


Assessment and Validation of a  
Partitioned CFD-CA Coupling  
Framework for Rotor Load Prediction  
Master of Science Thesis

Lennard Paul Frederik de Graaf



*This page was intentionally left blank.*

# Assessment and Validation of a Partitioned CFD-CA Coupling Framework for Rotor Load Prediction

by

Lennard Paul Frederik de Graaf

to obtain the degree of Master of Science  
at the Delft University of Technology,  
to be defended publicly on July 3, 2026 at 9:00 AM.

Student number: 4875842  
Project duration: May 15, 2025 – July 3, 2026  
Thesis committee: Dr. Wei Yu TU Delft, Committee chair  
Dr.ir. Alexander van Zuijlen TU Delft, supervisor  
Dr.ir. Matteo Pini TU Delft, Independent examiner  
Faculties: Faculty of Aerospace Engineering, Delft  
Company: Netherlands Aerospace Centre (NLR), Amsterdam

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.



# Preface

What better way to start this document than with a cliché. From a young age, I was fascinated by aircraft and my desire to know how they functioned. Even though I did not spend my childhood memorising aircraft types, begging my parents for flying lessons or visiting air shows, I was certain of a future in the aerospace industry and I knew that studying in Delft would be the first step toward that goal. Reflecting on my time at TU Delft, I realise how much I have grown as a person and as an engineer, and I am proud of my younger self for sticking with his instincts. At the same time, I can now say that he truly did not know what he was getting himself into.

This project would not have been possible without the support of many people. First, I would like to thank Onno Bartels and Stevie-Ray Janssen. After my initial e-mail, they went out of their way to help facilitate this thesis project at NLR and to make me feel part of the AVFP department. I am especially grateful to Stevie for helping create this opportunity and supervising me throughout this project. I would also like to thank Richard Bakker and Stefan van 't Hoff for their support. A word of thanks also goes out to my former colleagues at the ASTC department, and in particular Christoph Mertens, for giving me the opportunity to join and get to know NLR in my previous internship.

Second, I would like to thank Alexander van Zuijlen, who introduced me to the field of fluid-structure interaction during his course at TU Delft. At the time, I had not imagined that this would later become the topic of my thesis project. I am grateful for his excellent supervision, guidance, and feedback throughout this thesis.

Of course, I also want to thank my family. I am grateful to my mother, who never failed to confront me with my progress, and who generously lent me her time, perspective and project management skills. I also want to thank my father and my sister, who were a great source of unconditional help and support. And then there is Jip, our dog, who never cared much about I felt about my thesis or my progress. Instead, she always met me with excitement, reminding me to appreciate the small things and to enjoy moments away from everything else.

A special thanks goes out to my friends and team members from the Brunel Solar Team. Being a part of that project brought some of the highest highs and the lowest lows of my student life in the best and most memorable sense. Above all, I will forever cherish the memory of driving a self-built solar car, Nuna 11S, across South Africa. I have no doubt that the Brunel Solar Team gave me friends for life.

Lastly, I want to thank my girlfriend, Luciana. At a time when I felt lost and unsure about my next endeavours, she supported me, challenged me and helped me refine my beliefs. I am immensely grateful for her help in setting this all in motion, and for her unwavering support throughout this project.

*Lennard Paul Frederik de Graaf*  
Delft, July 2026

# Abstract

This thesis presents the implementation, verification, and assessment of a partitioned CFD-CA coupling framework between the high-fidelity flow solver ENSOLV and the comprehensive rotorcraft analysis environment FLIGHTLAB, based on the delta-airloads methodology. The coupling was implemented in FLIGHTLAB through a dedicated AERODELTA component, which applies an aerodynamic correction, derived from the difference between the CFD and baseline lifting-line airloads, during the trim process, while the updated rotor motion and deformation are fed back into the subsequent CFD solution. The numerical interface was first verified to confirm consistent transfer of reference-frame information, blade motion and deformation, and aerodynamic loads and moments between the two solvers. The framework was then assessed using a test matrix of rigid and flexible Mach-scaled BO105 rotor configurations, covering hover, forward flight, and descending flight, with both steady and unsteady CFD. For the rigid rotor, the coupled simulations converged stably across all flight conditions, and the predicted mean airloads and approximate blade-vortex interaction (BVI) locations compared reasonably well against HART-II experimental data, although the coarse CFD mesh limited the accuracy of the detailed sectional load fluctuations associated with BVI. In contrast, the flexible rotor simulations revealed a critical limitation: inconsistencies in the predicted aerodynamic pitching moment propagated into the elastic torsional response, producing an unstable aeroelastic feedback loop that caused chaotic blade flapping and feathering and, in several cases, prevented convergence of the trim solution. These results indicate that the ENSOLV-FLIGHTLAB coupling is a verified and reliable framework for rigid rotor load prediction, but that its current implementation is not yet suitable for flexible aeroelastic rotor analysis. The dominant limitations are the coupling between the aerodynamic pitching moment and the structural torsional response, and insufficient resolution and preservation of the rotor wake. Recommendations to address these limitations are provided to guide future development of the framework.

# Contents

<b>Preface</b>	<b>i</b>
<b>Abstract</b>	<b>ii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Background and Motivation . . . . .	1
1.2 Industrial Context and NLR . . . . .	2
1.3 Literature Background . . . . .	3
1.4 Research Objective and Scope . . . . .	4
1.5 Thesis Outline . . . . .	5
<b>2 Literature</b>	<b>6</b>
2.1 Fundamentals of Fluid-Structure interaction . . . . .	6
2.1.1 Principle of the coupled problem . . . . .	6
2.1.2 Reference Frames (Lagrangian, Eulerian, and ALE) . . . . .	7
2.1.3 Monolithic and Partitioned . . . . .	7
2.1.4 Time-integration Methods . . . . .	8
2.1.5 Fluid-Structure Interface . . . . .	9
2.2 Fundamentals of Rotorcraft . . . . .	11
2.2.1 Design of Rotorcraft . . . . .	11
2.2.2 Rotor and Blade Design . . . . .	11
2.2.3 Rotor Kinematics . . . . .	13
2.2.4 Rotor Control and Blade Pitch Variation . . . . .	14
2.3 Fundamentals of Rotor Aeromechanics . . . . .	15
2.3.1 The Field of Rotor Aeromechanics . . . . .	15
2.3.2 Blade Section Aerodynamics . . . . .	16
2.3.3 Forces acting on the Rotor Blade . . . . .	17
2.3.4 Structural Coupling . . . . .	18
2.4 Fundamentals of Rotor Aerodynamics . . . . .	19
2.4.1 Wake Structure and Induced Flow . . . . .	19
2.4.2 Transonic flow . . . . .	21
2.4.3 Blade-Vortex Interaction . . . . .	22
2.4.4 Dynamic Stall . . . . .	23
2.5 Rotorcraft prediction Methodologies . . . . .	23
2.5.1 History of Rotorcraft Prediction Methodology . . . . .	23
2.5.2 Formulation of the Delta-Airloads Method . . . . .	26
2.6 Review of Rotor Load Prediction Using partitioned CFD-CA Coupling . . . . .	27
2.6.1 Applications of partitioned CFD-CA Coupling Using the Delta-Airloads Method . . . . .	28
2.6.2 Synthesis of Findings, Results, and Limitations . . . . .	29
<b>3 Methods</b>	<b>32</b>
3.1 Coupling Strategy . . . . .	32
3.1.1 Coupling Philosophy . . . . .	32
3.1.2 Global Coupling Workflow . . . . .	32
3.1.3 Global Timing Scheme . . . . .	34
3.2 Solver Methods . . . . .	35
3.2.1 Comprehensive Analysis Method: FLIGHTLAB . . . . .	35
3.2.2 Aerodynamic Solver: ENSOLV . . . . .	37
3.3 Reference Frames and Solver Consistency . . . . .	39
3.3.1 Reference Frame Definitions: FLIGHTLAB . . . . .	39

3.3.2	Reference Frame Definitions: ENSOLV	40
3.3.3	Solver Consistency	41
3.4	Grid Interface	42
3.4.1	Grid Interface Coupling	42
3.4.2	Mapping Method: Volume Spline Interpolation	43
3.4.3	Transfer of Grid Motion	43
3.4.4	Transfer of Aerodynamic Loading	45
3.5	Implementation of the Delta-Airloads in FLIGHTLAB	46
3.5.1	Introduction of the AERODELTA Component	46
3.5.2	Computation of Delta-Airloads	47
3.5.3	Modified Trim Procedure	47
3.5.4	Relaxation Strategy	49
<b>4</b>	<b>Models</b>	<b>50</b>
4.1	Rotor geometry	50
4.2	FLIGHTLAB Model Set-up	51
4.2.1	Structural discretisation model	51
4.2.2	Aerodynamic Model	52
4.2.3	Numerical discretisation	52
4.3	ENSOLV	53
4.3.1	Flow and turbulence modelling	53
4.3.2	Temporal discretisation	53
4.3.3	Domain	54
4.3.4	Boundary conditions	54
4.3.5	Computational grid	54
<b>5</b>	<b>Results</b>	<b>57</b>
5.1	Design of Experiment	57
5.1.1	Analysis Strategy	58
5.1.2	Verification Test Cases	58
5.1.3	Extension to Unsteady CFD	59
5.1.4	Validation Case Selection	59
5.1.5	Summary of Test Matrix	60
5.2	Verification	60
5.2.1	Transfer of aerodynamic loading from CFD to CA	61
5.2.2	Alteration for Steady CFD	62
5.2.3	Verification of Delta-Airloads Implementation	63
5.3	Rigid Rotor Simulations	64
5.3.1	C100: Hover (3 kN), Steady CFD, Rigid Rotor	64
5.3.2	C102: Hover (6 kN), Steady CFD, Rigid Rotor	66
5.3.3	C100: Hover (3 kN), Unsteady CFD, Rigid Rotor	67
5.3.4	C851: Forward Flight, Unsteady CFD, and Rigid Rotor	69
5.3.5	Summary of the Rigid Rotor	73
5.4	Flexible Rotor Simulations	74
5.4.1	C400: Hover (3 kN), Steady CFD, Flexible Rotor	74
5.4.2	C402: Hover (4 kN), Unsteady CFD, Flexible Rotor	77
5.4.3	C851: Forward Flight, Unsteady CFD, and Flexible Rotor	80
5.4.4	Summary of the Flexible Rotor	83
5.5	Validation	83
5.5.1	C929: Descending Flight, Unsteady CFD, and Rigid Rotor	84
5.5.2	C929: Descending Flight, Unsteady CFD, and Flexible Rotor	88
5.5.3	Comparison with Experimental Data	91
5.5.4	Validation Summary	94
<b>6</b>	<b>Conclusion</b>	<b>95</b>
<b>7</b>	<b>Recommendations</b>	<b>98</b>
	<b>References</b>	<b>100</b>

---

<b>A</b>	<b>Appendix: Solver Verification</b>	<b>102</b>
A.1	FLIGHTLAB Verification . . . . .	102
A.1.1	Periodicity of the Rotor Motion . . . . .	102
A.1.2	Rigid and Flexible Behaviour . . . . .	103
A.1.3	Trim Targets and Control States . . . . .	104
A.2	Verification of the CFD solution . . . . .	105
A.2.1	Residuals . . . . .	105
A.2.2	Steady CFD solution behaviour . . . . .	106
A.2.3	Unsteady CFD solution behaviour . . . . .	107
<b>B</b>	<b>Appendix: Pre-Processing of CFD solution for Data Exchange</b>	<b>110</b>

# List of Figures

1.1	Schematic representation of rotor blade forces and motions, including flapping, lead-lag, and torsion. [1]	1
1.2	Overview of aerodynamic phenomena occurring in rotorcraft operation, including dynamic stall, blade-vortex interaction, and transonic flow effects. [2]	2
2.1	Serial staggered scheme. [8]	9
2.2	Serial staggered scheme with fluid sub-cycling. [8]	9
2.3	Parallel coupling scheme. [8]	9
2.4	Mid-point staggered scheme. [8]	9
2.5	Non-Matching meshes on fluid-structure interface. [7]	10
2.6	Transfer of interface quantities between non-matching meshes. Fluid loads are projected onto the structural mesh through interpolation of the structural shape functions [7].	11
2.7	Rotorcraft configurations. [9]	12
2.8	Rotor convention in the rotor disk plane. [9]	13
2.9	Rotor blade degrees of freedom.	13
2.10	(a) Force balance on a helicopter during forward flight (b) dissymmetry of the corresponding in-plane velocity distribution. [10]	15
2.11	Conceptual representation of rotor aeromechanics as the interaction between aerodynamics, structural dynamics, and rotor control. Adapted from Smith [11].	16
2.12	Rotor blade section aerodynamics. [9]	17
2.13	Rotor blade side view. Adapted from [9].	18
2.14	Rotor blade top view. Adapted from [9]	18
2.15	Elastic modal blade deformations: flap $u$ , torsion $\theta_t$ and lag $w$ . [12]	19
2.16	Summary of aerodynamic phenomena occurring in the flow field of a helicopter rotor. [13]	20
2.17	Rotor wake formation. [14]	21
2.18	Schematic rotor wake geometries at different flight velocities. [14]	21
2.19	Vortex-induced blade loading. [9]	22
2.20	Schematic representation of the different stages of the dynamic stall process: (a) attached flow, (b) flow reversal, (c) shear layer roll-up, (d) stall onset, (e) full stall, and (f) flow reattachment. [16]	23
2.21	Historical development of comprehensive analysis systems [1]	25
2.22	Airloads Workshop CFD-CA loose coupling breakthrough, UH-60A high-speed airloads, $r/D = 0.965$ , $\mu = 0.368$ . [4]	26
2.23	CFD-CA loose coupling procedure, $F/M=$ forces and moments. [3]	27
3.1	Flowchart describing NLR's implementation of the loose coupling based on the delta-airloads method.	34
3.2	Loose coupling staggered scheme.	35
3.3	FLIGHTLAB rotor coordinate systems and convention.	40
3.4	Definition of the inertial, body and wind-axis system in ENSOLV. [35]	40
3.5	Comparison of free-stream and angle of attack conventions between FLIGHTLAB and ENSOLV. Rotor blades shown in grey.	41
3.6	Geometric coupling interface between FLIGHTLAB and ENSOLV. (ENSOLV surface grid in grey and FLIGHTLAB points in red)	42
3.7	Initial, undeformed grid slice near the blade tip.	44
3.8	Deformed grid slice near the blade tip.	45
3.9	Aerodynamic loading distribution on the rotor blade. Force distribution from ENSOLV in red. Force representation in FLIGHTLAB in blue.	46

