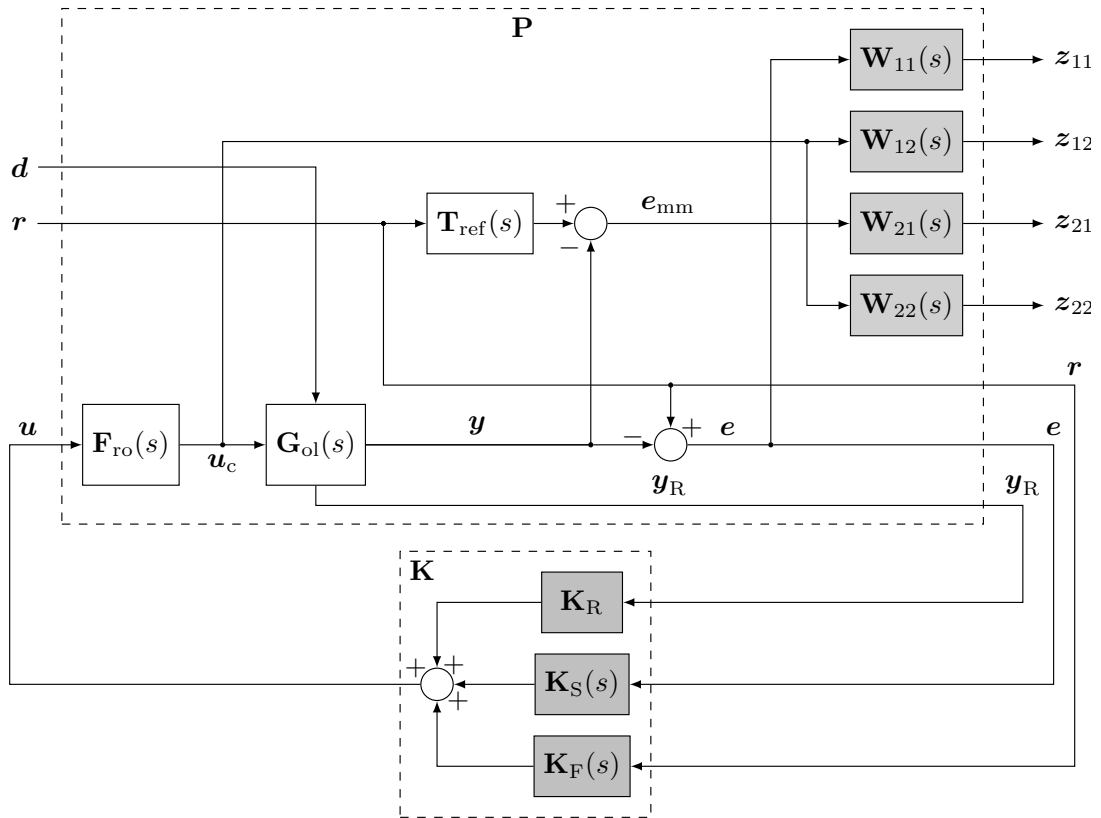


Figure 1.15 – Multi-objective \mathcal{H}_∞ framework

Controller synthesis framework

The synthesis of the fixed-structure controller described above is carried out thanks to a \mathcal{H}_∞ mixed-sensitivity formulation [Doy+89; AN06b] depicted in fig. 1.15. This scheme has been inherited from Strub's previous work [SB16] and adapted for attitude control.

Disturbance rejection and reference tracking are the two main requirements that are being addressed. For each objective, a set of two transfers are constrained by weighting functions in order to obtain satisfactory control performance without exceeding the actuator limitations.

To ensure disturbance rejection properties, a perturbation input d is added to the actuator output before the projectile model. The closed-loop transfer function $T_{d \rightarrow e}$ conditions the rejection performance of the plant, while the transfer $T_{d \rightarrow u_c}$, from the disturbance to the control signal, limits the amplitude and bandwidth of the actuator inputs. In the reference tracking case, the closed-loop transfer $T_{r \rightarrow y}$ is compared to the second order reference model T_{ref} that describes the expected tracking performance. The resulting model matching error e_{mm} is constrained so that the closed-loop time response matches the reference model output. As above, the transfer $T_{r \rightarrow u_c}$ is constrained to preserve the actuators.

The weighting filters are chosen as diagonal arrays of first order lead-lag transfer functions. The parameters $k_{i,LF}$, $k_{i,HF}$ and ω_i of Table 1.7 respectively define the low-frequency gain, hi-frequency gain and -3 dB bandwidth of the inverse filter W_i^{-1} pictured on Figure 1.16.




Figure 1.16 – Parametrization of the inverse filter W_i^{-1} [Str16]

$W_{11} (T_{d \rightarrow e})$		$W_{12} (T_{d \rightarrow u_c})$		$W_{21} (T_{r \rightarrow e_{mm}})$		$W_{22} (T_{r \rightarrow u_c})$	
$k_{11,LF}$	0.001	$k_{12,LF}$	2	$k_{21,LF}$	0.001	$k_{22,LF}$	2
$k_{11,HF}$	1.7	$k_{12,HF}$	0.001	$k_{21,HF}$	0.12	$k_{22,HF}$	0.001
$\omega_{11}(\phi)$	3 rad/s	ω_{12}	15 rad/s	$\omega_{21}(\phi)$	3 rad/s	ω_{22}	15 rad/s
$\omega_{11}(\theta)$	3 rad/s			$\omega_{21}(\theta)$	3 rad/s		
$\omega_{11}(\psi)$	2 rad/s			$\omega_{21}(\psi)$	2 rad/s		

Table 1.7 – Parameters of the weighting filters

The synthesis framework allows an accurate trade-off between the various objectives [SB16] but may be challenging to setup for non-expert users. Thus, a comprehensive procedure is proposed in order to provide a straightforward determination of the weighting filter parameters. The first step of the tuning procedure consists in specifying the reference model by taking into account the desired bandwidth and time response of each channel. The damping ratio is uniformly set to 0.78 for short settling time and limited overshoot while the bandwidth is specified on a channel-by-channel basis in order to take into account the open-loop dynamics. Because of the large gimbal frame inertia, the yaw axis resonance frequency of 2.56 rad/s is significantly lower than in pitch. Thus the required yaw bandwidth is capped at 2 rad/s whereas the closed-loop pitch bandwidth is set to 3 rad/s. Due to the limited roll authority, the specified roll bandwidth is chosen equal to the pitch bandwidth even though the roll axis behaves like a first order lag in series with an integrator and thus does not present any resonance.

The weight on the disturbance rejection transfer $T_{d \rightarrow e}$ plotted on fig. 1.17 constrains the sensitivity in the system bandwidth and cancels out any steady-state error. As the sensitivity rolls off at high frequencies, the plot of $T_{d \rightarrow e}$ presents a peak: $k_{11, HF}$ conditions the magnitude of this peak and thus should be minimized without exceeding the actuator constraints and degrading the model-matching transfer. The low frequency gain $k_{11, LF}$ is set to a small non-zero value in order to suppress the steady-state error while avoiding numerical issues. For each channel, ω_{11} corresponds to the respective reference model bandwidth. The template enforced on $T_{r \rightarrow e_{mm}}$ presented on fig. 1.18 attempts to minimize the difference between the closed-loop reference to output transfer and the reference model. This transfer also presents a peak at medium frequencies, therefore the upper bound of the high frequency transfer $k_{21, HF}$ is chosen as low as possible while keeping in mind the disturbance rejection objective and the actuator limitations. $k_{21, LF}$ and ω_{21} are set according to the same guidelines as in the disturbance rejection case. The weights on the actuator transfers $T_{d \rightarrow u_c}$ and $T_{r \rightarrow u_c}$ aim at limiting the amplitude and slew rate of the control signal. The parameters $k_{12, LF}$ and $k_{22, LF}$ are chosen so that the canards deflection limits are not reached in normal operation while ω_{12} and ω_{22} are fixed at one third of the actuator controller bandwidth to avoid reaching maximum slew rate. $k_{12, HF}$ and $k_{22, HF}$ are set to a small non-zero value to cut off the high frequency content of the control signal.

The synthesis results presented in figures 1.17 and 1.18 show that all three objectives are fulfilled for all channels, except for a small violation of the roll model-matching con-




Figure 1.17 – Disturbance rejection transfers $\mathbf{T}_{\mathbf{d} \rightarrow \mathbf{e}}$ (left) and $\mathbf{T}_{\mathbf{d} \rightarrow \mathbf{u}_c}$ (right) *The roll (green), pitch (blue) and yaw (red) channels are compared to the templates (dashed for roll and pitch, dash-dotted for yaw); frequency is in rad/s*




Figure 1.18 – Reference tracking transfers $\mathbf{T}_{\mathbf{r} \rightarrow \mathbf{e}_{mm}}$ (left) and $\mathbf{T}_{\mathbf{r} \rightarrow \mathbf{u}_c}$ (right) *same legend as above*

straint. Fig. 1.17 reveals that the disturbance rejection constraint is only active for the yaw channel at the edge of the desired closed-loop bandwidth. That behaviour is explained by the open-loop yaw resonance close to the desired bandwidth. The control signal is constrained for both roll and pitch axes around $\omega = 11$ rad/s. In fig. 1.18, the model-matching constraint is acting on the roll response at low frequencies up to 0.5 rad/s because of the integrator in roll. The yaw channel is constrained at the edge of its bandwidth due to the nearby open-loop resonance. At around 6 rad/s, all transfers show a peak with the roll channel slightly above its constraint. None of the actuator constraints are active in the reference tracking case. Additionally, the multivariable modulus margin is computed with the \mathcal{H}_∞ norm of the sensitivity at the output of the system [SP07]. The current synthesis gives a value of 0.66 that is above the common standard of 0.5.

1.4.2 Experimental verification of closed-loop performance

In order to assess the suitability of both modeling and control aspects, the controller structure described in the previous subsection has been implemented on the ACHILES.

The practical issues associated with controller discretization and computational delays were handled according to the guidelines of Strub et al. [Str16, pp. 111-112]. The control performance was assessed with respect to disturbance rejection and reference tracking objectives: the closed-loop model predictions were compared to the attitude measurements in order to verify its accuracy and confirm the validity of the decoupling assumptions.

Reference tracking test case




Figure 1.19 – Attitude tracking experiment *reference (dashed), simulation (blue) and measurements (red)*

The reference tracking performance is evaluated with a series of step inputs applied sequentially on roll, pitch and yaw channels. The amplitude and width of the steps are




Figure 1.20 – Actuator deflection during tracking *simulation (blue) and measurements (red)*

chosen so that no actuator saturation occurs on neither of the three axes.

On all three axes, the attitude measurements of Figure 1.19 exhibit a very good match with the model outputs although the measured response exhibit more overshoot than expected (3.0% instead of 0.5% in roll, 22% instead of 7.5% in pitch and 25% instead of 5.8 % in yaw). Cross-coupling effects are mostly noticed on the roll axis even though they are strongly attenuated by the feedback loop. Interestingly, on fig. 1.19, the pitching maneuver generates greater disturbances on the roll angle than the yawing maneuver. This is because the pitch and roll dynamics are much faster than the yaw dynamics due to the added yaw inertia of the gimbal frame.

The control signals corresponding to each axis are plotted on Figure 1.20. In pitch and

yaw, the actuator deflections conform well with the model predictions. However, there are discrepancies on the roll channel due to unmodeled cross-coupling effects that are rejected by the controller. The offset between predicted and actual steady-state values has different causes on the three axes. In pitch and yaw, it can be attributed to the limited accuracy of the trimming procedure. On the roll axis, the small radial offset between the CG and the rotation axis generates a disturbing moment that has to be compensated by the actuators.

Disturbance rejection test case




Figure 1.21 – Pitch (left) and yaw (right) disturbance rejection *disturbance (dashed), simulation (blue) and measurements (red)*

On both pitch and yaw channels, an artificial disturbance step is injected at the actuator input in order to emulate an attitude disturbance of known amplitude. The attitude response of the projectile is then recorded and compared to the simulated response of the closed-loop projectile model. Figure 1.21 shows that the measured response matches well with the simulated outputs. The small discrepancies in the yaw angle response reveal that closed-loop yaw damping is slightly lower than expected.

1.5 Conclusion

This chapter presented the legacy experimental setup and its evolution: the mechanical redesign of the gimbal structure made the roll axis functional, allowing the entire attitude dynamics of the projectile to be excited and controlled. Later on, the configuration of the mock-up was adapted to promote similarities with an innovative long-range

guided projectile concept and provide the capability to adjust one of its critical geometric parameters over a large range of motion.

The attitude dynamics model based on the non-linear equations of motion was declined to the 3 DoF case, then simplified to focus on the most significant dynamics and bring the number of estimated parameters down to 13 (including the noise model). Data collection experiments were performed in accordance with the decoupling assumptions: the identified model was found to adequately capture the roll, pitch and yaw dynamics of the mock-up while the parameter estimates were consistent with the a-priori knowledge of the projectile aerodynamics and its symmetry properties. The structure of the skid-to-turn autopilot was updated with an additional roll channel and a new procedure was introduced to determine the frequency templates of the weighting filters. The frequency-domain analysis confirmed that the closed-loop system could provide adequate performance and stability without exceeding the actuator limitations.

Finally, the performance of the attitude controller was assessed through reference tracking and disturbance rejection experiments. Local discrepancies between the measured data and the model predictions were observed on the roll actuator channel and attributed to unmodeled cross-coupling effects. However, their effect on the attitude time histories is minimal, showing that such disturbances are well rejected by the controller. Overall, the reference tracking and disturbance rejection performance of the autopilot matches very well with the model predictions, proving that the controller design is satisfactory and the HIL methodology is sound.

PARAMETER-VARYING MODELS OF THE PROJECTILE DYNAMICS

Generally speaking, the purpose of a model is to capture the input-output behavior of a system, process or transformation. A model can typically be described by its parameters and its structure: the former are physical or abstract quantities that appear in the computation of the model output, whereas the latter defines the relationship between the parameters and the input or output variables.

The nature of the model structure can be very diverse, ranging from a very flexible array of tunable parameters (i.e neural networks) to a set of explicit mathematical equations from which all terms are known a-priori. In the former case, the model could be seen as a "black box" where parameters have little physical meaning and can be varied freely to match the system dynamics. In the latter case, the model is said to be a "white box" which is used to depict a well-known physical phenomenon that dictates the sign and value of each parameter. For this study, "grey box" model structures are used to describe the projectile dynamics. It means that, while some parts of the model structure and parameters are prescribed by theoretical laws such as rigid-body mechanics, some relations can be arbitrarily defined where coefficients are tuned in order to fit input-output sequences measured on the actual system.

This chapter presents the modeling effort in the context of airframe-controller optimization for agile projectiles and proposes a supplement to the attitude model described in the previous chapter. It has been demonstrated that the angular dynamics of the ACHILES projectile could be assumed to be decoupled with negligible impact on the controller performance during the validation runs. Since the flight scenario considered in the next chapter focuses on manoeuvres in the pitch plane, the dimension of the study has been reduced from three to one degree-of-freedom. This allows some mechanical simplification of the experimental setup which can then collect data over a greater range of geometric parameters to fully take advantage of airframe-controller optimization. Also,

since the design of a gain-scheduled autopilot adapted to variable-air-speed operation has already been presented by Strub [Str+18], the wind speed is fixed to 30 m/s for the rest of the study.

The first section states the expected purposes of the model-building process as well as the various constraints and technical limitations that must be considered. The second part describes the protocols that have been adopted in order to gather experimental data and characterize the projectile pitch dynamics. More precisely, the aim is to determine the values of the unknown parameters which provide the best match between the model output and the recorded time histories. The third section presents a Linear-Parameter-Varying (LPV) model of the pitch dynamics while the fourth one justifies the need for non-linear structures. The parts five to seven describe the build-up of the non-linear model structures, starting from the very generic fluid mechanics equations down to the engineering methods implemented into the MATLAB code. Finally, the last section benchmarks the model structures against experimental data and checks their behavior for a wide range of stable and unstable projectile configurations. The suitability of the LPV model for controller synthesis is assessed and a non-linear structure is chosen for the validation of the closed-loop projectile.

2.1 Introduction to Modeling

2.1.1 Model purposes and requirements

Purposes

Given that models are only simplified representations of the actual physics, their interest should be assessed in the light of their fitness for their intended purpose rather than their absolute fidelity to the observed phenomena. Here, the end goal is to perform concurrent optimization of the guided projectile geometry and controller gains. Hence, the closed-loop behavior of the airframe-controller design conditions the success of the modeling and optimization framework. As a result, a degree of error in the prediction of the projectile open-loop dynamics is acceptable as long as the closed-loop system demonstrates adequate performance and stability characteristics.

However, in order to improve their agility, it would be desirable to obtain guided projectile designs that perform well over a wide range of angles of attack. To do so, the model structure must be able to capture the non-linearities that arise when operating

at the edge of the projectile flight envelope. Then, in the ideal case, the controller can be tailored to compensate for the change in open-loop dynamics. If these phenomena cannot be fully captured by the synthesis framework, a higher fidelity model should still be developed to assess the suitability of the controller design with respect to the previously overlooked phenomena.

Requirements

As the projectile geometry is to be part of the guided projectile design space, the model structure must implement the dependency of the pitch dynamics to the geometric parameters of the projectile configuration. This dependency is also most likely to be non-linear as both aerodynamic forces and inertia properties are usually not proportional to the dimensions of the lifting surfaces.

Finally, the formulation of the model has to comply with the limitations of the currently available synthesis and optimization methods. In particular, in order to use most of these tools, one should provide a model structure which is linear with respect to its states, inputs and outputs. There is also a constraint on the parametric dependency that can only be expressed using simple mathematical functions.

2.1.2 Modeling strategy

For the first step which is airframe-controller design and optimization, a linear-parameter-varying (LPV) model is developed for use in the controller synthesis framework. The second step consists in assessing the closed-loop response of the tuned airframe-controller design while including the effects of canard stall. To do so, a non-linear model must be used to more closely predict real-world performance and pinpoint possible stability issues.

In case the response of the non-linear system proves to be unsatisfactory, the linear synthesis model can be adapted by introducing some degree of uncertainty into the model coefficients. This solution can be combined with an adjustment of the frequency templates in the synthesis framework to improve the robustness of the tuned controller and mitigate the effects of canard stall. Validation of the closed-loop non-linear model is provided by wind tunnel experiments covered in the next chapter.

Finally, the flight envelope of the guided projectile can be estimated. This domain is defined as the range of pitch angles that can be reached with satisfactory transient behavior (e.g low settling time, small overshoot, well-damped oscillations). The amplitude

of the flight envelope, given by a metric named the *maximum control AoA* (c.f. section 3.3.4), can be seen as a figure of merit for comparing airframe-controller designs and assessing the benefit of the co-design methodology.

2.2 Experimental Data Collection

2.2.1 Trim map recording

In aviation, the trimming process consists in adjusting the pilot inputs (control surfaces deflection and power) so that the aircraft current flight parameters (altitude, heading and airspeed) can be maintained without further controls actions. This implies that forces and moments applied on the airframe are balancing each other and the system is at an equilibrium condition. In our case, since the degrees of freedom of the ACHILES projectile are restricted to its pitching motion, the set of equilibrium parameters consists of the pitch attitude θ and the virtual pitch control deflection δ_m . For a stable projectile flying at a constant airspeed and altitude, the trim attitude corresponding to a given control deflection is unique.

Hence, the trim map recording procedure consists of a sweep of the δ_m parameter space, starting from -16° to $+16^\circ$ in 1° increments. The sequence starts at zero up to the maximum, then down to the minimum and up again back to zero: this way, all trim points are visited twice (once in each direction) and the consequences of any hysteresis can be averaged. The step duration is adjusted according to the damping of the open-loop dynamics so that the pitch attitude transients can be eliminated in post-processing. Also, since the projectile admits an horizontal symmetry, it is expected that the trim map would be symmetric with respect to the origin and would include this point. These properties are leveraged to mitigate any offset due to a miscalibration of the canards actuators.

Figure 2.1 shows that the aspect of the trim map depends on the projectile geometry: with its high canards aspect ratio, the ACO configuration exhibits a strong non-linearity in its steady-state control response, the effectiveness of its control surfaces being significantly reduced for deflection amplitudes larger than 5° . The linear slope is the result of the competition between the static stability and the control authority of the projectile. Thus, around the origin, the trim map slope can be compared to the ratio of the identified aerodynamic derivatives Cm_{δ_m} over Cm_α .

The recorded trim maps serve multiple purposes: since they are the only datasets

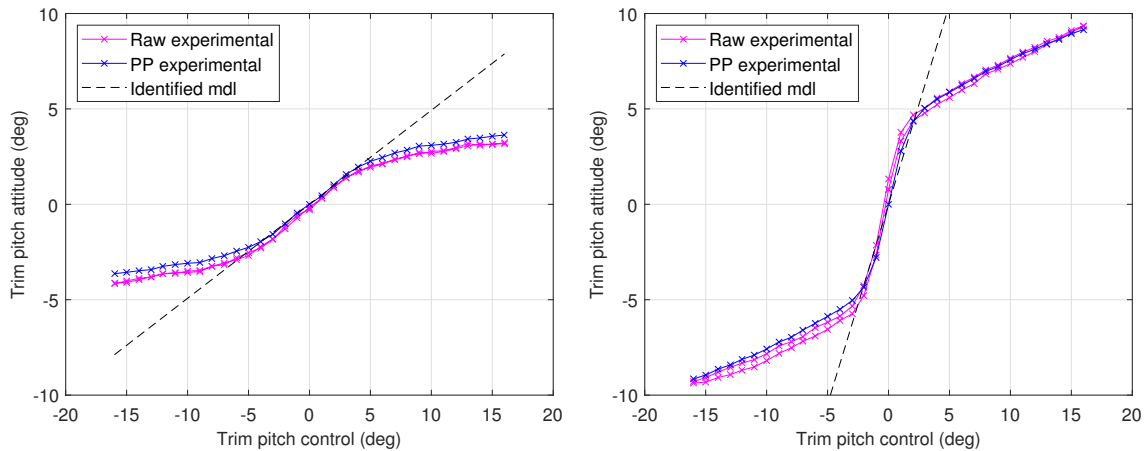


Figure 2.1 – Trim maps of the ACO projectile for different values of the geometric parameter: $x_F = 0.313$ m (left) and $x_F = 0.175$ m (right)

which capture the non-linear behavior of the projectile at large control deflections, they are a crucial input for all of the non-linear model structures (c.f. section 2.7) and can also be used to assess the fidelity of the model candidates. One major drawback of this data source is that trim maps can only be collected for configurations which exhibit a sufficient degree of open-loop stability. Else, the steady-state response requires too much time to be obtained or is excessively sensitive to process noise. Note that, theoretically, the trim map of a neutrally stable projectile (c.f. section 3.3) collapses to the Y-axis while unstable configurations do not admit any trim map at all.

2.2.2 System identification of the variable-stability projectile

Although some of the aerodynamic derivatives of the projectile configuration may be evaluated analytically (c.f. section 2.5.6), experimental work may provide more comprehensive and accurate results when available. However, since the projectile dynamics now depend on a geometric parameter, the system identification procedure introduced in section 1.3 needs to be updated to support the unstable or marginally stable configurations that may result from some regions of the parameter space.

Open-loop identification

In case where the configuration possesses a sufficient degree of open-loop stability, the existing methodology presented in section 1.3 does not need to be extensively modified.

To account for the restriction of the mock-up degrees of freedom, the model structure is adapted to the pitch-only case by restricting the state and output vectors to $\mathbf{x} = [\theta, q]^T$ and $Vy = \theta$. Similarly, the input vector Vu is limited to the virtual pitch control deflection δ_m , so that the new model is SISO. The parameters to be estimated are the aerodynamic derivatives Cm_α , Cm_q and Cm_{δ_m} , and the noise model terms $K_{\dot{\theta}\theta}$ and $K_{\dot{q}\theta}$.

The state-space form of this model is described by the equations (1.33) and (1.34) with the following matrices:

$$\mathbf{A} = \begin{bmatrix} 0 & 1 \\ A_{m\theta} & A_{mq} \end{bmatrix} \quad (2.1)$$

$$\mathbf{B} = \begin{bmatrix} 0 \\ B_{m\delta_m} \end{bmatrix} \quad (2.2)$$

$$\mathbf{C} = [1 \ 0] \quad (2.3)$$

$$\mathbf{D} = 0 \quad (2.4)$$

$$\mathbf{K} = \begin{bmatrix} K_{\dot{\theta}\theta} \\ K_{\dot{q}\theta} \end{bmatrix} \quad (2.5)$$

Rationale for closed-loop identification

When the stability and damping of the projectile decrease, the attitude transients become increasingly prominent in terms of amplitude and duration. These excursions do not agree with the assumption of small perturbations around $\theta = 0^\circ$ on which the identification model was built. Moreover, the prolonged transients require greater experiment run times which are impractical given that the battery life of the projectile mock-up is limited. As a result, the open-loop projectile dynamics cannot be identified “as it is” and a feedback loop must be implemented to keep the attitude in check and reduce the transients duration.

To explain some important concepts related to closed-loop identification, let Figure 2.2 represent our system driven by the reference signal r and composed of the plant G and the controller K . The controller output u is affected by the disturbance d while the system output y is corrupted by the measurement noise n . The transfers between the




Figure 2.2 – Example of a closed-loop system

different inputs and outputs are as follows:

$$u = \frac{K}{1 + KG}r - \frac{KG}{1 + KG}d - \frac{K}{1 + KG}n \quad (2.6)$$

$$y = \frac{KG}{1 + KG}r + \frac{G}{1 + KG}d + \frac{1}{1 + KG}n \quad (2.7)$$

Choice of the identification approach

Two families of closed-loop identification methods have been developed [Lju86]: the direct identification consists in using the input-output data of the plant without considering the effect of the feedback loop. This approach is the simplest as the estimate of the plant is directly found, however it breaks one key assumption of the estimation algorithms: due to the feedback, the output noise is no longer uncorrelated from the plant input. As a result, the prediction error minimization scheme is not guaranteed to produce unbiased estimates of the model parameters [Lju97]. Therefore, the bias should be quantified by the means of an alternative method, such as Monte Carlo simulations, to ensure the quality of the parameter estimates.

The indirect identification determines the transfer function of the closed-loop system. Information about the plant dynamics is retrieved in a second step, assuming the exact knowledge of the controller structure. In that case, the input data is either the reference or the disturbance signal of the closed-loop, and the uncorrelated noise assumption is fulfilled, leading to theoretical guaranties on the unbiased nature of the parameters estimates. Because of this property, the indirect approach has been retained for the identification of the projectile pitch dynamics.

Selection of the I/O signals and controller structure

In order to maximize the quality of the identified model, close attention should be paid to the relationship between the controller performance and the effect of output noise on the plant estimates. The transfers (2.6) and (2.7) can be expressed in the frequency domain to find the relationships between the Fourier transforms of the various signals.

Consider injecting the excitation signal on d and recording the system output y . By setting r to zero in (2.7) we have:

$$y = \frac{G}{1 + KG}d + \frac{1}{1 + KG}n \quad (2.8)$$

Note $\hat{T}_{d \rightarrow y}$ the estimate of the transfer $T_{d \rightarrow y} = \frac{G}{1 + KG}$ from d to y . The estimate of G is given by:

$$\hat{G} = \left(\hat{T}_{d \rightarrow y}^{-1} - K \right)^{-1} \quad (2.9)$$

Express the impact of measurement noise on $\hat{T}_{d \rightarrow y}$:

$$\hat{T}_{d \rightarrow y} = \frac{Y}{D} = \frac{G}{1 + KG} + \frac{1}{1 + KG} \frac{N}{D} = T_{d \rightarrow y} + \frac{1}{1 + KG} \frac{N}{D} \quad (2.10)$$

Substituting into (2.9):

$$\hat{G} = \frac{G + N/D}{1 - KN/D} = G + \frac{(1 + GK)N/D}{1 - KN/D} \quad (2.11)$$

Assume that the noise term is much smaller than the excitation signal ($KN \ll D$), so that:

$$\hat{G} = G + (1 + KG) \frac{N}{D} \quad (2.12)$$

Thus, at at frequency ω , the estimation of G would be the most accurate if:

$$N(\omega) \ll \frac{G}{1 + KG} D(\omega) \quad (2.13)$$

Following the same process, if the excitation signal is injected on the reference r , the output equation deduced from (2.7) is:

$$y = \frac{GK}{1 + KG}r + \frac{K}{1 + KG}n \quad (2.14)$$

With the same notations as before, the impact of measurement noise on $\hat{T}_{r \rightarrow y}$ is:

$$\hat{T}_{r \rightarrow y} = \frac{Y}{R} = \frac{KG + N/D}{1 + KG} \quad (2.15)$$

Substituting into (2.9) and assuming that the noise term is much smaller than the excitation signal ($KN \ll D$) gives:

$$\hat{G} = G + (K^{-1} + G) \frac{N}{D} \quad (2.16)$$

Thus, at at frequency ω , the estimation of G would be the most accurate if:

$$N(\omega) \ll \frac{KG}{1 + KG} R(\omega) \quad (2.17)$$

By comparing the conditions given by Equations (2.13) and (2.17), it can be noted that the excitation on d will give best results for lower controller gains and/or narrower bandwidth while the excitation on r favors larger gains and/or greater controller bandwidth.




Figure 2.3 – Layout of the identification controller

The starting point for the controller structure was a classical PID layout. For practical reasons related to controller wind-up during the wind tunnel spool-up phase, the integral action has been excluded from the structure. Thus, the identification controller presented on fig. 2.3 has been chosen as a simple proportional-derivative that requires high gains for fast control and low steady-state error. In accordance with the conclusions of the previous paragraph, the reference r has been selected to inject the excitation signal into the closed-loop system.

Controller tuning and model fit

A simplified \mathcal{H}_∞ synthesis scheme based on the output sensitivity function has been applied to a pair of projectile configurations suffering from open-loop stability issues. For both configurations the frequency template has been set up for a bandwidth of 2 rad/s, a low frequency gain of 0.3 and a high frequency gain of 1.3. At the actuator input, a gain margin of 9 dB and a phase margin of 45° have been specified, whereas at the plant output the gain margin requirement is 10 dB and the phase margin is set to 45° . The results of these syntheses are $k_{\theta ID} = 1.54$ and $k_{q ID} = 0.44$ for the most unstable configuration, versus $k_{\theta ID} = 1.26$ and $k_{q ID} = 0.48$ for the second most unstable projectile.




Figure 2.4 – Reference tracking of the most unstable projectile fitted with the identification controller

Figure 2.4 displays the identified model response on validation data along with the corresponding pitch attitude measurements. The dynamics of the closed-loop system are accurately reproduced and the steady-state error is well predicted. The identification controller provides a reasonably fast response with moderate error which satisfies the experimental constraints.

Parameter estimation results

The parameter estimates are found with the same algorithm as in section 1.3. The bootstrap methodology for uncertainty evaluation is also re-used. Results of both open-

and closed-loop identification campaigns are combined into data series covering the full range of fins position x_F allowed by the mock-up, as pictured in Figure 2.5.




Figure 2.5 – Estimates of Cm_α (left), Cm_q (center) and Cm_{δ_m} (right). *The two leftmost data points are from closed-loop identification.*

The closed-loop estimates of Cm_α line up very well with the trend observed on the open-loop data. Since the magnitude of this coefficient is small, so are the error bars. The estimation of Cm_q has proven to be more challenging due to convergence issues, thus one of the data points appears to be an outlier. As there is another outlier among the results obtained with the open-loop identification procedure, the cause of this issue may be related to the experimental setup and not the identification procedure. As for the Cm_{δ_m} , the worst-case uncertainty on the closed-loop parameter estimates is much greater than for the open-loop results and there may be a small offset between both datasets.

2.3 Linear-Parameter-Variable Model of the Projectile Pitch Dynamics

2.3.1 State-space structure

Given that the LPV pitch dynamics model shares the same purpose with the linear attitude model presented in section 1.2, its structure is very similar with a state-space form and estimated aerodynamic derivatives. The non-linear equation of motion (2.18) is derived under the same assumptions as in the 3 DoF case, except that the roll bearing and the yaw gimbal are locked and the projectile motion is restricted to pitch only.

$$I_{yy}(x_F) \ddot{\theta} = \bar{q}SD \left(Cm_\alpha(x_F) \theta + Cm_q(x_F) \frac{qD}{v} + Cm_{\delta_m}(x_F) \delta_m \right) \quad (2.18)$$

The definition of the state-space form is straightforward: the state vector x is composed of the pitch angle θ and the pitch rate q while the input is the virtual pitch control deflection δ_m and the output is the pitch attitude θ . As a result, the state and input matrices are given by:

$$A = \begin{bmatrix} 0 & 1 \\ A_{m\theta} & A_{mq} \end{bmatrix} \quad B = \begin{bmatrix} 0 \\ B_{m\delta_m} \end{bmatrix}$$

$$A_{m\theta} = \bar{q}Sd \frac{C_{m_\alpha}(x_F)}{I_{yy}(x_F)} \quad A_{mq} = \frac{\bar{q}Sd^2}{\bar{v}} \frac{C_{m_q}(x_F)}{I_{yy}(x_F)} \quad B_{m\delta_m} = \bar{q}Sd \frac{C_{m_{\delta_m}}(x_F)}{I_{yy}(x_F)}$$

The transfer function from δ_m to θ is then:

$$G_{\text{dynmdl}} = \frac{B_{m\delta_m}}{s^2 + (1 - A_{mq})s - A_{m\theta}} \quad (2.19)$$

2.3.2 Surrogate models of the aerodynamic and inertia properties

In order to enable airframe-controller design, the synthesis model must implement an explicit dependency to the projectile geometric parameters. This is achieved by expressing the aerodynamic coefficients C_{m_α} , C_{m_q} , $C_{m_{\delta_m}}$ and the moment of inertia I_{yy} as polynomial functions of the fins longitudinal position x_F . Previous work [Den+17] has shown that the aerodynamic coefficients associated to control surfaces could be tabulated as piecewise affine functions of Mach number. However, to the best of the author's knowledge, modeling of these coefficients as continuous functions of the airframe geometric parameters has not yet been attempted.

Aerodynamic coefficients

The motivation for expressing the aerodynamic coefficients as polynomial functions of the geometric parameter stems from the limitations of the software tools used to solve the synthesis problem. In fact, the *systeme* function is currently restricted to model terms defined as rational functions of tunable parameters [GA11].

The order of the polynomials should be as high as required to accurately capture the evolution of the coefficients as the parametric space is swept, but should also be as small as possible to avoid over-fitting if the number of experimental data points is

limited. Here, an original procedure based on a-priori knowledge of the projectile flight dynamics is used to determine the model order for each coefficient. The process leverages on analytical formulas commonly used in aircraft design, as well as experimental data. Under the assumptions that the projectile AoA remains small and the contribution of the canards and body to its pitch damping are negligible, the following conclusions can be inferred:

- The static stability coefficient Cm_α is given by the product of the fins lift curve slope $C_{L,\alpha,F}$ by the adimensional lever arm $(x_F - x_G)/D$. Thus, as $C_{L,\alpha,F}$ is constant at low AoA, Cm_α is an affine function of x_F
- As proposed by Cook [Coo12, pp. 377-380], the pitch damping coefficient Cm_q can be estimated from the damping moment generated by the tailplane incidence α_T induced by pitch rate (c.f. fig. 2.6). In our case the fins act as a tailplane and Cm_q can be approximated by Equation (2.20), which results in a second order polynomial in x_F .

$$Cm_q = -\frac{(x_F - x_G)^2 S_F}{D^2 S} C_{L,\alpha,F} \quad (2.20)$$

- The pitch control coefficient Cm_{δ_m} is mostly influenced by the airflow over the control canards located at the front of the projectile. Thus one may assume that it is independent from the position of the fins which are located further downstream. However, as the flow is subsonic, the pressure field around the fins extends upstream and may influence the canards. Experimental data of fig. 2.7 confirms this assumption and shows that the pitching moment induced by the deflection of the canards is reduced as the fins are in the vicinity of the canards. Thus, a first order polynomial expression has been retained for $Cm_{\delta_m} = f(x_F)$.

To sum up, the polynomial models of the aerodynamic coefficients are described by the following set of equations:

$$\begin{cases} Cm_\alpha(x_F) = P_{\alpha_0} + P_{\alpha_1} x_F \\ Cm_{\delta_m}(x_F) = P_{dm_0} + P_{dm_1} x_F \\ Cm_q(x_F) = P_{q_0} + P_{q_1} x_F + P_{q_2} x_F^2 \end{cases} \quad (2.21)$$

This a-priori analysis is confirmed by the trend of the parameter estimates obtained in section 2.2.2. Figure 2.7 show the evolution of the aerodynamic coefficients for fins positions ranging from forward to aft as well as the fitted polynomial models. The fit percentages based on normalized RMS errors are respectively 87.9% for Cm_α , 77.3% for




Figure 2.6 – Tailplane incidence due to pitch rate [Coo12]

Cm_q and 79.7% for Cm_{δ_m} indicating that the surrogate models fit well with the identified values of the aerodynamic coefficients.




Figure 2.7 – Surrogate models of Cm_α (left), Cm_q (center) and Cm_{δ_m} (right)

Moment of inertia

Because of large range of fins positions available on the mock-up, the pitch moment of inertia I_{yy} of the projectile is affected by the position of the fins. The parallel axis theorem can be applied to the finned sleeve and the body+canards assembly to express I_{yy} as a function of the MoI and CG of the canards and body (subscript B), the mass properties of the adjustable finned sleeve (subscript F) and the CG position of the whole projectile x_{GS} :

$$I_{yy} = I_{yy,B} + m_B \cdot (x_{G,B} - x_G)^2 + I_{yy,F} + m_F \cdot (x_{G,F} - x_G)^2 \quad (2.22)$$

Consequently, the moment of inertia of the projectile around its pitch axis has been represented by a second order polynomial in x_F given in Equation (2.23).

$$I_{yy}(x_F) = P_{iy_0} + P_{iy_1} x_F + P_{iy_2} x_F^2 \quad (2.23)$$

2.3.3 Minimum-size linear fractional representation

Once its coefficients have been expressed as polynomial functions of the tunable parameter, the model may be cast as a Linear Fractional Representation (LFR) composed of a linear time-invariant plant M interconnected with a diagonal matrix Δ containing only the varying parameter x_F .




Figure 2.8 – Linear Fractional Representation of the parameter-dependent plant

The size of Δ should be minimized in order to improve the performance of the synthesis routine and reduce the conservativeness of its results [RHB14]. Using default MATLAB tools gives an LFR of size 14 which seems excessively large given that the geometric parameter only appears 6 times in total: one in the affine Cm_α model, one in the Cm_{δ_m} model, two in the second-order Cm_q model and two in the I_{yy} model. In order to reduce the size of the LFR, the set of equations describing the pitch dynamics has been cast into the block diagram of Figure 2.9 in such a way that the number of repetitions of x_F is minimal. Then, each occurrence of the geometric parameter is replaced by an output z_i and an input v_i . Finally, the M and Δ matrices are respectively populated by the polynomial coefficients and the geometric parameter itself.

The set of equations describing the pitch dynamics model is given by (2.24). Since the state derivative $\ddot{\theta}$ is a function of the inverse of the moment of inertia, I_{yy}^{-1} has been

modeled to avoid introducing rational coefficients in the M matrix.

$$\begin{cases} I_{yy}(x_F) \ddot{\theta} = \bar{q}SD \left(Cm_{\alpha}(x_F) \theta + Cm_q(x_F) \frac{qD}{v} + Cm_{\delta_m}(x_F) \delta_m \right) \\ Cm_{\alpha}(x_F) = P\alpha_0 + P\alpha_1 x_F \\ Cm_{\delta_m}(x_F) = Pdm_0 + Pdm_1 x_F \\ Cm_q(x_F) = Pq_0 + Pq_1 x_F + Pq_2 x_F^2 \\ I_{yy}(x_F)^{-1} = Piyi_0 + Piyi_1 x_F + Piyi_2 x_F^2 \end{cases} \quad (2.24)$$

The block diagram of fig. 2.9 shows the occurrences of x_F and the associated inputs v_i and outputs z_i . The Δ matrix is simply $\mathbb{1}_6 \cdot x_F$ while M is given by Equation 2.8.

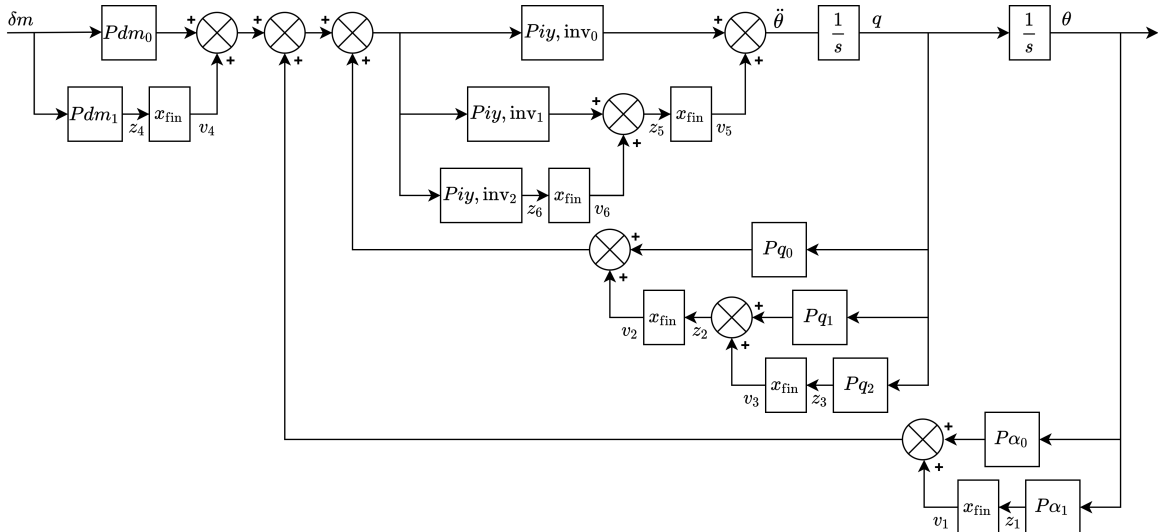


Figure 2.9 – Block diagram of the pitch dynamics model

$$\begin{bmatrix} \ddot{x} \\ \dot{x} \\ z_1 \\ z_2 \\ z_3 \\ z_4 \\ z_5 \\ z_6 \\ y \end{bmatrix} = \begin{bmatrix} Piyi_0Pq_0 & Piyi_0P\alpha_0 & Piyi_0 & Piyi_0 & 0 & Piyi_0 & 1 & 0 & Piyi_0P\delta m_0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & P\alpha_1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ Pq_1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ Pq_2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & P\delta m_1 \\ Piyi_1Pq_0 & Piyi_1P\alpha_0 & Piyi_1 & Piyi_1 & 0 & Piyi_1 & 0 & 1 & Piyi_1P\delta m_0 \\ Piyi_2Pq_0 & Piyi_2P\alpha_0 & Piyi_2 & Piyi_2 & 0 & Piyi_2 & 0 & 0 & Piyi_2P\delta m_0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \dot{x} \\ x \\ v_1 \\ v_2 \\ v_3 \\ v_4 \\ v_5 \\ v_6 \\ u \end{bmatrix} \quad (2.25)$$

Thanks to this process, the number of occurrences of x_F has been reduced from 14 to 6. Favorable impact on the controller synthesis process is to be expected but has not been accurately quantified.

2.4 Introduction to Non-Linear Aerodynamics Modeling

2.4.1 Rationale and means

The previously-mentioned LPV models include a linear description of the projectile aerodynamics by the means of the aerodynamic derivatives Cm_α , Cm_q and Cm_{δ_m} which are constant for a given value of the geometric parameter. However, if the local angle of attack of a canard increases beyond a given value, the flow on the upper surface of the plate separates, creating a recirculation area that leads to a net loss of lift [And17, p 329]. This phenomenon named canard stall implies that the generated aerodynamic moment is no longer proportional to the canards deflection angles and the model structure with constant coefficients becomes inaccurate.

Different solutions to this issue can be investigated: The value of the aerodynamic derivatives could be varied according to the projectile angle of attack. This choice has the advantage of turning the model into a quasi-LPV form for which robust control methods may offer theoretical stability guarantees [Tôt10]. However, most LPV models are built from look-up tables and require a comprehensive database of aerodynamic derivatives

that requires extensive CFD simulations or numerous wind tunnel experiments. Moreover, scheduling the model with respect to some of its fast-varying states comes with a number of shortcomings [Mac19]. Alternatively, non-linear models can reduce the volume of required data by implementing semi-empirical correlations based on experimental observations or use analytical formulas derived from fluid mechanics theory. The drawback is that they are much more challenging to use for controller synthesis since fewer tools are available for this type of structure. In this study, the main purpose of the higher-fidelity model (c.f. section 2.1.1) is simulation and controller validation, thus a non-linear form has been retained. Three different model structures have been developed and are described in section 2.7.

2.4.2 Fundamentals of the projectile pitch dynamics

In order to emulate the projectile pitching motion, three effects should be quantified: The pitch static stability moment generated in response to an angle of attack disturbance, the pitch damping moment generated in response to a non-zero pitch rate and the pitch control moment generated by a change of the canards deflection angles.

If the projectile configuration is stable and the controls deflection is non-zero, the control moment will tend to push the projectile away from its equilibrium position whereas the stability moment will try to return it to zero AoA. This competition results in the emergence of a new equilibrium position which is shifted away from the origin. As both stability and control moments depend on the trim AoA and controls deflections, they can be quantified by steady-state i.e. time-invariant measurements or simulations. One can either attempt to evaluate each effect independently or focus on the outcome of the competition between the two. The pitch damping coefficient, however, requires an off-trim condition in order for the pitch rate to build up, leading to a counteracting damping moment. Thus, it is generally more difficult to obtain as it requires transient measurements or unsteady simulations.

2.5 Theoretical Groundwork for Analytical Aerodynamics Predictions

The objective of this section is to review existing theories that could be able to predict the aerodynamic forces acting on the projectile. However, since the projectile is a combination of slender body and lifting surfaces, analytical models may only be applied at the

component level and mutual interactions need to be accounted for.

Under some fundamental assumptions, forces around simple geometrical shapes may be directly predicted by the potential flow theory without the need for discretization of the flow field or the geometry. This approach has the advantage of providing analytical expressions that are explicit functions of the projectile geometric parameters, including the ones that may be tuned as part of the co-design process. Such knowledge could also be leveraged to postulate more complex grey-box model structures in case the accuracy of the theoretical estimates prove to be insufficient.

That is why this section is not intended to provide a comprehensive derivation of the analytical results, but instead focuses on the fundamental principles and assumptions leading to the formulas. It must also be kept in mind that the results given by these theories may not be directly applicable to the ACHILES II projectile as some of the underlying assumptions may be violated. The theoretical results presented herein provide the foundation of the engineering methods presented in section 2.6 and implemented into the model structures described in section 2.7.

2.5.1 Simplification of the fluid mechanics equations

Starting from a very generic set of equations describing the fluid dynamics, a number of simplifications consistent with the scope of the study are made in order to cast the equations into a problem that admits analytic solutions. The theory of fluid mechanics is based on three fundamental principles applied to an infinitesimal fluid element [And17, p 106]:

- the conservation of mass, giving rise to the continuity equation
- Newton’s second law, cast into the momentum equation
- the conservation of energy, leading to the energy equation

One of the most comprehensive description of the behavior of fluids can be obtained thanks to the Navier-Stokes equations. They directly translate the above-mentioned principles with the only addition of a behavior law for the fluid viscosity, which is assumed to remain constant no matter the amount of shear stress it is subjected to. As this is the case with air, the Navier-Stokes equations can be considered for atmospheric flight.¹

Moreover, as the scope of this modeling effort is to estimate the lift forces applied on a projectile flying at low subsonic speeds, the viscous effects and the heat transfer inside

1. For the sake of brevity, the equations will not be described herein but can be found in most aerodynamics coursebooks such as [And17].

the flow field can be neglected. The rationale behind that choice is that viscous effects mostly affect drag predictions instead of lift [HC03, p 330] and significant heat transfer usually occurs at high Mach number i.e for supersonic and hypersonic flow conditions.

Additionally, as the HIL approach relies heavily on experimental data, no attempt is made to obtain aerodynamic models which remain valid far away from the freestream conditions achievable by the wind tunnel. Thus, the freestream velocity w.r.t the speed of sound is restricted to low Mach number and the flow can be considered incompressible.

Finally, we can assume that, at the initial conditions, the flow stream is free of any vorticity. This can be justified by the design of the wind tunnel which tends to minimize such phenomena. With that, and all of the above assumptions, the flow can be considered irrotational and thus can be studied in the frame of the potential flow theory.

2.5.2 Potential flow theory

Let $\mathbf{V} = [u \ v \ w]^T$ be the velocity vector of the fluid particle considered, the irrotationality condition means that the curl of \mathbf{V} is zero anywhere in the fluid domain. Thus, there exist a scalar function ϕ named the velocity potential function that is solution of the equation (2.26).

$$\mathbf{V} = \nabla\phi \tag{2.26}$$

Similarly, as the flow is incompressible, the volume of the fluid element is constant and the divergence of \mathbf{V} is zero anywhere in the fluid domain. Thus, there exist a vectorial function ψ named the stream function that verifies the equation (2.27).

$$\mathbf{V} = \nabla \times \psi \tag{2.27}$$

The magnitude of the stream function is constant along any streamline, which is defined as a curve passing through a given point $M(x, y)$ and tangent to the velocity vector of each fluid particle in the domain. In a given plane where z is constant, both the velocity potential and the magnitude of the stream function are connected to the local velocity field by the following relations

$$u = \frac{\partial\phi}{\partial x} \quad v = \frac{\partial\phi}{\partial y} \tag{2.28}$$

$$u = \frac{\partial\psi}{\partial y} \quad v = -\frac{\partial\psi}{\partial x} \tag{2.29}$$

By combining the expression of the velocity potential with the incompressibility con-

dition, ϕ can be found as a solution of the Laplace's equation (2.30).

$$\nabla^2 \phi = 0 \tag{2.30}$$

This second order partial differential equation is linear and admits well known harmonic functions as solutions [And17, p 238]. Thus, the principle of superposition can be used to decompose the flow features of a complex problem into a combination of elementary flows that come with simple expressions for potential functions. However, those solutions are only available for canonical planar flows, so the real-life tridimensional problem should be reduced into a two-dimensional problem while capturing the dominant flow features. The assumption on the orientation of the plane containing the dominant features leads to a split between various developments of the potential flow theory that will be detailed herein.

Elementary flows

Before attempting to reduce the dimension of our real-life problem, let us devise a strategy for representing complex 2D flows. The flow field around an arbitrary shaped body may be replicated by the superposition of an uniform flow representing the freestream conditions and multiple elementary flows originating from discrete sources arranged across the plane.

- The uniform flow of velocity V_∞ oriented in the positive x direction induces a velocity potential $\phi(x, y) = V_\infty x$ and a stream function $\psi(x, y) = V_\infty y$. The orientation of the flow is consistent with the equation of the streamlines which is $y = \text{constant}$.
- The two-dimensional flow generated by a source (or sink) placed at the origin where the fluid is appearing (or disappearing) at a uniform rate of $m \text{ m}^2\text{s}^{-1}$ induces a velocity potential $\phi(r, \theta) = \frac{m}{2\pi} \ln \left(\frac{r}{r_0} \right)$ and a stream function $\psi(r, \theta) = \frac{m\theta}{2\pi}$. Thus, the streamlines are radial lines extending from the origin.
- If a source and a sink are placed at an infinitesimal distance from each other, the resulting singularity is called a doublet. It induces a stream function $\psi(r, \theta) = -\frac{\mu}{2\pi r} \sin \theta$ and a velocity potential $\phi(r, \theta) = \frac{\mu}{2\pi r} \cos \theta$.
- The two-dimensional flow around a vortex of circulation Γ located at the origin induces a velocity potential $\phi(r, \theta) = \frac{\Gamma}{2\pi} \theta$ and a stream function $\psi(r, \theta) = \frac{\Gamma}{2\pi} \ln \left(\frac{r}{r_0} \right)$. Here, the streamlines are concentric circles centered on the origin.

Application: cylinder in an uniform flow

Let's first study the flow field produced by a source set upstream of a sink of equal magnitude m at a respective distance c and $-c$ of the origin. The stream function induced by this combination can be shown to be equal to $\psi = -\frac{m}{2\pi}\beta$ which gives in cartesian coordinates $\psi(x, y) = \frac{m}{2\pi} \arctan\left(\frac{2cy}{x^2+y^2-c^2}\right)$. For a constant ψ_B , the streamline equation becomes:

$$x^2 + y^2 - 2c \cot\left(\frac{2\pi\psi}{m}\right)y - c^2 = 0 \tag{2.31}$$




Figure 2.10 – Geometry of a source upstream of a sink (reproduced from [HC03])




Figure 2.11 – Streamlines induced by a source-sink pair (reproduced from [HC03])

Equation (2.31) describes a set of circles of radius $r = c\sqrt{\cot^2(2\pi\psi/m) + 1}$ centered on the y-axis at the ordinate $y = \pm c \cot(2\pi\psi/m)$. Adding the stream function of the uniform flow according to the principle of superposition gives $\psi = \frac{m}{2\pi} \arctan\left(\frac{2cy}{x^2+y^2-c^2}\right) + V_\infty y$ which for $\psi = 0$ is the equation of an oval that is symmetrical about the horizontal and vertical axes.

As the streamlines are not permeable to the flow, they can be replaced by solid boundaries representing the shape of the flying body under the assumption that the flow remains irrotational, which has already been made previously. Thus, if we want to model the flow around a cylinder, it seems reasonable to think that the distance between the source and sink must be made infinitely small. This leads to a new type of singularity named a doublet, for which the potential and stream functions are known. By superposition, the stream function of a cylinder in an uniform flow is then $\psi(r, \theta) = -\frac{\mu}{2\pi r} \sin \theta + V_\infty r \sin \theta$. Noting that the streamline $\psi = 0$ is a circle of radius $a = \sqrt{\frac{\mu}{2\pi V_\infty}}$, the expressions of the stream function and the velocity potential can be simplified into equations (2.32) and (2.33).

$$\psi(r, \theta) = -V_\infty \sin \theta \left(\frac{a^2}{r} - r \right) \quad (2.32)$$

$$\phi(r, \theta) = V_\infty \cos \theta \left(r - \frac{a^2}{r} \right) \quad (2.33)$$

Complex potential and conformal mapping

An alternative to the superposition of singularities is to use a class of functions named conformal transforms to deduce the complex potential of the flow around a given geometry from the solution of the uniform flow around a circular cylinder. The complex potential W is defined as a complex-valued function of the velocity potential and the magnitude of the stream function according to eq. (2.34).

$$W = \phi + i\psi \quad (2.34)$$

In the previous example, substituting the expressions of ϕ and ψ gives:

$$W(r, \theta) = V_\infty r \left(1 - \frac{a^2}{r^2} \right) (\cos \theta + i \sin \theta) \quad (2.35)$$

Introducing the complex variable $z = r(\cos \theta + i \sin \theta)$, the complex velocity potential of a cylinder in a uniform flow is now:

$$W(z) = V_\infty z \left(1 - \frac{a^2}{|z|^2} \right) \quad (2.36)$$




Figure 2.12 – Circular mid-wing airframe section into a uniform cross flow (reproduced from [Nie88])

Here the conformal transform is used to find the velocity potential around circular mid-wing combination shown in Figure 2.12, which is of particular interest as it allows to study the mutual interactions between the wing and the body. Let z be the complex variable in the plane where the flow is to be found and σ be the complex variable in the plane where the flow is known. It is required that the transformation cause no distortion of the complex plane towards infinity i.e. $z(\sigma) \rightarrow \sigma$ at infinity. According to Nielsen [Nie88, p. 27], a suitable mapping is produced by the following function:

$$z(\sigma) = \frac{A}{4} \left(\frac{\sigma}{r_0} + \frac{r_0}{\sigma} \right) \pm \left[\left(\frac{A}{4} \right)^2 \left(\frac{\sigma}{r_0} + \frac{r_0}{\sigma} \right)^2 - a^2 \right]^{1/2} \quad (2.37)$$

Note that r_0 is the radius of the cylinder in the σ -plane while a is the radius of the body in the z -plane. $A = 2r_0 = s + \frac{a^2}{s}$ gives the relation between the semi-span of the wing-body and the other geometrical parameters. The velocity potential of the wing-body combination in an uniform cross-flow oriented positively along the imaginary axis is then given by the author [Nie88, p. 30] as a function of a, s and the freestream velocity V_∞ :

$$W(z) = -iV_\infty \left[\left(z + \frac{a^2}{z} \right)^2 - \left(s + \frac{a^2}{s} \right)^2 \right]^{1/2} \quad (2.38)$$

2.5.3 Slender body theory

The slender body theory, originally developed by Munk [Mun24] to predict aerodynamic forces on airships, assumes an elongated body of revolution which radius is small compared to its length. In our case, this theory may be applied to the cylindrical body of the projectile. Since the radial velocity of the flow in any cross-section is assumed to be zero, the flow features can be described in any radial plane containing the body's axis of revolution. The basic set of equations includes the Laplace equation in polar coordinates and both boundary conditions on the body surface and in the far field. The surface condition expresses that there is no in- or outflow through the body surface while the infinity condition prescribes that the perturbation velocities induced by the presence of the body decay to zero towards infinity.




Figure 2.13 – Principle of superposition: lifting (bottom) and non-lifting (top) problems (reproduced from [HC03])

The case of predicting the forces on a slender body flying at an angle of attack is decomposed into a non-lifting problem which aims at capturing the deflection of the streamlines due to the presence of a slender body aligned with the flow, and a lifting problem which aims at quantifying the asymmetrical effects due to the cross flow component (c.f. fig. 2.13). As seen in section 2.5.2, the axial flow around an oval can be represented by a distribution of sources and sinks along the axis of revolution of the body, while the cross-flow around the cylindrical section can be modeled by a distribution of doublets oriented perpendicular to the cross-flow and along the body axis.

Luneau and Bonnet [LB89] give an expression for the normal force acting on a slender body as a function of the angles of attack and sideslip. For a body with no boat tail




Figure 2.14 – Disposition of singularities for the lifting problem (reproduced from [HC03])

(the section of the base equals the largest cross section) at zero sideslip, the normal force coefficient is simply:

$$C_{N,B} = \sin(2\alpha) \quad (2.39)$$

The derivation of the integrals will not be detailed here but can be found in Karamcheti [Kar80]. However, a couple of significant assumptions have to be mentioned, namely the magnitude of the perturbation velocities is considered small w.r.t the freestream velocity and the change in body radius by unit length is much less than unity. Both of these hypotheses may not hold at both ends of the body because of stagnation points or a rounded nose, thus the local pressure distribution may be off and the total force may be affected.

2.5.4 Lifting line theory

The lifting line theory formulated by Prandtl [Pra25] predicts the lift distribution over a wing of finite span but large aspect ratio i.e the wing span is large in comparison to the mean chord of the wing. Even if the flow around each wing section is affected by the downwash of the wingtip vortices, it is assumed to remain two-dimensional in the sense that no cross flow occurs between neighboring sections.

This theory differs from the previous one in the sense that the singularities are now lines disposed in a plane instead of points aligned an axis and the flow field is now tridimensional instead of being planar. The distribution of singularities is composed of three main elements: a bound vortex filament spreading along the wing span, a vortex sheet trailing from the wing downstream towards the far field, and a starting vortex filament which is shed from the wing at the initial instant and travels downstream with the flow.

The bound vortex models the circulation around the airfoil that generates lift, the vortex sheet captures the effect of the wingtip vortices and the starting vortex closes the vortex sheet so that the Helmholtz theorem is verified². While the starting vortex has negligible effect on the flow around the wing, the vortex sheet strongly influence the circulation around the bound vortex and gives rise to a non-uniform spanwise lift distribution. The counter-rotating pairs of vortices also generate a net downwards velocity component named the downwash that may affect any other lifting surface downstream of the wing.




Figure 2.15 – Prandtl lifting-line model of the finite wing (reproduced from [HC03])

In the case of the elliptic spanwise lift loading, the lift curve slope of the finite wing can be derived from the airfoil lift curve slope and the wing aspect ratio AR . Considering a flat plate in potential flow, the thin airfoil theory gives a lift curve slope of 2π [HC03, p. 412] and the lifting line theory extends this result of the finite wing according to Equation (2.40).

$$C_{L,W} = \frac{2\pi\alpha}{1 + \frac{2}{AR}} \quad (2.40)$$

2.5.5 Wing-body interference factors

The problem of evaluating the mutual interference between the fins and the body has been addressed by Nielsen [Nie88] on the basis of the potential flow and slender body theories. A similar approach has been chosen, splitting the problem into a lifting and

² The Helmholtz theorem states that the circulation around a vortex tube is constant and implies that a vortex line cannot end in the fluid [HC03, p. 455]

a nonlifting scenario. In the nonlifting case, the angle of attack is set to zero and the thickness of the fins is neglected, so that the velocity potential corresponds to the body alone. The lifting problem due to the cross flow generated by body angle of attack is treated by integrating the velocity potential around a circular mid-wing section (cf. eq. (2.38)) over the length of the finned body.

Body-on-wing interference

Initially, Bernoulli's law is used to derive the pressure loading on the fin due to the body. Two integrations are performed in sequence, first along the fin chord and then across the fin span, showing that the result of the integrals only depend on the body radius and the fin maximum span. The interference factor K_{WB} is defined as the ratio of the lift of the fin in the presence of the body compared to the lift of the fin alone. The result is expressed as a function of the body radius r_B and the fin maximum span b_F , using the ratio $\lambda = \frac{r_B + b_F}{r_B}$.

$$K_{WB} = \frac{1}{\pi(\lambda - 1)^2} \left[\frac{\pi}{2} \left(\frac{\lambda^2 - 1}{\lambda} \right)^2 + \left(\frac{\lambda^2 + 1}{\lambda} \right)^2 \arcsin \left(\frac{\lambda^2 - 1}{\lambda^2 + 1} \right) - \frac{2(\lambda^2 - 1)}{\lambda} \right] \quad (2.41)$$

Wing-on-body interference

Similarly, the lift of the body due to the presence of the fin is compared to the lift of the fin alone to derive the the interference factor K_{BW} . As the slender body theory provides expressions for the lift on the cruciform wing-body configuration and the lift on the fin alone, the expression of K_{BW} can be deduced from the previous results:

$$K_{BW} = \left(1 + \frac{1}{\lambda} \right)^2 - K_{WB} \quad (2.42)$$

Applying these formulas to the front canards and the rear fins of the ACO projectile layout, it can be estimated that the body carries over more than 60% of the undeflected canards normal force and around 45% of the fins lift.

Body interference with deflected control surfaces

As the control surfaces are deflected, they no longer form a simple circular mid-wing cross-section when combined with the body. Thus, the conformal transform given by eq.

(2.37) does not stand and the velocity potential must be induced by a more complex distribution of singularities. However, the method is essentially unchanged and the reader is referred to Nielsen for the derivation of the potentials and forces [Nie88, pp. 214-217].