3.10 Inner loop: Flowchart of the computation and application of the delta-airloads in FLIGHTLAB. . . . .	48
4.1 BO105 model main rotor geometry and blade twist distribution.[37] . . . . .	51
4.2 Definition of the structural and aerodynamic elements on the rotor blade in FLIGHTLAB. . . . .	52
4.3 Spherical computational domain with radius $R = 50$ m. . . . .	54
4.4 Local grid resolution near the blade tip region. . . . .	55
4.5 Top view of the rotor and structured multi-block grid topology. . . . .	56
5.1 Comparison of the aerodynamic solution from CFD and the spline-mapped aerodynamic loading on the CA discretization. Case C851: Forward flight & rigid rotor. . . . .	61
5.2 Discrete force $F_z$ per ACP, determined by CFD and transformed to the representation used by FLIGHTLAB. . . . .	62
5.3 Azimuthal variation of the aerodynamic force components obtained by rotating the steady CFD solution. Case C102, chosen location: ACP and at 75% span. . . . .	63
5.4 Global iteration history of the total vertical force $F_z$ for the C100 rigid hover case(3 kN). . . . .	64
5.5 Spanwise distribution of the vertical force $F_z$ in the first and final global iteration for the C100 rigid hover case. . . . .	65
5.6 Global iteration history of the total vertical force $F_z$ for the C102 rigid hover case (6 kN) with outer relaxation. . . . .	66
5.7 Spanwise distribution of the vertical force $F_z$ in the first and final global iteration for the C102 rigid hover case. . . . .	67
5.8 Global iteration history of the mean total vertical hub force for Case C100-UNS-RGD. . . . .	68
5.9 Hub force $F_z$ as a function of azimuth at iteration 1. . . . .	68
5.10 Hub force $F_z$ as a function of azimuth at iteration 6. . . . .	68
5.11 Spanwise distribution of $F_z$ during the first and final coupling iterations for Case C100. Comparison of steady and unsteady CFD. . . . .	69
5.12 Evolution of the mean hub forces and moments during the coupling procedure for Case C851. . . . .	69
5.13 Blade pitch distribution derived from the control states. . . . .	70
5.14 Total blade pitch distribution including blade twist. . . . .	70
5.15 Azimuthal distribution of hub forces for Case 851 during the first coupling iteration. . . . .	71
5.16 Azimuthal distribution of hub forces for Case 851 after convergence (iteration 10). . . . .	71
5.17 Azimuthal distribution of hub moments during the first coupling iteration. . . . .	71
5.18 Azimuthal distribution of hub moments in the tenth coupling iteration. . . . .	72
5.19 Spanwise and azimuthal distribution of $F_z$ from the CFD solution. Represented as discrete forces at the ACP's in FLIGHTLAB. . . . .	72
5.20 Spanwise and azimuthal distribution of baseline $F_z$ from FLIGHTLAB. Represented as discrete forces at the ACP's in FLIGHTLAB. . . . .	72
5.21 Q-criterion visualisation of the rotor wake for Case C851. . . . .	73
5.22 Global iteration history of the total vertical force $F_z$ for the flexible hover case C400 (3 kN). . . . .	75
5.23 Comparison of the flexible blade motion for Case C400 in the initial and converged coupling iterations. . . . .	76
5.24 Initial azimuthal loading and correction for the flexible hover case C400. . . . .	76
5.25 Converged azimuthal loading and correction for the flexible hover case C400. . . . .	77
5.26 Global iteration history of the mean total vertical hub force for Case C402. . . . .	78
5.27 Initial blade motion and azimuthal thrust distributions for Case C402. . . . .	79
5.28 Flexible blade motion of the converged solution (C402). . . . .	80
5.29 Evolution of the mean total hub forces for Case C851 over the coupling iterations. . . . .	80
5.30 Evolution of the azimuthal hub-force distributions for Case C851. . . . .	81
5.31 Evolution of the azimuthal hub-moment distributions for Case C851. . . . .	82
5.32 Comparison of the CFD blade-normal force distribution for Case C851 at the first and seventh coupling iterations. Represented on the aerodynamic FLIGHTLAB grid. . . . .	83
5.33 Evolution of the mean hub forces and moments during the coupling procedure for Case C929. (Rigid rotor) . . . . .	84

5.34	Azimuthal distribution of hub forces for Case 929 during the first coupling iteration (rigid rotor blade). . . . .	85
5.35	Azimuthal distribution of hub forces for Case 929 during the 7th coupling iteration (rigid rotor blade). . . . .	86
5.36	Azimuthal distribution of hub moments for Case 929 during the first coupling iteration (rigid rotor blade). . . . .	86
5.37	Azimuthal distribution of hub moments for Case 929 during the seventh coupling iteration (rigid rotor blade) . . . . .	86
5.38	Blade-normal force distribution from CFD. . . . .	87
5.39	Q-criterion iso-surface coloured by Mach number for Case C929 with a rigid rotor. . . . .	88
5.40	Evolution of the mean hub forces and moments versus coupling iterations for Case C929 (flexible rotor). . . . .	89
5.41	Azimuthal distribution of hub forces for Case 929 during the first coupling iteration (flexible rotor blade) . . . . .	90
5.42	Azimuthal distribution of hub forces for Case 929 during the ninth coupling iteration (flexible rotor blade) . . . . .	90
5.43	Azimuthal distribution of hub moments for Case 929 during the ninth coupling iteration (flexible rotor blade) . . . . .	90
5.44	Azimuthal distribution of hub forces for Case 929 during the ninth coupling iteration (flexible rotor blade) . . . . .	91
5.45	Blade-normal force distribution from CFD for case C929 in loop 7 (flexible rotor blade). . . . .	91
5.46	Normal force predictions at 87 % radial station. Comparison between HART-II experiment, and simulated rigid and flexible rotor. . . . .	93
5.47	Pitching moment predictions at 87 % radial station. Comparison between HART-II experiment, and simulated rigid and flexible rotor. . . . .	94
A.1	Assessment of the periodicity of the rotor motion over two consecutive revolutions . . . . .	102
A.2	Vertical rigid blade position in hover for iterations 1 and 8. Case C100 . . . . .	103
A.3	Relative blade pitch deformation in the ACPs versus span and azimuth. . . . .	104
A.4	Rotor blade flapping and pitch deformation in the ACPs versus span and azimuth. Case C400. . . . .	104
A.5	Hub loads for the trimmed rigid hover solution (Rigid hover case C100). . . . .	105
A.6	RMS- and max-norm residuals of the continuity, momentum, and energy equations vs. the iteration index. Case C100-RGD-STD initial CFD simulation. . . . .	106
A.7	Total lift for the steady C100 hover case. . . . .	107
A.8	Residual history for the unsteady simulation during coupling loop 1 (Case C100). . . . .	108
A.9	Residual history for the unsteady simulation during coupling loop 8 (Case C100). . . . .	108
A.10	Lift force distributions for unsteady CFD (Case 100) at different coupling loops. . . . .	109
B.1	Left: cell face-centred grid. Middle: cell vertex grid. Right: overlap of both. . . . .	111

# List of Tables

2.1	Overview of rotor blade structural coupling mechanisms (cause → response → effect). Adapted from Johnson [9]. . . . .	20
2.2	Summary of Key Findings for Loose CFD-CA Coupling. . . . .	31
4.1	Mach-scaled BO105 main rotor key data . . . . .	51
5.1	Test matrix for verification and validation cases used in the CFD-CA coupling study. . .	60
5.2	Comparison of aerodynamic loading, represented by total force $F_z$ , between CFD and FLIGHTLAB for C100 case. $F_x$ and $F_y$ are approximately 0 and thus omitted. . . . .	61
5.3	Verification of delta-airloads for global coupling iterations. Total rotor thrust in $z$ for a steady hover case with trim target of 3 kN. . . . .	63
5.4	Swashplate collective pitch $\theta_0$ and $\Delta$ -airloads correction per global iteration (C100). . . .	65
5.5	Collective pitch $\theta_0$ and delta-airloads correction per global iteration (C102). . . . .	66
5.6	Collective pitch $\theta_0$ and delta-airloads correction per global iteration (C100-UNS-RGD), with final steady CFD values for comparison. . . . .	67
5.7	Evolution of the control states for Case C851 over the global coupling iterations. . . . .	70
5.8	Collective pitch $\theta_0$ for the first five global iterations (flexible C400), with steady C100 results for comparison. . . . .	75
5.9	Coupling loop convergence history (Case C402). . . . .	78
5.10	Approximate control-state evolution for Case C851 over the first seven coupling iterations. . .	81
5.11	Experimental trim control angles of the HART-II experiment . . . . .	83
5.12	Evolution of the control states for Case C929 over the global coupling iterations for the simulated rigid rotor. . . . .	84
5.13	Case C929 — Control state evolution across coupling (flexible blade) . . . . .	88
A.1	Trimmed control states for Case C100 rigid hover. . . . .	104
B.1	Verification of force conservation during face-centre to vertex redistribution. . . . .	111

# Nomenclature

## Abbreviations

Abbreviation	Definition
ACP	Aerodynamic Control Point
ALE	Arbitrary Lagrangian-Eulerian
AOA	Angle of Attack
BVI	Blade-Vortex Interaction
CA	Comprehensive Analysis
CFD	Computational Fluid Dynamics
CFL	Courant–Friedrichs–Lewy (number)
CSD	Computational Structural Dynamics
CSS	Conventional Serial Staggered
DOF	Degree of Freedom
EARSM	Explicit Algebraic Reynolds Stress Model
FEM	Finite Element Method
FSI	Fluid-Structure Interaction
HART-II	Higher Harmonic Control Aeroacoustic Rotor Test II
NLR	Netherlands Aerospace Centre
RANS	Reynolds-Averaged Navier–Stokes
RBF	Radial Basis Function
RMS	Root-Mean-Square
URANS	Unsteady Reynolds-Averaged Navier–Stokes

## Symbols

Symbol	Definition	Unit
$A$	Finite-volume cell discretisation matrix	[-]
$C_s$	Structural damping matrix	[-]
$D$	Rotor diameter	[m]
$f_s$	Structural external force vector	[N]
$F, M$	Forces and moments	[N], [Nm]
$F^{LL}, F^{CFD}$	Lifting-line and CFD sectional airloads in the delta-airloads formulation (Chapter 2)	[N]
$F^F, F^S$	CFD-derived and FLIGHTLAB (structural baseline) aerodynamic loading at inner trim-loop iteration $k$ , used in the delta-airloads implementation (Chapter 3)	[N]
$\mathcal{F}$	Fluid-solver operator in the partitioned coupling scheme	[-]
$F_x, F_y, F_z$	Force components in the reference frame; $F_z$ denotes vertical force / rotor thrust	[N]
$i$	Global coupling iteration index	[-]
$K_s$	Structural stiffness matrix	[-]
$M_s$	Structural mass matrix	[-]
$M_x, M_y, M_z$	Moment components about the coordinate axes	[Nm]
$n$	Interface normal vector; also time-step index depending on context	[-]
$p$	Pressure or stress quantity	[Pa]


Symbol	Definition	Unit
$r_k$	Relaxation factor applied at inner trim-loop iteration $k$	[-]
$r_0$	Initial relaxation factor	[-]
$R$	Rotor radius; also source-term vector depending on context	[m] / [-]
$\mathcal{S}$	Structural-solver operator in the partitioned coupling scheme	[-]
$t$	Time	[s]
$u$	Displacement or deformation	[m]
$W$	Vector of conservative flow variables	[-]
$w$	Lag deformation	[m]
$x$	Position of moving computational mesh	[m]
$\dot{x}$	Mesh velocity	[m s <sup>-1</sup> ]
$\alpha$	Angle of attack	[rad]
$\beta$	Flapping angle	[rad]
$\Gamma$	Fluid–structure interface	[-]
$\mu$	Advance ratio	[-]
$\rho$	Density	[kg m <sup>-3</sup> ]
$\sigma_f, \sigma_s$	Fluid and structural stress tensors	[Pa]
$\tau$	Viscous stress tensor	[Pa]
$\theta$	Pitch angle / torsion angle	[rad]
$\theta_t$	Torsional deformation	[rad]
$\theta_0$	Collective pitch	[rad]
$\theta_{1c}$	Lateral cyclic pitch	[rad]
$\theta_{1s}$	Longitudinal cyclic pitch	[rad]
$\Omega$	Rotor angular velocity	[rad s <sup>-1</sup> ]
$\Omega_f$	Fluid domain	[-]
$\Omega_s$	Structural domain	[-]
$\Phi_c$	Convective flux term in ALE formulation	[-]
$\psi$	Blade azimuth angle	[rad]
$\epsilon$	Convergence tolerance	[-]
$f$	Subscript denoting fluid	[-]
$s$	Subscript denoting structure / structural	[-]
LE, TE	Subscripts denoting leading-edge and trailing-edge quantities	[-]
$x, y, z$	Cartesian coordinate directions	[-]
0	Collective component	[-]
1c	First-harmonic cosine component	[-]
1s	First-harmonic sine component	[-]
$\dot{(\cdot)}$	First time derivative	[-]
$\ddot{(\cdot)}$	Second time derivative	[-]
$(\cdot)^{(k)}$	Coupling sub-iteration index	[-]
$(\cdot)^n, (\cdot)^{n+1}$	Time-level notation	[-]

# Introduction

## 1.1. Background and Motivation

Rotorcraft are a class of aircraft distinguished from fixed-wing aircraft by their use of large-diameter rotating wings, referred to as rotors, to generate lift and control. A rotor consists of multiple blades attached to a central hub and rotates to produce aerodynamic loading. As the blades move through the air, the local inflow conditions vary continuously along the span of the blade and throughout each rotation, leading to a non-uniform loading distribution over the rotor disk.

A defining feature of rotor systems, described by the field of rotor aeromechanics, is the interaction between aerodynamics, structural motion, and control. Typically, rotor blades are long, slender structures with high aspect ratios that are designed to be flexible and low in weight. When subjected to aerodynamic and inertial forces, these blades may deform, leading to motions such as flapping, lead-lag, and torsion, as illustrated in Figure 1.1. These motions alter the local inflow conditions experienced by the blade and, consequently, the forces acting on it. In addition to the aerodynamic loading and structural response, rotor behaviour is inherently linked to control. During operation, the pitch of each blade is continuously adjusted using collective and cyclic pitch inputs, allowing the rotor to produce the required forces and moments for different flight conditions. As a result, rotor behaviour is governed by the interaction of aerodynamic loading, structural motion, and control inputs.



**Figure 1.1:** Schematic representation of rotor blade forces and motions, including flapping, lead-lag, and torsion. [1]

The flow environment around a rotor is characterised by a wide range of aerodynamic phenomena that occur simultaneously, as illustrated in Figure 1.2. Near the rotating blade, the flow can become compressible and unsteady, particularly on the advancing side, where transonic or locally supersonic




Figure 2 A summary of specific flow problems which occur on a helicopter. From Caradonna (1992).

Figure 1.2: Overview of aerodynamic phenomena occurring in rotorcraft operation, including dynamic stall, blade-vortex interaction, and transonic flow effects. [2]

conditions and shocks may occur. On the retreating side, lower inflow velocities combined with higher angles of attack can result in flow separation and dynamic stall. In addition, the rotating blades may shed tip vortices that interact with a following blade, a phenomenon known as blade-vortex interaction (BVI). The wake downstream of the rotor consists of these vortex structures, which persist over multiple rotor revolutions and influence the inflow conditions experienced by the blades.

The coupled interaction between aerodynamics, structural motion, and control governs rotor behaviour and directly influences key performance aspects, including efficiency, aeroelastic stability, vibration levels, and handling characteristics. To perform rotorcraft analysis, it is therefore essential to accurately predict the key quantities that describe the aerodynamic blade loading, rotor motion and deformation, and control states. Moreover, inaccurate prediction methods can lead to increased development costs, extended project timelines, and elevated risk during rotorcraft development. This motivates the development and assessment of modelling methodologies capable of accurately resolving this behaviour.

## 1.2. Industrial Context and NLR

The research presented in this thesis was carried out as part of ongoing development at the Netherlands Aerospace Centre (NLR) into improving rotorcraft prediction capabilities. As a national research institute, NLR's primary goal is to support the growth and development of the aerospace industry in the Netherlands. Hence, NLR actively researches and develops new modelling, simulation, and analysis techniques. In the field of rotorcraft, NLR is interested in the accurate prediction of rotor blade loads, in particular for applications such as fatigue analysis, where reliable load estimates are critical for the structural design and lifetime assessment of rotor blades.

To simulate rotorcraft, NLR employs the comprehensive analysis tool FLIGHTLAB. This simulation environment models rotorcraft aerodynamics, flight mechanics, structures, and control systems within a single framework, allowing for accurate and efficient rotorcraft analyses during all stages of the design process. To model the aerodynamics, FLIGHTLAB uses a lifting-line approach, which provides computational efficiency but often fails to accurately represent complex flow phenomena.

To improve the prediction of rotor loads, NLR utilises the in-house flow solver ENSOLV. The high-fidelity flow modelling in ENSOLV allows for a detailed and accurate simulation of the three-dimensional flow around a rotor and is able to capture aerodynamic phenomena that low-order models cannot. However, the high computational cost associated with CFD simulations makes them unsuitable for routine use within comprehensive analyses like FLIGHTLAB.

To combine the strengths of both approaches, NLR has implemented a loose coupling between ENSOLV and FLIGHTLAB based on the delta-airloads method, originally proposed by Potsdam et al. [3]. While this framework has been developed and applied in initial studies, its predictive performance for practical rotorcraft applications, such as fatigue load prediction, has not yet been established. This motivates the need for a systematic evaluation of the coupling framework's performance, with the aim of improving its applicability for rotorcraft analysis and load prediction.

### 1.3. Literature Background

Traditionally, rotorcraft aeromechanics analyses have relied on comprehensive analysis (CA), an interdisciplinary methodology, for predicting rotorcraft characteristics. These tools combine models for aerodynamics, structural dynamics, flight dynamics, and control into a single code. A key strength of comprehensive analyses is their broad applicability. These codes are capable of calculating rotorcraft performance, blade motion and loading, structural loads, and flight dynamics for all operating conditions and rotorcraft configurations at every stage of the design process [4].

The design and development of rotorcraft requires the capability to calculate rotor performance and aerodynamic loading. To support these calculations, comprehensive analyses rely on lifting-line and vortex wake models to account for the airloads. While these approaches are efficient, they rely on simplified assumptions and are limited in their ability to accurately represent the complex three-dimensional flow field of rotors. As a result, effects such as flow reversal, dynamic stall, and blade-vortex interaction (BVI) are often only approximated or neglected by these models. As a consequence, severe discrepancies can arise in the computed airloads and rotor response for high-speed, high-thrust, and manoeuvring conditions, where complex aerodynamic effects become dominant.

The simplifications applied in CA tools also affect the modelling of rotorcraft structural dynamics. Generally, CA tools rely on reduced-order structural models that are designed to efficiently capture the dominant blade motions, such as flapping, lead-lag, and torsion, but are unable to resolve the full three-dimensional structural behaviour of the blade. While higher-fidelity approaches, such as the three-dimensional finite element method (FEM), could provide more structural information, a compromise is made based on the need for routine simulations with reduced computational cost and complexity.

Modern CFD methods offer the opportunity to overcome the aerodynamic limitations of comprehensive analyses. By directly resolving the complexities of flow separation, blade stall, three-dimensional unsteady transonic flow, and rotor wake, CFD is able to compute a high-fidelity aerodynamic solution. However, the computational cost associated with CFD makes it impractical to directly replace the aerodynamic models in CA for routine applications.

This has led to the development of coupled CFD–CSD approaches, in which a CFD solver is combined with a structural dynamics model, allowing for the simultaneous prediction of blade deformation, rotor motion, and aerodynamic loading. However, coupling CFD to a comprehensive analysis is not trivial. In tightly coupled problems, aerodynamic airloads, blade deformation, and rotor trim are fully interdependent, and the aerodynamics and structural dynamics must be integrated simultaneously at every time-step. Here, the interdependence becomes a downside, making this approach computationally expensive and impractical for routine applications.

The alternative is to use a loosely coupled approach in which information is transferred between the CFD and CSD on a periodic, per-revolution basis. The most effective coupling strategy in this class is the delta-airloads method, originally proposed by Potsdam et al. [3]. Rather than fully replacing the aerodynamic loads computed by the low-fidelity model of the CA, this approach determines a correction based on the CFD results. This correction, referred to as the delta-airloads, is based on the difference between the airloads of the CA and those of the CFD and is added to the predicted airloads of

the low-fidelity model. After several coupling iterations, the CA airloads are completely replaced by the CFD airloads, and the result is a converged, coupled, and trimmed aerodynamic and structural solution.

The potential of loosely coupled CFD–CSD approaches for rotor load prediction has been shown in a number of studies. Most importantly, these studies demonstrated that coupling CFD to a CA framework leads to an improvement in rotor airload prediction, particularly for cases involving unsteady aerodynamics such as BVI and transonic flow.

In addition, a key finding in these studies is that the benefit of the CFD correction is not uniform, but is strongest when correcting the dominant aeroelastic mechanisms. This was demonstrated by Datta et al. [5], [6] for the sectional pitching moment. Physically, the sectional pitching moment was shown to cause elastic blade torsion, which in turn affected the airloads and structural response. This work established the causal relationship between aerodynamic moment prediction and rotor loading, and underscored the importance of accurate moment prediction for the loose coupling.

Moreover, these studies have shown that accurate structural modelling is a prerequisite for successful coupling. The structural model must represent the correct blade dynamics for the blade to provide an appropriate aeroelastic response to the applied aerodynamic loading. Moreover, the modelling and inclusion of the wake in the aerodynamic solution were found to be essential, as blade loading is strongly influenced by wake interactions.

While these studies demonstrate that loosely coupled CFD–CSD methods can improve the prediction of rotor loading, they also highlight that the effectiveness of these methods depends strongly on modelling choices and implementation details. This further reinforces the need for structured evaluations in the development of loosely coupled frameworks. This need is reflected in the industrial context at NLR, where the predictive performance of the coupling between ENSOLV and FLIGHTLAB has not been fully established.

## 1.4. Research Objective and Scope

The objective of this thesis is to implement, verify, and assess a loosely coupled CFD–CSD framework between ENSOLV and FLIGHTLAB based on the delta-airloads methodology, and to identify the requirements and limitations affecting its ability to successfully predict rotor loads. The focus of this thesis is on evaluating the coupling implementation, the numerical behaviour, and the capability to predict rotor blade loading, rotor motion, and trim for a selected number of flight conditions and rotor configurations.

As part of this research, a series of simulations is performed with progressively increasing complexity. The analysis starts with simplified hover cases to verify the coupling methodology and evaluate the behaviour of the numerical interface. Subsequently, more complex conditions, including unsteady CFD, forward flight, and descending flight, are used to extend the analysis. Both rigid and flexible rotor configurations are considered to assess the influence of structural blade modelling on the coupled solution.

In this work, a distinction is made between computational robustness and physical credibility. The former refers to the ability of the coupled framework to obtain stable, converged, and repeatable solutions for the chosen modelling choices and numerical set-up. The latter refers to whether the converged solutions are consistent with the expected rotor behaviour and, if available, experimental data. This distinction is made because successful implementation and numerical convergence have been shown not to automatically imply predictive accuracy, particularly when simulating flexible rotor blades.

The main research question addressed in this thesis is:

*What is required for the successful implementation, verification, and assessment of a loosely coupled CFD–CSD framework based on the delta-airloads methodology for rotor-load prediction?*

This main research question is addressed through the following sub-questions:

- To what extent does the implemented coupling approach represent the proposed delta-airloads methodology and the rotor aeromechanical problem?

- For which selected modelling choices, numerical settings, flight conditions, and rotor configurations does the coupled framework produce stable, converged, and predictively reliable rotor-load simulations?
- What limitations prevent the framework from being used reliably for flexible aeroelastic rotor-load prediction?

## **1.5. Thesis Outline**

This thesis is structured as follows. Chapter 2 presents the theoretical background on fluid-structure interaction, rotorcraft aeromechanics, and coupling methodologies. Chapter 3 describes the coupling methodology and numerical approach. Chapter 4 details the implementation of the models and the simulation set-up. Thereafter, Chapter 5 presents the results of the simulation cases and their assessment. Lastly, Chapter 6 discusses the findings and provides conclusions and recommendations for future work.

# 2

## Literature

This chapter introduces the fundamental concepts required to understand rotorcraft aeromechanics. It first presents the principles of fluid-structure interaction, followed by the basic design, operation, and control of rotorcraft. The chapter then describes the coupled aerodynamic and structural behaviour of rotor blades, including blade motion, loading, and key aerodynamic phenomena such as wake effects, blade-vortex interaction, transonic flow, and dynamic stall. Together, these topics provide the physical foundation for the rotor prediction and fluid-structure interaction methods used in this work.

### 2.1. Fundamentals of Fluid-Structure interaction

#### 2.1.1. Principle of the coupled problem

As the name suggests, Fluid-Structure interaction (FSI) concerns multi-physics problems where the interactions between fluid and structure cannot be ignored and the two systems must be considered together. Typically, a fluid-structure interaction system consists of two domains: a fluid domain  $\Omega_f$  and a structural domain  $\Omega_s$ , which share a common interface  $\Gamma$ . At the interface, the two physical systems can exchange information in the form of the displacement (structure on fluid) and forces (fluid on structure). The exchange of information between the two systems is enforced through interface conditions. These conditions allow for the governing equations of the fluid and structure models to be coupled. They can be expressed as a dynamic condition and two kinematic conditions.

The dynamic condition states that the tractions of the surface of the structure are in equilibrium with those on the fluid side. The general form states that the stress exerted by the fluid on the structure must be balanced by the traction exerted by the structure on the fluid:

$$\boldsymbol{\sigma}_s \mathbf{n}_s = \boldsymbol{\sigma}_f \mathbf{n}_f \quad \text{on } \Gamma, \quad (2.1)$$

where  $\boldsymbol{\sigma}_f$  and  $\boldsymbol{\sigma}_s$  denote the fluid and structural stress tensors, respectively, and  $\mathbf{n}_f$  and  $\mathbf{n}_s$  are the outward unit normal vectors of the fluid and structural domains at the interface. The fluid stress tensor contains both the pressure and viscous contributions, and can be written as

$$\boldsymbol{\sigma}_f = -p\mathbf{I} + \boldsymbol{\tau}, \quad (2.2)$$

where  $p$  is the fluid pressure,  $\mathbf{I}$  is the identity tensor, and  $\boldsymbol{\tau}$  is the viscous stress tensor. This formulation therefore includes both pressure loading and wall shear stress at the interface. The kinematic conditions ensure compatibility of displacement and velocity at the interface:

$$\mathbf{u}_f = \mathbf{u}_s \quad \text{on } \Gamma, \quad (2.3)$$

$$\dot{\mathbf{u}}_f = \dot{\mathbf{u}}_s \quad \text{on } \Gamma, \quad (2.4)$$

where  $\mathbf{u}_{f,s}$  is the displacement, and  $\dot{\mathbf{u}}_{f,s}$  is the velocity. These conditions ensure that the fluid boundary moves together with the structural surface, and that no gaps or overlaps occur between the two domains. In numerical simulations, these interface conditions are implemented through boundary conditions and coupling algorithms connecting the fluid and structural solvers [7].

### 2.1.2. Reference Frames (Lagrangian, Eulerian, and ALE)

Typically, different reference frames for the fluid and structural domains are used to formulate the governing equations of FSI problems. These descriptions reflect the physical nature of each system and how their motion is best represented.

Structural dynamics are generally formulated using a Lagrangian approach. In this reference frame the computational mesh moves with the material points of the structure. Hence, the governing equations of motion are written with respect to the position of the structure  $\mathbf{u}_s$ , and the mesh deforms together with the body. The structural dynamics equations are typically formulated using a discretized finite-element representation as

$$\mathbf{M}_s \ddot{\mathbf{u}}_s + \mathbf{C}_s \dot{\mathbf{u}}_s + \mathbf{K}_s \mathbf{u}_s = \mathbf{f}_s, \quad (2.5)$$

where  $\mathbf{M}_s$ ,  $\mathbf{C}_s$ , and  $\mathbf{K}_s$  are the structural mass, damping, and stiffness matrices respectively, and  $\mathbf{f}_s$  represents the external forces acting on the structure.

In contrast, fluid dynamics are usually formulated in an Eulerian reference frame, in which the computational mesh remains fixed in space and the fluid flows through it. The governing equations are then written in terms of flow variables such as the velocity  $\mathbf{u}_f$ , pressure  $p$ , and density  $\rho$ , which are evaluated at fixed spatial locations.

For numerical simulations of FSI problems, the fluid equations are often discretized using a finite-volume approach and then expressed in an Arbitrary Lagrangian–Eulerian (ALE) formulation. In semi-discrete finite-volume form, the ALE Navier–Stokes equations can be written as

$$\frac{d}{dt} [A \mathbf{W}] + \Phi_c (\mathbf{W}, \mathbf{x}, \dot{\mathbf{x}}) = \mathbf{R} (\mathbf{W}, \mathbf{x}), \quad (2.6)$$

where  $\mathbf{W}$  is the vector of conservative flow variables and  $A$  is the matrix associated with the discretization of the finite-volume cells in the computational domain. The term  $\Phi_c$  arises from the discretization of the convective fluxes, and accounts for the relative motion between the fluid and the computational mesh. The vector  $\mathbf{R}$  represents the contributions of the source and viscous terms. Lastly,  $\mathbf{x}$  and  $\dot{\mathbf{x}}$  denote the position and velocity of the moving computational mesh. The ALE formulation therefore allows the fluid equations to be solved using a moving reference volume on a mesh that is moving in response to the movement of the structure. [7]

### 2.1.3. Monolithic and Partitioned

To model the coupled FSI system, the governing equations of the fluid and structure must be solved simultaneously at a given time step. Two principal approaches exist to solve the coupled system: monolithic and partitioned formulations.