To distinguish between the interference of body on undeflected canards and the additional effect of canards deflections, a new factor k_{WB} is introduced. It has the same meaning as K_{WB} except that, in that case, the body is fixed at zero AoA and lift is generated by deflecting the canards. As for the previous case, Equation (2.43) gives the interference factor as a function of the geometric ratio λ .

$$k_{WB} = \frac{1}{\pi^2} \left[\frac{\pi^2}{4} \left(\frac{\lambda + 1}{\lambda} \right)^2 + \pi \left(\frac{\lambda^2 + 1}{\lambda(\lambda - 1)} \right)^2 \arcsin \left(\frac{\lambda^2 - 1}{\lambda^2 + 1} \right) - \frac{2\pi(\lambda + 1)}{\lambda(\lambda - 1)} + \frac{(\lambda^2 + 1)^2}{\lambda^2(\lambda - 1)^2} \left(\arcsin \left(\frac{\lambda^2 - 1}{\lambda^2 + 1} \right) \right)^2 - \frac{4(\lambda + 1)}{\lambda(\lambda - 1)} \arcsin \left(\frac{\lambda^2 - 1}{\lambda^2 + 1} \right) + \frac{8}{(\lambda - 1)^2} \log \left(\frac{\lambda^2 + 1}{2\lambda} \right) \right] \quad (2.43)$$

Mutual interference between deflected control surfaces

The last source of interference that will be quantified with the help of the slender body theory is the interaction between deflected canards. This phenomenon is due to the pressure field around a deflected control canard spilling over to the adjacent surfaces. For instance, during a pitch-up command, all the canards are deflected trailing edge down. Thus, the suction region above the bottom canards will be contaminated by the overpressure generated at the lower surface of the top canards, reducing the total lift of the four canards. This effect is accounted for in the formula of the equivalent canard angle of attack (eq. (2.47)) by the factor Λ_{ij} that quantifies the interference of the canard j on the canard i . No derivation is provided for the estimation of these coefficients, the implemented values being extracted from a report by Nielsen and Hensch [NHS77, p. 242].

2.5.6 Apparent-mass formula for pitch damping

Two types of analytical methods can be used for pitching moment estimation: integral methods such as the apparent-mass formula attributed to Bryson [Bry53], or pointwise methods like the tail damping estimation proposed by Cook [Coo12, pp. 343-346]. For

the latter, the damping moment stems from the angle of attack induced by pitch rate at the tailplane center of pressure(c.f. fig. 2.6). In reality, the pitch rate generates a vertical velocity field with an intensity proportional to the distance from the CG. Thus, pointwise methods neglect the contribution of each cross section of the projectile body and underestimate the pitch damping coefficient. Being an integral method, the apparent-mass method is much better suited to airframes where the body significantly contributes to the aerodynamic forces, as the contribution of each cross-section can be taken into account. Moreover, if the cross section is a finned body, the interaction effects between components are included.

The method of apparent masses presented by Nielsen [Nie88] provides estimates for all forces and moment coefficients of a given slender missile configuration, except for the drag force. It is based on the evaluation of the kinetic energy of the fluid along the missile axis, from which the normal forces and rolling moment on each cross-section can be obtained by differentiation. The pitch damping moment is expressed as a function of the apparent mass coefficients which only depend on the section geometry. For a non-spinning projectile ($p = 0$), the pitch damping moment estimate is computed by integrating the apparent mass coefficient of each section from the projectile base to its nose, as in Equation (2.44).

$$Cm_q = -4 \left[A_{22,\text{base}} \left(\frac{x_{\text{base}}}{\lambda_{\text{base}}} \right)^2 + \int_{(x/\lambda)_{\text{base}}}^{(x/\lambda)_{\text{nose}}} A_{22} \left(\frac{x}{\lambda} \right) d \left(\frac{x}{\lambda} \right) \right] \quad (2.44)$$

If the cross-section is constant, the apparent mass coefficient can be brought out of the integral and the whole process is greatly simplified. In our case, the coefficient of a cylindrical section fitted with four cruciform rectangular fins is given as a function of the local body radius to fin span ratio λ by Equation (2.45).

$$A_{22} = \lambda^2 \left(1 - \frac{1}{\lambda^2} + \frac{1}{\lambda^4} \right) \quad (2.45)$$

2.6 Estimation of Aerodynamic Forces on Projectile Components

The previous section shows that the relevance of the aerodynamic theories is conditioned by some assumptions on the geometry of the flying body. This is why, in the case of long-range projectile designs, determining the adequate method may be challenging as the assumptions of both slender-body or lifting-line theories may not be strictly verified.

This is why engineering methods have been developed to relax some of the geometric assumptions on the projectile components.

Also, some key physical phenomena, such as canard stall, are driven by viscous effects which have explicitly been neglected in the whole theoretical framework. Since, to the best of the author's knowledge, such viscous effects cannot be added back into the analytical models, semi-empirical correlations based on experimental data have been introduced as remedies to the theories inaccuracies, bridging the gap between the ideal flow around canonical bodies and the behavior of the actual projectile in wind tunnel. This section presents the engineering methods that have been applied to the problem of estimating the aerodynamic forces and moments over the ACHILES projectile up to the canard stall condition.

2.6.1 Equivalent angle of attack method for wing-body interactions

The prediction of lift forces on the fins and canards might not be as straightforward as in the case of an aircraft wing or even a missile fin due to the peculiar wing-body geometry. The experimental mockup features an 80 mm caliber body with four rectangular canards and fins of respective chords 30 and 45mm and aspect ratios of 2 (see section 1.1.2 for more details). These dimensions are a trade-off between the representativity of the mockup w.r.t long range projectile configurations and the technical limitations of the wind-tunnel setup.

However, this makes for an unusually small wing chord to body diameter ratio which raises the question of wing-fuselage interactions. Two modeling options can be explored: either a pair of opposite fins can be considered to form an ideal finite wing of aspect ratio 4, or each fin is assumed to be a slender wing of aspect ratio 2. Both assumptions are equally debatable: on one hand, the Prandtl lifting line theory is commonly admitted to be valid for wings of aspect ratios greater than 5. Additionally, in order to feature an ideal elliptic loading, the planform of the untwisted wing should be elliptical and not rectangular. On the other hand, Jones slender wing theory is has only been derived for pointed wings even if recent works have shown that it could be extended to rectangular wings [DM17, p 10] of aspect ratios below one. Also, the vicinity of the fuselage may strongly influence the flow field on the fin as one of its tips is coincident with the wall surface instead of being in free air.

The solution to this problem has been found into the theory of equivalent angle of attack developed by Hemsch and Nielsen [HN83]. Its aim is to expand the results of section 2.5.5, derived for slender wing-body configurations, to bodies with non-slender fins at moderate angles of attack. This approach allows to deduce the forces produced by a finned body from the lift polar of a wing formed by two fins joined together, without any assumption on the fins aspect ratios. According to Nielsen [Nie88, pp. 134-136], once the lift polar of the wing alone is known, the interference factors K_{CB} , k_{CB} and Λ_{ij} can be applied to non-slender fins. Since there is no requirement for this lift polar to be linear, any type of function can be used to model canard lift.

A few definitions should be introduced before the expression for the equivalent AoA of the lifting surfaces is given: The local angle of attack α_{loc} is the AoA seen by a lifting surface as a result of the projectile flying at an incidence angle α and an angle of sideslip β . It also depends of the orientation of the lifting surfaces on the projectile i.e the respective X and + patterns for the canards and the fins. The i -th canard deflection angle δ_i is the rotational motion around the canard mid-chord swivel provided by its servo actuator. For any control canard, the total angle of attack is the sum of the local angle of attack and the canard deflection angle. Thus, the equivalent angle of attack of a deflected control canard $\alpha_{C,eq,i}$ is decomposed into the contribution of the undeflected canard $\alpha_{C,eqz,i}$ at its local AoA and the contributions of all the canards deflections δ_j assuming that the local AoA is zero:

$$\alpha_{C,eqz,i} = K_{CB} \alpha_{C,loc,i} \quad (2.46)$$

$$\alpha_{C,eq,i} = \tan^{-1} \left(\tan(\alpha_{C,eqz,i}) + k_{CB} \left[\tan(\alpha_{C,eqz,i} + \delta_i) - \tan(\alpha_{C,eqz,i}) + \sum_{j \neq i} \tan(\Lambda_{ij} \delta_j) \right] \right) \quad (2.47)$$

On this projectile, the fins cannot be deflected. As a result, the equivalent angle of attack of fin i reduces to:

$$\alpha_{F,eq,i} = K_{FB} \alpha_{F,loc,i} \quad (2.48)$$

In this framework, the influence of the body on the fins and canards centers of pressure is neglected. This is supported by the remark of Nielsen [Nie88, p. 131] who states that the forward shift of the fins CP due to the body is at most a few percent of the chord.

2.6.2 Data-based non-linear canard lift model

As the model aims at capturing complex non-linear effects, sometimes the underlying physical phenomena violate the assumptions of the chosen theoretical framework. This is precisely the case with canard stall: at high angle of attack, the flow field is dominated by viscous effects as the boundary layer over the canard upper surface tends to separate, leading to a large wake or recirculation area [And17, p 63]. This is inconsistent with the assumption of flow irrotationality that governs potential flows and their applications to slender bodies and finite wings. In that case, a suitable alternative may be found by reconstructing the canard lift polar from experimental data that captures the real flow physics.

Canard lift reconstruction routine

Various means of implementing the non-linear canard lift polar have been investigated, including the use of wind-tunnel data obtained from the literature. However, the resulting behavior of the model showed inaccuracies in terms of maximum canard lift and canard lift curve slope. Thus, a routine was designed to reconstruct the canard lift polar from the trim map data of the actual HIL projectile. The reader is referred to section 2.2 for a description of the data collection procedure.

For each pair of recorded trim angle of attack α_t and canards deflection $\delta_{m,t}$, the following procedure is applied:

1. The raw trim AoA data is averaged for negative and positive canards deflections and the offset at the origin is removed.
2. The equivalent canard AoA is computed (c.f. section 2.6.1) and the position of the canards center of pressure is interpolated from aerodynamics prediction data.
3. The individual canard normal force is deduced from the analytic condition for pitching moment equilibrium at the trim point, given by:

$$Cm_{\delta_m} \delta_{m,t} + Cm_{\alpha} \alpha_t = 0$$

The normal force coefficient generated by one canard can then be related to the total moment applied by all four canards, assuming an equal contribution from

each of them:

$$Cm_{\delta_m} \delta_{m,t} = 2\sqrt{2} C_{N,C} K_{CB} \left(\frac{S_C}{S} \right) \left(\frac{x_G - x_{CP,C}}{D} \right)$$

Finally, the average canard normal force coefficient is:

$$C_{N,C} = \frac{-Cm_\alpha}{2\sqrt{2}K_{CB} \frac{S_C}{S} \frac{x_G - x_{CP,C}}{D}} \alpha_t$$

4. After projection in the local aerodynamic frame, the canard lift coefficient is tabulated as a function of the local equivalent angle of attack. This forms the canard-alone lift polar, corrected for the body interference effects.

It should be noted that, during the AoA sweep, the pitch stability coefficient Cm_α is assumed to be constant and equal to its estimated value around zero AoA. However, as the trim AoA and canards deflection increase, the canards experience stall and their destabilizing contribution to the projectile pitch stability decreases, increasing the total magnitude of the pitch stability coefficient. Thus, the model accuracy for large canards deflections and/or forward fins positions may be limited.




Figure 2.16 – FAST diagram of canard lift reconstruction routine

The canard lift is plotted as a function of local AoA on Figure 2.17 for different positions of the fins. One can notice that the linear region close to the origin extends up to less than five degrees, then the slope of the C_L curve reduces. Note that the slope of the linear region is significantly smaller for forward fins positions ($0.175 \text{ m} \leq x_F \leq 0.215 \text{ m}$). Even though some destructive canard-fins interactions may be present, the magnitude of this effect and the absence of a lift plateau at large AoA (c.f. [ASS15]) cast doubt on the

representativity of the obtained results. Some other non-linear effects such as the increase of pitch stability with AoA may have been lumped with the canard stall phenomena. Also, as the canard model is based on trim data, the influence of dynamic effects on canard stall cannot be captured by this model.

Non-linear canard lift polar




Figure 2.17 – Reconstructed data and polynomial models of the canard lift polar for different fins positions x_F (m)

Once the canard lift has been reconstructed from experimental data over a range of local angles of attack, a polynomial model is used to create a continuous function that fits the reconstructed data. It aims at capturing the non-linearity associated with canard stall and should provide satisfactory asymptotic behavior at the origin and at large local angles of attack. Preliminary analysis of reconstructed data shows a linear trend for low AoA and an oblique asymptote at large AoA. As a result, the following model formulation has been selected:

$$C_{L,C}(\alpha_C) = \text{sgn}(\alpha_C) \left(C_{L,\text{aoff}} + C_{L,\text{aslp}}|\alpha_C| + \frac{1}{p_0 + p_1|\alpha_C| + p_2\alpha_C^2} \right) \quad (2.49)$$

$C_{L,aslp}$ and $C_{L,aoff}$ are the respective slope and offset of the oblique asymptote of the function, while p_0 , p_1 and p_2 are the coefficients of the inverse polynomial part of the function. p_0 and p_1 can be expressed as functions of the other parameters and the slope at the origin $sl_{0,CLC}$ thanks to the following properties of the lift polar:

- $C_{L,C}(0) = 0$ implies that $p_0 = \frac{-1}{C_{L,aoff}}$
- $\frac{\partial C_{L,C}}{\partial \alpha_{eq}}(0) = sl_{0,CLC}$ implies that $p_1 = \frac{C_{L,aoff} - sl_{0,CLC}}{C_{L,aslp}^2}$

The remaining unknowns are $C_{L,aslp}$, $C_{L,aoff}$, $sl_{0,CLC}$ and p_2 . The slopes $C_{L,aslp}$ and $sl_{0,CLC}$ are directly extracted from the reconstructed data and do not need to be adjusted during the tuning process. In order to facilitate surrogate modeling and improve the predictive capabilities of the canard lift model, p_2 is also kept constant, so that $C_{L,aoff}$ is the only parameter tuned by the least squares error minimization algorithm in order to fit the reconstructed data. Overall, the polynomial models provide an excellent fit over the full set of reconstructed data, with an average value of 95.7% based on the normalized root-mean-square error (NRMSE) criterion.

Surrogate models of lift polar parameters




Figure 2.18 – Surrogate models of $C_{L,aslp}$ (left), $C_{L,aoff}$ (center) and $sl_{0,CLC}$ (right)

Applying the canard lift reconstruction process to a dataset covering multiple fins positions showed that canard lift was affected by this geometric parameter and followed a clear trend. Thus, surrogate models of the parameters $C_{L,aslp}$, $C_{L,aoff}$, $sl_{0,CLC}$ and p_2 were introduced in order to capture the parametric dependency of the lift polar. In a similar fashion as in section 2.3.2, polynomial representations were retained to provide a continuous description of the parameters. Their respective orders have been chosen to capture the trend of the parametric variation of each coefficient while minimizing the impact of measurement and estimation error:

$$C_{L,aslp} = m_{0,aslp} + m_{1,aslp}x_F \quad (2.50)$$

$$C_{L,aoff} = m_{0,aoff} + m_{1,aoff}x_F + m_{2,aoff}x_F^2 + m_{3,aoff}x_F^3 \quad (2.51)$$

$$sl_{0,CL,C} = m_{0,sl0} + m_{1,sl0}x_F + m_{2,sl0}x_F^2 + m_{3,sl0}x_F^3 \quad (2.52)$$

$$p_2 = m_{0,sqr} \quad (2.53)$$

Figure 2.18 confirms that all of the parameters decrease when the fins are shifted forwards. This shows the value of expressing the model as a function of physical parameters whose variations are monotonous and do not require high order surrogate models. As a result, the canard lift polar model can be extrapolated to forward fins positions ($x_F \leq 0.17$) where the projectile stability is too low to record accurate trim maps.




Figure 2.19 – Canard lift polars for specified x_F (left) and models fit (right)

Figure 2.19 shows that the degradation of the lift polar model fit is very limited when its parameters are replaced by their respective surrogate models. Even though the effect is a bit more pronounced for forward fins configurations, the fit remains adequate with NRMSE indexes in excess of 0.86.

2.6.3 Analytical formula for fins lift

The complexity of the lift polar is determined by the range of local angle of attack seen by the fins. As the fins remain aligned with the body, they share the same geometrical angle of attack which will be restricted to less than 10 degrees in this study. Thus, a linear

fins lift polar can be used and only the lift curve slope needs to be determined.

Potential flow theories give the lift curve slope of either slender wings ($AR < 2$) or large aspect ratio wings ($AR > 5$). In our case, the aspect ratio of the equivalent wing composed of two fins joined at their root chords is 4. For such wing geometries, flying at Reynolds number close to the test conditions ($Re \approx 90,000$), Ananda and Seiling [ASS15] show that the Helmbold equation best matches the experimental results. This formula predicts the lift curve slope of moderate to large aspect ratio finite wings as a function of the slope of the airfoil and the wing aspect ratio AR . For a flat plate in a potential flow, the thin airfoil theory gives a slope of 2π and the formula reduces to:

$$C_{L,\alpha} = \frac{\pi AR}{1 + \sqrt{1 + \left(\frac{AR}{2}\right)^2}} \quad (2.54)$$

For the ACO airframe, $C_{L,\alpha,F} = \frac{4\pi}{1+\sqrt{5}}$ and the fins lift coefficient is simply:

$$C_{L,F} = \frac{4\pi}{1 + \sqrt{5}} \alpha_F \quad (2.55)$$

Note that this representation neglects the flow field perturbations that may be induced by the upstream canards, such as downwash or vorticity. This assumption is supported by the arrangement of the four fins which are shifted 45 degrees with respect to the canards orientation (c.f. section 1.1.2).

2.6.4 Semi-empirical body lift and center of pressure

Body normal force and CP position

As seen in section 2.5.3, the slender body theory can provide estimates of the normal force acting on a fuselage-like shape. However, it is also well known that the hemispherical nose violates the assumptions of this theory as its cross-section varies too quickly along the revolution axis. The experimental work conducted by Darling [Dar73] on hemisphere-cylinder bodies of various length-to-diameter ratios clearly shows that the theory systematically underestimates the normal force slope of such geometries. Thus, the results of the slender body theory have been discarded in favor of an affine regression based on Darlings wind-tunnel data [Dar73, Fig. 2]:

$$C_{N,\alpha,B} = 2 + 0.05 \left(\frac{L_B}{D_B} - 2 \right) \quad (2.56)$$

A similar method has been applied for the center of pressure (CP) location described in the figure 9 of the same report. An attempt has been made to include the CP shift due to angle of attack by using the method of Barth [Bar69] even if the geometry and freestream conditions of this study are marginally outside of the domain of validity given by the author. Thus, the formula for the CP position is:

$$x_{CP,B} = D_B \left[0.656 + 0.1426 \left(\frac{L_B}{D_B} \right) + 2\sqrt{\alpha^2 + \beta^2} \right] \quad (2.57)$$

Darling mentions that such experimental results are valid for angles of attack up to 10 degrees, which is consistent with the scope of this study.

Effect of lifting surfaces interference

The lift carried over by the body can be quantified using the interference factors given in section 2.5.5. According to Hensch [Hem+76, p. 22] the normal force carried over from the lifting surfaces can be assumed to be independent from the means of lift production i.e local surfaces AoA or surfaces deflections. Thus, the normal force carried over from the canards is equal to the sum of the canard forces projected in the Body frame multiplied by K_{BC}/K_{CB} . Similarly, the carry-over force from the fins is obtained by multiplying the fins forces by K_{BF}/K_{FB} .

Additionally, it is assumed that the body center of pressure location remains unchanged in the presence of the canards and fins, except for the fact that the carry-over forces are applied on the body at the respective fins and canards locations. According to Nielsen [Nie88, p. 132], this CP shift exists for transonic configurations where the Mach waves on the panels interact with the body. In the low subsonic regime, none of these effects shall be encountered.

2.7 Parameter-Dependant Models of the Aerodynamic Moment

2.7.1 Panorama of the non-linear modeling approaches

Different modeling approaches have been explored in order to both capture the dependency to the geometric parameter and the canard stall phenomenon. Three non-linear models have been created, each with a different design philosophy. As the experimental

database was expanded, the focus has shifted from a mostly analytical structure to a data-based model in order to improve fidelity at the expense of versatility and predictive capabilities.

The component-based structure makes use of the local lift models obtained in section 2.6 to reconstruct the steady-state aerodynamic forces and moments from the contributions of each part of the projectile and their mutual interactions. The rationale behind this design paradigm is that a mostly analytical model structure may offer more versatility than data-based formulations and could more easily be extended to less stable fins positions. At the early stage of this work, the outcome of the closed-loop identification procedure described in section 2.2.2 was still uncertain, so having a model structure that could be extended beyond the range of stable fins positions for which experimental data could be acquired was considered very desirable. Since it was expected that the body and canard lift would be independent from the fins position, such component models could be based on available experimental data and remain valid for unstable fins positions. However, the results of the canard lift reconstruction process were later found to be x_F -dependent and additional experimental data was obtained, limiting the benefit of this approach.

The semi-local model is an attempt at improving the fidelity of the component-based structure by leveraging on identified aerodynamic derivatives that cover the full range of fins positions. The apparent-mass method has been replaced by the Cm_q surrogate model and the local lift models for the fins and body have been discarded. The global pitch stability and pitch control moments are modeled in place of the local component forces. The pitch stability moment is estimated from the Cm_α surrogate model augmented with the contribution of canard stall. The pitch control moment is still based on the local canard forces but only includes the effect of canards deflections and is rescaled using the Cm_{δ_m} surrogate model.

The data-based structure gets rid of the local canard model and instead captures the non-linear stall effects by mapping the trim attitude of the projectile as a function of the controls deflection. The surrogate Cm_q model is retained for the damping contribution.

2.7.2 Component-based model

As its name implies, this architecture relies on component buildup to predict the total forces and moments acting on the configuration. First, the equivalent angle of attack of each lifting surface is computed according to the formula of section 2.6.1, including

body-on-wing effects and control surface interference. The remaining interaction effects are captured by the component interference factors of section 2.5.5. Then, local lift models are used to evaluate the component forces and moments: the canard forces are found with the non-linear local lift polar of section 2.6.2, the body forces or moment are estimated with the semi-empirical model of section 2.6.4 and the fins lift is given by the Helmbold formula of section 2.6.3. Finally, the contributions from the projectile body and all fins and canards are summed at the airframe center-of-gravity to determine the total lift and pitching moment of the configuration. Figure 2.20 describes the implementation of the model, red boxes being specific to the canards, green boxes to the fins, blue boxes to the body while white boxes apply to any of those components.

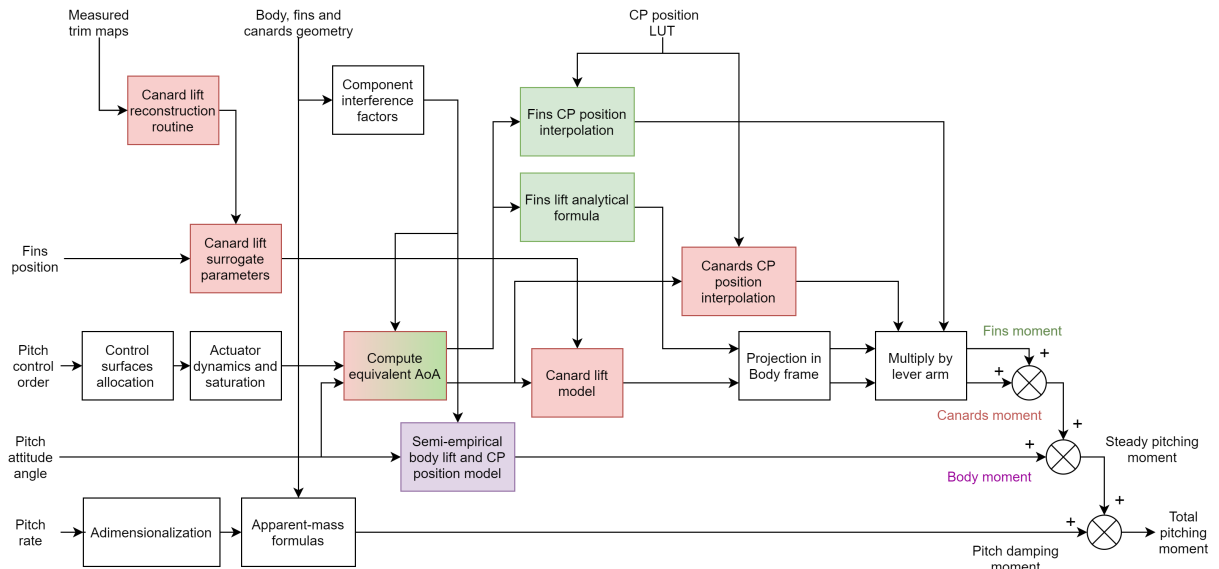


Figure 2.20 – FAST diagram of component-based model

2.7.3 Semi-local model

This layout presented on Figure 2.21 is the combination of the local canard model of the previous structure and the surrogate models of the aerodynamic derivatives described in section 2.3.2. Thus, the pitch stability and damping coefficients are obtained from these models which are based on the experimental data collected in section 2.2.2. Conversely, the pitch control authority is computed from individual canards forces by using the same lift polar as above, but only with the part of the equivalent AoA that represents the controls deflection. An effort has been made to capture the effect of canard stall on the projectile

pitch stability by augmenting the Cm_α coefficient with an additional term described in the next paragraph. Also, the pitch control moment stemming from the canard lift polar has been scaled to match the identified Cm_{δ_m} in the linear domain.

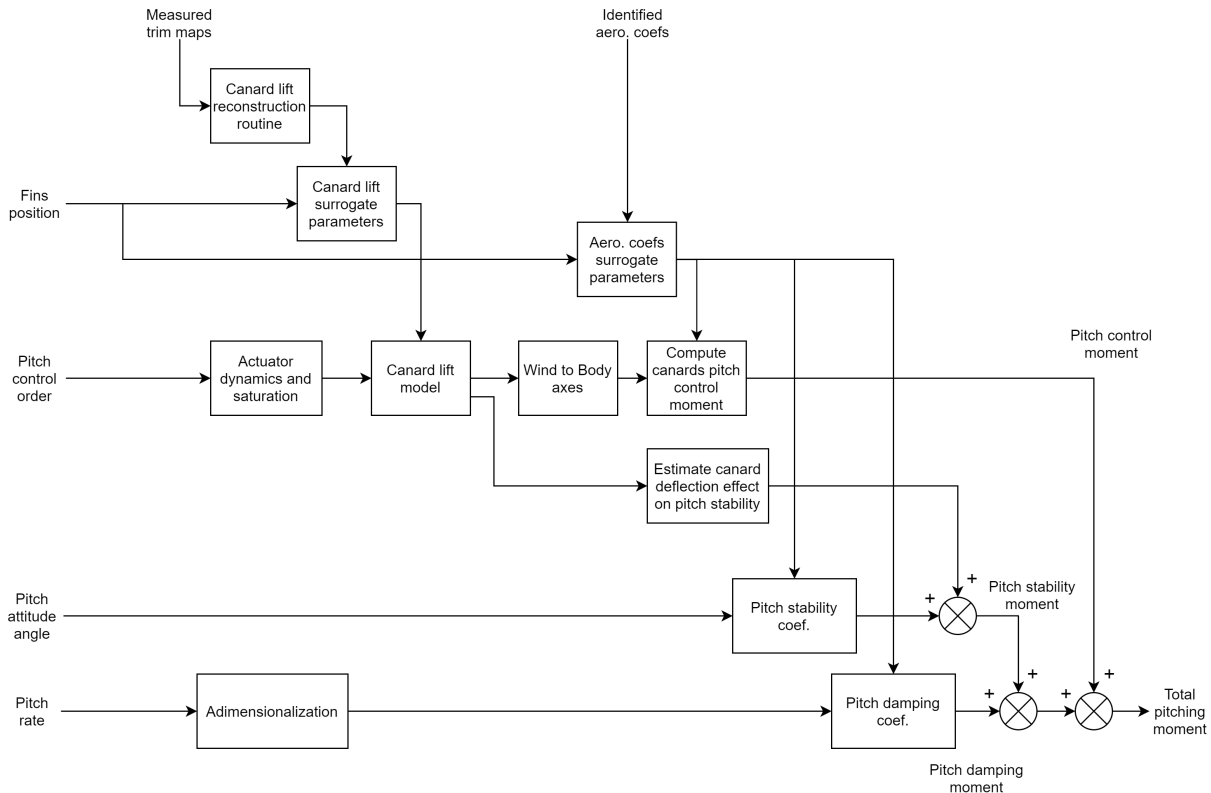


Figure 2.21 – FAST diagram of semi-local model

Estimation of the canards deflection effects on pitch stability

The pitch stability of the projectile is the result of the contributions of the fins, body and canards. While the body and horizontal fins only see the projectile angle of attack, the canards are subjected to local angles of attack which are a combination of the projectile AoA and sideslip as well as the control surfaces deflections. Given that a trailing-edge-down deflection is required to create a nose-up pitching moment, the local canards AoA are usually greater than the projectile AoA. Thus, the canards stall earlier than the fins, producing a rearwards shift of the projectile center of pressure and increasing its static stability. Since the static stability coefficient Cm_α is identified in the vicinity of $\alpha = 0^\circ$, it does not capture non-linear effects. As a result, an additional term proportional to

the ratio of the non-linear canard lift curve slope to the linear slope is appended to the stability coefficient to model the effect of canard stall:

$$Cm_{\alpha,cor}(x_F) = Cm_{\alpha}(x_F) - (x_G - x_{CPC})K_{scl} \left(1 - \frac{sl_{CL,C}(\alpha_{C,eq})}{sl_{0,CLC}} \right) \quad (2.58)$$

At the origin ($\delta_m = 0$), the pitch stability coefficient is equal to its identified value. As the canards deflections increase, the local lift curves slope $sl_{CLC}(\delta_m)$ drops below its maximum value $sl_{0,CLC}$ and the magnitude of the stability coefficient increases.

The scaling factor K_{scl} is used to tune the sensitivity of pitch stability to canards lift saturation. It is adjusted so that the steady-state model response fits well with the experimental trim maps across the range of stable fins positions for which trim maps could be recorded: $0.175 \leq x_F \leq 0.313$.

Scaling of the canards pitch control moment

The pitch control moment $Cm_{C,defl}$ is proportional to the total canard lift induced by the canards deflection. At the component level, this quantity is found by subtracting the lift of undeflected canard $C_{L,C,undefl,i}$ from the total canard lift $C_{L,C,i}$. The former quantity can be determined by substituting the equivalent AoA of the undeflected canards (2.46) in the non-linear canard lift equation (2.49). Similarly, evaluating this equation with the (deflected) canard equivalent AoA (2.47) gives the total canard lift.

$$C_{L,C,undefl,i} = C_{L,C}(\alpha_{C,eqz,i}) \quad (2.59)$$

$$C_{L,C,i} = C_{L,C,i}(\alpha_{C,eq,i}) \quad (2.60)$$

The proportionality factor is chosen so that the slope $sl_{0,CLC}$ of the linear region of the pitch control moment matches with the identified pitch control coefficient $Cm_{\delta_m,exp}$. The factor $2\sqrt{2}\cos(\alpha)$ comes from the control allocation and the projection from Wind to Body frame.

$$Cm_{C,defl} = 2\sqrt{2}\cos(\alpha) \frac{Cm_{\delta_m,exp}}{sl_{0,CLC}} \sum_{i=1}^4 C_{L,C,i} - C_{L,C,undefl,i} \quad (2.61)$$

2.7.4 Data-based model

This structure pictured on fig. 2.22 favors system identification and non-linear regression to deduce the projectile pitching moment from a comprehensive experimental database. The pitch trim map model captures the competition between control authority and pitch stability and predicts the projectile trim attitude for a given pitch control deflection. In addition, the surrogate models of Cm_α and Cm_q are employed to include the dynamic effects and compute the total pitching moment. The structure of the trim map model is designed to generate maps for arbitrary fins positions, its building process is detailed in the next paragraph.