In a monolithic approach, the fluid and structural governing equations are assembled into a single system of equations and are solved simultaneously. This formulation enforces the interface conditions implicitly within the global solution procedure, and therefore ensures a strong coupling between the two physical domains. As a result, monolithic methods are generally robust and stable, particularly in problems where the interaction between the fluid and structure is strong [7].

At the same time, monolithic approaches are often difficult to implement in practice. Fluid and structural solvers are typically developed as separate specialised software programmes, and each is optimised for solving their respective governing equations. Building a unified system requires the integration of these tools. In practice, this requires substantial modifications to the existing codes and leads to complex software frameworks that are difficult to maintain. Additionally, the numerical characteristics and the time-integration schemes of fluid and structural problems often differ significantly. Solving both sets of equations simultaneously may therefore reduce the solver efficiency and lead to unnecessary computational cost. [7]

For these reasons, partitioned approaches are widely used in engineering practice. Partitioned methods treat the fluid and structure solvers as separate entities that exchange information at the fluid-structure interface. For a given structural shape, the fluid solver determines the loads on the structure and exports these to the structural solver. Conversely, the structural solver determines the structural motion and deformation for the given loads and consequently shares the updated motion with the fluid solver. Since the solvers only provide the physical quantities of interest at the interface, they are often referred to as 'black-boxes'. Partitioned methods are also referred to as staggered schemes.

Partitioned methods offer significant advantages from a software engineering perspective. Most importantly, they allow the use of existing CFD and structural solvers without the need for extensive modifications. This allows for each solver to be numerically optimised for its respective physical problem. However, a downside to partitioned approaches is that they may suffer from numerical instabilities, particularly when strong coupling exists between the fluid and structural dynamics. [7]

#### 2.1.4. Time-integration Methods

In terms of time-integration methods for the numerical modelling of FSI problems, partitioned approaches can be further divided based on how the two solvers are coupled at each time step. There are two variables which control the interaction between the solvers: whether the solvers are run in serial or in parallel, and whether the time-stepping of the coupled simulation occurs explicitly or implicitly. In general, the explicit coupling of solvers is referred to as loosely coupled, and the implicit coupling of solvers is referred to as strongly coupled.

##### Loose and Tight Coupling Method

In partitioned FSI simulations, the coupling method determines when information is exchanged between the fluid and structural solvers. In a loosely coupled approach, the solvers exchange information only once per time step, or on a periodic, per-revolution basis in the case of rotorcraft. While this explicit coupling strategy is computationally efficient and simple to implement, it introduces a lag of one time step between solvers. As a result, loosely coupled simulations may become unstable for strongly coupled problems or when large time steps are used.

In contrast, tightly coupled approaches make use of an iterative procedure within each time step. Thus, the fluid and structural solvers repeatedly exchange information until the conditions are satisfied simultaneously at the interface. This approach means the coupled system is implicitly integrated in time. Although this approach significantly improves numerical stability and accuracy, it increases computational cost because multiple sub-iterations are required at each time step [7].

##### Serial and Parallel Schemes

In a serial coupling scheme, the fluid and structural solvers are executed sequentially. An example of a serial scheme is the Conventional Serial Staggered (CSS) scheme. For a time step from  $t_n$  to  $t_{n+1}$ , the procedure consists of the following steps:

1. Update the fluid mesh according to the structural displacement at time  $t_n$ .
2. Solve the fluid equations on the updated mesh to obtain the new flow state  $W_{n+1}$ .
3. Transfer the resulting loads to the structural solver.
4. Advance the structural solution to obtain the updated displacement  $u_{n+1}$ .

The idea behind this serial procedure is that one solver completes its computation before the other solver advances, and the choice can be made as to which solver advances first. This procedure is illustrated in Figure 2.1. Because the fluid and structural solvers are executed in sequence without sub-iterations within the same time level, this scheme corresponds to a loosely coupled serial approach. Since the fluid solver generally requires a smaller time step than the structure, the serial scheme can be adapted with added sub-cycling for the fluid solver, as shown in Figure 2.2. By sub-cycling the fluid solver, the structural model is only solved when necessary.

The serial scheme can be made tightly coupled by adding an iterative procedure within each time step. In this case, the fluid and structural solvers exchange information repeatedly until convergence of the interface conditions is achieved. Denoting the sub-iteration index by  $k$ , the iterative coupling process

can be written as

$$\mathbf{p}_f^{(k)} = \mathcal{F}(\mathbf{u}_s^{(k)}), \quad (2.7)$$

$$\mathbf{u}_s^{(k+1)} = \mathcal{S}(\mathbf{p}_f^{(k)}). \quad (2.8)$$

These sub-iterations continue until the interface residual satisfies a convergence criterion such as

$$|\mathbf{u}_s^{(k+1)} - \mathbf{u}_s^{(k)}| < \epsilon, \quad (2.9)$$

where  $\epsilon$  is a set tolerance. This iterative procedure ensures that both the dynamic and kinematic interface conditions are satisfied at the same time level, improving the numerical stability.




Figure 2.1: Serial staggered scheme. [8]




Figure 2.2: Serial staggered scheme with fluid sub-cycling. [8]

The parallel execution of the solvers can be obtained by altering the procedure as shown in Figure 2.3. The resulting scheme facilitates the parallel exchange of information between solvers by limiting the exchange to one moment in time. Often the parallel execution is improved upon by discretizing in time using the second-order midpoint rule. In this computation, the fluid is solved for half a time step out of sync. These approaches can also become tightly coupled by adding sub-iterations, as mentioned previously, as shown in Figure 2.4.




Figure 2.3: Parallel coupling scheme. [8]




Figure 2.4: Mid-point staggered scheme. [8]

### 2.1.5. Fluid-Structure Interface

Besides the time-integration methods that control when information is exchanged across the interface, how information is exchanged across the interface is equally relevant in FSI simulations. At the interface between the fluid and structural domains, the physical conditions must be satisfied. In partitioned approaches, this further requires matching the fluid and structural meshes at the interface, and matching the deformation of the fluid mesh to the ALE formulation.

#### Consistent and Conservative Transfer

For the transfer of quantities between meshes, it is important that the transfer operation ensures a physically correct coupling. Mathematically, the transfer should be consistent and conservative.

Consistency refers to the ability of the transfer operation to accurately reproduce the transferred field across the interface. For example, the structural displacement must be transferred consistently. For a

given structural displacement, the transfer operation should reproduce the same displacement at the fluid interface without introducing artificial deformation.

Conservation, on the other hand, ensures that the mechanical work exchanged between the structure and the fluid domains is preserved during the transfer. One example of a quantity that must be conserved is the fluid load. If conservation is not satisfied, artificial energy may be added or removed from the system, which may lead to non-physical results.

Hence, many interpolation schemes are designed to satisfy both the consistent and conservative transfer of information across the fluid-structure interface. Common approaches include the weighted residual method and radial basis function (RBF) interpolation. [7]

### Non-matching interface meshes

Typically, the fluid and structural meshes are generated independently, as the structural dynamics can usually be represented by a coarse mesh, whereas the fluid behaviour can only be computed accurately using a fine mesh. Hence, the structural and fluid meshes are non-matching, meaning that there is no one-to-one matching of interface nodes. This situation is illustrated in Figure 2.5.





Figure 2.5: Non-Matching meshes on fluid-structure interface. [7]

Due to this mismatch, the boundary conditions cannot be directly imposed at the corresponding interface nodes. Instead, interpolation or projection methods are required to transfer the information between the two meshes. The structural displacement must be interpolated onto the fluid mesh to update the fluid interface, and the fluid loads must be transferred back to determine the structural response. Commonly, the displacement of the fluid interface can be expressed using structural shape functions. The displacement of a given fluid node is then a weighted combination of the surrounding structural nodes. Conversely, the fluid forces are applied as concentrated loads at the structural elements. This is illustrated in Figure 2.6.

### Mesh Motion

There are many ways of deforming the fluid volume mesh following the consistent transfer of the structural deformation to the fluid interface. A simple approach is to regenerate the fluid mesh after each coupling iteration. Although robust, this approach is rarely used due to the computational costs, continuous need for user interaction, and the lack of automatic conservation of physical quantities. More common are mesh deformation methods where the displacement of the structure is redistributed into the interior of the fluid mesh. Examples of these methods include spring analogy methods, Trans-Finite Interpolation, Laplacian smoothing, and radial basis function interpolation. These techniques allow the mesh to follow the structural motion while preserving mesh quality and avoiding excessive element distortion. [7]



**Figure 2.6:** Transfer of interface quantities between non-matching meshes. Fluid loads are projected onto the structural mesh through interpolation of the structural shape functions [7].

## 2.2. Fundamentals of Rotorcraft

### 2.2.1. Design of Rotorcraft

The broad class of aircraft known as rotorcraft is distinguished from fixed-wing aircraft by their use of large-diameter rotating wings to generate lift. These rotating wings are referred to as rotor blades and are attached to a central hub to form a rotor. A rotorcraft may employ one or more rotors, each of which is made up of at least two rotor blades. Because lift is generated through rotation rather than forward translation, a helicopter is capable of vertical flight, vertical landing, and sustained hover.

The most common rotorcraft configuration is the single main rotor helicopter. To provide lift, propulsion, and control, the rotor of a helicopter rotates about a vertical axis. As the blades rotate, they sweep out the rotor disk surface, and their relative motion with respect to the surrounding air generates aerodynamic forces that support the weight of the aircraft. In addition to the classic helicopter, numerous alternative rotorcraft configurations exist, including tandem rotor helicopters, coaxial rotor systems, tiltrotor aircraft, and compound helicopters. These designs mainly differ in the number and placement of their rotors and how these are used to achieve lift, propulsion, and torque balance. Figure 2.7 illustrates several common rotorcraft configurations.

Rotorcraft operate in four flight regimes. In hover, rotor thrust balances the aircraft weight, and the rotorcraft remains stationary at some point off the ground. The second flight regime is vertical climb, in which additional thrust propels the rotorcraft upward. Third, there is vertical descent, which is more complicated due to the presence of both upward and downward flow around the rotor. Moreover, to achieve forward flight, the rotor disk is tilted in the flight direction to generate a thrust component. Lastly, landing is a combination of forward flight and vertical descent.

### 2.2.2. Rotor and Blade Design

For rotorcraft to achieve vertical flight, the thrust force supplied by the rotor should support the vehicle's weight. Conservation of momentum requires that the rotor thrust is achieved by accelerating air downward, and is accompanied by an equal and opposite reaction of the rotor blades against the air. Kinetic energy is thus imparted on the air in the wake that needs to be supplied by the power source of the rotor in order to sustain level flight. This energy is often denoted as the induced power  $P_i$ , a property of both fixed and rotating wings, and constitutes the absolute minimum power required for equilibrium flight.

The efficiency of the rotor in vertical flight can be expressed in terms of the power loading  $T/P$ , the ratio of rotor thrust to rotor power required, where a higher power loading indicates a higher vertical flight efficiency. Additionally, the performance of the rotor can be expressed in terms of the disk loading  $T/A$ , the ratio of rotor thrust to rotor disk area. For rotorcraft in hover, a low disk loading is critical to achieving a high power loading. This can be explained by the fact that the induced power loading  $P_i/T$  is inversely proportional to the square root of the rotor disk loading. Hence, as the efficiency of rotor thrust generation increases the disk loading decreases. [9]



single main rotor helicopter (UH-1)



single main rotor helicopter (SH-60)



single main rotor helicopter (Bo-105)



tandem helicopter (CH-47D)



compound helicopter (EC X3)



coaxial helicopter (KA-32)



tiltrotor (XV-15)



tiltrotor (V-22)

Figure 2.7: Rotorcraft configurations. [9]

$$P_i/T \propto \sqrt{T/A} \quad (2.10)$$

Furthermore, this implies that induced power is inversely proportional to the radius of the rotor. As a result, helicopters are characterised by large-diameter rotors and have a disk loading in the range of 240–720  $[N/m^2]$ . Within the aerospace industry, the high disk loading rotating wings, such as propellers and turbofan engines, are often smaller and used as propulsive devices. Their operating conditions, high axial velocities and significantly lower thrust to weight, make them more suitable as opposed to large-diameter rotors. [9]

The requirement for a large rotor diameter has several consequences for rotor blade design and rotor dynamics. Most importantly, the large rotor diameter implies long and slender blades with high aspect ratios, which are required to maintain aerodynamic efficiency. In addition, the blades are designed with a geometric twist distribution, which gradually reduces the blade twist toward the tip. This is required to maintain a more uniform spanwise aerodynamic loading since the tangential flow velocity increases linearly with radius, causing the outer sections to experience higher relative velocities.

Finally, the geometry of the blades makes them relatively flexible structures. To withstand the significant centrifugal and aerodynamic forces during flight, rotor blades must therefore be designed with strong internal structures. Modern rotors are therefore commonly manufactured with composite materials to achieve high strength-to-weight ratios. However, this flexibility also allows the blades to respond dynamically to the aerodynamic loads through motions such as flapping, lead-lag and torsion. While this flexibility is beneficial to alleviate aerodynamic loads and maintain equilibrium, it also results in complex aeroelastic interactions between the aerodynamic loads and the structural response of the rotor blades.

### 2.2.3. Rotor Kinematics

The general convention for motion in the rotor disk is illustrated in Figure 2.8. The positive direction of rotation is defined as counterclockwise when viewed from above. The position of a blade in the disk is specified by the azimuth angle  $\psi$ , measured in the plane of rotation and increasing in the direction of rotation. In forward flight, the freestream velocity is aligned with the azimuth  $\psi = 0^\circ$ . The right side of the rotor disk is referred to as the advancing side, where the rotational velocity of the blade adds to the freestream velocity. The left side is called the retreating side, where the rotational velocity subtracts from the freestream velocity. This velocity asymmetry across the rotor disk plays a fundamental role in rotor aerodynamics.




Figure 2.8: Rotor convention in the rotor disk plane. [9]




Figure 2.9: Rotor blade degrees of freedom.

In addition to the general rotor motion, rotor blades exhibit three primary motions, illustrated in Figure 2.9. The first is blade flapping, which is the out-of-plane deflection of the blade about a flap hinge, or, in

the absence of a hinge, the bending of the blade at the root and along the span. This motion produces a vertical displacement of the blade relative to the rotor disk. The second motion is lead-lag, which is the in-plane motion about a lag hinge or through in-plane bending of the blade. This motion results in the fore-aft movement of the blade within the rotor disk. The third motion is blade feathering, which is the rotation of the blade about the blade's spanwise axis. Feathering changes the aerodynamic angle of attack with respect to the airflow and provides the primary means of controlling rotor thrust and moments.

The extent to which these motions are possible depends on the design of the rotor hub. Since rotor blades are typically long and slender, significant bending moments can occur at the blade root. Therefore, several design solutions exist to reduce the loads transmitted to the hub. Allowing the blades to flap freely, for example, reduces bending loads by enabling the blade to align more closely with the aerodynamic forces acting on it. In rotors where the blades are sufficiently aeroelastically flexible, these load-relieving motions can occur through structural deformation rather than through mechanical hinges.

Several hub configurations have been developed to accommodate these blade motions. One of the earliest designs is the articulated rotor, in which each blade is connected to the hub through both a flapping hinge and a lead-lag hinge. These hinges allow the blade to move relatively freely, significantly reducing bending loads at the root. Another simplified configuration is the teetering rotor, which is typically used on two-bladed helicopters. In this design, the blades are connected through the hub and share a single flapping hinge, allowing the rotor to teeter like a seesaw. While this configuration removes the need for individual flapping hinges, lead-lag motion is generally not accommodated in teetering rotors.

Modern rotor systems often employ hingeless or bearingless hubs. In hingeless rotors, the blades are attached rigidly to the hub, and the flapping and lead-lag motions occur through elastic bending of the blade root and along the blade span. In bearingless rotors, this concept is extended further by eliminating mechanical pitch bearings as well. Instead, flexible composite structures accommodate the required blade motions. Although these designs reduce the mechanical complexity and maintenance requirements of the rotor system, they generally transmit higher loads to the hub structure.


Regardless of the specific design, the fundamental purpose of the rotor hub is to accommodate blade motion in a way that reduces the structural loads while allowing the rotor to generate the forces and moments required for flight and control.

#### **2.2.4. Rotor Control and Blade Pitch Variation**

Unlike fixed-wing aircraft, where lift, propulsion, and control are provided by separate systems, a helicopter rotor must perform all three functions simultaneously. The rotor therefore requires a mechanism that allows the aerodynamic loading of the blades to be adjusted continuously during operation. Control is achieved through feathering motion, which rotates the blade about its spanwise axis, thus resulting in a change in angle of attack. By varying the blade pitch, the rotor can generate the forces and moments required to support the aircraft weight, provide propulsion, and control the attitude of the rotorcraft.

Two primary types of pitch control are employed in helicopter rotors: collective pitch and cyclic pitch. Collective pitch refers to a simultaneous and uniform change in the pitch angle of all blades, independently of their azimuth position. Increasing the collective pitch increases the blade pitch angle over the entire rotor disk and results in an increase in the total rotor thrust. Since the collective pitch affects all azimuthal positions, this control input is used primarily to regulate the lift produced by the rotor and to control the vertical motion of the aircraft, such as climb or descent.

While the collective pitch alters the overall thrust of the rotor, it does not provide directional control. In forward flight, the rotation of the blade means that the inflow conditions experienced by the rotor are inherently asymmetric, illustrated in Figure 2.10. On the advancing side of the rotor disk, the blade moves into the flow and the relative velocity experienced by the blade is the sum of the rotational velocity and the forward flight velocity. Conversely, on the retreating side of the rotor disk, these two velocities are aligned and subtract. Near the root of the retreating blade, this can lead to a region of reversed flow.



**Figure 2.10:** (a) Force balance on a helicopter during forward flight (b) dissymmetry of the corresponding in-plane velocity distribution. [10]

Since the lift generated by a blade section is proportional to the square of the relative velocity, the advancing blade would produce significantly more lift than the retreating blade if the blade pitch were constant throughout the rotation. This phenomenon is commonly referred to as the dissymmetry of lift of rotors. The result of this imbalance would be a rolling moment on the rotor that would make stable forward flight impossible.

To overcome this inherent aerodynamic asymmetry, the pitch of the rotor blades is varied periodically during each revolution and is known as cyclic pitch. Through cyclic pitch control, the angle of attack of the blades is reduced on the advancing side of the rotor disk and increased on the retreating side. In this way, the lift distribution over the rotor disk can be balanced and large rolling moments can be prevented, allowing the rotor to remain in equilibrium.

In general, the blade pitch angle can be expressed as a periodic function of the azimuth angle  $\psi$ ,

$$\theta(\psi) = \theta_0 + \theta_{1c} \cos \psi + \theta_{1s} \sin \psi, \quad (2.11)$$

where  $\theta_0$  represents the collective pitch, and  $\theta_{1c}$  and  $\theta_{1s}$  represent the cyclic pitch components. These cyclic terms correspond to the lateral and longitudinal pitch angles of the rotor blade, and allow the rotor thrust vector to be tilted in the desired direction. The combination of collective and cyclic pitch allows the rotor to be trimmed to a state in which the aerodynamic forces and moments acting on the rotorcraft are balanced.

## 2.3. Fundamentals of Rotor Aeromechanics

This section provides a description of the rotor from an aeromechanical perspective. The rotor is treated as a problem in which aerodynamic forces, structural mechanics, stability, flight dynamics, and noise interact to determine the behaviour of the rotor.

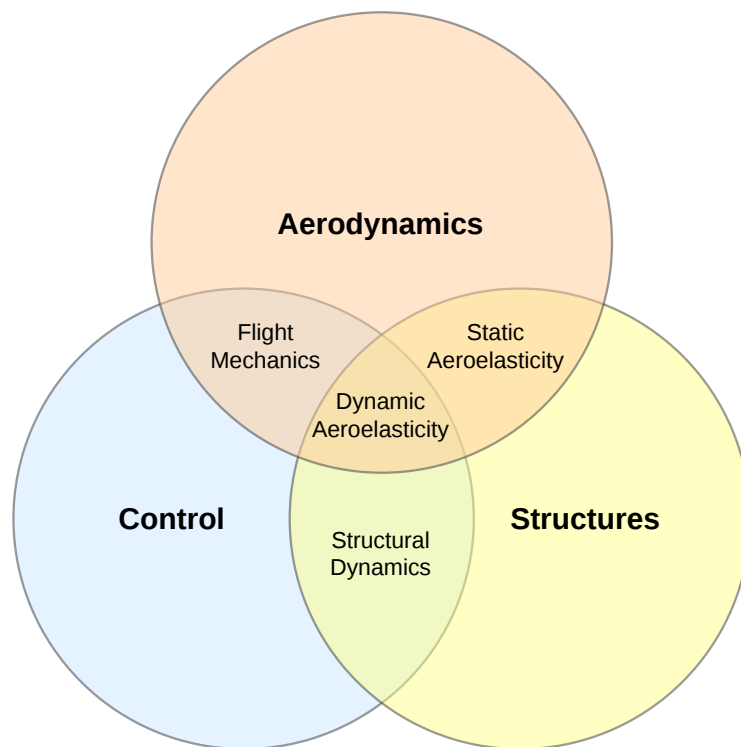
### 2.3.1. The Field of Rotor Aeromechanics

From a physical perspective, the rotor consists of slender blades attached to a central hub that rotates about the rotor shaft. As the blades rotate, aerodynamic forces are generated along the span of each blade. These loads cause the blades to deform and move relative to the rotor disk, producing structural reaction forces internally and at the blade root. The combined effect of these forces determines the loads at the hub and aircraft that ultimately govern the rotor thrust and torque.

The analysis of this system therefore requires the combined consideration of several disciplines. Aerodynamics determines the airloads on the blade sections, structural mechanics describes the

structural response to these loads, and rotorcraft flight mechanics relates the forces and moments to the motion and control of the rotorcraft. The analysis of the rotor inherently couples these disciplines since changes in one domain directly influence the others.

The multidisciplinary field of rotor aeromechanics is illustrated in Figure 2.11. The interaction between aerodynamics and rotor control is the basis of rotorcraft flight mechanics, whereas the interaction between rotor control and structural mechanics describes the dynamic response of the rotor system. Moreover, the coupling between aerodynamics and structural dynamics leads to static aeroelastic behaviour of the rotor blades. Finally, when all three domains interact simultaneously, the resulting behaviour is referred to as dynamic aeroelasticity, which governs the time-dependent response of the rotor system and the loads transmitted to the aircraft structure [1].



**Figure 2.11:** Conceptual representation of rotor aeromechanics as the interaction between aerodynamics, structural dynamics, and rotor control. Adapted from Smith [11].

In this framework, the field of rotor aeromechanics is the foundation for the prediction of blade loads, structural response, and the total force and moment balance of the rotor. These quantities and more are central to the analysis of rotor performance, loads, stability, and flight dynamics, and therefore play a critical role in the design and analysis of rotorcraft systems.

### 2.3.2. Blade Section Aerodynamics

To start describing the behaviour of a rotor, the aerodynamic forces acting on the rotor can be analysed at the blade section level. For this analysis, each blade section can be treated as a two-dimensional airfoil, where the local flow field is a combination of the rotational velocity, forward flight velocity, and induced velocity. Figure 2.12 shows the aerodynamic forces at the rotor section for a general local flow field.

In the rotating hub plane reference frame, the velocity components are split into the tangential  $u_T$ , perpendicular  $u_P$ , and radial  $u_R$  components. The in-plane resultant velocity  $U$ , inflow angle  $\phi$ , and angle-of-attack  $\alpha$  are defined as,

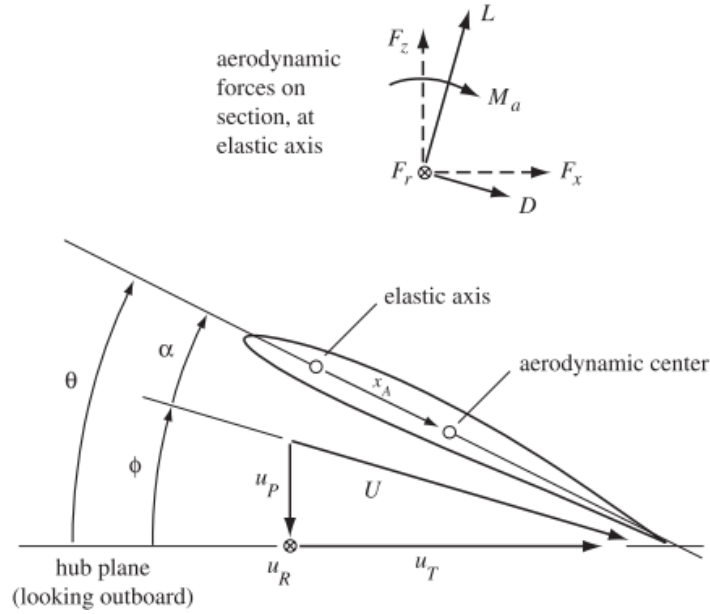


Figure 2.12: Rotor blade section aerodynamics. [9]

$$U = \sqrt{u_T^2 + u_P^2} \quad (2.12) \quad \phi = \tan^{-1}\left(\frac{u_P}{u_T}\right) \quad (2.13) \quad \alpha = \theta - \phi \quad (2.14)$$

where  $\theta$  includes the collective and cyclic pitch as well as the inherent blade twist. The aerodynamic forces generated by the blade section are the lift and drag forces  $L$  and  $D$ , given by

$$L = \frac{1}{2}\rho U^2 c c_l \quad (2.15) \quad D = \frac{1}{2}\rho U^2 c c_d \quad (2.16)$$

where  $\rho$  is the air density,  $c$  is the blade chord, and  $c_l(\alpha, M)$  and  $c_d(\alpha, M)$  are the sectional lift and drag coefficients, dependent on angle of attack and Mach number.

To find the contributions of the loads in the hub plane, the components of the sectional lift and drag can be resolved into the hub-plane axes  $F_z$  and  $F_x$ . In addition, the blade also produces a radial force  $F_r$  that consists of a radial drag force and an in-plane component of the blade lift due to the flapped orientation of the blade. Lastly, the sectional aerodynamic pitching moment  $M_a$  is included, which is positive for nose-up pitching. To simplify the effects of the loading on blade deformations, the forces and moment are represented about the elastic axis, the axis about which the blade bends and twists.

### 2.3.3. Forces acting on the Rotor Blade

Extending the section aerodynamics to the entire rotor blade shows that the rotor blades also experience inertial, centrifugal and Coriolis forces in addition to the aerodynamic loading. A side view and a top view of a rotor blade during rotation are shown in Figures 2.13 and 2.14.

In the rotating reference frame (Figure 2.13), the vertical motion of the blade is dependent on the equilibrium between the aerodynamic force, the centrifugal force, and the inertial force. The aerodynamic force generated by the blade section is equal to the section lift and is generally much larger than the other forces as it is required to balance the aircraft weight. This force tends to cause out-of-plane blade deformation (flapping), shown here by the flap angle  $\beta$ . Additionally, the blade element experiences a centrifugal force that acts radially outward. As the rotor spins, every mass element of the blade tends to move in a straight line, while the rotation constrains it to a circular path. The centrifugal force thus stiffens the blade and opposes blade flapping. In the case of blade flapping, the blade element also experiences an inertial force associated with its vertical acceleration away from its previous position, which further opposes the flap motion.




Figure 2.13: Rotor blade side view. Adapted from [9].




Figure 16.3. Rotor blade lagging moments.

Figure 2.14: Rotor blade top view. Adapted from [9]

In the rotor disk plane (Figure 2.14), the blade is subject to aerodynamic drag as well as the radial centrifugal force. Where the aerodynamic drag can cause the blade to lead or lag its position, the centrifugal force restores this motion by stiffening the blade. When in-plane deformation (lead-lag motion) occurs, additional inertial forces arise that oppose this motion. Furthermore, Coriolis forces may develop due to the coupling between the lead-lag motion and rotation. Together, these forces balance the in-plane motion of the rotor blade.

Lastly, as a result of the forces, each blade element is also subject to moments. The most influential is the pitching moment, which acts around the radial blade axis, and is responsible for the torsional deformation of the blade section, leading to an altered blade pitch. The two other moments, often smaller in magnitude, lead to in-plane and out-of-plane motion of the blade.

### 2.3.4. Structural Coupling

To analyse the behaviour of the rotor from an aeromechanical standpoint, the rotor is modelled as a rotating, flexible structure subjected to the aerodynamic and inertial loading. Additionally, control is applied through collective and cyclic pitch at the hub. A hingeless rotor configuration is used as the reference model. In this configuration, the blades are rigidly attached to the hub without discrete flapping or lag hinges. Hence, all blade motion is accommodated through elastic deformation of the blade structure.