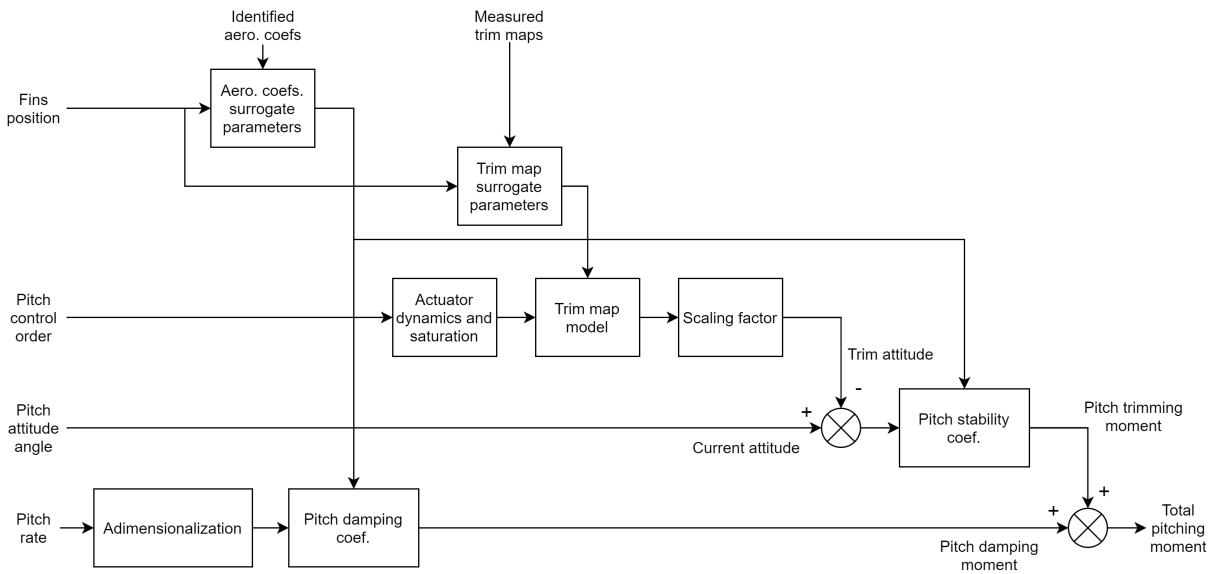


Figure 2.22 – FAST diagram of data-based model

Non-linear trim map model

The raw trim maps collected according to the protocol detailed in section 2.2.1 are post-processed before being used for modeling purposes. The post-processing step leverages on the symmetry properties of the projectile in order to mitigate the effect of process noise. First, the mean offset of the trim AoA series is removed, then a central symmetry w.r.t the origin is applied to the negative part of the trim map and the AoA measurements corresponding to the same controls deflections are paired together and averaged. Finally, only the positive part of the trim map is kept and its remaining offset is removed so that the projectile can be trimmed at zero AoA with undeflected control surfaces.




Figure 2.23 – Post-processed trim maps and polynomial models for different fins positions x_F (m)

The measured trim maps exhibit a similar trend as the reconstructed canards lift polars with a linear region around the origin followed by an oblique asymptote. Thus, the polynomial structure introduced in section 2.6.2 is adapted for the purpose of trim map prediction:

$$\alpha_{\text{trim}}(\delta_m) = \text{sgn}(\delta_m) \left(\alpha_{\text{aoff}} + \alpha_{\text{aslp}} |\delta_m| + \frac{1}{a_0 + a_1 |\delta_m| + a_2 \delta_m^2} \right) \quad (2.62)$$

The model coefficients α_{aoff} , α_{aslp} , a_0 , a_1 and a_2 are determined using the same methodology as in section 2.6.2. Similarly, surrogate models are used to capture the parametric dependency and express α_{aoff} , α_{aslp} , $sl_{0,\alpha}$ and a_2 as polynomial functions of the fins position x_F . However, Figure 2.23 shows that the spacing between the different trim maps increases greatly as the fins are brought forwards and x_F decreases. This is because, as the static stability of the projectile decreases, the trim map asymptotically expands towards infinity: in theory, a neutrally-stable projectile could be trimmed at any angle of attack without any steady-state controls deflections. Thus, the slopes of the trim map, both at the origin and at large controls deflection, diverge to infinity as shown by the trend of the

parameters α_{aslp} and $sl_{0,\alpha}$ plotted on fig. 2.24.

$$\alpha_{\text{aoff}} = n_{0,\text{aoff}} + n_{1,\text{aoff}}x_F \quad (2.63)$$

$$\alpha_{\text{aslp}} = n_{0,\text{aslp}} + n_{1,\text{aslp}}x_F + n_{2,\text{aslp}}x_F^2 + n_{3,\text{aslp}}x_F^3 \quad (2.64)$$

$$sl_{0,\alpha} = n_{0,sl0} + n_{1,sl0}x_F + n_{2,sl0}x_F^2 + n_{3,sl0}x_F^3 \quad (2.65)$$

$$a_2 = n_{0,\text{sqr}} + n_{1,\text{sqr}}x_F + n_{2,\text{sqr}}x_F^2 + n_{3,\text{sqr}}x_F^3 \quad (2.66)$$

On the range of stable fins position considered, the polynomial model fits the measured trim map very well with an average NRMSE-based fit index of 96,8%. The model based on surrogate parameters remains accurate with an average fit index of 87,1%. However, given the asymptotic divergence of three of the parameters, care should be taken when extrapolating the trim map model to less stable configurations as it may become increasingly inaccurate.

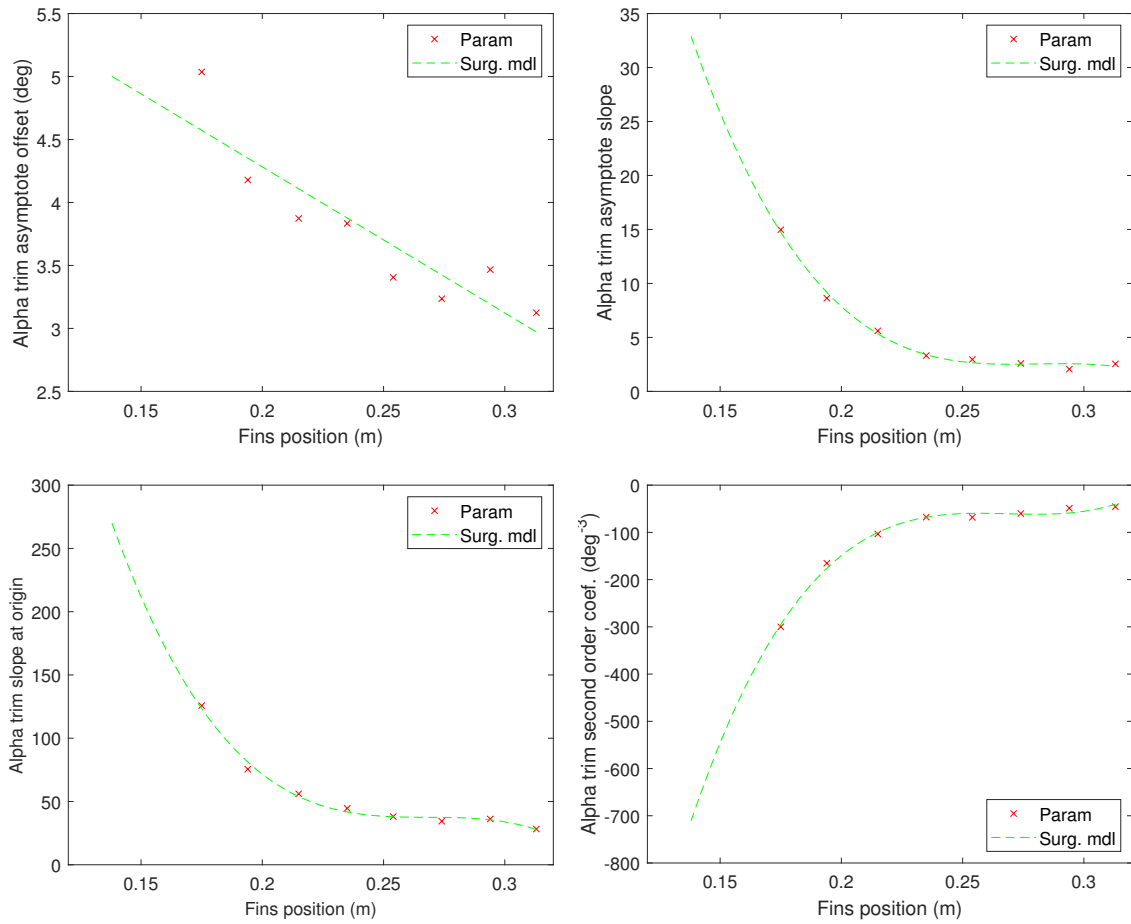


Figure 2.24 – Surrogate models of α_{aoff} , α_{aslp} , $sl_{0,\alpha}$ and a_2




Figure 2.25 – Trim maps for specified x_F (left) and models fit (right)

2.8 Model Assessment and Selection

This section aims at quantifying the ability of the model structures to accurately predict the pitch dynamics of a family of projectiles with varying degrees of static stability. The LPV model must be able to capture the projectile behavior sufficiently well so that the synthesis framework produces an airframe-controller design that shows satisfactory performance in wind tunnel. As mentioned in the introductory section of this chapter, the prediction of the actual closed-loop response of the projectile is the ultimate purpose of the non-linear models.

However, before the controller design stage, closed-loop response measurements are not available so model assessment must be performed on open-loop data. First, the model responses to small control inputs are compared to pitch attitude time sequences acquired after the system identification experiments. Then, the equilibrium manifolds given by the non-linear structures are matched up to trim maps recordings. The LPV model of section 2.3 must be able to reproduce the open-loop pitch dynamics to provide a good starting point for controller synthesis. The non-linear models of section 2.7 should also be able to do so, while capturing the effect of canard stall on the projectile trim map. The component-based, data-based and semi-local structures are benchmarked in order to find the most suitable model formulation for assessing control performance at large angles of attack, up to the non-linear domain. The model should also be able to accurately describe the projectile pitch dynamics over a wide range of fins positions, including statically unstable projectile configurations.

2.8.1 Open-loop fidelity assessment

Step response

As part of the system identification procedure, validation sequences have been collected in order to check the LPV model fit for different fins positions. The responses of the non-linear structures have also been added in order to assess their transient-matching abilities.

The open-loop projectile is excited with a series of steps on the pitch control channel. The sequence alternates between positive and negative deflections in order to limit the magnitude of the pitch angle excursions. The amplitude is kept constant while the duration of the step is varied in order to excite the plan on a larger bandwidth. The system response exhibits a typical under-damped second order behavior, featuring decaying oscillations as the new equilibrium position is reached.




Figure 2.26 – Pitch attitude responses comparison for $0.254 \leq x_F \leq 0.313$

The time histories of Figure 2.26 show that all models are able to capture the steady-state and transient response of the open-loop projectile relatively well, with the exception of the component-based model being excessively damped with oscillations decaying too fast. The trim attitude predicted by the semi-local structure seems to have the greatest offset with respect to the steady-state measurements, however the transients are well reproduced. The LPV and data-based models offer the best fit, with well-matched amplitude and period in the oscillatory regime and accurate trim conditions.

As the fins position decreases towards less stable configurations, the amplitude of the oscillations increase and model fit seems to degrade. This is especially true for the component-based structure which predict much smaller trim attitudes than the experimental data. For the least stable fins position ($x_F = 0.175$ m), the amplitude of the oscillations is overestimated by the LPV model and underestimated by the other structures.




Figure 2.27 – Pitch attitude responses comparison for $0.175 \leq x_F \leq 0.235$

Also, the trim attitude measured on the positive side appears to be slightly smaller than the opposite of the negative trim attitude. Such asymmetric effects cannot be predicted by any of the model structures and degrade their respective fits. Figure 2.28 confirms that the overall fit of the models is correct, with values above 0.6 except for the component-based structure that sharply declines for the most forwards fins position. The data-based model offers the best overall fit with an average of 81.9%, followed by the LPV model with 78.5%. The semi-local model has the worst average fit with 68.2%, while the component-based structure manages 68.6%.

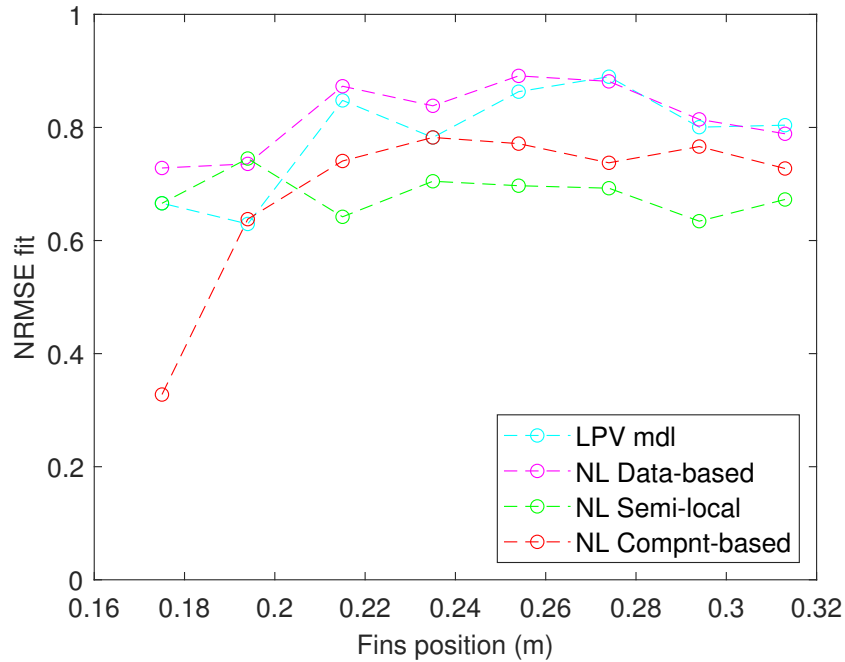


Figure 2.28 – Open-loop response fit for different fins positions

Trim map prediction

The equilibrium manifold of a given stable projectile configuration, also named trim map, is recorded as part of the data collection procedure detailed in section 2.2. In the pitch-only case, it consists of a graph of the trim pitch attitude θ (which is also the angle of attack α) as a function of the virtual pitch control deflection δ_m . This map is used for various purposes, such as finding the control deflection corresponding to a reference trim attitude, as well as predicting the maximum AoA that could be achieved by the configuration for a given control deflection limit. In the frame of model assessment, the accuracy of the predicted trim map gives some insight on the ability to capture the effect of non-linear phenomena on the steady-state response of the model.

For the most rearwards fins positions covered in Figure 2.29, the trim maps remain relatively similar, with an early transition away from linear behavior for controls deflection amplitudes of less than 5° . The black dashed line shows that the linear portion of the trim maps are consistent with the identified aerodynamic coefficients, if only for the component-based model which slope is slightly overestimated. As the fins are shifted forwards, the linear slope of the trim map increases and so does the range of motion of the projectile mockup. According to Figure 2.30, for the most forward fins position, the projectile reaches





Figure 2.29 – Trim maps comparison for $0.254 \leq x_F \leq 0.313$

an AoA of more than 9° with only 16° of virtual pitch control deflection. Moreover, pitch angles up to 5° can be achieved without encountering any significant non-linearity in the pitch control response. These observations clearly demonstrate the benefit of reducing the projectile static stability in order to mitigate canard stall.

As far as model fidelity is concerned, all of the model structures become less accurate, especially the component-based one whose fit drops below 60%. This is mainly due to an underestimation of the slope in the linear domain, which also offsets the oblique asymptote in relation to the measured data. The semi-local structure is able to predict the linear domain well but starts to overestimate the asymptotic slope as the fins abscissae decrease. The data-based model gives the best results with only a little offset visible for the most forwards fins position.

Results of Figure 2.31 are in agreement with the previous observations, showing de-

Figure 2.30 – Trim maps comparison for $0.175 \leq x_F \leq 0.235$

creasing model fits for fins positions forwards of 0.215 m. It also quantifies the degradation and gives the trends for fins positions ahead of the range covered by the trim maps. Even if the overall fits remain very good with a respective 95.3, 89.2 and 85.1 percents for the data-based, semi-local and component-based structures, the plots underline that the latter model may not be suitable for less stable configurations. Further investigation is required in order to find which model could perform best over an extended range of fins position.

2.8.2 Extension to statically-unstable projectiles

As for any other optimization problem, the larger the parameter space is, the better the quality of the optimum may be. Given that the longitudinal position of the projectile fins

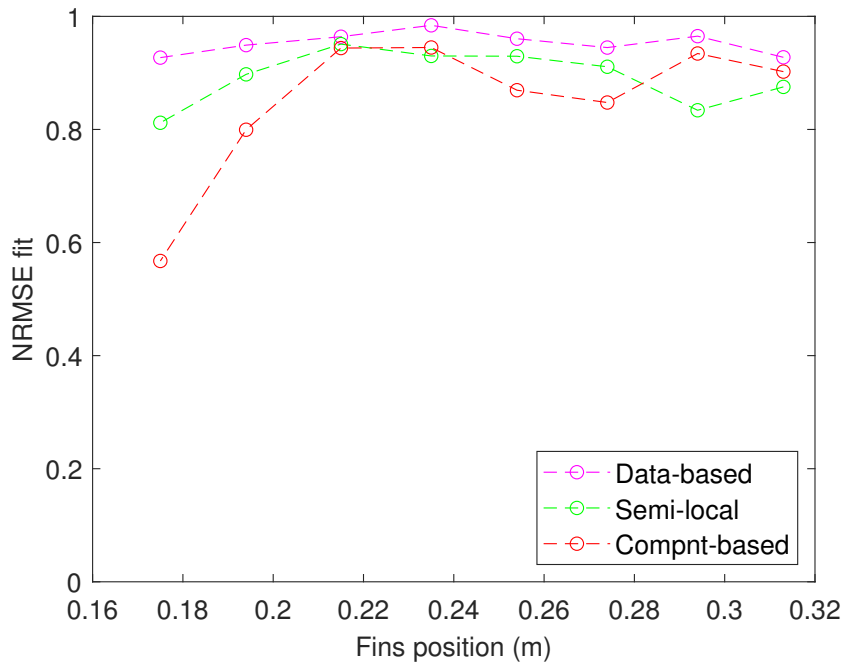


Figure 2.31 – Open-loop response fit for different fins positions

is tuned inside the synthesis framework, the range of x_F on which the aerodynamics model is valid must be as large as possible to maximize the benefits of the airframe-controller optimization approach. Moreover, the literature mentions that air vehicles can benefit from relaxed static stability to improve their maneuvering performance [ABH73]. As a result, an effort has been made to extend the modeling capabilities of the HIL framework to statically-unstable projectile configurations.

First, the LPV model would need to be extended, implying that the surrogate models of the identified aerodynamic coefficients would have to be extrapolated to more forwards fins positions or additional pitch identification experiments should be performed. Looking at the trend of the Cm_q graph of fig. 2.7, it is difficult to assess whether the second order surrogate model would remain accurate for lower x_F . Moreover, the non-linear models also need to be improved in order to provide closed-loop validation over the extended parameter range. As mentioned in sections 2.7.1 and 2.8.1, the attempt at developing a model structure that could rely on limited experimental data to predict the pitch dynamics of unstable projectiles has been met with poor results. The only model candidate that seems capable to reproduce the projectile behavior over the whole range of fins positions relies heavily on identified aerodynamic coefficients and requires additional data. Thus,

a new closed-loop pitch identification technique detailed in section 2.2.2 has been introduced to provide the required aerodynamic coefficients for unstable or marginally stable configurations.

Linearized aerodynamic coefficients

This data can also be leveraged to obtain additional insight on the behavior of the models for more extreme fins positions: the estimated values of Cm_α , Cm_q and Cm_{δ_m} for the fins positions of 0.156 m and 0.138 m are appended to the set of previously identified coefficients (c.f. section 2.2.2) and plotted on Figure 2.32. Then, all of the four models are linearized with respect to the state variables α and q and the input δ_m . Results displayed on Figure 2.32 are obtained using the small-perturbations approach around the operating point where states and inputs are zero.




Figure 2.32 – Comparison of Cm_α (left), Cm_q (center) and Cm_{δ_m} (right) from system identification and models linearization

Figure 2.32 indicate that the LPV model provides an excellent fit with the identified coefficients, confirming the good open-loop results obtained for the step response. As such, it can be inferred that this model structure correctly predicts the projectile pitch dynamics in the linear region and is adequate for controller synthesis and airframe optimization.

The semi-local model is the only non-linear structure whose linearized behavior is able to match the reference data over the full range of fins positions. It remains very close to the LPV results for all three coefficients. The large errors in the linearized component-based response explain its deficiencies observed on the trim maps and time-response plots. The excessive damping found for aft fins positions (fig. 2.26) stems from the inability of the apparent-mass method to accurately predict the pitch damping coefficient. The large steady-state offset observed on both trim maps (fig. 2.30) and time responses (fig. 2.27)

for forward x_F is due to the excessive pitch stability and greatly underestimated control power shown on the Cm_α and Cm_{δ_m} plots of fig. 2.32. The response of the linearized data-based model to small α and q perturbations is consistent with the identified coefficients. However, the rightmost plot of fig. 2.32 show that the linearized pitch control coefficient of the model goes to zero as the configuration gets closer to instability. This non-physical behavior is the result of a modeling artifact shown in the output equation (2.67) of the data-based model given by the diagram of fig. 2.22:

$$Cm = Cm_\alpha (\alpha - \alpha_{\text{trim}}(\delta_m)) + Cm_q \frac{qD}{\bar{v}} \quad (2.67)$$

As Cm_α goes to zero when the configuration gets closer to neutral static stability and the predicted trim AoA α_{trim} remains finite, the resulting Cm is zero for all δ_m and the control surfaces are totally ineffective, which does not agree with the physics. Thus, the formulation of the data-based model is not suitable for unstable or marginally stable projectiles configurations.

2.8.3 Model comparison and selection

After comparing experimental data with the predictions of the different model structures, one non-linear formulation must be selected in order to perform the closed-loop validation of the controller response and quantify its performance. Ideally, it should be able to accurately capture the effects of canard stall and predict the projectile pitch dynamics over a wide range of fins positions. As stated in the previous section, this last consideration is especially important in the context of airframe-controller optimization. The following table compares the model fit percentages on open-loop trim maps and step response measurements. The trends with respect to decreasing fins positions and the behavior of the linearized models are leveraged to qualitatively assess the potential for model extension to statically unstable projectile configurations.

Model structure	Step responses fit (%)	Trim maps fit (%)	Step response fit trend	Trim map fit trend	Linearized mdl trend
Data-based	81.9	95.3	+	+	--
Semi-local	68.2	89.2	+	-	+
Component-based	68.6	85.1	--	--	--

Table 2.1 – Non-linear models comparison

If it was not for its inability to correctly predict the response of neutral or unstable configurations, the data-based structure would have been the best option as it has the

highest fit for both types of open-loop experiments conducted with stable configurations. Unfortunately, the investigation in section 2.8.1 has shown that there is little perspective for extending the domain of validity of this model formulation. Thus, this structure should only be used if the fins position is known in advance and sufficiently aft (i.e $x_F \geq 0.175$).

The fidelity of the component-based structure can be deemed acceptable for the most aft fins positions, but quickly degrades as x_F decreases, even if the projectile configuration remains stable. Here, part of the model deficiencies could be attributed to the analytical formula used to estimate the pitch damping moment. Also, the evolution of the linearized pitch control coefficient may reveal an issue with the modeling of the control surfaces. It may be possible to improve the model fidelity by introducing semi-empirical correction factors in order to account for unmodeled phenomena, such as the aerodynamic interactions between fins and canards, but this option is left for future work.

As a result, the semi-local model has been selected to assess the performance of the airframe-controller designs in simulation. However, even if the linearized coefficients are consistent with the experimental results over the whole range of stable and unstable fins positions, the mediocre fit on the open-loop response and the distortion of the trim map at large AoA for low x_F indicate that further model validation should be performed. The closed-loop wind-tunnel tests presented in the next chapter (section 3.5) cover the test conditions at which the model predictions may be inaccurate and provide a definitive answer on the suitability of the model structure.

2.9 Conclusion

In this chapter, one LPV and three non-linear model structures have been investigated to perform and validate airframe-controller optimization. The purpose of the modeling process and the function of each type of model have been defined: the LPV structure is designed for controller synthesis while the non-linear model is used to assess the closed-loop performance. Both formulations have to be dependent on the fins position in order to predict the behavior of statically stable and unstable projectiles.

Experimental data has been collected to capture the non-linearities of the equilibrium manifold and the transient response to small input disturbances. These measurements are used to determine the unknown parameters of the "grey box" model structures. The range of fins positions for which the projectile models are valid has been expanded thanks to a new closed-loop system identification method.

The LPV model structure is described with a focus on surrogate models of the aerodynamic and inertia terms as polynomial functions of x_F . The order of these expressions is chosen according to a-priori knowledge on the configuration aerodynamics. Then, the model is cast into a linear fractional representation and reformulated in order to reduce its size and improve computational performance.

Non-linear models are introduced to circumvent the limitations of the linear structure with respect to canard stall. The fluid mechanics theory is presented, using potential flows to derive analytical lift formulas for wings and slender bodies, as well as estimating the wing-body interactions and the pitch damping effect. Engineering methods are used to bridge the gap between the theoretical framework and the actual projectile geometry and flight domain. The canard lift polar is reconstructed from the measured trim map up to the stall region and the equivalent angle of attack allows to better predict the component interference on the definitive projectile geometry.

Three non-linear models are presented, ranging from a mostly analytical approach to a data-based structure. Model fidelity is assessed on open-loop trim maps and step responses of stable configurations, then closed-loop experiments are employed to forecast the model ability to predict the dynamics of unstable projectiles. The data-based structure provides the best fit on stable projectile responses but is unsuitable for projectiles with relaxed static stability. The component-based structure is hampered by its reliance on inaccurate analytical models and shows large errors for forwards fins positions. The semi-local model is the only structure that appears to perform evenly across the range of fins positions, thus it has been chosen for closed-loop validation of the airframe-controller designs.

INTEGRATED DESIGN OF PROJECTILE AERODYNAMICS AND CONTROL

In this chapter, a new methodology is proposed as a means to improve the projectile maneuverability despite stringent control system limitations. First, the interest of plant-controller optimization in the frame of guided projectile design is probed, then the second section develops the literature review about this particular concept. A novel design framework is introduced in section three to account for the effects of canard stall early at the design stage, speeding up the development process and reducing the risk of uncovering performances deficiencies at the validation stage. Plant-controller optimization is applied to the projectile aerodynamics to maximize the control performance and mitigate the detrimental effects of the non-linearity of the open-loop plant response. In the fourth section, two distinct controller structures are implemented into the new design framework and evaluated in simulation. A control scheme tailored to a Relaxed Static Stability (RSS) configuration is also tuned and simulated. Finally, section five focuses on wind tunnel testing to provide experimental validation of the models used in the previous section and check the control performance of the guided projectile candidates.

3.1 Motivations and Use Case

3.1.1 Example of a flight scenario

As mentioned in the introduction, innovative projectile designs [Dec+18] [Vas+20] have been developed in order to achieve increased range through a glide phase performed at a shallow flight path angle. Thus, in urban warfare or non-line-of-sight firing scenarios, a top attack may be required in order to clear obstacles and terrain features such as elevated buildings or treelines. Moreover, it is also beneficial to warhead lethality as some targets like armored vehicles are more vulnerable when struck from above. Thus, a

scenario combining a shallow glide phase and a top attack, as pictured in fig. 3.1, is of particular interest.




Figure 3.1 – Combined gliding flight and top attack scenario

Two maneuvers are required in order to follow the above-mentioned flight profile. First, at the apogee (see fig. 3.1, label I), the projectile should transition from ballistic to gliding flight. Due to the low dynamic pressure stemming from the combination of high altitude and low airspeed, the projectile may need a substantial lift coefficient to follow the desired trajectory. At the beginning of the terminal phase, it is necessary to swiftly transition from a glide to a steep dive by the means of an aggressive pitching maneuver (see fig. 3.1, label II). A considerable load factor is required in order to minimize the turn radius and avoid obstacles.

3.1.2 Properties and limitations of gliding projectiles

Emerging gliding projectile configurations feature large lifting surfaces in order to maximize their lift-to-drag ratios. Compared to classical guided projectiles such as artillery shells fitted with course correction fuses, these additional surfaces enable the gliding projectiles to generate significant lift by increasing their angle of attack [Dec+18] instead of relying solely on the forces produced by their control surfaces. As a result, precise pitch attitude control is crucial to ensure maximum maneuvering performance and accurate flight path tracking. In order to do that, the available control moment should be sufficient to generate adequate pitch rate and angle of attack. This poses a challenge as available

control force and actuator bandwidth are usually very limited for this type of airframe [Fre11] due to the design constraints on the control and actuation system (CAS) such as G-hardening, packaging space, power usage and unit cost. Thus, actuator rate and deflection limits are susceptible to be reached in normal operation which must be avoided as, without proper anti-windup schemes, it may degrade controller performance and even lead to closed-loop instability [Tha+19].

In order to mitigate the actuator limitations, the trade-off between pitch authority (which conditions the amplitude of control surfaces deflections) and stabilization effort (that drives the requirements on actuator bandwidth) is of primary interest. This compromise is driven by the static stability coefficient of the projectile: the Cm_α not only quantifies the magnitude and direction of the pitching moment generated in response to an angle of attack disturbance, it also conditions the size of the trim map. This is because, for a given control order δ_m , the equilibrium pitch attitude depends on the ratio of Cm_α and Cm_{δ_m} . Legacy methods for sizing the static stability of the projectile are based on open-loop stability criteria and thus may be excessively conservative in the frame of closed-loop control [McC99]. More recently, Fresconi et al. [FCF12] conducted a parametric study of the influence of various actuation schemes, control laws and geometric parameters on projectile range. Results showed that center of gravity (CG) position, which determines static stability, had a significant influence on performance. However, further research by the same author has been conducted with a stable configuration featuring a static margin in agreement with standard design practices [Fre+18] and no attempt at tuning this parameter has been found in the literature.

3.1.3 Purpose of plant-controller optimization

This work proposes an original application of the Plant-Controller Optimization (PCO) process to concurrently tune the static stability and the controller gains of a fin-stabilized guided projectile. The legacy approach to guided projectile design implies sizing the lifting surfaces according to open-loop static criteria derived from empirical knowledge and mission requirements (i.e. open-loop stability margin and required course correction capability), and only then tuning the gains to perform best with the given airframe.

The expected benefit of the integrated design of aerodynamics and control over this sequential “design then control” process is twofold. First, as shown in the pioneering work of Fathy [Fat+01], solving the design and control problem successively does not guarantee optimal system performance: this could only be achieved by using nested or

simultaneous optimization strategies. In our case, it means that the sequential design process will likely lead to worse closed-loop performance than the simultaneous tuning of both aerodynamic and controller parameters. Concurrent design of projectile geometry and controller structure also allows to mitigate risks and avoid unexpected redesigns of the aerodynamic configuration. This could be the case if the airframe proves to be excessively stable and requires too much control force to trim or, alternatively, if its lack of open-loop stability results in excessive actuator bandwidth requirements.