The rotor blade is modelled as a rotating cantilever beam. Its motion can be described in terms of the fundamental deformation modes about the elastic axis, illustrated in Figure 2.15. These are out-of-plane bending (flap), in-plane bending (lead-lag), and torsion. These motions arise as structural responses to the loads acting on the blade and affect the inflow conditions of the rotor blade, leading to the coupled interaction between aerodynamics, structure and control.

Flap motion is primarily driven by the aerodynamic lift along the blade span and produces out-of-plane bending of the blade. Conversely, the centrifugal force generated by the rotation reduces blade flapping, giving the blade an increased effective stiffness and limiting the amplitude of the motion. The local effect of the flapping motion is a change in the orientation of the blade relative to the flow, which modifies the local inflow angle and angle-of-attack. As a result, the blade lift and drag are modified, leading to both a blade flap and lag response. Flapping and lag are further coupled as flapping also introduces radial motion of the blade elements, which leads to in-plane Coriolis forces that cause the blade to lead-lag. The influence of flapping is especially pronounced in forward flight, where the azimuthal variation in lift due to the cyclic pitch also leads to azimuthally varying blade flapping and aids in altering the orientation of the rotor thrust vector.

Lead-lag motion of the blade occurs in the plane of rotation and is produced by aerodynamic drag, centrifugal forces, and Coriolis forces. Comparable to blade flapping, lead-lag motion alters the local inflow conditions and leads to increased in-plane hub forces, changes in relative phase of the blades, rotor vibration, and rotor stability. In addition, blade flap and lag motion are coupled through the aerodynamic loading: increased flapping alters lift and drag, which may increase blade lag, whereas blade lag affects the inflow and velocity distribution, leading to the flap response.




Figure 2.15: Elastic modal blade deformations: flap  $u$ , torsion  $\theta_t$  and lag  $w$ . [12]

Torsional blade deformation arises from the aerodynamic pitching moment and from structural moments generated by blade bending. The elastic axis is generally offset from the aerodynamic centre. As a result, bending deformation produces torsional moments and blade pitching is often coupled to flap and lag motions. Torsional deformation modifies the blade pitch and thereby the aerodynamic loading on the blade, strongly influencing both flap and lag motion further.

Examples of further influences are the control states, which control blade feathering. The collective pitch changes the blade pitch across the entire revolution, leading to increased blade deformation, coning, and rotor thrust. The effect of the cyclic pitch is similar, introducing periodic variations in loading and motion and contributing to the dynamic response of the blade. Due to the dynamic response of the blade, the flapping motion lags in phase relative to the aerodynamic loading.

The behaviour of the rotor can therefore only be understood by considering the coupled interaction between the loading, structural response, and control inputs. An overview of the main coupling mechanisms, in terms of cause, response, and resulting effect, is summarised in Table 2.1.

## 2.4. Fundamentals of Rotor Aerodynamics

Unlike fixed-wing aircraft, the aerodynamics of rotorcraft are characterised by unsteady and three-dimensional vortical flows. Due to the rotation of the rotor, each blade moves through regions of varying velocity and inflow conditions, and may interact with previous blade wakes. Consequently, these variations cause the aerodynamic loads to vary.

An overview of the aerodynamic phenomena that occur during rotorcraft flight, and tend to cause difficulties such as noise, vibrations and performance limitations, is illustrated in figure 2.16. These include dynamic stall on the retreating blade, transonic flow on the advancing blade, blade-vortex interaction (BVI) generated by the rotor wake and vibrations and acoustic noise.

### 2.4.1. Wake Structure and Induced Flow

Most of the aerodynamic phenomena illustrated in figure 2.16 arise from the vortical wake that is created by the rotor blades as they rotate and translate through the air.

As each blade rotates, the shed vorticity from the rotor blades rolls up into discrete concentrated tip vortices, which form the dominant structures in the wake as shown in figure 2.17. However, the wake is not merely a downstream feature of the rotor flow. Instead, it directly influences the aerodynamic loading through the induced velocity field. While the concentrated tip vortices temporarily remain in the trace of the blade tips, they induce local variations in airflow for the incoming rotor blades. Since the blade loading is dependent on the local inflow angle, and the local inflow angle depends on the wake geometry, this creates a nonlinear feedback mechanism between the blade aerodynamic and the wake

**Table 2.1:** Overview of rotor blade structural coupling mechanisms (cause → response → effect). Adapted from Johnson [9].

Cause	Rotor/ Structural Response	Effect / Coupling
Aerodynamic lift variation	Flap motion (out-of-plane bending)	Redistribution of lift, tilt of the thrust vector, and bending-induced torsional moments
Radial motion during flapping	Coriolis forces (in-plane inertial forces)	Induces lead-lag motion (flap → lag coupling)
Inertial forces, Coriolis forces, and aerodynamic drag	Lead-lag motion (in-plane bending)	Generates in-plane hub forces, modifies blade phase, and contributes to vibration and stability
Blade bending (flap and lag)	Torsional deformation (twist)	Produces pitch-flap and pitch-lag coupling through bending-induced moments
Aerodynamic pitching moment	Torsion	Modifies blade pitch and aerodynamic loading (torsion → flap/lag coupling)
Centrifugal forces (rotation)	Axial tension and structural stiffening	Increases natural frequencies and limits deformation
Collective pitch input	Uniform load increase	Increases coning, torsion, and overall structural loading
Cyclic pitch input	Periodic lift variation	Drives 1/rev flapping response and dynamic blade motion




Figure 2 A summary of specific flow problems which occur on a helicopter. From Caradonna (1992).

Figure 2.16: Summary of aerodynamic phenomena occurring in the flow field of a helicopter rotor.[13]

evolution. Additionally, these vortices may interact with themselves and collide with rotor blades or other components of rotorcraft before evolving into a wake system downstream of the rotor increasing the complexity of the rotor aerodynamics further. [9], [13]

In hover, the vortices form a spiral pattern behind the blade tips and beneath the rotor disk as illustrated in figure 2.18. Since the rotational velocity of the rotor is much higher than the convection of the vortices away from the rotor, the tip vortices interact with the following blades. The induced velocity generated




Figure 2.17: Rotor wake formation. [14]




Figure 2.18: Schematic rotor wake geometries at different flight velocities. [14]

by these vortices alter the local inflow angle at each affected blade section and therefore the aerodynamic load along the blade. Because the rotor operates within its own wake, the inflow conditions encountered by the blades depend greatly on the strength, position and evolution of the wake.

In forward flight, the wake is no longer axisymmetric as the freestream velocity causes the wake to be directed downstream. The tip vortices are then convected both downward and rearward, which leads to the distorted wake as shown in figure 2.18. Although the wake is convected away from the path of the rotor blades more rapidly than in hover, the vortices still influence the blade loading during following revolutions.

### 2.4.2. Transonic flow

In forward flight, the flow velocity relative to a rotor blade is the vector sum of the blade rotational velocity and the forward velocity of the aircraft. On the advancing side of the rotor disk these velocities add, leading to higher local velocities, especially near the blade tip. At sufficiently high forward speeds, this may cause the local Mach number to approach or exceed the critical Mach number and lead to

transonic flow conditions.

Under these conditions, shock waves may form on the blade surface, causing shock-boundary layer interaction, local flow separation and strong pressure gradients. These compressibility effects modify the pressure distribution over the rotor blade, increase the aerodynamic drag and introduce oscillations in the pitching moment. Because the rotor blades periodically pass through regions of different relative velocity during each revolution, the position and strength of the shock is dependent on the blade azimuth. As a result, the aerodynamic loads often are highly unsteady, which leads to increased rotor vibration and noise. Transonic flow on the advancing blade therefore is one of the primary limitations on the forward speed of helicopters. [2], [15]

### 2.4.3. Blade-Vortex Interaction

Blade-vortex interaction (BVI) occurs when a rotor blade encounters a vortex previously shed from the tip of another blade. As each blade generates lift, strong tip vortices are formed that convect downstream to form the rotor wake. Due to the high rotational velocity of the blades, the rotor blades repeatedly pass through this wake, and the vortices may pass close to or directly collide with a following blade.

When a blade encounters a vortex, the vortex induces a local flow field that alters the effective inflow velocity and angle of attack experienced by the blade sections. Since the tip vortices are strong and the distance between the blade and vortex is small, the induced velocities can be large. This leads to localized and rapid variations in the aerodynamic loading along the span. Typically, this aerodynamic disturbance appears as a sharp change in lift and moment, leading to impulsive loads on the rotor blade structure. The loads propagate through the rotor hub and into the airframe, causing increased rotor vibration and structural loading. [9]

The flow field induced by a vortex leads to a characteristic loading distribution along the blade. Near the vortex, regions of increased and decreased local lift are produced as a result of the induced vertical velocity across the blade span. The resulting lift distribution typically shows two peaks on either side of the vortex core, corresponding to the largest induced upwash and downwash. Figure 2.19 illustrates the induced loading distribution following the interaction between a rotor blade and the vortex.




Figure 2.19: Vortex-induced blade loading. [9]

The characteristics of BVI depend on the relative orientation between the vortex and the blade. The limiting cases are: almost perpendicular interactions, where the vortex approaches the blade span approximately perpendicular, and almost parallel interactions, where the vortex path is aligned with the blade span. The perpendicular case is characteristic of the blade interaction with the tip vortex from the preceding blade, on the advancing and retreating sides of the rotor disk in forward flight, and leads to vibration and structural loads. Due to the orientation of the vortex and blade, this encounter is a steady, three-dimensional aerodynamic problem. On the other hand, the parallel case is characteristic of the blade interaction with vortices that are one or more revolutions old, and often occurs in the first quadrant of the rotor disk ( $0^\circ - 90^\circ$ ). This encounter is an unsteady, two-dimensional problem and is a factor in BVI noise generation. [9]


Overall, BVI is a major source of aerodynamic disturbances in rotorcraft as these interactions lead

to localized blade loads, rotor vibrations and rotor noise. Therefore, BVI must be considered in the aerodynamic and structural analysis of rotors.

#### 2.4.4. Dynamic Stall

Dynamic stall is an unsteady aerodynamic phenomenon that can occur when rotor blades experience rapid increases in angle of attack. In helicopters, this phenomenon is common on the retreating side of the rotor disk in forward flight, where the local relative velocity is reduced and the blade must operate at higher angles of attack in order to maintain lift. Because the angle of attack changes rapidly, the boundary layer does not separate at the static stall limit. Instead, flow separation is delayed and the section lift continues to increase. [9]

The sequence of flow events during the evolution of dynamic stall is illustrated in Figure 2.20. The process begins with the formation of a separation bubble near the leading edge of the airfoil. As the angle of attack increases, the boundary layer near the leading edge separates and rolls into a leading-edge vortex that develops on the blade surface. While this vortex remains attached to the rotor blade, it temporarily increases the sectional lift and moment coefficients beyond the static stall values. However, the vortex eventually convects downstream and detaches from the airfoil leading to a breakdown of the flow. This results in a collapse in lift, a sharp increase in drag and a large negative pitching moment as the center of pressure shifts aft. [9]



**Figure 2.20:** Schematic representation of the different stages of the dynamic stall process: (a) attached flow, (b) flow reversal, (c) shear layer roll-up, (d) stall onset, (e) full stall, and (f) flow reattachment. [16]

For the rotor the consequences of dynamic stall are related to the formation and shedding of the dynamic stall vortex, which produces large transient variations in lift, drag and pitching moment on the blade sections. These rapidly changing aerodynamic loads propagate through the blade structure and into the rotor hub, increasing rotor vibration and structural loading. In rotorcraft, dynamic stall often occurs cyclically on the retreating side of the rotor disk, leading to periodic fluctuations in aerodynamic loading. As a result, dynamic stall can limit rotor performance, lead to increased acoustic noise and increase the loading on the rotor blades.

## 2.5. Rotorcraft prediction Methodologies

### 2.5.1. History of Rotorcraft Prediction Methodology

Rotorcraft prediction methodologies are organized the following three strategies: : Comprehensive Analysis (CA), Computational Fluid Dynamics (CFD), and a CFD-CA coupling that integrates high-fidelity aerodynamics with structural dynamics. This section discusses how the emergence of each

strategy was driven by the limitations and technical challenges encountered in rotorcraft development.

### **Early Foundation**

The foundation of rotorcraft prediction was Glauert's work on autogyros and early helicopter rotors in the 1930's. These early prediction methods combined blade element theory for sectional lift and drag, momentum theory for induced flow effects and linear beam theory to include the blade flapping dynamics. In addition, they represented the effects of the wake using prescribed geometries or simplified inflow models. Critically, this early work established the physical decomposition of the rotor aerodynamics - the relation between the local blade section loading and the global induced flow - that remains fundamental to rotorcraft analysis.[1]

However, a lack of a consistent coupling between the aerodynamics, structural deformation and flight dynamics limited the capabilities of these early methods. Typically, aerodynamics, structural dynamics and rotor control were solved independently. For instance, these approaches considered the blade motion independently of the unsteady aerodynamics and solved the wake by imposing or approximating it rather than as a dynamic system. While this separation was adequate for hover and low-speed forward flight, these methods proved increasingly inaccurate for flight conditions characterized by high advance ratio, strong blade-vortex interaction (BVI), transonic flow, and dynamic stall. Moreover, as rotor systems became larger, stiffer and operated at higher speeds the interdependence between the aerodynamic loads and the structural deformation and the lack of a reliable prediction framework become more apparent.[1]

### **Rise of Comprehensive Analyses**

A major conceptual change occurred in the 1960s and 1970s with the development of Comprehensive Analysis (CA), an integrated multidisciplinary rotorcraft prediction methodology. Previous technical challenges encountered in rotorcraft development programs, most notably for the AH-56A Cheyenne, demonstrated that the behaviour of rotorcraft could not be reliably predicted using separated physical models. As discussed by Ormiston [4], the Cheyenne program highlighted the direct link between the deficiencies in prediction capabilities and the risks associated during development. During this project, inaccurate prediction of aeroelastic led to stability problems, conservative design compromises, increased development cost, and ultimately cancellation of the production program. As a consequence, the need for a unified prediction framework for the design and development of rotorcraft was recognized.

The ability to predict the coupled aeromechanic behaviour of rotors is essential for effective research and development. First, virtually all aspect of rotorcraft performance are influenced by aeromechanics, including efficiency, aeroelastic stability, vibrations and handling characteristics. Second, in the absence of reliable prediction methodologies designers are forced to adopt conservative designs, compromising on performance and innovative concepts. Third, inaccurate prediction methods can result in increased development costs, extended project timelines and increased risks during the development process. For these reasons, the research and development of prediction methodologies was established as a primary research subject within the rotorcraft aeromechanics community. [4]

Comprehensive Analysis addressed the earlier shortcoming by utilising lifting-line aerodynamics together with nonlinear beam structural dynamics that was capable of representing the flap, lag, and torsional degrees of freedom of the rotor blade using finite-element formulations. To incorporate flight dynamics and the control system, the aerodynamic and structural models were coupled through trim algorithms that simultaneously solved for the blade motion and rotor control inputs. The resulting unified computational environment fundamentally transformed rotorcraft design and development by allowing consistent prediction of rotor performance, sectional airloads, blade structural loads, and aeroelastic stability. An overview of comprehensive analysis and its developers is illustrated in figure 2.21.[17].

Despite its improved prediction capabilities, CA still relied on low-fidelity aerodynamic models. As rotorcraft approached higher speed and higher load flight conditions, the aerodynamic deficiencies of CA became apparent. In particular, conventional CA methods introduced errors in the sectional pitching moment phase and magnitude near the blade tip under transonic conditions. The incorrect pitching moment lead to erroneous elastic blade twist, thereby altering the blade lift and the wake development. Further analysis confirmed that the structural modelling within CA was largely correct, the limitation remained the aerodynamic modelling. [18]

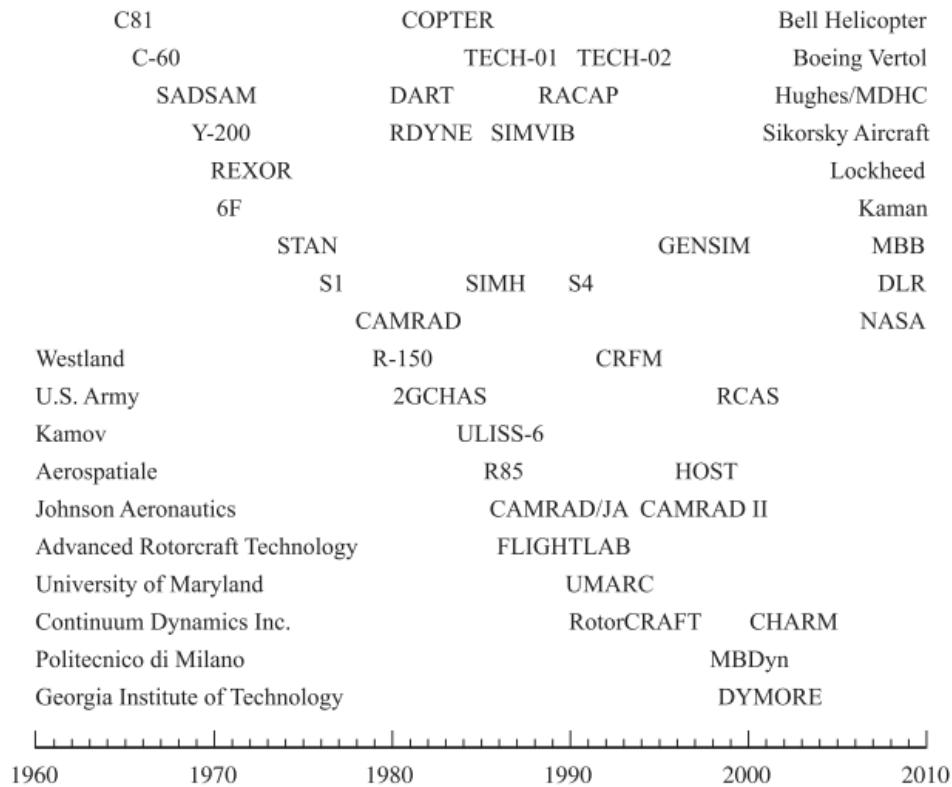


Figure 2.21: Historical development of comprehensive analysis systems [1]

### Rise of Computational Fluid Dynamics (CFD)

Computational Fluid Dynamics (CFD) was introduced to overcome the aerodynamic limitation of the low-fidelity aerodynamics of CA by directly solving the three-dimensional, viscous, compressible flow equations. In contrast to lifting-line methods, CFD is capable of resolving nonlinear phenomena such as blade-vortex-interactions, dynamic stall and transonic flow. The correct modelling of these effects is critical for accurate prediction of section airloads.

However, CFD alone is unable to account for the structural dynamics or control. For the rotor to achieve equilibrium, the consistent coupling between the aerodynamic loads, structural deformation and control inputs is required. While tightly-coupled solver frameworks, where information between CFD and structural solvers are exchanged at every time-step, were widely regarded as the more accurate approach, hardware limitations forced the development of methodologies that were computationally cheaper and more practical for routine application. This limitation motivated the development of loose CFD-CA coupling as an extension of Comprehensive Analysis.

### Loose Coupling Breakthrough

Initial attempts to integrate CFD into rotor prediction methodologies focused on directly replacing the internal aerodynamic model within CA, hereby implicitly treating the problem as solely an aerodynamic problem. In early coupling attempts, this approach destabilized the trim process and lead to convergence problems due to the fact that achieving rotor equilibrium required the simultaneous satisfaction of trim targets, structural deformation and control inputs. Forcing the CFD solution directly into the nonlinear trim loop proved numerically fragile and computationally too expensive. [4].

The breakthrough came following the analysis of the NASA/ Army UH-60A Airloads Flight Test Program during the Airloads workshop[19]. Datta et al. [18] showed that CFD could accurately resolve the essential unsteady aerodynamic physics when removing the 1/rev airloads, and thus that the difficulty primarily concerned the prediction of the steady 1/rev blade motion. In other words, the main difficulty was achieving a consistent aeroelastic equilibrium between aerodynamic loads, blade deformation, and control inputs otherwise known as the rotor trim problem.

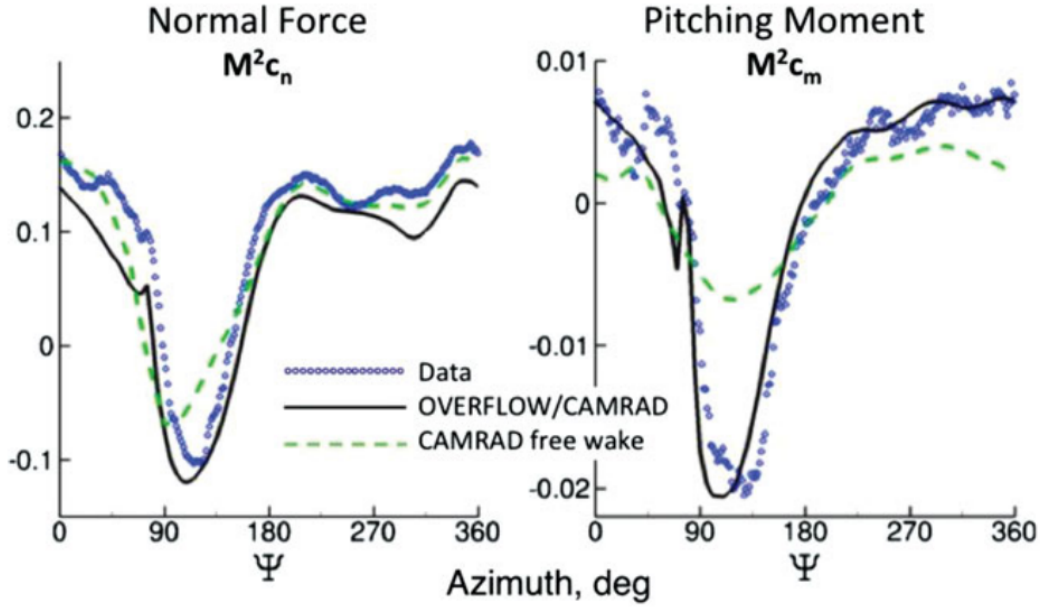


Figure 2.22: Airloads Workshop CFD-CA loose coupling breakthrough, UH-60A high-speed airloads,  $r/D = 0.965$ ,  $\mu = 0.368$ . [4]

A reformulation of the coupling strategy proved to be the solution and was met with immediate success. Instead of replacing the internal aerodynamic model of CA with the CFD solution, Potsdam et al. proposed to introduce the CFD solution as an incremental correction to the aerodynamic loading. [3]. This approach retains the CA trim framework while utilising the higher fidelity CFD. The close agreement of CFD-CA normal force and pitching moment airloads compared to measured data is shown in figure 2.22

This approach provided the first practical and robust realization of loose CFD-CA coupling for trimmed steady rotor flight. It demonstrated that high-fidelity aerodynamics and efficient trim enforcement could coexist within a partitioned framework. The method reframed rotor aeroelastic coupling as a stabilized correction problem and established the foundation for modern partitioned fluid-structure interaction formulations.

### 2.5.2. Formulation of the Delta-Airloads Method

The loose coupling methodology developed by Potsdam et al. is used to couple a CA code to a CFD solver on a periodic, per-revolution basis [3]. Rather than embedding the CFD loading directly within the trim loop, the method utilises an incremental correction strategy, commonly referred to as the delta-airloads method. In this approach, the aerodynamic loads used in the CA at coupling iteration  $i$  are written as

$$\vec{F}_i = \vec{F}_i^{LL} + \left( \vec{F}_{i-1}^{CFD} - \vec{F}_{i-1}^{LL} \right), \quad (2.17)$$

where  $\vec{F}^{LL}$  denotes the lifting-line solution from the CA and  $\vec{F}^{CFD}$  denotes the sectional airloads computed by CFD. The correction term is therefore the difference between the CFD and lifting-line solution from the previous iteration. The formulation can also be rewritten as

$$\vec{F}_i = \vec{F}_{i-1}^{CFD} + \left( \vec{F}_i^{LL} - \vec{F}_{i-1}^{LL} \right), \quad (2.18)$$

thus illustrating that the CA determines only the loading required to retrim the rotor, while CFD supplies the bulk of the aerodynamic loading.

The coupling workflow is best described via Figure 2.23. First, a trimmed rotor solution is obtained using the lifting-line aerodynamics in the CA without CFD data. The resulting blade motion is transferred to the CFD solver, which then computes the aerodynamic loading data  $F_0^{CFD}$  and  $M_0^{CFD}$ . In the case of

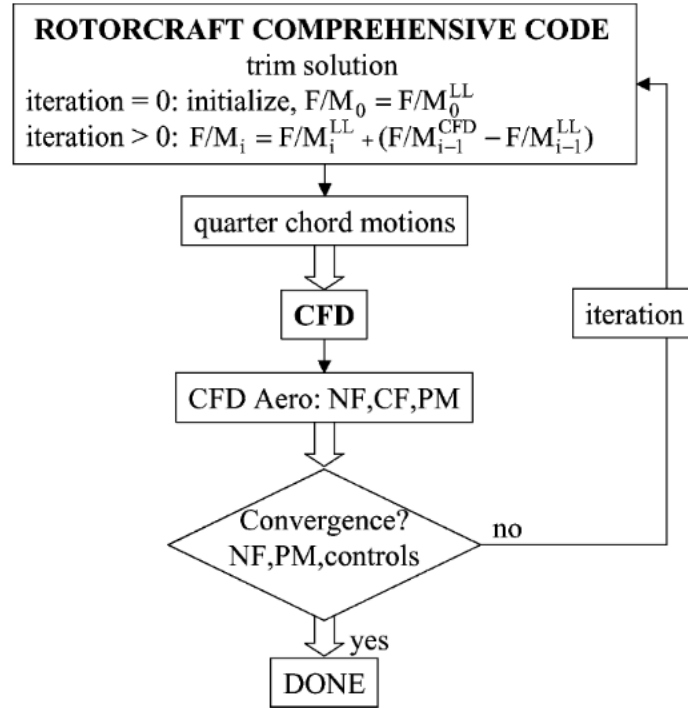


Figure 2.23: CFD-CA loose coupling procedure, F/M= forces and moments.[3]

Potsdam's formulation, these include the sectional normal force, pitching moment, and chord force, and only include the pressure components of the force. These CFD airloads are post-processed, used to determine the delta-airloads correction and the rotor is subsequently retrimmed using the updated aerodynamic loads. This process is repeated until the control angles and aerodynamic loads converge between successive coupling iterations.

Upon convergence, the lifting-line loads no longer change between iterations, such that

$$\vec{F}_i^{LL} = \vec{F}_{i-1}^{LL}, \quad (2.19)$$

meaning that the correction term vanishes. The total aerodynamic loads used in the CA therefore reduce to

$$\vec{F}_i = \vec{F}^{CFD}, \quad (2.20)$$

indicating that the final aeroelastic equilibrium is fully CFD-consistent while satisfying all trim constraints within the comprehensive framework.

To ensure that the CFD solutions are consistent between iterations the CFD must be run for sufficient rotor revolutions between coupling iterations to eliminate aerodynamic transients and establish periodicity. Failure to establish a consistent and periodic solution leads to unstable or non-convergent coupling solutions.

## 2.6. Review of Rotor Load Prediction Using partitioned CFD-CA Coupling

Following the historical development of rotorcraft prediction methodologies and the formulation of the delta-airloads methods, the application of loosely coupled CFD-CA methodologies has been investigated across literature. This section reviews these applications focussing on how partitioned coupling frameworks have improved rotor load prediction, the modelling choices commonly employed, and the limitations that persist across different studies. Lastly, the findings are summarized to outline what considerations are critical for the development of a loosely coupled CFD-CA methodology.

### 2.6.1. Applications of partitioned CFD-CA Coupling Using the Delta-Airloads Method

The first practical demonstration of the loose coupling formulation was achieved by Potsdam et al. using the delta-airloads formulation [3]. In that study, the OVERFLOW Navier-Stokes solver was coupled with CAMRAD-II to predict UH-60A rotor airloads in forward flight. The study showed that stable, trimmed solutions could be obtained for forward-flight conditions, with improved sectional normal force and pitching moment compared to stand-alone CA. However, it also showed that the convergence of the coupled solution did not guarantee complete physical accuracy of the solution, as the result still depended on the fidelity of the CFD.

The work of Potsdam was extended by Lim et al. to BVI-dominated flight cases based on the UH-60A, HART I, and HART II configurations [20]. The intent of this study was to show that the same coupling strategy could be applied to a regime that is dominated by vortex interaction. Their results showed that the loose coupling could handle unsteady aerodynamic conditions and improved the higher-harmonic airloads compared to the reduced-order methods, in particular in descending flight. However, discrepancies remained in the coupling solution that were sensitive to the temporal resolution and structure of the flow-field and wake. The study therefore reinforced that the loose coupling can improve BVI prediction, but only to the extent that the correct physics are actually captured by the aerodynamic model.

While early studies focused on demonstrating feasibility and stability, others began to identify the physical mechanisms responsible for the observed improvements. Datta et al. [21], applied a loosely coupled CFD-CA to predict vibratory loads in high-speed forward flight. Critically, their work highlighted an important physical mechanism underlying the improved performance of CFD-CA coupling, namely to accurately capture of the unsteady pitching moments, particularly in transonic flow regions near the blade tip. These pitching moments drive the elastic twist deformation, which in turn affects the aerodynamic loading distribution. They showed that this aeroelastic feedback loop is essential for accurately predicting the higher harmonic loading and rotor vibration.