3.2 State-of-the-Art of Plant-Controller Optimization

3.2.1 Concept

Plant-controller optimization (PCO) consists in tuning both plant parameters and controller gains to find the most beneficial design trade-off. This concept, also known as integrated design and control or more simply co-design, first appeared in the field of structure mechanics [HLD85] [OH87] along with active control of flexible bodies. Nowadays, there exist a number of control systems for which some physical parameters of the plant have significant influence on the controller design problem and thus on the performance of the closed-loop system. This approach has been successfully applied to aircraft and airships [KH00] [YL03] [DVT15], electric and hybrid vehicles [Lv+18] [Sil+16], robotic manipulators [RWH06], electric motors [RP99], elevators [Fat+02] and automotive suspensions [Fat+03].

The benefits of PCO are manifold: for instance, performing both plant sizing and controller tuning simultaneously may provide the expected level of closed-loop performance at a minimal cost. Alazard [Ala+13] used co-design to find the maximum transmission delay that could be tolerated by a satellite attitude control system, Niewoehner [NK96] minimized the control authority requirement of a fighter jet in approach phase, and Denieul [Den+17] sized the elevons of a blended-wing-body airliner under handling qualities constraints. Alternatively, one could take advantage of both process and feedback dynamics in order to obtain the best possible performance index. Morris [Mor92] proposes a procedure to design a family of tail-less aircraft and their controller for minimum drag with or without handling qualities constraints.

3.2.2 Optimization strategies

The coupling between the plant and the controller optimization problems was studied by Reyer [RP00] and Fathy [Fat+01]. An effort was made to assess whether or not the proposed strategies would lead to a global optimum for both plant and controller design. The latter paper lists four different strategies for conducting plant-controller optimization:

- The *sequential* strategy embodies the "design then control" paradigm where the controller gains are chosen to achieve the best possible performance for a given plant design. As the coupling between plant and controller design problems is not explicitly accounted for, this strategy often leads to sub-optimal solutions [Fat+01].
- The *iterative* strategy alternatively solves the plant and controller design sub-problems one after the other until convergence criteria are reached. Note that, as the problems are interdependent, a change in the plant parameters will degrade the optimality of a given controller and vice versa. Thus, in the general case, no mathematical proof of convergence towards the system-wide optimum exists. Nevertheless, this method has been widely used as, in some cases, each sub-problem could be solved individually at a moderate computational cost using linear matrix inequalities (LMI) [NK96; LLW05a].
- The *nested* strategy is bi-level: while the outer loop optimizes the cost function by altering only the plant design, the inner loop generates an optimal controller for each iteration of the plant. In the case where the parameters of the plant do not depend on the controller design, the global optimality of this approach can be guaranteed [Fat+01].
- The *simultaneous* strategy searches for an optimum by varying both plant and controller parameters at the same time. In principle, the system-wide optimum is reachable but this approach introduces additional mathematical complexity as the combined optimization problem is very often non-convex [Fat+01] or even non-smooth [Ala+13; Per+16; Den+17].

In our case, the simultaneous approach has been investigated. Thanks to the non-smooth optimization routine Systune, tunable plant parameters could be included into the controller synthesis problem. All the design and controller variables were concurrently adjusted in order to minimize the \mathcal{H}_∞ norms of specified transfers in a multi-objective synthesis setup.

3.2.3 Optimization problem setup

In a general sense, an optimization problem can be defined by a set of parameters to be varied (the parameter space), a metric that should be minimized (the cost function) and a set of constraints that must be fulfilled by every iteration of the plant-controller design. The constraints and the parameters are often process- and controller-dependent: the degrees-of-freedom of the optimization problem can be as diverse as the number of trays in a distillation column [LF94], the gain of the acceleration controller of an electric vehicle [Lv+18] or the size of an aircraft control surfaces [Den+17]. The cost function is usually related to higher level requirements such as the efficiency over a driving cycle [Lv+18] or the sizing of some critical components of the system which may have an impact on its overall performance and cost [Ala+13; Den+17].

The selection of the cost function minimization algorithm could also be investigated as part of the optimization strategy. However, in this study, the PCO problem is solved using a commonly-used controller synthesis routine for which substantial information is available in the literature [AN06b; GA11].

Cost function

The value of the cost function is used to quantify the degree of fitness of a given design with respect to its intended purpose. That way, the algorithm can discriminate between different designs and select the best one according to the chosen criterion.

In several papers [Mor92; RP99; KH00; DVT15], the formulation of the cost function was inspired by optimal control and includes a weighted sum of squared states or input variables, allowing for a tradeoff between disturbance rejection and control effort. Penalty functions may be included to implement constraints on the parameter values that may not otherwise be supported by the optimization scheme. Morris [Mor92] penalized flying-wing designs that required excessive control deflections to maintain trimmed flight while Deodhar [DVT15] handicapped airship candidates that required large ballast in order to account for the associated manufacturing challenge. The quadratic cost function may also be augmented to consider non-control-related objectives: Morris [Mor92] included a weighting term on aerodynamic drag of his designs while Reyer [RP99] added a penalty on the maximum speed requirement for an electric motor as well as its mass.

Alternative formulations were also introduced: Ravichandran [RWH06] used the integral of time multiplied by absolute error to quantify the tracking performance of a

robotic arm design. Lv [Lv+18] used a set of time-domain performance indexes obtained on benchmark cases in simulation to find the best controller and plant design. Similarly, Nikpoorparizi [NDV17] compared the energy efficiency of airborne wind energy systems for a given wind profile in simulation.

In our case, the structured \mathcal{H}_∞ synthesis framework provides options for defining both cost function and constraints: the norm of a given transfer can be declared as a "soft" constraint that will be minimized under the condition that other "hard" constraints are met. Modern synthesis tools based on non-smooth optimization also support additional metrics such as closed-loop stability margins and time-domain requirements for soft or hard constraints [Apk12].

Constraints

Optimization constraints may be enforced on the system performance indexes or on the value of its tunable parameters. This can be done for a variety of purposes such as finding a compromise between different design requirements or enforcing critical aspects of system performance: Niewoehner [NK96] uses \mathcal{H}_∞ norms to limit the flight path disturbance and the controls deflections of an aircraft struck by a wind gust while Lv [Lv+18] specifies lower bounds on top speed, maximum grade ability and braking effectiveness of his electric vehicle designs.

Constraints may also be used to restrict the parameter space in order to reflect physical limitations or stay within the domain of validity of the plant model: Kajiwara [KH00] imposed lower and upper bounds on the geometrical dimensions of its UAV fuselage in order to ensure feasible designs. They can also be applied to variables which are not degrees-of-freedom of the optimization problem but are indirectly affected by the design process: Denieul [Den+17] limits the amplitude and rate of control surfaces motions to avoid actuator saturation and unmodeled non-linearities in the control response while Kajiwara prescribes the sign of an aerodynamic coefficient in order to obtain a minimum-phase system.

Numerically, constraints are frequently expressed as inequalities on the numerical values of parameters or variables but the \mathcal{H}_∞ norm may also be used to bound the frequency response to a given excitation signal [NK96; Den16]. In that case, the value of the \mathcal{H}_∞ norm achieved by the optimizer quantifies the distance to the constraint bound: the smaller the norm, the further away the metrics are from their respective bounds, given that a norm greater than one implies that one or more constraints have been violated.

One can also leverage an useful property of the \mathcal{H}_∞ norm which is that all constraints on the norms of the mono-variable transfers are met if the bound on the multi-variable transfer is respected. The drawback is that the condition is only sufficient and introduces some conservatism on the synthesis results [YL03].

Also, the \mathcal{H}_∞ bounds do not directly constrain the amplitude of the time response of the system. Therefore, as suggested by Denieul [Den16], it would be interesting to include explicit time-domain criteria in the PCO problem. Apkarian [ARN11] and Chambron [CBA18] propose different approaches to ensure that the outputs of a linear system remain inside arbitrary bounds in the presence of an unknown disturbance.

This work proposes a means of circumventing the latter limitation by relying on a non-linear time-domain approach to complement the synthesis constraints. That way, frequency requirements on the actuator transfers can be relaxed as the effect of aerodynamic non-linearities on the controller output will be accounted for in the subsequent validation phase.

3.2.4 Control design

Controller structure

In order to solve the control side of the PCO problem, the designer must first decide on a structure for the controller. These considerations include the order of the controller, its centralized or decentralized scheme and the layout of its interconnections with the plant. For instance, Alazard [Ala+13] retains a decentralized proportional-derivative scheme for its attitude controller while Niewoehner [NK96] assumes a full state feedback.

The controller architecture may be constrained by the hardware or software limitations of the computer embedded on the system such as the maximum controller order [ANF] as well as the availability of state vector measurements. Also, some of the mathematical tools used to solve the co-design problem such as linear matrix inequalities may dictate the use of a full-order controller [Fat+02] [YL03].

The controller layout can also be tailored to the dynamics of the open-loop plant such as strong resonances in the frequency band of interest, non-minimum phase or integrator behavior. For instance, Strub [SB16] implements a roll-off filter to prevent excitation of unmodeled flexible dynamics of his experimental setup. Alternatively, the structure can be adjusted according to the performance objectives set for the closed-loop system: Lv [Lv+18] alternates between two controller layouts, each designed to provide best perfor-

mance according to a different figure of merit.

Synthesis methods

In the last two decades, significant progress has been made on the numerical methods available for solving the controller synthesis problem. Non-smooth optimization algorithms [AN06b] provided computationally efficient solutions to ill-conditioned problems such as fixed-structure \mathcal{H}_∞ controller synthesis. This methodology allows greater flexibility in the definition of the controller structure or the problem formulation compared to standard LQR or full-order \mathcal{H}_∞ problems [GA11].

Early efforts on plant-controller optimization [NK96; KH00; YL03; LLW05b] favored LMI-based algorithms to solve for the optimal set of controller gains under \mathcal{H}_2 and/or \mathcal{H}_∞ constraints. Since then, structured \mathcal{H}_∞ synthesis has gained popularity in the control community and was successfully applied to several PCO problems in aerospace [Ala+13; Per+16; Den+17].

The choice of a synthesis method is conditioned by the specific features of the plant-controller optimization problem, i.e the type of constraints and cost function, as well as the structure of the controller and the nature of the plant parameters. In most cases, the PCO problem is obtained by augmenting the controller synthesis problem with plant-related parameters and constraints [Ala+13; Per+16; Den+17]. However, there are examples in the literature where the controller gains are tuned outside of a dedicated controller synthesis framework. Morris [Mor92] uses a LQR criterion based on a reference model for step tracking, but this metric is augmented with the trimmed drag of the aircraft and additional penalty functions to form the cost function that is minimized.

The LQR synthesis allows a trade-off between disturbance rejection and control action minimization by minimizing the weighted sum of the control and disturbance energy. A variant of this method has been applied by Fathy [Fat+02] to improve the ride quality of an elevator fitted with an active control system.

The \mathcal{H}_∞ synthesis constrains the higher bound of the transfer function greatest singular value over the frequency domain. The mixed-sensitivity formulation [Kwa93] enforces such frequency templates on the system sensitivity functions. Alazard [Ala+13] used such templates to impose a roll-off on the controller response and reject sensor noise as well as unmodeled dynamics. Denieul [Den+17] applied the same method to enforce model-matching constraints on the flight dynamics of a blended-wing-body aircraft. \mathcal{H}_∞ synthesis can also provide stability and performance guarantees w.r.t parametric uncer-

tainty. Alazard [Ala+13] applies an iterative multi-model approach to find the worst-case parametric configurations and includes them in the next controller design step. Robust stability is guaranteed by evaluating the upper bound of the structured singular value.

Synergies between synthesis and optimization

Theoretical guarantees on solutions existence, optimizer convergence and system-level optimality can be obtained by associating a synthesis framework with an optimization strategy. For instance, Yang [YL03] derives necessary and sufficient conditions for the existence of PCO solutions under \mathcal{H}_∞ constraints. The convergence properties of two LMI-based iterative optimization algorithms are established, however system-level optimality is not guaranteed. Conversely, Fathy [Fat+01] shows that, under certain assumptions, a nested optimization scheme, including a LTI plant and a LQR controller, reduces to a static optimization problem and leads to global optimality.

3.3 Airframe-Controller Design Methodology

In this section, an original design framework is proposed in order to account for the effects of canard stall at the design stage by integrating a non-linear model of the projectile pitch dynamics into the tuning process. A plant-controller optimization scheme is implemented in order to limit the control effort by tailoring the static stability of the projectile. The end goal being to find an airframe-controller configuration that performs adequately over the widest possible range of angles of attack. The traditional multi-objective \mathcal{H}_∞ synthesis framework is augmented with airframe stability tailoring thanks to the LPV model of the projectile dynamics. Parametric uncertainties are used to penalize excessive actuator usage and promote configurations which are more robust to canard stall. After being checked against frequency templates and stability margins, the closed-loop performance of the co-designed configuration is systematically assessed in the non-linear domain using the semi-local aerodynamic model developed in the previous chapter. This workflow allows to pinpoint the eventual control issues, such as actuators saturation or controller wind-up, and adjust the synthesis constraints to mitigate them.

Quantification of airframe static stability

The degree of static stability of an aerodynamic configuration is defined by its static margin m_S which corresponds to the distance between the projectile center of pressure (i.e. the point of application of the sum of all aerodynamic forces) and its center of gravity, projected on its longitudinal axis and counted positively when the CP is aft of the CG. According to Kessler [Kes66, p. 53], this quantity can be expressed as a function of the aerodynamic derivatives by $m_S = -Cm_\alpha/Cz_\alpha$. Since the domain of angle of attack explored in this study is restricted to $\pm 10^\circ$, it is expected that the projectile lift curve slope Cz_α will remain strictly positive. Thus, the static stability can be deduced from the sign of the pitch stability coefficient Cm_α : negative values imply a restoring moment in response to an AoA disturbance and a stable airframe. The case $m_S = 0$ corresponds to an absence of restoring moment and is referred to as a neutrally stable projectile. A positive Cm_α indicates a diverging response to AoA disturbances and an unstable aerodynamic configuration.

3.3.1 Selection of the tunable geometrical parameter

The projectile layout described in chapter one, section 1.1.2 can be described by the geometric parameters pictured on Figure 3.2. The chosen degree-of-freedom should significantly alter the airframe static stability while having little influence on other design criteria which are outside the scope of this study, such as maximum lift-to-drag ratio. Also, due to the experimental nature of this study, the configuration should be easy and quick to adjust and must comply with the internal volume limitations of the actual mock-up.




Figure 3.2 – Parametric projectile geometry (blue arrow shows the fins position w.r.t airframe datum section)

Since the aft fins represent approximately 70% of the lifting surfaces area and are located further away of the CG than the canards, their longitudinal position x_F has a significant influence on the location of the projectiles center of pressure and its static stability. Moreover, varying the fins position affects neither area nor the aspect ratio of the lifting surfaces so the projectile aerodynamic performance should not be significantly affected (c.f section 2.6.3). Finally, the new mock-up layout presented in section 1.1.2 features a 3D printed sleeve which allows the fins longitudinal position to be adjusted over a 175 mm range of motion. As a result, the most forwards setting produces a statically unstable projectile while the fully aft position gives a comfortable static margin of approximately 0.66 caliber, well above the common standard of 0.5 cal [Kes66]. Since this tuning range appears to be adequate, x_F has been chosen as the geometrical parameter to be varied in order to adjust the projectile static stability.

Influence of fins position on open-loop pitch dynamics

In order to get more insight on the outcome of the co-design, the parameter space of the fins position has been swept during the model identification test campaign, allowing to plot the dynamics and frequency response of the open-loop projectile for increasing values of x_F .

Figure 3.3 depicts the root locus of the LPV pitch dynamics model G_{dynmdl} . For fins position between 130 and 149 mm, the dynamics consists of two real poles, one on each side of y-axis. The system which has its right hand pole in zero ($x_F = 0.149$ m) corresponds to the neutrally stable projectile ($Cm_\alpha = 0$). As the finned sleeve is brought back, the poles converge towards each other, merging into a doublet of stable poles at $\omega = 0.876$ rad/s for $x_F = 0.150$ m. Then, they split into a pair of complex conjugate poles which grow further apart as the magnitude of their imaginary parts increase. The natural frequency of the poles increase as the fins are shifted aft: this could be expected as the transfer function of the model presented in the previous chapter (eq. (2.19)) shows that the pulsation only depends on the pitch stability coefficient which magnitude increases with x_F . Conversely, the pole damping is a function of both aerodynamic damping and stability coefficients. Thus, its variations are more complex, with a decreasing ratio between 160 mm and 220 mm followed by an increase from 250 mm onward.

According to the Bode plot of Fig. 3.4, the systems featuring a RHP pole ($x_F < 0.149$ m) have moderate static gain, no resonance and a phase close to -180 degrees. While the peak gain increases as the fins are shifted aft, the static gain decreases, which is sensible




Figure 3.3 – Root locus of the plant for various fins positions

given that the restoring moment rapidly outweighs the control surfaces moment (c.f 2.4.2). An exception being the configuration with a pole in zero which has infinite static gain. Note that all of the projectiles with stable complex conjugate poles ($x_F > 0.149$ m) present a gain resonance that shifts to higher frequencies as the fins position is increased.

From these plots, one may wonder which fins position offers the most favorable dynamics for closed-loop control. A high static gain is definitely desirable as it allows to limit the canards deflection amplitude and delay the onset of aerodynamic non-linearities. Plants with lower resonance magnitude and higher damping ratio may also present less of a challenge for the control system, allowing the actuator bandwidth requirement to be relaxed. For all those reasons it is expected that the co-design will favor neutral or marginally stable configurations. Also, given its outstanding characteristics, the neutral-stability airframe may be the object of a dedicated study aimed at assessing its potential in the frame of agile projectile design.




Figure 3.4 – Bode plot of the plant for various fins positions

3.3.2 Models of the adjustable-stability projectile

The layout of the projectile lifting- and control surfaces can be described by a set of geometric parameters among which the fins position is the driver of the projectile static stability. In this section, the LTI model of the open-loop pitch dynamics is built-up and included into the synthesis scheme, using parametric uncertainties to represent the non-linearities affecting the control surfaces.

Synthesis model build-up

The plant model implemented into the airframe-controller optimization framework is composed of the LPV model G_{dynmdl} presented in the previous chapter (c.f. Equation (2.19)), augmented with a model of the control and actuation system G_{CAS} which is introduced herein. The combination of the two models is of order 9 and features a single input δ_m and output θ .

The CAS model includes the dynamics of the servomotors driving the canards and accounts for the delay induced by the on-board processing chain. The computational

delay corresponding to one sampling period $T_s = 0.01$ s is modeled by a first-order Pade approximation, while the zero-order hold behavior is captured by a first-order low-pass filter with a time constant of $0.5T_s$. The servo-actuators implement a position control loop using the feedback from an absolute magnetic sensor. The servo dynamics are represented by a third-order model which closed-loop transfer function is given by Strub [SB16]:

$$G_A(s) = \frac{k_{\text{pos}}K(k_p s + k_i)}{\tau s^3 + (1 + Kk_p)s^2 + K(k_i + k_{\text{pos}}k_p)s + Kk_{\text{pos}}k_i} \quad (3.1)$$

Thus, the CAS model is represented by the following equation:

$$G_{CAS}(s) = \frac{1 - 0.5T_s}{(1 + 0.5T_s)^2} \frac{k_{\text{pos}}K(k_p s + k_i)}{\tau s^3 + (1 + Kk_p)s^2 + K(k_i + k_{\text{pos}}k_p)s + Kk_{\text{pos}}k_i} \quad (3.2)$$

The servo controller gains k_{pos} , k_p and k_i are found with \mathcal{H}_∞ synthesis while the servo motor parameters K and τ are estimated by system identification. The values of all these coefficients are provided in Table 3.1.

Parameter	K	τ	k_{pos}	k_p	k_i
Value	326.2 deg.s ⁻¹ V ⁻¹	0.0182 s	29.4 s ⁻¹	0.052 V.s.deg ⁻¹	0.3941 V.deg ⁻¹

Table 3.1 – Actuator model parameters

Modeling of parametric uncertainties and canard stall

Two limitations of the LPV projectile model should be accounted for in the synthesis process. First, the static stability, pitch damping and pitch control moments are estimated using surrogate models based on identified aerodynamic coefficients. A degree of uncertainty associated with the estimation of these parameters from measured time responses is present, adding up to the fit error of the surrogate models. Although an attempt was made to quantify the uncertainty on the parameter estimates (c.f section 1.3.3), its results were considered to be overly optimistic as the some of the disturbing phenomena fall outside the assumption of random noise. Thus, a parametric uncertainty of +/- 10% was assigned to the pitch stability and pitch damping coefficients.

The second limitation applies to the modeling of the aerodynamic moment generated by the control canards. The LPV structure imposes a linear behavior which is only valid for small controls deflections at low angles of attack. Outside of this domain, the canards

experience stall and the aerodynamic control moment is reduced. The approach that has been implemented to account for those non-linear effects consists in assigning large unbalanced parametric uncertainty to the pitch control coefficient: the upper bound of +10% is retained but the lower bound is reduced to -50% so that part of the non-linear domain of the trim maps falls into the region covered by the uncertain model, as shown on Fig. 3.5.




Figure 3.5 – Comparison of uncertain pitch control model with measured trim maps for different fins positions

The value of this lower bound is a compromise between covering the part of the trim map which is most likely to be explored during normal operation and being excessively conservative while operating in the linear range. Note that both considerations are affected by the choice of the geometric parameter, as less stable configurations show a more pronounced break between the linear and the non-linear domain. For $x_F = 0.175$ m (right plot), the uncertain model covers all operating points up to AoAs of approximately 7° , which is close to the forecasted flight domain of the projectile (see section 3.4.2).

3.3.3 Multi-objective synthesis and optimization

The synthesis scheme pictured in fig. 3.6 is arranged according to the different objectives and constraints that may be considered for airframe-controller design. In the pitch control case, all the transfers are mono-variable so that the input and output sensitivity functions are equivalent. Thus, the whole synthesis scheme is very similar to the classical S/KS mixed sensitivity approach, with an additional weighting template on GS to

improve robustness and disturbance rejection properties.




Figure 3.6 – Multi-objective synthesis framework

Three objectives are considered for tuning the controller: the performance requirement is expressed as a disturbance rejection constraint using the output sensitivity function $T_{\theta_d \rightarrow \theta_{tot}}$ multiplied by the weighting filter W_S to form the performance channel $T_{\theta_d \rightarrow z_{DR}}$. The actuator limitations are translated by the closed-loop transfer from the pitch attitude reference to the actuator input $T_{\theta_r \rightarrow \delta m}$ multiplied by the filter W_{KS} to form the control effort channel $T_{\theta_r \rightarrow z_{act}}$. The transfer from the actuator input disturbance δm_d to the output θ_{tot} is constrained by the filter W_{GS} to improve the low-frequency disturbance rejection properties and prevent exact simplification of the system flexible modes which could lead to robustness issues. Hence, this transfer is referred to as the robustness channel $T_{\delta m_d \rightarrow z_{rob}}$.

As in section 1.4.1, W_S and W_{KS} are chosen as first order lead-lag filters characterized by their respective low-frequency gain, high-frequency gain and -3 dB bandwidth (c.f. fig 1.16). The frequency template $1/W_S$ is shaped like a high-pass filter in order to provide good reference tracking and disturbance rejection performance. Conversely, the template $1/W_{KS}$ features low-pass behavior to attenuate the effect of measurement noises, preserve the actuators from high frequency content and limit the control energy. Depending on the controller structure, $1/W_{GS}$ can be a static gain bound to limit the GS peaks associated with pole-zero cancellations [Apk09, p 73], or a high-pass filter to focus on the disturbance rejection aspects. Note that the KS template provides a means to limit the control effort but is not intended to accurately represent the CAS limitations, some of which can only be captured with the non-linear model presented in the next section. The time-domain performance of the control system in the non-linear region should also be accounted for during the tuning of the frequency templates.

3.3.4 Non-linear simulation of closed-loop performance

During the tuning process, systematic assessment of each iteration of the airframe-controller design is necessary to ensure that its performance will not be excessively degraded by non-linear aerodynamic phenomena. Existing non-linear simulators based on look-up tables fed by CFD [VTS22] assume the control response to be in the linear domain and may not capture the effect of canard stall. Thus, a new environment depicted in Figure 3.7 has been developed around the semi-local aerodynamic model presented in section 2.7.3.

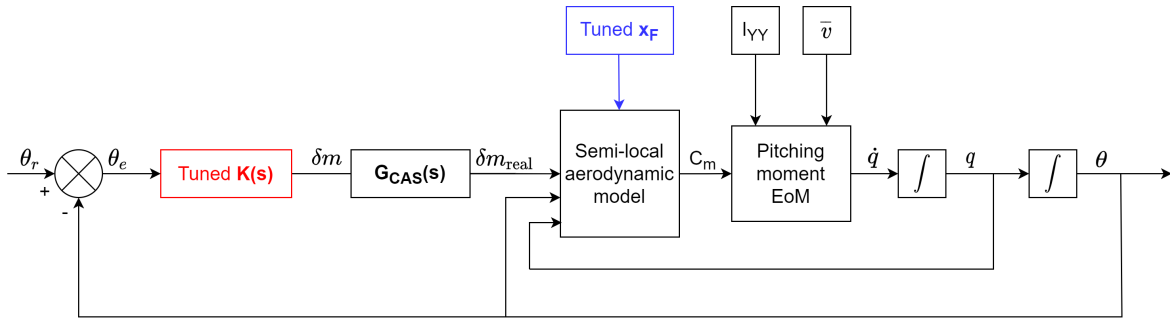


Figure 3.7 – Non-linear pitch dynamics simulator

This framework differs from the synthesis model in that, for a given fins position, the state-space representation of the pitch dynamics based on constant aerodynamic derivatives is replaced by a non-linear representation of the aerodynamic moment coupled with Euler’s equation of pitching motion. The equation of pitching motion is implemented as introduced in the first chapter (section 1.2), using the actual pitching moment as an input instead of its linearized decomposition factoring aerodynamic derivatives. The aerodynamic moment Cm is computed from the pitch attitude, rate and control deflection according to the flowchart pictured in Fig. 3.8. Details of the implemented equations are available in appendix A.1.

A new figure of merit based on the time response of the control system is defined to assess its performance in the face of aerodynamic non-linearities. The chosen criterion is based on the size of the projectile trim map, with the additional constraint that the projectile dynamics must be quick and well damped (upper bound on the settling time). This metric, referred to as the projectile *max control AoA*, is defined as the greatest angle of attack that can be reached (within a $\pm 5\%$ band) from zero pitch angle in less than two seconds. Figure 3.9 shows that, at max. control AoA, the time response of the system




Figure 3.8 – Flowchart of the semi-local model

fits into the $\pm 5\%$ band from 2 seconds onward. This metric is evaluated iteratively by running non-linear simulations with increasing target AoAs until the computed settling time reaches two seconds.

To sum up, this integrated synthesis and simulation environment allows to rapidly converge to an airframe-controller design which satisfies traditional linear performance, robustness and stability criteria but is also much more likely to conserve its performance in spite of non-linear aerodynamic phenomena. A performance metric is proposed to assess the design performance in the non-linear domain.




Figure 3.9 – Graphical interpretation of the max. control AoA

3.4 Synthesis and Simulation of Agile Projectiles

This section presents an application of the co-design methodology to the maximization of the projectile angle of attack control domain under stringent actuator limitations. Two different controller structures are implemented into the co-design scheme, resulting in distinct airframe geometries with dissimilar performance levels. The most promising configuration is benchmarked against a projectile designed according to current best practices in order to emphasize the benefits of ACO. Finally, controller synthesis is performed with an neutrally stable projectile so that the performance of this peculiar aerodynamic configuration can be evaluated.

Definition of the optimization problem

In order to describe the optimization problem, one should define its cost function, degrees-of-freedom and optimization constraints. Here, the parameters to be adjusted are the controller gains as well as the longitudinal position of the fins x_F . The \mathcal{H}_∞ norms of the transfers $T_{\theta_d \rightarrow z_{DR}}$, $T_{\theta_r \rightarrow z_{act}}$ and $T_{\delta m_d \rightarrow z_{rob}}$ defined in the previous section give an insight on the performance and actuator usage of the airframe-controller designs. The co-design problem aims at improving the performance of the flight control system in spite of aerodynamic and mechatronic limitations. Therefore, the norm of the performance transfer represents the cost function and shall be minimized without any lower bound,

under the condition that the actuator constraints and stability requirements are met.

As mentioned in the paragraph on multi-objective synthesis, the output sensitivity acts as a performance index for disturbance rejection and reference tracking. Thus, it has been declared as a "soft" constraint in the synthesis framework. All the other metrics prevent unsatisfactory behavior of the closed-loop system and often conflict with the performance criterion. That is why minimizing the control effort, stability and robustness criteria beyond their maximum admissible levels can reduce system performance. Hence, these requirements have been set as optimization constraints and marked as "hard" constraints for the synthesis algorithm. The problem was solved with the MATLAB Systune routine which relies on non-smooth optimization algorithms [AN06a]. This function was selected for its ability to implement both soft and hard constraints, and mix different constraint types such as frequency templates and stability margins. In order to circumvent the lack of global optimality guarantee of this method, eleven runs were conducted with different sets of randomized initial parameter values that were automatically generated by the synthesis routine.

Controller structure candidates

Two controller structures have been trialed: the first attempts were made with a full order array of feedback gains that was cast in state-space form similarly to Berard [BBS12, p 159]. The rationale behind this complex structure was to let a lot of degrees of freedom for the structured synthesis algorithm to tune, and attempt to bring down the \mathcal{H}_∞ norms of the frequency transfers as close as possible to the theoretical minimum. However, in order to improve the numerical tractability of the synthesis process, the state matrix is restricted to a tridiagonal form, reducing the number of tunable parameters without affecting the order of the compensator. Because of this choice, the controller structure differs from a canonical full-order compensator and the optimality of the \mathcal{H}_∞ norm is no longer guaranteed.

The controller G_{comp} is SISO with θ as the input and δ_m as the output. Given that the augmented plant is of 9th order, the respective sizes of A_{comp} , B_{comp} , C_{comp} and D_{comp} are 9-by-9, 9-by-1, 1-by-9 and 1-by-1, for a total of 100 gains which has been reduced to 44 by shaping the A_{comp} matrix. Its implementation into the feedback loop is shown on fig. 3.10.

Due to its the large numbers of degrees-of-freedom, this setup proved to be quite challenging to tune. Moreover, the tuned controller response turned out to be very similar




Figure 3.10 – State-space controller structure

to the one of a classical PID layout (c.f fig. 3.12). Thus, a less complex dual-loop controller was introduced. The layout pictured in fig. 3.11 is based on cascaded pitch attitude and rate feedback loops, with a P-IP controller in the direct chain. This structure has been commonly used for aerospace control applications [MR05] [TP21] and features only three control gains. The rate feedback gain k_q provides a stabilizing effect while the attitude gain k_θ and the integral gain k_{qi} ensure reference tracking without steady-state error.




Figure 3.11 – Dual-loop controller structure

3.4.1 Co-design with state-space controller

Tuning of the frequency templates and stability margins

The determination of the frequency templates often requires multiple iterations even though guidelines do exist [Kwa93] and knowledge of the open-loop behavior can be leveraged to streamline the process [RRL20]. Depending on the controller structure, some of the closed-loop properties such as disturbance rejection due to integral action may be prescribed by the control layout or have to be enforced by the adequate transfers. Thus, the shape of the weighting filters differ greatly and can only be chosen on a case-by-case basis.

For the state-space controller, the initial set-up converged to a neutrally stable air-frame paired with a controller that showed very poor disturbance rejection properties at low frequencies. Moreover, it was noted that the controller zeros tended to cancel out the poles of the pitch dynamics model, which could be undesirable since the open-loop plant featured uncertain flexible modes. The definitive synthesis set-up features increased stability requirements as well as a more conservative performance target. The robustness filter W_{GS} was ultimately chosen as a +16 dB static gain bound shown on fig. 3.16. The bandwidth of the closed-loop sensitivity template (fig. 3.14) has been relaxed to 2.5 rad/s, with a low-frequency asymptote of 0.001 and a high-frequency gain of 1.3. The margins at system output and the actuator input are specified at 12 dB gain and 30 degrees phase.

The characteristics of the control effort weighting template W_{KS} shown on fig. 3.15 were based on Strub's reference study [SB16] which featured the exact same servo-actuator. The only minor change is that the maximum bandwidth requirement was increased by 50% in order to cope with the potential demand of less stable projectile configurations. Yet, the specified bandwidth is only half of the theoretical value attainable by the servo controller, so there is still comfortable performance headroom. Thus, the low-frequency and high-frequency gains retain their respective values of 0.001 and 2 while the filter bandwidth is expanded to 15 rad/s.

$W_S (T_{\theta_d \rightarrow \theta_{tot}})$		$W_{KS} (T_{\theta_r \rightarrow \delta_m})$		$W_{GS} (T_{\delta m_d \rightarrow \theta_{tot}})$	
$k_{S,LF}$	0.001	$k_{KS,LF}$	0.001	k_{GS}	6.310
$k_{S,HF}$	1.3	$k_{KS,HF}$	2		
ω_S	2.5 rad/s	ω_{KS}	15 rad/s		

Table 3.2 – Parameters of the weighting filters (state-space controller)

Frequency-domain results

The frequency response of the tuned state-space controller pictured on Figure 3.12 shows integral action at low frequencies up until 7 rad/s where derivative action takes over. This last effect is band-limited as roll-off starts around 20 rad/s with a slope of 20 dB per decade. Overall, this behavior is very similar to a PID controller combined with a roll-off filter, save for the steeper gain drop before the gain minimum which contributes to stability margins with a steep phase increase between 6 and 8 rad/s.

Figure 3.13 presents the root locii of the tuned plant and the state-space controller.




Figure 3.12 – Bode plot of the tuned state-space controller

It reveals that the pair of dominant poles of the adjusted open-loop pitch dynamics is located at 7.58 rad/s with a damping ratio of only 0.113. Their resonant frequency is computed at 7.48 rad/s, close to the dip in both actuator transfer and loop gain. The controller attempts to cancel these poles with a pair of zeros located at 6.78 rad/s, hence GS is the highest in this area. There appears to be another pole-zero cancellation within the controller itself at around 2.57 rad/s, even though the robustness transfer is not as affected. The next pair of controller poles lies at 5.14 rad/s with a damping ratio of 0.944 and is free of nearby zeros. For this controller layout, the fins position x_F corresponding to the optimal plant dynamics has been found 0.211 m behind the reference datum, which corresponds to a static margin of approximately 0.30 calibers. This value, close to the middle of the parameter range, can be interpreted as the outcome of a trade-off between conflicting effects as none of extreme fins positions gives the best results.

The performance template of fig. 3.14 reveals that the sensitivity target is met on all frequencies but a small band between 2 and 3 rad/s. At low frequencies, the sensitivity is the inverse of the loop gain, thus it is the most affected by parametric uncertainty. The resonance is marginal with a peak gain of 1.31 between 10 and 20 rad/s. At higher frequencies, the open-loop gain drops and the transfer converges to a value of one. The -3 dB bandwidth of the closed-loop sensitivity is 1.84 rad/s which is lower than the template bandwidth because of the different high-frequency gain asymptotes.

The reference to actuator transfer pictured in fig. 3.15 is constrained by its template




Figure 3.13 – Root locus of optimized plant (blue) and tuned controller (red)

above 20 rad/s. At low frequencies, this transfer behaves like the inverse of the plant and is thus very sensitive to parametric uncertainties. At high frequencies, $KS(s)$ is the image of the controller, enforcing the roll-off properties of the compensator. The transfer bandwidth remains much lower than its prescribed value at only 2.84 rad/s for the nominal system. This is due to a steep notch in the frequency response, which reaches an attenuation of -25 dB at 6.95 rad/s. The low frequency gain on the actuator channel also stays under the template, with a nominal value of 0.910 and a worst-case magnitude of 1.43.

The peak of the robustness transfer illustrated on fig. 3.16 is very close to the chosen maximum bound on the frequency band between 7 and 8 rad/s, highlighting the limitations of this controller structure. $GS(s)$ matches the inverse of the controller at low frequencies but behaves like the plant at high frequencies. The gain margin at the system output given by fig. 3.17 is saturated at its minimum specification: the critical frequency band spreads on one decade from 2 to 20 rad/s, with the exception of the vicinity of 8 rad/s where there is a strong dip in the loop gain. The constraint on the phase margin is not active with more than 30 degrees of excess margin at the system output. The delay margin is 0.462 s, which is comfortable since the embedded control loop runs at 100 Hz on the projectile mock-up.




Figure 3.14 – Disturbance rejection transfer S and associated template




Figure 3.15 – Control effort transfer KS and associated template

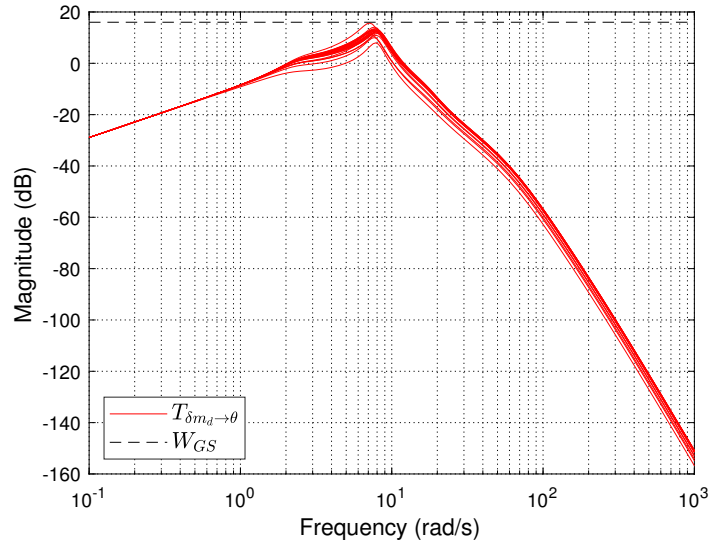


Figure 3.16 – Robustness transfer GS and associated template

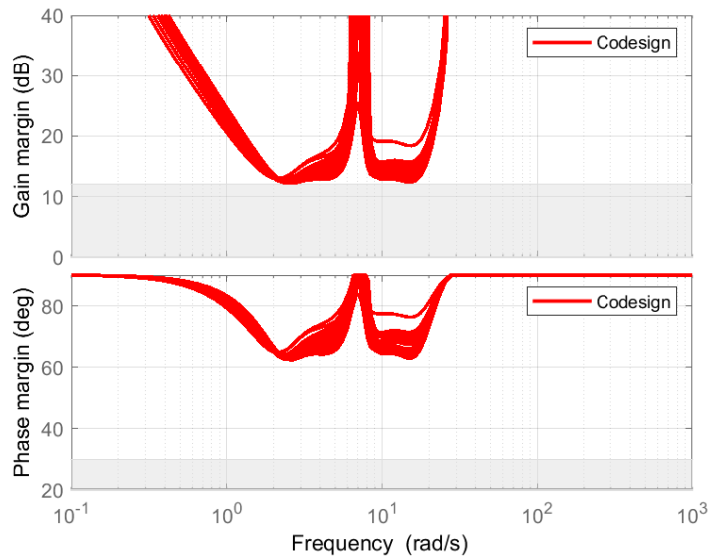


Figure 3.17 – Stability margin at the output θ

Time response analysis

Once the linear performance, stability and robustness requirements are satisfied, the non-linear response of the guided projectile configuration is assessed thanks to the simulator presented in section 3.3.4. Two objectives are pursued: first, to pinpoint the effects of canard stall on the closed-loop response and then, to quantify the ability of the co-design to mitigate the performance limitations due to this phenomenon. The former goal is met by injecting a series of step inputs of increasing amplitudes. The pitch response and the actuator time histories are displayed in fig. 3.18. The latter aim is achieved using the maximum control AoA criterion defined in section 3.3.4. To that end, the amplitude of the reference input is increased until the response time of the non-linear step response exceeds two seconds. The maximum pitch angle target and the associated linear and non-linear closed-loop responses are plotted on fig. 3.19.




Figure 3.18 – Step tracking of the state-space controller for increasing amplitudes

The set of step responses of fig. 3.18 reveals that significant differences between the linear and non-linear pitch attitude histories develop as the step amplitude is increased. The first thing that can be noted is the discrepancy between the rising- and the falling edge of the response: the non-zero trim condition can be met after an over-damped transient while the return to zero pitch always features an overshoot. This shows that the system behavior at large AoA differs from the prediction of the linear model, which is consistent with the physics of canard stall.

For a 2 degree step, both non-linear pitch attitude and controls deflection match well with their linear counterpart. As the target attitude increases, the control deflection required to reach steady-state turns out to be several times greater than the linear predic-




Figure 3.19 – Projectile step response at maximum control AoA

tions: it is 1.8 times larger for the 4° step and 2.9 times bigger for the 6° step. Therefore, the pitch response of the higher fidelity closed-loop projectile model is more sluggish: the 95% rise time increases from 0.8 s to 2.1 s for the 4° step and 3.6 s for the 6° step. The non-linearities also affect the controller response around zero, with a slower descent time and a more pronounced overshoot: the 67% descent time increases from 0.49 s to 0.87 s for a 4° step and 1.37 s for a 6° step while the overshoot grows from 9.4% to 17.8% for the 4° step and 20.2% for the 6° step.

For pitch steps of 8 degrees and up, the commanded deflection reaches its saturation bound, capping the trim attitude at 6.72 degrees. Because of that, the state-space controller suffers from wind-up as the unsaturated control action constantly increases for as long as the attitude error persists. Thus, when the reference switches to zero, the control action needs time to decrease beyond its maximum value, introducing lag in the system response. The severity of that effect depends on the magnitude of the residual attitude error and the duration of the saturated condition. Consequently, for practical use cases, the target attitude needs to be bounded below 6.7 degrees in order to prevent this condition from occurring.

Despite promising frequency-domain results, the edge-of-envelope performance of the state-space airframe-controller design appears to be underwhelming with a maximum control AoA of only 3.83 degrees. Fig. 3.19 shows that the time response to this command presents high-frequency oscillations affecting the pitch attitude predicted by the non-linear model. The magnitude of the actuator deflection seems reasonable with a steady-state non-linear value just above 6 degrees, which should ensure that only moderate

control moment non-linearities should be encountered. Yet, this design turns out to be too slow and badly damped, and could not be significantly improved through the numerous attempts at frequency template tuning. Thus, the focus has shifted to the more simple and easier to tune dual-loop controller structure.

3.4.2 Co-design with dual-loop controller

Tuning of the frequency templates and stability margins

Since the dual-loop structure explicitly includes an integrator, the low-frequency disturbance rejection properties turned out to be much easier to obtain. Moreover, this layout requires less stringent stability margins to prevent excessive oscillations in both linear and non-linear time responses. Therefore, the gain and phase margin requirements were relaxed to 8 dB and 30 degrees. The bound on the robustness transfer was kept at 16 dB, even though this controller seems less prone to quasi-exact pole-zero cancellations, probably due to its lower number of degrees-of-freedom.

As for the performance target, the weighting filter bandwidth is set to 3.5 rad/s which represents a more ambitious objective compared to previous work [RRL22], while the high-frequency asymptote is kept to its precedent value of 1.3 that turned out to be a good compromise between performance and control effort. As usual, the low-frequency bound on S is set to a very small nonzero value in order to suppress steady-state error. Given that the servo control loop and projectile optimization range have been left unchanged, there is no reason to adjust the control effort template found in the previous example.

$W_S (T_{\theta_d \rightarrow \theta_{tot}})$		$W_{KS} (T_{\theta_r \rightarrow \delta_m})$		$W_{GS} (T_{\delta m_d \rightarrow \theta_{tot}})$	
$k_{S,LF}$	0.001	$k_{KS,LF}$	0.001	k_{GS}	6.310
$k_{S,HF}$	1.3	$k_{KS,HF}$	2		
ω_S	3.5 rad/s	ω_{KS}	15 rad/s		

Table 3.3 – Parameters of the weighting filters (dual-loop controller)

Frequency-domain results

The frequency response of the tuned dual-loop compensator presented on Figure 3.20 is typical of a PID controller. The low-frequency integral action is present up until $\omega = 6.75$ rad/s where the derivative action takes over. The maximum attenuation is 6.49 dB, much




Figure 3.20 – Bode plot of the tuned dual-loop controller

less pronounced than for the state-space structure. There is no controller roll-off since the derivative effect is obtained by a rate feedback which is already filtered by the plant dynamics (i.e. the pitch damping effect).

The pole-zero map of the tuned plant and controller is depicted on fig. 3.21. The tuned airframe dominant poles exhibit a natural frequency of 4.25 rad/s and a damping ratio of 0.197, slower but better damped than the previous airframe-controller design. These dynamics are obtained for an optimal fins position of 0.168 m which corresponds to a static margin of approximately 0.22 calibers. The compensator has one pole at the origin and two conjugate left-hand-plane zeros that have a frequency of 3.46 rad/s and a damping of 0.474.

The performance template of fig. 3.22 reveals that the system bandwidth is too low to fit into the template and the output sensitivity spikes above its upper bound. For a nominal plant, the bandwidth of the performance transfer is only 1.56 rad/s, while the worst-case sensitivity peaks at 1.61. The maximum is reached at the frequency of 4.3 rad/s, close to the resonance of the tuned plant dominant poles computed at 4.1 rad/s. These metrics indicate that the reference tracking and disturbance rejection properties of the new controller structure may be slightly worse than for the state-space layout but, given the limitations of the synthesis model in capturing the system non-linearities, the simulation results presented in the next paragraph are more likely to depict actual performance levels.




Figure 3.21 – Root locus of optimized plant (blue) and tuned controller (red)




Figure 3.22 – Disturbance rejection transfer S and associated template

The diagram 3.23 displays comfortable margins between the control system KS transfer and the template representing the actuator limitations. Compared to the previous synthesis results, there is significant improvement across the board: the magnitude of the static gain has been reduced threefold down to 0.290 while the transfer bandwidth has dropped to 2.18 rad/s. But perhaps the most striking difference is that the gain asymptote now remains 25 dB under the prescribed template, which indicates that the




Figure 3.23 – Control effort transfer KS and associated template

actuator bandwidth requirements could be significantly relaxed. The 11 dB gap between the closed-loop GS transfer and the constraint of fig. 3.24 also reveals that the structured controller has much better robustness properties than the state-space layout presented in the previous section.

Among all the optimization constraints, only the stability margins requirements of fig. 3.25 are active. Both input and output gain margins are saturated at their minimum value of 8 dB, while the phase margins stay above 45 degrees. The respective gain minima are reached around 14.6 rad/s at the input and 4.0 rad/s at the output. The latter frequency is very close from the peak of the sensitivity plot, suggesting that both stability and control performance requirements are critical around the open-loop resonance. The delay margin is evaluated at 0.0945 s, significantly less than in the previous case but still approximately one order of magnitude higher than the embedded control loop time period.




Figure 3.24 – Robustness transfer GS and associated template




Figure 3.25 – Stability margins at the actuator input δ_m (left) and at the output θ (right)

Time response analysis

The step sequence time histories of fig. 3.26 indicate that the alterations in the projectile response are similar in nature to those observed with the state-space controller. However, the magnitude of all of those effects is significantly attenuated: for instance, the 95% rise time increases from 0.78 s to only 0.83 s at 4° and 1.56 s at 6° . Favorable effect of the non-linearities can even be observed, with the falling edge overshoot being reduced from 7.3 % to 2.4 % for the 10° step.

No wind-up is encountered as the control deflection remains below 10° even for the




Figure 3.26 – Step tracking of the dual-loop controller for increasing amplitudes




Figure 3.27 – Projectile step response at maximum control AoA

greatest step command. This is due to the lower actuator transfer gain of the linear system, combined with what seems to be a lower sensitivity to the non-linearities of the control surfaces response. It could be inferred that the new controller structure and the optimizer choice of reducing the airframe static stability have led to a more robust design, as forecasted by the lower GS peak in fig. 3.24.

These improvements translate into a much greater flight envelope, with a maximum AoA of 7.30° , obtained with slightly smaller controls deflection and no oscillation of the pitch response. Finally, the comparison between the different airframe-controller designs at the edge of their respective AoA envelopes (fig. 3.28) demonstrates that the rapidity and damping of the guided projectile have been greatly improved in comparison to the state-space co-design.





Figure 3.28 – Comparison of the airframe-controller designs steps to their respective max. control AoA

3.4.3 Simulation benchmark against stable configurations

Now that a promising airframe-controller design has been found, it can be leveraged to assess the benefits of co-design with respect to conventional design practices. For that, a reference configuration based on a fixed static margin of 0.5 calibers is introduced. This value, determined from best practices commonly used in unguided ammunition [Kes66] or guided projectile [Fre+18] design, corresponds to a fins position of 0.277 m which is greater than the optima found in both co-designed configurations. As controller is synthesized for this configuration using the same dual-loop layout as in the last co-design setup.

However, the synthesis templates have been adapted to the dynamics of more stable airframes. Notably, due to the increased control effort required to overcome the projectile natural stability, the low frequency asymptote of the reference to actuator gain template pictured in fig. 3.30 has been raised to a magnitude of 4. The output sensitivity objective has also been relaxed, with a high frequency bound of 1.2 and a bandwidth of 2 rad/s. This is to account for the "waterbed effect" [RLD10] where overly strict templates on the performance channel could lead to the control effort constraint becoming unfeasible.

Figure 3.29 – Output sensitivities S comparison

Frequency-domain comparison

The results of the most promising dual-loop ACO have been overlaid with the transfers of the previously mentioned stable projectile configuration. The sensitivity plot of fig. 3.29 indicates that the standard projectile features similar, if not slightly better, disturbance rejection and reference tracking properties than the co-designed one. However, the example of section 3.4.1 has shown that frequency domain results could be misleading as nonlinearities may have a significant influence on control performance.

The advantages of co-design are more apparent in Figure 3.30 which describes the actuator usage of both configurations. The steady-state gain and bandwidth of the optimal airframe-controller design are significantly lower than the ones of the more stable configuration, with the nominal values respectively coming at 0.290 and 2.43 rad/s versus 1.77 and 3.40 rad/s. For the worst combination of uncertain parameters, the standard configuration requires amplification factors up to 3.8 between the attitude reference and the controls deflections ($KS(s)$ static gain), well above the limit of 2 that has been introduced by Strub [SB16] and conserved for the previous tuning examples. Thus, it appears very likely that this controller will command large actuator inputs and exceed the linear range of the canards pitching moment, as verified in the next paragraph. Thus, the performance estimates based on the linear model are most likely to be excessively optimistic and, in practice, the optimal projectile may outperform the standard design.


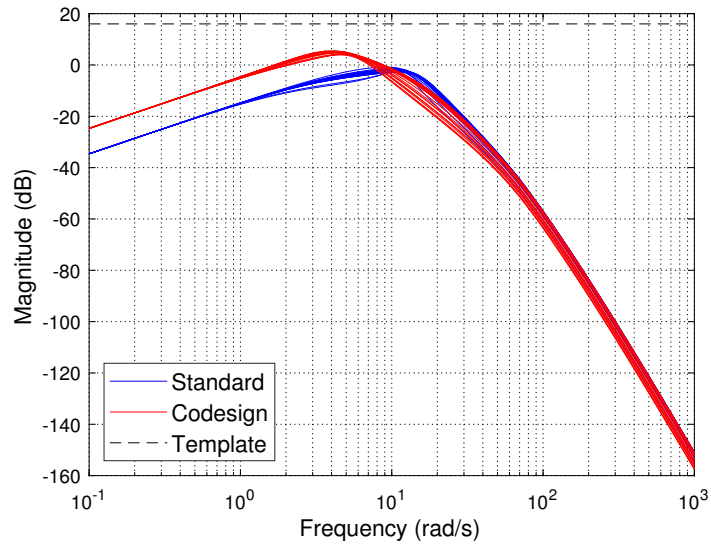


Figure 3.30 – Reference to actuator transfers KS comparison

In any case, robustness to input perturbation and exact pole-zero cancellations appears not to be a significant concern, as the responses of both systems provide comfortable margins with respect to the maximum gain requirement of fig. 3.31.

Figure 3.31 – Actuator to output transfer GS comparison

Time domain comparison

The responses of the two configurations computed with the non-linear model are overlaid in both a reference tracking and a disturbance rejection scenario. Figure 3.32 compare the step responses for an amplitude of 5 degrees while fig. 3.33 shows the effect of a disturbance generated by a 5 degree step on the pitch control channel.




Figure 3.32 – Reference tracking of standard and optimized airframes

The reference tracking case confirms the suspicions regarding the excessive actuator usage of the stable configuration: the pitch control reaches its saturation limit in less than two seconds, having achieved an attitude of only 4.1 degrees. Although its rise time is




Figure 3.33 – Disturbance rejection of standard and optimized airframes

not as fast, with 0.57 s instead of 0.36 s, the co-designed projectile is able to meet the 5° attitude command with ease, having used less than 2.3 degrees of controls deflection. The non-linearities also degrade the aspect of the standard projectile response with attitude oscillations persisting for 1.5 seconds while the optimal configuration is devoid of such flaw. In terms of performance, the maximum control AoA of the co-designed projectile is unsurprisingly much larger than the more stable one. Due to these canard stall issues, the standard configuration only manages 2.50 degrees, much smaller than the 7.30 degrees of the dual-loop ACO.

The second figure reveals that the increased control authority of the optimal airframe comes with the drawback of being more sensitive to actuator input disturbance. For this configuration, the peak attitude excursion reaches 5.5 degrees and require 1.91 second to decay beyond 0.1° , while the more stable design peaks at 2.7 degrees and decays in less than 0.96 second.

3.4.4 Control of projectiles with relaxed static stability

The open-loop dynamics study has shown that neutral static stability configurations may be particularly well suited for agile projectiles as they feature the highest amplification between the steady-state control deflection and the trim pitch attitude. As a result, it is expected that such designs feature a greater maximum AoA which would ensure the

best lift performance of the projectile and eventually open the way to additional maneuvers similar to those performed by modern fighter jets, such as aerodynamic braking obtained from extreme AoA excursions.

To put this assumption to the test, the fins position has been set to 0.1494 m which corresponds to neutral static stability according to the LPV synthesis model. The dual-loop controller structure has been retained to provide the adequate reference tracking and disturbance rejection performance.

Frequency-domain results




Figure 3.34 – Disturbance rejection transfer S and associated template

As the controller layout is left unchanged with respect to the co-design scenario, the synthesis has been attempted with similar frequency templates and stability margins requirements. However, the forecasted performance level differs from the previous case: the output sensitivity function of fig. 3.34 has a narrower bandwidth of 0.964 rad/s while the worst case sensitivity peaks at 1.50 for a frequency of 2.67 rad/s. Compared to the previous two airframe-controller designs, this system is predicted to be the slowest but its maximum sensitivity is in-between the values of the two other controllers.

The reference to actuator transfer of fig. 3.35 has a bandpass shape, with asymptotic decay at low frequencies that matches the physical interpretation of zero control deflection required to maintain any trim attitude. The predicted actuator usage is exceptionally




Figure 3.35 – Control effort transfer KS and associated template

low, with a worst-case gain of only 0.126 and more than 32 dB separation between the response and its prescribed template at high frequencies. However, it should be kept in mind that such responses are unlikely to be observed in practice as the fins position cannot be adjusted to give exactly zero static margin. Moreover, such results are only valid in the vicinity of zero AoA where the linear aerodynamic model of the projectile has been obtained.

The gap between the robustness transfer and its maximum bound pictured on fig. 3.36 remains greater than 8 dB, meaning that no pole-zero cancellation issues are to be expected. As for the stability margins of fig. 3.37, the gain constraints are just met or very close to their lower bounds while the magnitude of the phase margins are comparable to the dual-loop co-design example.




Figure 3.36 – Robustness transfer GS and associated template




Figure 3.37 – Stability margins at the actuator input δ_m (left) and at the output θ (right)

Time response analysis




Figure 3.38 – Comparison of reference tracking of co-designed and RSS configurations

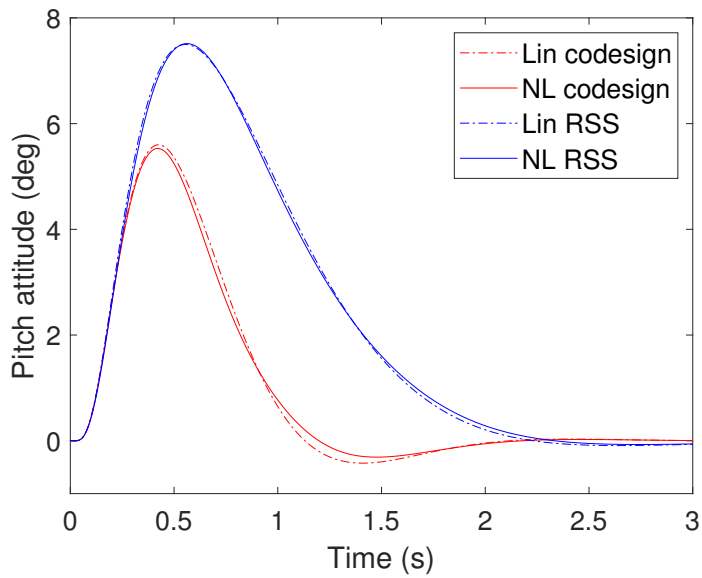


Figure 3.39 – Comparison of disturbance rejection of co-designed and RSS configurations




Figure 3.40 – Step tracking of the RSS projectile for increasing amplitudes

Figures 3.38 and 3.39 overlay the responses of the Reduced Static Stability (RSS) projectile and the dual-loop co-design to step tracking and disturbance rejection scenarios similar to those of the previous section. The baseline linear response of the neutral stability projectile is slower but better damped compared to the dual-loop co-design, which is consistent with the sensitivity plots. The 95% rise time is up from 0.78 s to 1.48 s but the overshoot has dropped from 7.3 % to 1.2 %. The disturbance rejection plot unveils that the RSS configuration is more sensitive to actuator input disturbance, with a peak attitude excursion of 7.5° and 2.17 seconds needed to decay beyond 0.1° , versus 5.5° and 1.91 seconds for the co-design. Given that such disturbance magnitude is quite extreme for actuator input noise, this behavior may be tolerable but should be investigated during upcoming wind-tunnel validation tests.

However, that does not mean that the RSS projectile is always slower than the co-design: in fact, Figure 3.40 shows that the impact of the non-linearities on control performance has been further reduced, with the simulated rise times being 1.54, 1.65, 1.82 and 2.0 seconds versus 0.83, 1.56, 2.08 and 2.41 seconds for the respective 4, 6, 8 and 10 degrees steps. Comparison between the control deflections of both configurations is provided by Figure 3.41: the very low actuator usage of the RSS projectile allows to reduce the canards local AoA, leading to better simulated maneuvering performance than the co-design with a maximum control AoA of around 10 degrees. However, it could also be a potential drawback since the actuators positions must be controlled with great precision, requiring high-resolution servo controllers and tight backlash specifications.

The simulated actuator deflections of Fig. 3.40 also show that the open-loop RSS

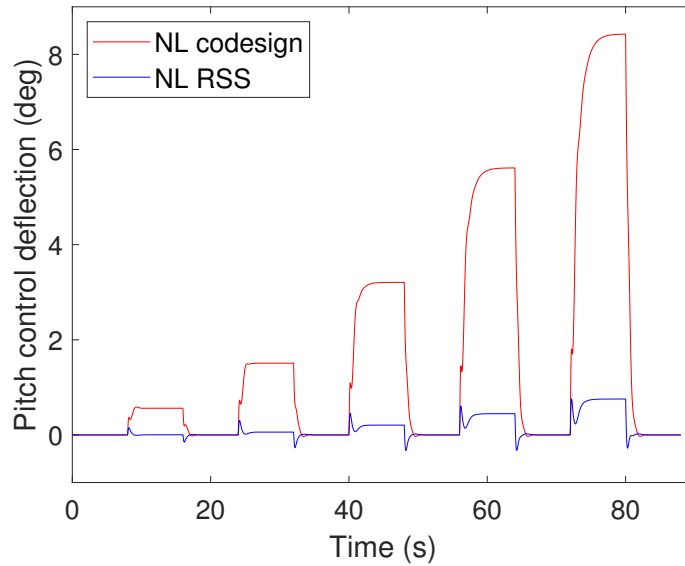


Figure 3.41 – Comparison of actuator usage of co-designed and RSS configurations

projectile does not remain neutrally stable on the whole range of AoA: the strictly positive steady-state control deflections predicted for any non-zero trim attitude reveal that the static gain of the plant is finite unlike what was predicted in section 3.3.1 and the non-linear projectile model is in fact statically stable. This can be due to two factors: first, due to the small discrepancies between the LPV and the non-linear models, the neutral fins position is actually slightly different for each structure. Secondly, the canard stall may alter the projectile static stability by shifting its center of pressure aft at larger angles of attack. This observation explains why the term "relaxed static stability" has been preferred to "neutral static stability" for describing this aerodynamic configuration, since the static stability of the projectile shifts between neutral, stable and unstable depending on the trim condition.

3.5 Experimental Testing and Validation of Design Candidates and Methodology

This section covers the experimental investigations that have been performed in the ISL subsonic wind tunnel in order to assess the actual performance of the various airframe-controller designs and the accuracy of the model predictions. The end goal is to determine how closely the non-linear simulator is able to predict the reference tracking performance

in order to validate the flight domain metric used to select the best configuration.



Figure 3.42 – Overview of the HIL projectile on its gimbal in wind tunnel [Strub, 2016]

The ACHILES projectile is pictured on Figure 3.42: a comprehensive description of the experimental setup can be found in section 1.1. In the frame of performance validation, both co-design candidates and the RSS projectile have been submitted to two different tests, one being focused on reference tracking and the other on disturbance rejection. In the tracking case, the guided projectile was tasked with following successive steps of increasing amplitude while returning to zero in-between each step, following the same command as what was simulated in the previous section. For disturbance rejection, this signal was injected as a disturbance at the actuator input, the reference being kept to zero.

3.5.1 Wind-tunnel testing of the state-space co-design

This aerodynamic configuration is the most stable of all three, with a fins position of 0.211 m giving an estimated static margin of 0.30 calibers. This design point is well within the parametric range on which the semi-local aerodynamic model has been validated during open-loop trials (c.f section 2.8). Figures 3.43 and 3.44 compare the attitude

and controls deflection measurements with the predictions of both linear and non-linear models.

Reference tracking




Figure 3.43 – Pitch attitude tracking of the state-space co-design




Figure 3.44 – Control action during reference tracking experiment

For this configuration, the experimental data confirm that the actuator saturation,

predicted by the non-linear model, indeed occurs. The effect of canard stall are visible for reference steps of 4° and more, with pitch control saturation being reached as early as 6° . Time histories of the linear model differs greatly from the recorded attitude, with no saturation present and a much faster response from 4° onwards. This example proves that the actual projectile dynamics are poorly predicted by such simplistic model structure on most of the AoA range. Conversely, the non-linear model is able to predict the closed-loop pitch dynamics fairly well up to around 6 degrees. In the saturated range, the trim angle of attack is overestimated and the lag due to controller wind-up is milder than what is shown by the experimental setup. With regard to the quality of fit obtained by the semi-local structure for a neighboring fins position of 0.215 m and controls deflection up to 16° (c.f fig. 2.30), the gap between the non-linear model and the measured response is somewhat surprising. However, the extrapolation of the polynomial canard model to higher control deflections may have induced some degree of error, while the experimental setup could also be part of the issue (c.f section 3.5.4).

Disturbance rejection




Figure 3.45 – Disturbance rejection of the state-space co-design

As for disturbance rejection, some similarities can be noted in regard to the models fit. According to fig. 3.45 the linear representation predicts excessively large oscillations, especially for greater disturbances. This is because the disturbance is injected at the actuator input and thus it is being affected by canard stall as well. Figure 3.47 proves that,




Figure 3.46 – Control action during disturbance rejection experiment

in this case, the non-linear structure accurately captures the projectile pitch dynamics unlike the linear model. Figure 3.46 shows that actuator saturation is not reached, probably because this configuration has high open-loop gain in the medium frequency range due to the plant resonance (c.f. section 3.3.1). The amplitude of the control action is well predicted by all model structures, the linear response being slightly faster than both non-linear simulation and measurements.

One may notice that this test is more demanding than the 5° step that was used in the previous part, and may be deemed a bit excessive considering that the aerodynamic models are designed to predict the projectile AoA up to 10° only. However, in the absence of a mechanical device that would allow to repeatably alter the projectile pitch angle, injecting a disturbance signal at the actuator input is the only option to get an insight on the controller ability to recover from an unexpected attitude excursion.




Figure 3.47 – Zoom on the 8° disturbance step rejection

3.5.2 Wind-tunnel testing of the RSS projectile

Contrary to the previous configuration, this airframe layout is the least stable of the three, with a fins position of 0.149 m well beyond the lower limit of 0.175 m for open-loop validation of the semi-local model. Thus, it could be expected that this case study would put the modeling and control framework to the test, eventually leading to the failure of the non-linear simulator to adequately predict reference tracking performance. As a result, some adjustments have been made to the model structure in order to improve model fit on the lower end of the parametric range. Figures 3.48 and 3.49 compare the respective attitude and pitch control deflection predicted by both the original and the adjusted model to measured data.

Non-linear model adjustment procedure

The model tuning procedure is based on closed-loop actuator time histories: the objective is to correct the linear part of the projectile static stability as well as the scaling factor accounting for canard stall effect on projectile stability. Recall Equation (2.58) which describes the pitch stability contribution to the semi-local model of the aerodynamic pitching moment:

$$Cm_{\alpha,cor}(x_F) = Cm_{\alpha}(x_F) - (x_G - x_{CPC})K_{scl} \left(1 - \frac{sl_{CL,C}(\alpha_{C,eq})}{sl_{0,CLC}} \right) \quad (3.3)$$




Figure 3.48 – Pitch attitude tracking of the RSS projectile




Figure 3.49 – Control action during reference tracking experiment

The surrogate model $Cm_\alpha(x_F)$ captures the linear behavior of the pitch stability coefficient. The polynomial interpolation of this surrogate models adds a degree of uncertainty to the parametric representation. Moreover the source data of the model is plagued by higher estimation error in this part of the x_F range as closed-loop parameter identification had to be performed on relaxed stability designs (c.f. section 2.2.2). Also, the scaling factor K_{scl} has been manually adjusted based on open-loop results obtained for more stable

configurations and may no longer be adapted.

The actuator deflections during reference tracking give some insight on the projectile behavior and the model deficiencies. Figure 3.49 reveals that the controller require increasingly negative pitch control to trim the projectile at 2 and 4 degrees AoA, which implies that Cm_α is positive on this part of the AoA range and the projectile is unstable. For greater angles of attack, the steady-state control action changes sign as positive defections are commanded for 8 and 10 degrees AoA, meaning that Cm_α is now negative and the projectile is stable. These variations illustrate that the RSS projectile is actually unstable at small AoA and becomes increasingly stable as the operating point moves away from the linear region.

The observed physics are consistent with the proposed model structure, but it seems that the scaling of these effects are off: first, the projectile model predicts positive stability in the linear region as $Cm_\alpha(x_F)$ is too small. Also, the increase in positive deflections with trim AoA seems to be milder than what is observed, indicating that the non-linear effects are under-estimated. Thus, an offset of 4.0 is added to $Cm_\alpha(x_F)$ and the value of K_{scl} is increased tenfold. The magnitude of the corrections are chosen so that the adjusted model is valid for the both the RSS projectile and the dual-loop co-design, covering the range of fins positions from 0.149 to 0.168 m, complementary to the interval $0.175 \text{ m} \leq x_F \leq 0.313 \text{ m}$ on which the original model is valid.

Reference tracking

The closed-loop response of the RSS configuration submitted to the step sequence is pictured on fig. 3.48. Both the linear and the legacy non-linear model fail to reproduce the large overshoots which are particularly visible at moderate AoA. The latter structure predicts a first-order like response unlike its re-tuned variant that is able to capture the oscillations. The adjusted model correctly represents the asymmetry in the measured time response, with the return to zero leading to greater overshoots than for the large reference inputs. However, some degree of discrepancy can still be noted on the amplitude of the peaks, which is slightly underpredicted, especially at large AoA.

Figure 3.50 compares the 10% response time of the different models as a function of the step amplitude. It reveals that the re-tuned model is the only one able to predict the variations of the closed-loop projectile response time with respect to the whole range of excitation amplitudes. The non-monotony of this metric is intriguing with a minimum time of 1.1 second being obtained for an 8° step command. This behavior is due to the




Figure 3.50 – Response time prediction and measurement

large overshoots which increase the projectile response time at low AoA. Note that, in the RSS case, the original non-linear model is very inaccurate and excessively optimistic with a settling time staying very close to the linear predictions. Since this structure has been used in the controller tuning framework, it may explain why the controller damping is lower than expected and the response time at low AoA is degraded.

To quantify the accuracy of each model structure, the relative error between the predicted 10% response time and the measured value has been computed for each of the five different steps. The errors are then averaged over the range of step amplitude to give the overall ability of the model to predict the response time of the closed-loop projectile in the 2° to 10° AoA range. Thus, the average relative error of each model with respect to the measured response times is respectively 41.2, 34.0 and 16.6 % for the linear, non-linear and re-tuned structures. Additionally, the maximum absolute error is given to illustrate how far the model predictions can deviate from the experimental data. The largest discrepancies are respectively 1.33 s, 1.25 s and 0.386 s for the linear, non-linear and re-tuned models.

Disturbance rejection

The disturbance rejection trial provides a valuable insight on the peculiar dynamics of the RSS configuration. Figures 3.51 and 3.52 present the complete time histories of the pitch attitude and the actuators deflection, while fig. 3.53 focuses on one of the step

responses.

The re-tuned model matches experimental measurements for step amplitudes up to 4° . After that, the amplitude of the attitude excursions quickly grows beyond the the expected validity domain of the aerodynamic model which is around 10 degrees: the maxima of 13.1° , 17.7° and 22.4° are reached for the respective disturbances of 6° , 8° and 10° . As a result, the inconsistencies in peak attitude and control action get more and more prominent, and the damping of the pitch dynamics starts to be significantly overestimated by all the models. Beyond 15° , it is highly probable that the projectile fins would stall, significantly decreasing the degree of open-loop stability and accelerating the unstable pole. Thus, the severity of this off-design condition is such that the control loop may not be able to provide the adequate stabilizing action, leading to a loss of control of the projectile which can be observed at the end of the measurement sequence. The zoom pictured on Figure 3.53 details how the actual response of the projectile differs from the model predictions. The discrepancy of the peak after the first zero crossing is very noticeable showing that, in this particular case, the model behavior in the moderate AoA domain may also be challenged.




Figure 3.51 – Disturbance rejection of the RSS projectile




Figure 3.52 – Control action during disturbance rejection experiment




Figure 3.53 – Zoom on the 8° disturbance step rejection

3.5.3 Wind-tunnel testing of the dual-loop co-design

The final set of experiments aims at validating the non-linear model used to assess the performance of the most promising co-designed configuration. Both original and re-tuned non-linear structures have been evaluated to determine which one was the most accurate and if the quality of their performance estimates were sufficient to confirm the interest of the ACO methodology.

Reference tracking




Figure 3.54 – Pitch attitude tracking of the dual-loop co-design

The pitch attitude and controls deflection time histories of the dual-loop configuration subjected to the usual step sequence are reproduced on figs. 3.54 and 3.55. The aspect of the attitude response is similar to the original non-linear model prediction, except for the mild overshoots observed on the 4° step and at each return to zero which are underestimated. The re-tuned model displays slightly excessive oscillations but captures the convergence to trim attitude quite accurately at large AoA, as confirmed by fig. 3.56. The linear structure is only able to correctly estimate the pitch attitude up to 4° AoA, after which discrepancies due to canard stall start to appear. Finally, the model fit percentages respectively come up at 88.5, 85.3 and 72.7 % for the re-tuned, original and linear model structures. The re-tuned model is the most accurate when it comes to actuator usage prediction, even though the quasi-absence of steady-state control action




Figure 3.55 – Control action during reference tracking experiment

for the 2° step shows that the static stability of the projectile is actually higher than the simulated one. The increase in actuator usage for moderate and large AoA is a bit too pessimistic, with larger peak deflections compared to actual measurements. The original non-linear structure shows an opposite trend, as the impact of canard stall on the controls deflection is underestimated, leading to smaller peak deflections.

Figure 3.56 quantifies the models ability to estimate the reference tracking performance of the co-designed configuration. It is apparent that none of the model structures is able to closely match the measured response time on the whole range of excitation amplitudes. The original non-linear model correctly estimates this metric up to the 6° step while the re-tuned one shows a local minimum typically associated with RSS configurations but accurately predicts the response time to the 8 and 10° steps. Overall, the adjusted NL model only has the second best average relative error with 33.8 % versus 21.9 % for the original NL model and 40.2 % for the linear structure. However, its predictions always remains fairly close from the measurements, with a maximum absolute error of 0.789 s versus 1.70 s and 2.73 s for the respective original NL and linear structures.




Figure 3.56 – Response time prediction and measurement

Disturbance rejection

The simulation results of the three model structures are in good agreement with the pitch response to the disturbance steps displayed on figs. 3.57 and 3.59. The re-tuned model lacks a bit of damping when returning to zero but, overall, the attitude peaks are well captured. It is worth mentioning that the amplitude of those excursions is significantly lower than for the RSS projectile: the maxima of 7.2 , 9.5 and 11.5° have been reached for the 6 , 8 and 10° steps instead of the 13.1 , 17.7 and 22.4° obtained with the neutrally stable configuration. The control action pictured in fig. 3.58 is also well predicted, confirming the models behavior observed on the pitch attitude time histories.

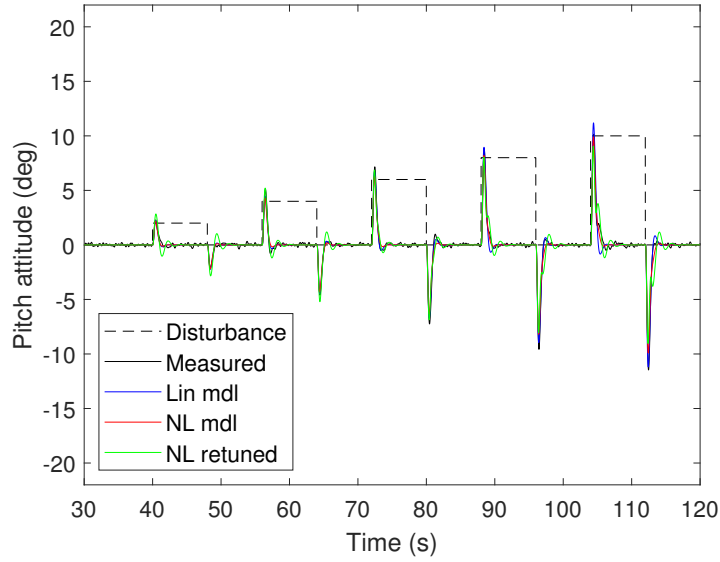


Figure 3.57 – Disturbance rejection of the dual-loop co-design




Figure 3.58 – Control action during disturbance rejection experiment




Figure 3.59 – Zoom on the 8° disturbance step rejection


Impact of non-linear model adjustment on maximum control AoA predictions

Figure 3.60 – Comparison of the responses of both variants of the non-linear model at their respective maximum AoA

The two non-linear formulations predict relatively distinct values of the projectile response time in the 6 to 8° amplitude range corresponding to the maximum angle of attack of the dual-loop co-design according to the original model (c.f. section 3.4.2). Thus, the max. control AoA estimates of the re-tuned model have been compared to those of the original one: the difference between the two values is significant, with the new model predicting only 4.62° versus 7.30° for the original one. These values are in agreement with the data of fig. 3.56 which shows that the response time predicted by the re-tuned structure rises more steeply than the estimates of the original model. Moreover, Figure 3.60 reveals that the aspect of the closed-loop time response differs between the two models, the new structure predicting oscillations that may slow down the system response and reduce the maximum control AoA of the co-designed configuration.

3.5.4 Conclusions on model fidelity and agile projectile design**Experimental assessment of agile projectile designs**

Based on the closed-loop test results of the three configurations, a number of conclusions regarding their fitness for purpose can be drawn. The state-space co-design is

the least sensitive to actuator input disturbances, however, its flight domain is limited by early actuator saturation due to canard stall (c.f figures 3.45 and 3.43). This configuration is the only one that required control action outside of the aerodynamic model domain of validity for pitch attitude targets as low as 6 degrees. Thus, it is definitely unsuitable for agile projectiles and should be discarded.

The relaxed static stability configuration was expected to be slower than the structured co-design for angles of attack smaller than 6° (see Figure 3.38) but, according to simulations, should have had a larger maximum control AoA. Unfortunately, the predictions of the non-linear model turned out to be inadequate as this projectile geometry was way outside its validation range. As a result, a methodology was proposed in order to adjust some of the parameters of the aerodynamic model. The re-tuned model showed better agreement with experimental data and was able to more accurately estimate the reference tracking performance of the configuration, as shown on Figure 3.50. Ultimately, this configuration is plagued with undesirable transient characteristics such as large overshoots and high sensitivity to actuator input disturbance (figures 3.48 and 3.51). Most of these effects are fairly well captured by the adjusted model structure but further study is required to understand their root causes and adapt the synthesis framework in order to mitigate them.

The dual-loop co-design gives balanced flight characteristics, with significantly enhanced performance compared to standard guided projectile design practices and a well-damped transient behavior (c.f. Figure 3.32). Until the high control authority of the RSS projectile can be harnessed, this optimal airframe-controller design remains the most advantageous prospect. Figure 3.61 confirms that this configuration gave the best results during the closed-loop wind tunnel validation campaign.

Model fidelity and accuracy of performance estimates

Experimental testing has revealed some limitations of the non-linear model structure that has been used to assess the performance of the various configurations developed in section 3.4. In spite of the modeling efforts described in the previous chapter, the retained model structure has proven unable to accurately predict the reference tracking performance of the guided projectile on the whole range of fins positions. The lack of open-loop validation data covering the parameter values corresponding to unstable and marginally stable configurations degraded the model accuracy on this part of the design space and prevented a comprehensive assessment of the model fidelity, leading to these

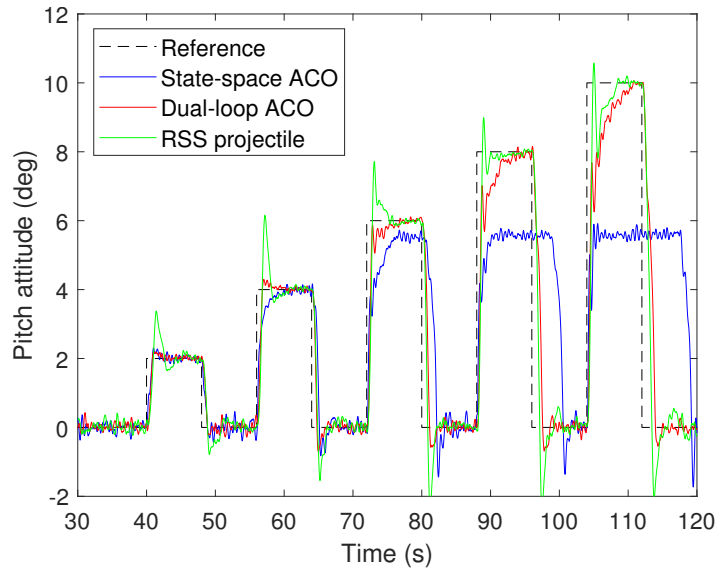


Figure 3.61 – Comparison of the measured step responses of the various agile projectile designs

shortcomings being discovered during wind tunnel validation. However, the information contained in the closed-loop actuator time histories was enough to overcome this issue by adjusting the model coefficients for this specific range of fins position.

Thus, the designer has to do with two variants of the same model structure, one for the forward range of fins positions ($x_F < 0.175$ m) and the other for the aft range. For configurations close to the switch point, such as the dual-loop codesign, the two models could be used to obtain a set of bounds on the maximum control AoA of the guided projectile. The values of 4.6° and 7.3° show that the uncertainty on the performance of this configuration is fairly large because of poor agreement between the model structures. However, the comparison with the standard configuration remains unequivocally favorable to the co-design, proving the benefits of this new methodology with respect to traditional design practices.

Shortcomings of the experimental validation process

Part of the discrepancies between the simulation outputs and the experimental results can be attributed to the limitations of the wind-tunnel setup: the mechanical design of the projectile mock-up is mostly based on bulky legacy components, leading to a stringent internal volume constraint. Thus, unlike contemporary setups [Nel16], the position of the mock-up center of gravity cannot be adjusted. Therefore, a random bias on the pitch axis

may be induced by the misalignment between the the CG and the center of the gibal frame, offsetting the trim map of the projectile. Moreover, the canards servo being fitted with relative position encoders, the zero of each control surface had to be set manually, introducing some error in the canards deflection readings. Due to the combination of both of these error sources, an offset as large as 0.92° was found on the pitch deflection time histories and removed during post-processing. This could contribute to the mismatch between the amplitudes of the experimental and model responses.

The accuracy of the attitude measurements is restricted by the limitations of the inertial measurement unit, which is affected by estimation errors and sensor dynamics due to onboard processing. According to manufacturer data, the typical attitude accuracy is given as $\pm 0.5^\circ$ when static and $\pm 2.0^\circ$ during motion. Experimentally, a noise amplitude averaging 0.5° was observed on the pitch attitude angle: for this reason, the 10 % settling time of the projectile response has been used for model validation instead of the 5 % criterion which was excessively affected by noise. Some of the disturbances could be attributed to the flexibility of the gibal assembly, especially the rods connecting the projectile mock-up to the outer frame that have to be as thin as possible in order not to interfere with the freestream: the circular section of those rods could lead to vortex shedding that would apply periodical forces on the structure. Another source of perturbations could be the tunnel itself: subsequent investigation of the freestream flow field revealed a pulsation in the axial velocity that was attributed to some intrusive measurement equipment upstream of the test section.

3.6 Conclusion

This chapter focused on the build-up of an integrated airframe-controller design and simulation framework in order to assess the potential of plant-controller optimization as a means to improve guided projectile agility. In order to do so, previous contributions on PCO theory and applications have been reviewed: this research contributed to the choice of an optimization strategy and synthesis method, as well as the formulation of the constraints and cost function.

An original design framework has been introduced, with the traditional multi-objective \mathcal{H}_∞ controller synthesis being augmented with airframe static stability tailoring thanks to a LPV pitch dynamics model. Closed-loop performance and stability assessment relies on the non-linear aerodynamic model developed in the previous chapter in complement to the

frequency templates and stability margins. That way, the effect of the non-linearities on the control authority due to aerodynamic stall of the control surfaces can be accounted for at the design stage to reduce the impact of this detrimental effect on the actual airframe-controller configuration.

The co-design framework has then been used to maximize the angle of attack at which the guided projectile could be controlled despite the limited effectiveness of its actuators and control canards. In order to show the flexibility of the methodology, two configurations have been developed, each one with a different controller structure. The most promising design has been benchmarked against a guided projectile derived from current best practices, showing much better performance in the face of aerodynamic limitations. In addition, the control of an airframe featuring relaxed static stability has been investigated to evaluate the potential of this outstanding configuration.

To validate the accuracy of the simulation results and the suitability of the developed airframe-controller configurations, a wind-tunnel test campaign consisting of closed-loop control experiments was carried out. Measurements revealed the deficiency of the non-linear aerodynamic model for configurations close to neutral static stability. A new methodology was proposed in order to adjust the model parameters on the basis of closed-loop time histories. The re-tuned model proved to be able to capture most of the system behavior and improve the quality of the tracking performance estimates on the forward range of fins positions. Overall, the combination of the two models allowed to predict the response time of the projectile to pitch attitude steps ranging from 2 to 10 degrees with approximately 30% error. Even though some discrepancies between the model predictions could be observed for configurations close to the model switch point, the simulated maximum control AoA of the dual-loop co-designed projectile remains superior to the one of the reference configuration. Hence the ACO methodology seems promising in the frame of agile projectile design. On the other hand, the RSS projectile showed promising simulation results but turned out to be plagued with undesirable transient characteristics and proved to be more affected by disturbances. Additionally, its extreme sensitivity to control surfaces deflections may be a concern due to the mechanical imperfections of the actuators drivetrains. Thus, this configuration is not yet mature and appears more challenging to control than the co-designed one.

CONCLUSION AND PERSPECTIVES

The aim of this work was to propose an alternative to the sequential (aerodynamics then control) guided projectile design methodology. In fact, the downside of the legacy approach is that critical control system deficiencies due to inadequate modeling or synthesis may only be discovered at the flight test stage, very late into the development process. In that case, the aerodynamic configuration of the projectile would have to be altered, adding considerable cost and time overruns.

The new methodology attempts to circumvent this issue by introducing airframe-controller optimization (ACO) schemes and mixed-fidelity modeling. A key idea of the proposed approach is to complement the linear model used for controller synthesis with a non-linear structure able to capture the dominant effects of canard stall of the projectile dynamics. This model was integrated into the ACO scheme where \mathcal{H}_∞ synthesis was used to concurrently tune the projectile geometry and the controller gains in order to mitigate the impact of canard stall on closed-loop performance and maximize projectile agility.

In the first chapter, the HIL approach proposed by Strub was leveraged to set up an attitude control law and verify its effectiveness. The experimental setup, which consists in a projectile mock-up mounted on a 3 DoF gimbal in wind tunnel, was adapted for this task. The model structure was derived from the non-linear equations of rotational motion, accounting for the additional degree of freedom in roll. A new system identification campaign was conducted to estimate the mock-up aerodynamic derivatives before including the model into the synthesis framework. This scheme, featuring a multi-objective \mathcal{H}_∞ setup, was adapted along with the controller structure. Notably, the weighting filters associated with the disturbance rejection, model matching and control effort attenuation objectives were adjusted according to the open-loop characteristics of the projectile. Finally, the performance of the attitude controller was assessed through reference tracking and disturbance rejection experiments, confirming that the measured autopilot performance matched well with the model predictions and that the extension to three degrees-of-freedom was successful.

The remaining two chapters of the thesis cover the issue of operation at larger angles of attack and the mitigation of control surfaces stall. This challenge was tackled with a

multi-disciplinary approach combining the methodology of control theory with detailed modeling of the projectile aerodynamics: the airframe static stability was tuned along with the controller gains in order to maximize performance while accounting for the actuators and control surfaces limitations. The scope of the study was reduced to pitch attitude control in order to focus on the most critical issue with respect to a forecasted operational scenario.

The second chapter describes the development of the aerodynamic models required for airframe-controller optimization and validation of the configuration candidates in non-linear domain. The experimental setup was employed to record the equilibrium manifold of the projectile and identify linear dynamic models for a range of open-loop stable and unstable aerodynamic configurations. To achieve that, a new closed-loop system identification protocol had to be developed and implemented on the experimental setup. A linear parameter-varying formulation was derived from this dataset and tailored for synthesis and optimization. The dependency to the geometric parameter was introduced by using polynomial surrogate models of the projectile aerodynamic derivatives. After reviewing the application of potential flow theory to the prediction of the projectile pitching moment, three non-linear model candidates were investigated in order to capture the effect of canard stall on the pitch dynamics. The model structures range from a formulation based on the analytical formulas of fluid mechanics to a purely data-based model structure relying on open-loop trim maps. Additional work was required to reconstruct the canard lift polar from a series of steady-state attitude measurements, feeding all but the data-based model. The models outputs were benchmarked against open-loop measurements in order to select the best candidate for assessing the closed-loop response of a family of projectile configurations. Model fidelity was assessed with trim maps and step responses for stable configurations, while closed-loop experiments were employed to forecast the model ability to simulate the dynamics of unstable projectiles. The semi-local model structure was preferred since it features the most consistent fidelity across the range of fins positions.

The final chapter presents the integrated airframe-controller design and simulation framework designed to improve the guided projectile agility by increasing the maximum angle of attack at which it can be controlled. The ACO methodology has been applied to two different concepts with distinct controller structures. Additionally, this framework has been used to tune a relaxed static stability (RSS) projectile concept. This original design framework augments the traditional multi-objective \mathcal{H}_∞ synthesis scheme with airframe static stability tailoring thanks to a LPV pitch dynamics model. The closed-

loop performance and stability of each iteration of the airframe-controller design was systematically assessed in non-linear domain with the semi-local model. A performance metric named the maximum control AoA was proposed to quantify the reference tracking performance and control authority in spite of canard stall conditions. Frequency- and time-domain simulation data was provided for each of the three configurations. The most promising co-design was benchmarked against a guided projectile derived from current best practices, showing much better performance when subjected to realistic aerodynamic limitations with a significant increase in maximum control angle of attack.

To validate the accuracy of the simulation results and the suitability of the developed airframe-controller configurations, a wind-tunnel test campaign consisting of closed-loop control experiments was carried out. Measurements revealed the limitations of the non-linear aerodynamic model for configurations close to neutral static stability. The closed-loop experimental data was leveraged to create a new variant of this model by adjusting some of its parameters. The re-tuned model proved to adequately capture the system behavior and was able to estimate the response time of the less stable configurations over a wide range of excitation amplitudes with an average 35 % error. Comparison between the max. control AoA predictions of both old and new model structures showed that large uncertainty was present on the performance estimates of the dual-loop codesign, while still being vastly superior to the traditionally designed projectile. The RSS projectile was found to be plagued with undesirable transient characteristics despite promising simulation results. Thus, this last configuration is not yet mature and appears more challenging to control than the co-designed one. All of these elements concur to show that the ACO methodology is relevant in the frame of agile projectile design.

Perspectives

This study unveiled the ability of the co-design to maximize the maneuvering performance obtained with a given control and actuation system (CAS). It demonstrated that it could salvage a CAS design that would otherwise be considered unsatisfactory: the canard planform used for this study was investigated as early as 2014 [Str+14] and quickly superseded by a more effective shape [Str+15] because of the canard stall issues.

Being able to extract significantly more performance from these flawed control surfaces proves the value of airframe-controller optimization in the frame of multi-disciplinary design. By adjusting the value of a single parameter of the projectile geometry, the control

designer is able to significantly improve the guided projectile maneuver performance without presumably altering critical performance indexes outside the scope of autopilot design, such as maximum range or payload capacity. This airframe tuning could even be made less intrusive by substituting the longitudinal fins position by the longitudinal CG position: although ill-suited to the ACHILES mock-up due to internal space constraints, such adjustment could be realized with an internal counterweight [Fre+18] with the guarantee of not impacting the aerodynamic performance of the projectile. Ultimately, the choice of the tunable parameter could be left to the designer depending on the application, as long as the static margin of the configuration is able to be adjusted. With regard to full-scale projectiles, it is even conceivable to alter the mass distribution of existing designs in order to improve their maneuvering performance and introduce new capabilities that rely on their improved agility, such as top attack flight profiles or moving target tracking. Alternatively, this methodology could allow to relax the actuator requirements for a given mission, allowing lower-cost or more rugged components to be used.

From an experimental standpoint, all of the building blocks for an attitude-controlled, high-performance airframe-controller design have been developed and tested. This objective could be achieved combining the 3 DoF modeling and control framework with the non-linear pitch model and ACO scheme, using the decoupling assumption and the tetrahedral symmetry of the projectile to simplify the problem. However, some limitations due to the design of the experimental setup, such as the excessive inertia and friction of the yaw gimbal or the lack of internal volume of the mock-up, prevent this design from being implemented. It must be noted that the suggestions for improved realism presented in Strub's PhD thesis [Str16] such as the redesign and integration into supersonic tunnels or the extension to pseudo 6-DoF through force balance measurements, are still relevant today. However, one of the strengths of this setup compared to other experimental means available at the ISL is its frugality and ease of operation. Embedding fragile sensors and increasing the freestream Mach number may increase the cost and time required to collect experimental data, undermining the ACHILES purpose in the frame of conceptual and preliminary design of projectile control schemes.

From a modeling standpoint, the limitations of the current approach are apparent on multiple aspects. The non-linear model structure used for simulation and early controller validation was unable to accurately reproduce the closed-loop pitch dynamics of the projectile over the entirety of the range of fins positions and excitation amplitudes: a second set of model parameters was necessary to cover less stable airframe configurations. Two

alternatives can be explored to remedy this shortcoming: either include more closed-loop test data to improve the model fidelity in the non-linear domain for fins positions where open-loop trim maps cannot be obtained, or increase the complexity of the analytical part of the model by implementing vortex tracking schemes as proposed by Nielsen [Nie88]. Additionally, the current model structure assumes that the aerodynamic derivatives identified around the zero pitch angle equilibrium condition are valid across the AoA envelope explored by the projectile: this assumption may induce significant error in the canard lift reconstruction process because of strong non-linearities in the pitch stability characteristics. Here too, two options could be considered: either include some higher order terms in the decomposition of the pitching moment, or map the aeroderivatives as a function of trim AoA and use local linear aerodynamic models. In either case, this requires an overhaul of the system identification procedure in order to estimate the additional parameters.

As for control theory, closed-loop simulation results may be improved by performing another synthesis of both RSS and dual-loop codesign configurations while considering the adjusted non-linear model. This could help to improve the unsatisfactory transient characteristics of the neutral-stability projectile and reduce the gap between the max. control AoA predictions for the dual-loop codesign. To go one step further, the treatment of the non-linear variations in the control authority due to canard stall could be altered. The present solution is to assign large parametric uncertainty to the pitch control aeroderivative, retaining a linear model for synthesis purposes. Yet, choosing the bounds of this uncertain coefficient is a difficult compromise between robustness to the worst-case scenario and mitigation of the degree of conservativeness in the synthesis results. Using multiple local synthesis models with different uncertainty bounds could alleviate this issue, but this type of LPV control requires extensive modeling work and large amounts of experimental data. Implementing control methods which are inherently designed for non-linear systems such as nonlinear dynamic inversion may also offer more options for preserving control effectiveness outside of the linear domain while maximizing performance in the linear range. However, robustness to parameter uncertainty must be kept in mind given the experimental nature of the model building process. Closed-loop testing also showed that some of the controller structures may be affected by windup effects near the limits of their flight domain. Various techniques such as integrator clamping, back calculation or model recovery could be implemented to mitigate this issue as part of future work.

BIBLIOGRAPHY

- [ABH73] DC Anderson, RL Berger, and JR Hess Jr, « Maneuver Load Control and Relaxed Static Stability Applied to a Contemporary Fighter Aircraft », *in: Journal of Aircraft* 10.2 (1973), pp. 112–120.
- [Aga99] Ramesh Agarwal, « Computational fluid dynamics of whole-body aircraft », *in: Annual review of fluid mechanics* 31.1 (1999), pp. 125–169.
- [AGB95] Pierre Apkarian, Pascal Gahinet, and Greg Becker, « Self-scheduled \mathcal{H}_∞ control of linear parameter-varying systems: a design example », *in: Automatica* 31.9 (1995), pp. 1251–1261.
- [Ala+13] Daniel Alazard, Thomas Loquen, Henry De Plinval, and Christelle Cumer, « Avionics/Control co-design for large flexible space structures », *in: AIAA Guidance, Navigation, and Control (GNC) Conference*, 2013, p. 4638.
- [Ale46] Leigh Alexander, « The Origin of Greek and Roman Artillery », *in: The Classical Journal* 41.5 (1946), pp. 208–212.
- [AN06a] Pierre Apkarian and Dominikus Noll, « Nonsmooth \mathcal{H}_∞ synthesis », *in: IEEE Transactions on Automatic Control* 51.1 (2006), pp. 71–86.
- [AN06b] Pierre Apkarian and Dominikus Noll, « Nonsmooth optimization for multi-disk \mathcal{H}_∞ synthesis », *in: European Journal of Control* 12.3 (2006), pp. 229–244.
- [And17] John David Anderson Jr, *Fundamentals of aerodynamics*, McGraw-Hill Education, 2017.
- [ANF] Pierre Apkarian, Dominikus Noll, and Alexandre Falcoz, *How we saved the Rosetta mission*, URL: <http://pierre.apkarian.free.fr/papers/rosettaProbe.pdf>.
- [Apk09] Pierre Apkarian, *Elements de la théorie de la commande robuste*, 2009.
- [Apk12] Pierre Apkarian, « Tuning controllers against multiple design requirements », *in: 2012 16th International Conference on System Theory, Control and Computing (ICSTCC)*, IEEE, 2012, pp. 1–6.

-
- [ARN11] Pierre Apkarian, Laleh Ravanbod-Hosseini, and Dominikus Noll, « Time domain constrained \mathcal{H}_∞ -synthesis », *in: International Journal of Robust and Nonlinear Control* 21.2 (2011), pp. 197–217.
- [AS53] Mac C Adams and WR Sears, « Slender-body theory-review and extension », *in: Journal of the Aeronautical Sciences* 20.2 (1953), pp. 85–98.
- [ASS15] GK Ananda, PP Sukumar, and Michael S Selig, « Measured aerodynamic characteristics of wings at low Reynolds numbers », *in: Aerospace Science and Technology* 42 (2015), pp. 392–406.
- [AW95] John David Anderson and John Wendt, *Computational fluid dynamics*, vol. 206, Springer, 1995.
- [Bar69] Barth, *Data for determining normal force moment and tangential force characteristics of slender nose-cylinder configurations in the transonic velocity range*, Messerschmitt-Buelkow-Blohm Corporation, 1969.
- [Bas92] Marshall J Bastable, « From breechloaders to monster guns: Sir William Armstrong and the invention of modern artillery, 1854-1880 », *in: Technology and Culture* 33.2 (1992), pp. 213–247.
- [BBS12] Caroline Bérard, Jean-Marc Biannic, and David Saussié, *La commande multivariable: Application au pilotage d'un avion*, Dunod, 2012.
- [Bec95] John Becklake, « The V2 Rocket—a Convergence of Technologies? », *in: Transactions of the Newcomen Society* 67.1 (1995), pp. 109–123.
- [Boh06] Torsten P Bohlin, *Practical grey-box process identification: theory and applications*, Springer Science & Business Media, 2006.
- [Bry53] Arthur E Bryson Jr, « Stability derivatives for a slender missile with application to a wing-body-vertical-tail configuration », *in: Journal of the Aeronautical Sciences* 20.5 (1953), pp. 297–308.
- [CBA18] Emmanuel Chambon, Laurent Burlion, and Pierre Apkarian, « Time-response shaping using output to input saturation transformation », *in: International Journal of Control* 91.3 (2018), pp. 534–553.
- [ÇLi11] Tayfun ÇLimen, « A generic approach to missile autopilot design using state-dependent nonlinear control », *in: IFAC Proceedings Volumes* 44.1 (2011), pp. 9587–9600.

-
- [Coo12] Michael V Cook, *Flight dynamics principles: a linear systems approach to aircraft stability and control*, Butterworth-Heinemann, 2012.
- [Cos95] Mark Costello, « Range extension and accuracy improvement of an advanced projectile using canard control », *in: 20th Atmospheric Flight Mechanics Conference, Baltimore, USA*, 1995, pp. 324–331.
- [Dar73] John A Darling, *Handbook of Blunt-body Aerodynamics: Volume 1-Static Stability*, vol. 73, 225, Naval Ordnance Laboratory, White Oak, 1973.
- [Dec+18] Cédric Decroq, Bastien Martinez, Marie Albisser, Simona Dobre, Patrick Gnemmi, Yannick Bailly, and Jean-Claude Roy, « Aerodynamic prediction of a projectile fitted with fins », *in: International Journal of Numerical Methods for Heat & Fluid Flow* 28.5 (2018), pp. 1218–1236.
- [Den+17] Yann Denieul, Joël Bordeneuve, Daniel Alazard, Clément Toussaint, and Gilles Taquin, « Multicontrol Surface Optimization for Blended Wing–Body Under Handling Quality Constraints », *in: Journal of Aircraft* 55.2 (2017), pp. 638–651.
- [Den16] Yann Denieul, « Preliminary Design of Control Surfaces and Laws for Unconventional Aircraft Configurations », PhD thesis, 2016.
- [Den98] P Denis, « Le code de prévision aérodynamique de l'ONERA : “MISSILE” », *in: RTO AVT Symposium on “Missile Aerodynamics”, Sorrento, Italy*, Cite-seer, 1998, pp. 11–14.
- [DM17] Adam C DeVoria and Kamran Mohseni, « A vortex model for forces and moments on low-aspect-ratio wings in side-slip with experimental validation », *in: Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 473.2198 (2017), p. 20160760.
- [Doy+89] John C Doyle, Keith Glover, Pramod P Khargonekar, and Bruce A Francis, « State-space solutions to standard \mathcal{H}_2 and \mathcal{H}_∞ control problems », *in: IEEE Transactions on Automatic Control* 34.8 (1989), pp. 831–847.
- [DVT15] Nihar Deodhar, Chris Vermillion, and Peter Tkacik, « A case study in experimentally-infused plant and controller optimization for airborne wind energy systems », *in: 2015 American Control Conference (ACC)*, IEEE, 2015, pp. 2371–2376.
- [ET94] Bradley Efron and Robert J Tibshirani, *An introduction to the bootstrap*, CRC press, 1994.

-
- [Fat+01] Hosam K Fathy, Julie A Reyer, Panos Y Papalambros, and AG Ulsov, « On the coupling between the plant and controller optimization problems », *in: Proceedings of the 2001 American Control Conference. (Cat. No. 01CH37148)*, vol. 3, IEEE, 2001, pp. 1864–1869.
- [Fat+02] Hosam K Fathy, Scott A Bortoff, G Scott Copeland, Panos Y Papalambros, and A Galip Ulsoy, « Nested optimization of an elevator and its gain-scheduled LQG controller », *in: ASME International Mechanical Engineering Congress and Exposition*, vol. 36290, 2002, pp. 119–126.
- [Fat+03] Hosam K Fathy, Panos Y Papalambros, A Galip Ulsoy, and Davor Hrovat, « Nested plant/controller optimization with application to combined passive/active automotive suspensions », *in: Proceedings of the 2003 American Control Conference, 2003*. Vol. 4, IEEE, 2003, pp. 3375–3380.
- [FB96] V Fleck and C Berner, « Increase of range for an artillery projectile by using the lift force », *in: Proceedings 16th International Ballistics Symposium, San Francisco, CA, USA, 1996*, pp. 23–27.
- [FCF12] Frank Fresconi, Ilmars Celmins, and Lois Fairfax, « Optimal parameters for maneuverability of affordable precision munitions », *in: 50th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition, 2012*, p. 254.
- [FLM21] Christian Franzmann, Friedrich Leopold, and Christian Mundt, « Experimental Determination of Pitch Damping Coefficients Using a Wire Suspension », *in: AIAA Scitech 2021 Forum, 2021*, p. 0836.
- [Fre+18] Frank Fresconi, Ilmars Celmins, James Maley, and Bryant Nelson, « Experimental Flight Characterization of a Canard-Controlled, Subsonic Missile », *in: 2018 Atmospheric Flight Mechanics Conference, Atlanta, USA, 2018*, p. 3154.
- [Fre11] Frank Fresconi, « Guidance and control of a projectile with reduced sensor and actuator requirements », *in: Journal of Guidance, Control, and Dynamics* 34.6 (2011), pp. 1757–1766.
- [GA11] P Gahinet and P Apkarian, « Structured \mathcal{H}_∞ Synthesis In MATLAB », *in: 18th IFAC World Congress, Milan, Italy, 2011*, DOI: 10.3182/20110828-6-it-1002.00708.

-
- [GL06] Alvin Gatto and Mark Lowenberg, « Evaluation of a three degree of freedom test rig for stability derivative estimation », *in: Journal of aircraft* 43.6 (2006), pp. 1747–1761.
- [HC03] Edward Lewis Houghton and Peter William Carpenter, *Aerodynamics for engineering students*, Elsevier, 2003.
- [Hem+76] Michael J Hemsch, Charles A Smith, Jack N Nielsen, and Stanley C Perkins Jr, *Calculation of Component Forces and Moments of Arbitrarily Banked Cruciform Missiles with Control Deflections*. Tech. rep., Nielsen Engineering and Research Inc Mountain View CA, 1976.
- [HLD85] Arthur L Hale, Ronald J Lisowski, and Wayne E Dahl, « Optimal simultaneous structural and control design of maneuvering flexible spacecraft », *in: Journal of Guidance, Control, and Dynamics* 8.1 (1985), pp. 86–93.
- [HN83] MJ Hemsch and JN Nielsen, « Equivalent angle-of-attack method for estimating nonlinear aerodynamics of missile fins », *in: Journal of Spacecraft and Rockets* 20.4 (1983), pp. 356–362.
- [Kar80] Krishnamurty Karamcheti, *Principles of Ideal-Fluid Aerodynamics (Reprint Edition)*, 1980.
- [Kes66] EL Kessler, « Design for Control of Projectile Flight Characteristics », *in: US Army Engineering Design Handbook* (1966).
- [KH00] Itsuro Kajiwara and Raphael T Haftka, « Integrated design of aerodynamics and control system for micro air vehicles », *in: JSME International Journal Series C Mechanical Systems, Machine Elements and Manufacturing* 43.3 (2000), pp. 684–690.
- [Kwa93] Huibert Kwakernaak, « Robust control and \mathcal{H}_∞ -optimization—tutorial paper », *in: automatica* 29.2 (1993), pp. 255–273.
- [LB89] James Luneau and Alain Bonnet, *Aérodynamique: théories de la dynamique des fluides*, Cépaduès-éditions, 1989.
- [LF94] ML Luyben and CA Floudas, « Analyzing the interaction of design and control—1. A multiobjective framework and application to binary distillation synthesis », *in: Computers & chemical engineering* 18.10 (1994), pp. 933–969.

-
- [Lju02] Lennart Ljung, « Prediction error estimation methods », *in: Circuits, Systems and Signal Processing* 21.1 (2002), pp. 11–21.
- [Lju86] Lennart Ljung, *System Identification: Theory for the User*, Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1986, ISBN: 0-138-81640-9.
- [Lju97] Lennart Ljung, « Identification in closed loop: Some aspects on direct and indirect approaches », *in: IFAC Proceedings Volumes* 30.11 (1997), pp. 155–160.
- [LLW05a] Fang Liao, Kai Yew Lum, and Jian Liang Wang, « An LMI-based optimization approach for integrated plant/output-feedback controller design », *in: Proceedings of the 2005, American Control Conference, 2005*. IEEE, 2005, pp. 4880–4885.
- [LLW05b] Fang Liao, Kai Yew Lum, and Jian Liang Wang, « Mixed $\mathcal{H}_2/\mathcal{H}_\infty$ Sub-Optimization Approach for Integrated Aircraft/Controller Design », *in: IFAC Proceedings Volumes* 38.1 (2005), pp. 103–108.
- [LM02] F Lawrence and B Mills, « Status update of the AEDC wind tunnel virtual flight testing development program », *in: 40th AIAA Aerospace Sciences Meeting & Exhibit*, 2002, p. 168.
- [Lv+18] Chen Lv, Xiaosong Hu, Alberto Sangiovanni-Vincentelli, Yutong Li, Clara Marina Martinez, and Dongpu Cao, « Driving-style-based codesign optimization of an automated electric vehicle: A cyber-physical system approach », *in: IEEE Transactions on Industrial Electronics* 66.4 (2018), pp. 2965–2975.
- [MA12] Mike Milner and PEO Ammunition, « Precision strike association Excalibur overview », *in: Combat Ammunition Systems* (2012).
- [Mac19] Dawid Machala, « Free-flight projectile behaviour: LPV modelling and global sensitivity analysis », PhD thesis, Université de Lorraine, 2019.
- [McC99] Robert McCoy, *Modern exterior ballistics: The launch and flight dynamics of symmetric projectiles*, Schiffer Pub., 1999.
- [ME81] Alfred E Magnus and Michael A Epton, *PAN AIR: A computer program for predicting subsonic or supersonic linear potential flows about arbitrary configurations using a higher order panel method. Volume 1: Theory document (version 1.1)*, tech. rep., NASA, 1981.

-
- [Mor92] Stephen Morris, « Integrated aerodynamics and control system design for tailless aircraft », *in: Astrodynamics Conference*, 1992, p. 4604.
- [MR05] Curtis Mracek and D Ridgely, « Missile Longitudinal Autopilots: Comparison of Multiple Three Loop Topologies », *en, in: AIAA Guidance, Navigation and Control Conference and Exhibit, San Fransisco, USA*, American Institute of Aeronautics and Astronautics, Aug. 2005, ISBN: 978-1-62410-056-7, DOI: 10.2514/6.2005-6380, URL: <http://arc.aiaa.org/doi/abs/10.2514/6.2005-6380> (visited on 03/18/2016).
- [MS77] Phillip H Morrison and David S. Amberntson, « Guidance and control of a cannon-launched guided projectile », *in: Journal of Spacecraft and Rockets* 14.6 (1977), pp. 328–334.
- [Mun24] Max M Munk, *The aerodynamic forces on airship hulls*, tech. rep., 1924.
- [NDV17] Parvin Nikpoorparizi, Nihar Deodhar, and Christopher Vermillion, « Modeling, control design, and combined plant/controller optimization for an energy-harvesting tethered wing », *in: IEEE Transactions on Control Systems Technology* 26.4 (2017), pp. 1157–1169.
- [Nel16] Bryant P Nelson, *Development and Evaluation of a Dynamic, 3-Degree-of-Freedom (DOF) Wind Tunnel Model*, tech. rep., Bennett Aerospace Cary, 2016.
- [NHS77] Jack N Nielsen, Michael J Hensch, and Charles A Smith, *A Preliminary Method for Calculating the Aerodynamic Characteristics of Cruciform Missiles to High Angles of Attack Including Effects of Roll Angle and Control Deflections*, tech. rep., Nielsen Engineering and Research Inc Mountain View CA, 1977.
- [Nie88] Jack Nielsen, *Missile aerodynamics*, American Institute of Aeronautics and Astronautics, Inc., 1988.
- [NK96] RJ Niewoehner and Isaac Kaminer, « Integrated aircraft-controller design using linear matrix inequalities », *in: Journal of guidance, control, and dynamics* 19.2 (1996), pp. 445–452.
- [OC08] Douglas Ollerenshaw and Mark Costello, « Model predictive control of a direct fire projectile equipped with canards », *in: Journal of Dynamic Systems, Measurement, and Control* 130.6 (2008), p. 061010.

-
- [OH87] Junjiro Onoda and Raphael T Haftka, « An approach to structure/control simultaneous optimization for large flexible spacecraft », *in: AIAA journal* 25.8 (1987), pp. 1133–1138.
- [Olm16] J. M. Olmos, *Guidance and control design for a class of long-range projectiles*, French-German Research Institute of Saint-Louis (ISL), 2016.
- [Per+16] J Alvaro Perez, Christelle Pittet, Daniel Alazard, and Thomas Loquen, « Integrated control/structure design of a large space structure using structured \mathcal{H}_∞ control », *in: IFAC-PapersOnLine* 49.17 (2016), pp. 302–307.
- [Pra25] Ludwig Prandtl, *Applications of modern hydrodynamics to aeronautics*, US Government Printing Office, 1925.
- [Pre16] Jamie Prenatt, *Katyusha: Russian multiple rocket launchers 1941–present*, Bloomsbury Publishing, 2016.
- [RHB14] Clement Roos, Georges Hardier, and Jean-Marc Biannic, « Polynomial and rational approximation with the APRICOT Library of the SMAC toolbox », *in: 2014 IEEE Conference on Control Applications (CCA)*, IEEE, 2014, pp. 1473–1478.
- [RLD10] Mike Ruth, Kenneth Lebsock, and Cornelius Dennehy, « What’s new is what’s old: Use of Bode’s integral theorem (circa 1945) to provide insight for 21st century spacecraft attitude control system design tuning », *in: AIAA Guidance, Navigation, and Control Conference*, 2010, p. 8428.
- [RM95] Clifford Ratliff and Edward Marquart, « An assessment of a potential test technique-Virtual Flight Testing (VFT) », *in: Flight Simulation Technologies Conference*, 1995, p. 3415.
- [RP00] Julie A Reyer and Panos Y Papalambros, « An investigation into modeling and solution strategies for optimal design and control », *in: ASME design engineering technical conferences*, Citeseer, 2000, pp. 10–13.
- [RP99] Julie A Reyer and Panos Y Papalambros, « Optimal design and control of an electric DC motor », *in: Proceedings of the 1999 ASME Design Engineering Technical Conferences*, Citeseer, 1999, pp. 85–96.

-
- [RRL20] Valentin Riss, Emmanuel Roussel, and Edouard Laroche, « Attitude Control of a Fin-Stabilized Projectile on a Three-Axis Gimbal in Wind Tunnel », *in: 2020 28th Mediterranean Conference on Control and Automation (MED)*, IEEE, 2020, pp. 37–43.
- [RRL22] Valentin Riss, Emmanuel Roussel, and Edouard Laroche, « Concurrent Airframe-Controller Optimization of a Guided Projectile fitted with Lifting Surfaces », *in: AIAA SCITECH 2022 Forum*, 2022, p. 0620.
- [RWH06] Thambirajah Ravichandran, David Wang, and Glenn Heppler, « Simultaneous plant-controller design optimization of a two-link planar manipulator », *in: Mechatronics 16.3-4* (2006), pp. 233–242.
- [SB16] Guillaume Strub and Michel Basset, « Skid-to-turn autopilot design and validation for an experimental guided projectile prototype », *in: AIAA Guidance, Navigation, and Control Conference, San Diego, USA*, 2016, p. 1875.
- [Sil+16] Emilia Silvas, Theo Hofman, Nikolce Murgovski, LF Pascal Etman, and Maarten Steinbuch, « Review of optimization strategies for system-level design in hybrid electric vehicles », *in: IEEE Transactions on Vehicular Technology 66.1* (2016), pp. 57–70.
- [SP07] Sigurd Skogestad and Ian Postlethwaite, *Multivariable feedback control: analysis and design*, vol. 2, Wiley New York, 2007.
- [Str+14] Guillaume Strub, Vincent Gassmann, Spilios Theodoulis, Simona Dobre, and Michel Basset, « Hardware-in-the-loop experimental setup development for a guided projectile in a wind tunnel », *in: 2014 IEEE/ASME International Conference on Advanced Intelligent Mechatronics, Besançon, France*, 2014, pp. 458–463.
- [Str+15] Guillaume Strub, Spilios Theodoulis, Vincent Gassmann, Simona Dobre, and Michel Basset, « Pitch axis control for a guided projectile in a wind tunnel hardware-in-the-loop setup », *in: Journal of Spacecraft and Rockets 52.6* (2015), pp. 1614–1626.
- [Str+18] Guillaume Strub, Spilios Theodoulis, Vincent Gassmann, Simona Dobre, and Michel Basset, « Gain-scheduled autopilot design and validation for an experimental guided projectile prototype », *in: Journal of Guidance, Control, and Dynamics 41.2* (2018), pp. 461–475.

-
- [Str16] Guillaume Strub, « Modeling, Identification and Control of a Guided Projectile in a Wind Tunnel », PhD thesis, Mulhouse, 2016.
- [Tha+19] Sovanna Thai, Spilios Theodoulis, Clément Roos, and Jean-Marc Biannic, « Robust design for the roll-channel autopilot of a canard-guided dual-spin projectile », *in: IFAC-PapersOnLine* 52.12 (2019), pp. 232–237.
- [Tip+20] Steven Tipàn, Spilios Theodoulis, Sovanna Thai, and Michael Proff, « Nonlinear dynamic inversion flight control design for guided projectiles », *in: Journal of Guidance, Control, and Dynamics* 43.5 (2020), pp. 975–980.
- [Tót10] Roland Tóth, *Modeling and identification of linear parameter-varying systems*, vol. 403, Springer, 2010.
- [TP21] Spilios Theodoulis and Michael Proff, « Robust Flight Control Tuning for Highly Agile Missiles », *in: AIAA Scitech 2021 Forum*, 2021, p. 1568.
- [TW17] Spilios Theodoulis and Philippe Wernert, « Flight dynamics & control for smart munition: the ISL contribution », *in: IFAC-PapersOnLine* 50.1 (2017), pp. 15512–15517.
- [Vas+20] Joseph D Vasile, Joshua Bryson, Benjamin C Gruenwald, Luisa Fairfax, Luke Strohm, and Frank Fresconi, « A Multi-Disciplinary Approach to Design Long Range Guided Projectiles », *in: AIAA Scitech 2020 Forum*, 2020, p. 1993.
- [VBF20] Joseph D Vasile, Joshua Bryson, and Frank Fresconi, « Aerodynamic Design Optimization of Long Range Projectiles using Missile DATCOM », *in: AIAA Scitech 2020 Forum*, 2020, p. 1762.
- [VJ84] S Vukelich and J Jenkins, « Missile DATCOM-Aerodynamic prediction of conventional missiles using component build-up techniques », *in: 22nd Aerospace Sciences Meeting*, 1984, p. 388.
- [VTS22] Gian Vinco, Spilios Theodoulis, and Olivier Sename, « Flight Dynamics Modeling and Simulator Design for a New Class of Long-Range Guided Projectiles », *in: EuroGNC (CEAS Conference on Guidance, Navigation and Control)*, 2022.
- [Wat07] Barry D Watts, *Six decades of guided munitions and battle networks: Progress and prospects*, Citeseer, 2007.

-
- [WH03] R Whyte and W Hathaway, *PRODAS—Projectile Design and Analysis Computer Aided Engineering Tool*, tech. rep., Arrowtech associates, Burlington, USA, 2003.
- [WP97] Eric Walter and Luc Pronzato, *Identification of parametric models from Experimental Data*, Springer Verlag New-York, 1997.
- [YL03] Guang-Hong Yang and Kai-Yew Lum, « An optimization approach to integrated aircraft-controller design », *in: Proceedings of the 2003 American Control Conference, 2003*. Vol. 2, IEEE, 2003, pp. 1649–1654.

APPENDIX

A.1 Equations of the Semi-Local Aerodynamic Model

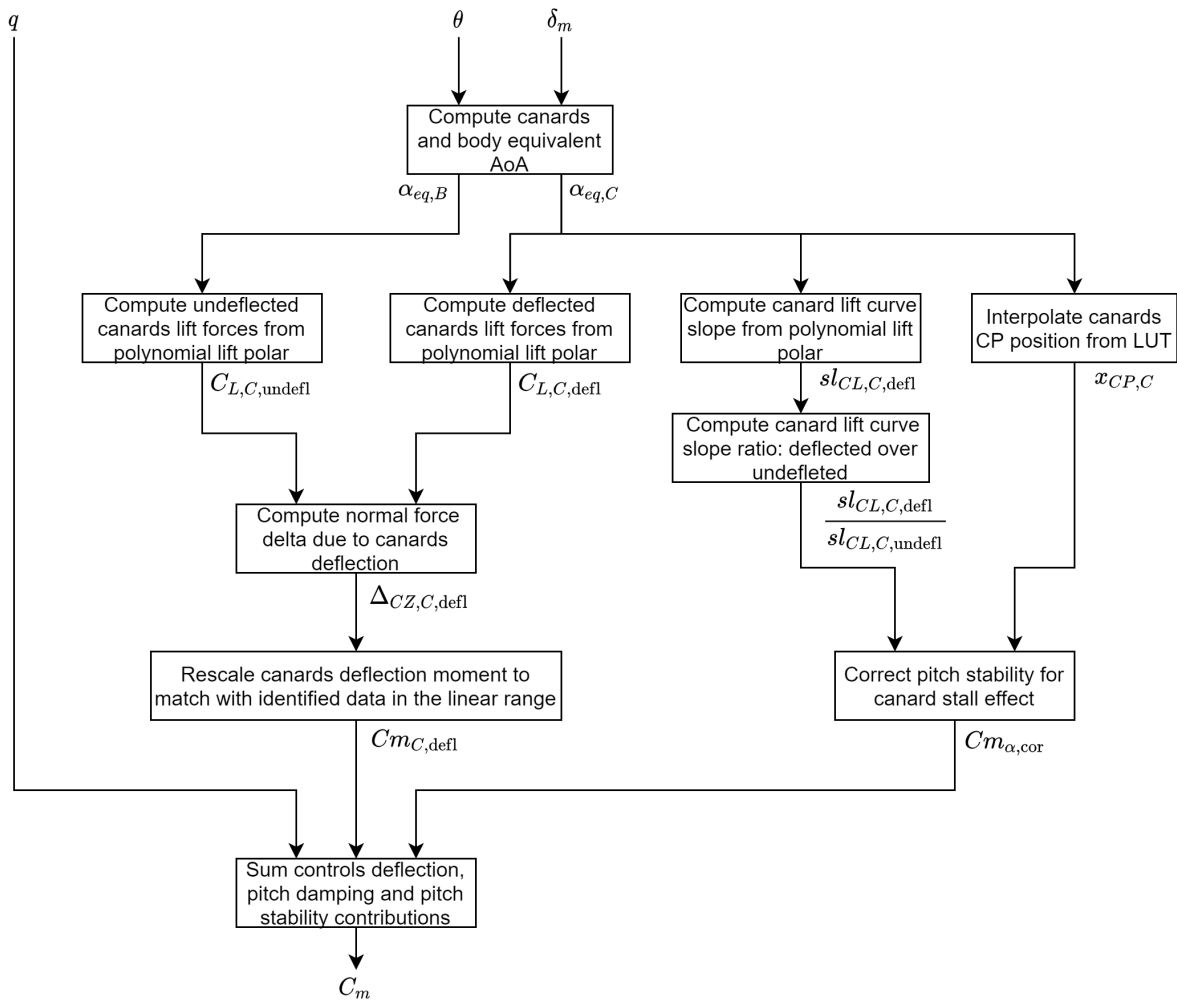


Figure 62 – Flowchart of the semi-local model

Equivalent AoA of canards and body

Assume that canards are in the X configuration and roll angle is zero to define their angular positions:

$$\boldsymbol{\phi}_C = \left[\frac{\pi}{4}, \frac{3\pi}{4}, \frac{5\pi}{4}, \frac{7\pi}{4} \right]^T$$

Use control allocation to compute canards deflections:

$$\boldsymbol{\delta}_C = \begin{bmatrix} 1 & 1 & 1 \\ 1 & -1 & 1 \\ 1 & -1 & -1 \\ 1 & 1 & -1 \end{bmatrix} \begin{bmatrix} 0 \\ \delta_m \\ 0 \end{bmatrix}$$

Compute canards local AoA:

$$\boldsymbol{\alpha}_C = \sin^{-1}(\cos(\boldsymbol{\phi}_C) \sin(\alpha))$$

Compute the equivalent AoA of the undeflected canards:

$$\boldsymbol{\alpha}_{C,eqz} = k_{CB} \boldsymbol{\alpha}_C$$

Estimate the canard-canard interactions with $\boldsymbol{\Lambda}$ interpolated from [NHS77, p. 242]:

$$\boldsymbol{\Lambda} = \begin{bmatrix} 0 & -0.112 & -0.046 & -0.112 \\ 0.112 & 0 & 0.112 & -0.046 \\ -0.046 & -0.112 & 0 & -0.112 \\ 0.112 & -0.046 & 0.112 & 0 \end{bmatrix}$$

Deduce the equivalent AoA of the deflected canards:

$$\boldsymbol{\alpha}_{C,eq} = \tan^{-1} \left(\tan(\boldsymbol{\alpha}_{C,eqz}) + k_{CB} \left[\tan(\boldsymbol{\alpha}_{C,eqz} + \boldsymbol{\delta}_C) - \tan(\boldsymbol{\alpha}_{C,eqz}) + \tan(\boldsymbol{\Lambda} \boldsymbol{\delta}_C) \right] \right)$$

Local lift forces of deflected and undeflected canards

$$C_{L,C,\text{defl}} = \text{sgn}(\alpha_{C,eq}) \left(C_{L,\text{aoff}} + C_{L,\text{aslp}} |\alpha_{C,eq}| + \frac{1}{p_0 + p_1 |\alpha_{C,eq}| + p_2 \alpha_{C,eq}^2} \right)$$

$$C_{L,C,\text{undefl}} = \text{sgn}(\alpha_{C,eqz}) \left(C_{L,\text{aoff}} + C_{L,\text{aslp}} |\alpha_{C,eqz}| + \frac{1}{p_0 + p_1 |\alpha_{C,eqz}| + p_2 \alpha_{C,eqz}^2} \right)$$

The parameters $C_{L,\text{aoff}}$, $C_{L,\text{aslp}}$, p_0 , p_1 and p_2 correspond to the non-linear model of the canard lift polar presented in section 2.6.2.

Normal force delta due to canards deflection

$$\Delta_{CZ,C,\text{defl}} = \cos(\alpha) \sum_C (C_{L,C,\text{defl}} - C_{L,C,\text{undefl}})$$

Rescaling of canards deflection moment

$$Cm_{C,\text{defl}} = \frac{Cm_{\delta m}}{2\sqrt{2} \cos(\alpha) sl_{0,CL,C}} \Delta_{CZ,C,\text{defl}}$$

As a reminder, Cm_α , Cm_q and $Cm_{\delta m}$ are the pitch areodynamic derivatives identified from experimental data (c.f section 2.2.2) while $sl_{0,CL,C}$ is the slope at the origin of the reconstructed canard polar (c.f. section 2.6.2).

Lift curve slopes of deflected and undeflected canards

$$sl_{CL,C,\text{defl}} = C_{L,\text{aslp}} + \frac{p_1 + 2p_2 |\alpha_{C,eq}|}{(p_0 + p_1 |\alpha_{C,eq}| + p_2 \alpha_{C,eq}^2)^2}$$

$$sl_{CL,C,\text{defl}} = \frac{1}{4} \sum_C sl_{CL,C,\text{defl}}$$

$$sl_{CL,C,\text{undefl}} = sl_{0,CL,C}$$

Correction of pitch stability for canard stall effect

$$Cm_{\alpha,\text{cor}} = Cm_{\alpha} - (x_G - x_{CP,C})K_{scl} \left(1 - \frac{sl_{CL,C,\text{def}}}{sl_{CL,C,\text{undef}}} \right)$$

Computation of the total pitching moment

$$Cm = Cm_{\alpha,\text{cor}} \alpha + Cm_q q \frac{D}{\bar{v}} + Cm_{C,\text{def}}$$