In parallel, international validation efforts provided a systematic assessment of loosely coupled methods across multiple solvers and frameworks. During the HART-II International Workshop [22], a major cooperative research program to investigate helicopter rotors, rotor wake and prediction, multiple solver combinations were assessed, compared and validated. The results showed that different loosely coupled frameworks were capable of capturing the primary features of BVI aerodynamics, the formation, convection, and interaction of rotor tip vortices. However, the study also showed that prediction accuracy remained highly dependent on turbulence modelling, wake resolution, and time-step selection. Hence, the dominant limitations of the loose coupling are associated with CFD modelling choices rather than the coupling method itself.

Research also focused on improving the fidelity of structural and aerodynamic modelling within loosely coupled frameworks. The exchange of all six-components of the sectional loads was introduced in a loosely coupled framework by Park et al. [23]. Their results showed good agreement with HART-II experimental data in terms of airloads, blade deformation and wake geometry, indicating that the increased fidelity was able to enhance the prediction accuracy. Jung et al. [24] further extended this work by including rotor-fuselage interaction effects and demonstrated that modelling the complete rotorcraft configuration improved the phase accuracy of BVI loads and trim predictions.

The influence of structural modelling on the prediction capabilities of loosely coupled CFD-CA was investigated by You et al. [25]. By comparing CAMRAD-II and DYMORE within a common loose coupling framework, they showed that while CAMRAD-II provided better agreement for low-frequency airloads and blade deformation, DYMORE offered improved prediction of BVI airloads and trim control inputs. This highlighted the importance of both aerodynamic and structural modelling choices in determining overall simulation accuracy.

The role of aerodynamic modelling fidelity was investigated by Amiraux et al. [26]. Their research compared the performance of lifting-line, hybrid RANS/free-wake and CFD within a loosely coupled framework. Their results indicated that while lifting-line methods are suitable for preliminary analysis, they are unable to capture key BVI phenomena. The hybrid approach showed to be a compromise between cost and accuracy, but most importantly it was concluded that fully resolved CFD was required to predict vortex interactions and associated phenomena with great accuracy.

Recent studies have also extended loosely coupled CFD-CA methods to advanced complex rotor systems. Passe et al. [27] investigated the interactional aerodynamics of a coaxial rotor using a loosely coupled CFD-CA framework. This study found that loosely coupled CFD-CA method significantly improved the prediction of pitching moments, in particular in regions of reverse and unsteady flows. The low-fidelity model of the CA was shown to produce discontinuities in these regions and incorrect stall behaviour, whereas the CFD provided smooth and physically consistent solutions. Moreover, Passe et al discussed that loosely coupled methods require sufficient CFD convergence between coupling iterations to ensure stability, typically in the order of multiple rotor revolutions.

Beyond the pairing of individual solver, loosely coupled CFD-CA methods have also incorporated during the development of frameworks, such as the Helios framework developed by Sitaraman et al.[28]. Helios introduced a multi-solver environment that integrates CFD solvers, wake solvers, and CA codes within a unified Python-based framework. Unlike earlier implementation of solvers, Helios focuses on modularity and interoperability between solvers. An important features of Helios is its heterogeneous meshing strategy, which combines unstructured near-body grids for resolving viscous flow around the blades with adaptive off-body grids for efficient wake capture. This allows Helios to accurately capture rotor wakes while maintaining computational efficiency. Validation efforts have demonstrated that Helios is able to accurately predict rotor airloads and wake structures in forward flight, while maintaining stable convergence. However, similar to other implementations, performance remains dependent on CFD resolution and modelling fidelity, and computational cost is still heavily dominated by the CFD simulation.

In addition, coupling libraries such as preCICE have been developed to provide a flexible and solver-independent approach to partitioned multi-physics coupling [29]. Unlike integrated frameworks such as Helios, preCICE behaves as an external coupling interface that provides the communication between independently developed CFD and CA solvers. Moreover, it provides data mapping algorithms, and supports both explicit and implicit coupling strategies, allowing users to tailor the coupling to their specific requirements. The modular design of preCICE lowers the barrier to implement loosely coupled methods, especially when existing solvers were not originally designed to work together. Coincidentally, the general-purpose nature of this coupling library introduces additional challenges, which must be addressed to ensure accurate aeroelastic simulations.

Overall, the literature demonstrates that loosely coupled CFD-CA methods have evolved from early proof-of-concepts into mature tools capable of predicting rotor airloads, blade deformation, and key aeroelastic phenomena. The choice to apply the CFD solution to a trimmed rotor proved to be successful and has led to significant improvements in capturing unsteady aerodynamic effects when applied across a range of flight conditions, including BVI-dominated and high-speed regimes. At the same time, the studies consistently indicate that the predictive capabilities of loosely coupled frameworks remains sensitive to modelling choices, particularly the aerodynamic and wake representation, as well as numerical resolution. This motivates a closer examination of the key considerations regarding the implementation and capabilities of the loose coupling.

### 2.6.2. Synthesis of Findings, Results, and Limitations

Based on literature several key findings emerged regarding the performance, implementation, and limitations of loosely coupled CFD-CA frameworks for rotorcraft aeroelastic prediction.

A consistent finding in the literature is that aerodynamic modelling remains the primary source of error in rotor load prediction. While the structural models within CA solvers are often generally adequate, the low-fidelity aerodynamic model fail to capture physical phenomena such as transonic effects, vortex dynamics, dynamic stall, and three-dimensional radial flow [6], [18]. The results of loosely coupled methods are therefore governed less by the coupling itself than the quality of the CFD solution that is used. In practice, this means that the prediction capabilities are strongly dependent on the fidelity of the CFD solution, including turbulence modelling, spatial resolution, and temporal discretization [26], [30].

Critically, this implies that the loose coupling do not remove aerodynamic modelling uncertainty, instead it transfers the uncertainty into the aeroelastic solution by responding in a more physically accurate way. For this reason, a coupled framework may converge smoothly and yet deliver an inaccurate result as the CFD solution is unable to preserve or resolve the relevant flow structures with sufficient accuracy[3],

[30]. This is particularly important for rotorcraft applications, where unsteady aerodynamic loads are depend on both magnitude and phase accuracy. Errors in the temporal resolution of the CFD solution could therefore lead to incorrect loading, rotor motion and trim states even when the mean airloads appear reasonable [20], [30].

A central finding across multiple studies is that the loose CFD-CA coupling does not improve rotorcraft aeroelastic quantities uniformly. Instead, these improvements are selective and strongest when correcting the aerodynamic quantities that dominate the structural response of the rotor blade. Datta et al. [5], [6] showed that the primary benefit of CFD coupling is not solely the improved section lift, but rather the improved prediction of the sectional pitching moment, particularly in the transonic outboard region of the rotor blade.

This improvement is highly relevant because the sectional pitching moments contribute to elastic torsional deformation of the blade. The resulting torsional motion modifies the local blade pitch and therefore the effective angle of attack. In turn, this alters the sectional normal force distribution, and leads to vibratory airloads that cause additional torsional and bending moments in the blade. The effectiveness of the coupled framework therefore depends on whether the CFD correction improves the aerodynamic quantities that drive this aeroelastic response, in particular the magnitude and phase of the sectional pitching moment.

As a result, improvements in lift prediction alone are insufficient to guarantee improved aeroelastic predictions. CFD should therefore be prioritised where it improves the aerodynamic quantities that drive the structural response, in particular the sectional pitching moments. This can be achieved by applying CFD, or locally increasing the mesh resolution, in blade regions where the low-fidelity methods have been shown to fail, such as the transonic outboard region identified by Datta et al [5]. In addition, this implies that the validation of a loosely coupled framework has a preferred hierarchy, first sectional pitching moments and torsional motion, then normal force and only then integrated quantities such as bending moments and control loads.

At the same time, accurate structural modelling is a prerequisite for the CFD loads to produce the correct aeroelastic response in a loosely coupled framework. Datta showed that after validating the structural model using measured flight-test airloads the remaining high-speed lift-prediction errors could be attributed to the aerodynamic and not structural modelling. Although the structural model is not often the dominant source of error, this work showed that the structural model must be sufficiently refined and validated to correctly respond to the aerodynamic corrections. This included the accurate representation of torsional stiffness, modal content, mass offsets, boundary conditions, damping, and control-system dynamics. Without this, the improvements in aerodynamic loading that CFD offers cannot be translated into the correct aeroelastic response.[5]

In addition, the literature highlights that the predictive accuracy is highly dependent on wake modelling in CFD. While improving the sectional pitching moment prediction can lead to accurate elastic torsion, it proved to be insufficient to capture the impulsive loading and the flap-bending response of the blade (approximately 75%R), which are governed by wake interactions/BVI [5]. Therefore, CFD should be able to accurately resolve the formation, convection, and preservation of the tip vortices, such that the rotor wake remains strong and coherent over multiple rotor revolutions.

This requirement introduces critical modelling considerations for loosely coupled simulations. In particular, accurate prediction of BVI requires the CFD solver to capture the initial vortex roll-up as well as the subsequent interactions with following blades at later azimuthal positions. In order to achieve such accuracy, stringent requirements are imposed on grid resolution, numerical dissipation and turbulence modelling as excessive diffusion may lead to the rapid decay of vortices and the loss or interactions. As reported by the HART-II studies, the ability of CFD to preserve the wake structure over several rotor revolutions is essential for capturing the magnitude and timing of the impulsive airloads and vibrations[30].

Overall, the literature indicates that the accuracy with which the dominant aeroelastic mechanism are captured, determines the effectiveness of loosely coupled CFD-CA methods. As summarized in Table 2.2, improvements to the predictive capabilities are selective and can mainly be achieved by enhanced prediction of sectional pitching moment given its effect on the structural response. However, wake preservation, numerical resolution and aerodynamic modelling fidelity also continue to be limiting

factors. These findings suggest the successful implementation requires the targeted application of CFD combined with a validated structural model, rather than simply coupling the two solvers.

**Table 2.2:** Summary of Key Findings for Loose CFD-CA Coupling.

<b>Aspect</b>	<b>Key Insight and Implication</b>
Aerodynamic Modelling	Primary source of error; CFD fidelity (turbulence, grid, time-step) governs accuracy [18], [30].
Coupling Behaviour	Convergence does not ensure correctness; errors in CFD propagate into aeroelastic response [3].
Core Mechanism	Improvement driven by pitching moment: $C_m \rightarrow \theta_e \rightarrow C_n \rightarrow$ loads [5].
Selective Improvement	Strongest improvements in pitching moment, torsion, and vibratory loads; weaker for wake-driven loads [5], [18].
Structural Modelling	Must be accurate and validated (torsion, modes, damping, controls) to correctly respond to CFD loads [5].
Wake Modelling	Critical for BVI and impulsive loads ( $\sim 75\%R$ ); requires vortex preservation over multiple revolutions [5], [30].
Numerical Sensitivity	Highly sensitive to turbulence model, spatial resolution, and temporal accuracy [26], [30].

# 3

## Methods

This chapter describes the methodology of the partitioned CFD-CA framework developed at NLR. The coupling strategy, solver methodologies, interface coupling, data exchange, and the modifications required for the implementation of the framework are presented in detail.

### 3.1. Coupling Strategy

This section defines the coupling strategy used to simulate the fluid-structure interaction in rotorcraft aeromechanics. This approach is based on a partitioned CFD-CA framework with a partitioned coupling strategy between a comprehensive analysis solver and a CFD solver. This section describes the coupling philosophy, the global coupling workflow, and the sequence of operations within each coupling iteration.

#### 3.1.1. Coupling Philosophy

To investigate the fluid-structure interaction problems of rotorcraft aeromechanics, NLR has chosen to construct a partitioned, loosely coupled simulation framework. In this framework, a rotorcraft comprehensive analysis code and a CFD solver are treated as independent codes that exchange information at set coupling intervals. During each coupling iteration, the comprehensive analysis solver determines the periodic motion of the rotor, which is fed back to the aerodynamic solver. The CFD simulation then determines the aerodynamic loads, which are fed back to the structural model as a corrective term.

NLR's choice for a partitioned framework with partitioned coupling is motivated by practical implementation considerations. This approach allows the CFD and comprehensive analysis solvers to utilise their own numerical settings, discretizations, and time-integration schemes, while reusing existing, well-validated tools. In contrast, developing an efficient monolithic solver with comparable CFD, structural-dynamics, trim, control-system, and rotorcraft-modelling capabilities would require substantial integration and development effort. Maintaining solver independence therefore allows NLR to leverage its in-house solvers and expertise.

Within this coupling strategy, the delta-airloads method is used to correct the low-fidelity aerodynamic loads computed by the comprehensive analysis solver. The correction is defined as the difference between the high-fidelity CFD airloads and the baseline airloads predicted internally by the comprehensive analysis code. This correction is then applied during the next trim calculation, so that the structural response is driven by aerodynamic loads that progressively approach the CFD solution. Iteratively, the CA airloads are replaced by the CFD airloads, while the comprehensive analysis solver continues to provide the trim and structural response. In this way, the method produces a consistent coupled aerodynamic and structural solution.

#### 3.1.2. Global Coupling Workflow

The delta-airloads method forms the underlying workflow that couples the structure and fluid models and provides the framework for the exchange of the blade motion, control inputs, and the aerodynamic

loads between both solvers. The workflow is therefore the starting point for describing the architecture of the partitioned coupling. The original formulation of the workflow by Potsdam et al. [3] was discussed in Section 2.5.2. The workflow employed by NLR is similar to that of Potsdam, and is shown in Figure 3.1.

The following notation is used: the superscripts  $S$  and  $F$  denote quantities obtained from the structural and the fluid (CFD) solver, respectively. The subscript  $i$  indicates the global coupling iteration, where an iteration refers to one periodic rotor revolution. The vectors  $F$  and  $M$  represent the aerodynamic forces and moments acting on the rotor blades.

The delta-airloads workflow is initiated by the structural model. Due to the lack of available aerodynamic CFD data at the initialisation, the comprehensive analysis solver employs a built-in, low-fidelity aerodynamic model to estimate the forces and moments  $F_0^S$  and  $M_0^S$  of the rotorcraft. Consequently, the comprehensive analysis solver performs a trim calculation of the rotorcraft. This results in a periodic, trimmed rotor motion as a function of azimuth. The blade motion is subsequently passed to the CFD simulation, which determines the aerodynamic loading  $F_0^F$  and  $M_0^F$  for the given rotor motion. In turn, the aerodynamic loading is transformed for use by the comprehensive analysis solver. At the subsequent global iteration  $i = 1$ , the aerodynamic loads  $F_0^F$  and  $M_0^F$  are applied as a correction to the low-fidelity aerodynamic method of the comprehensive analysis, yielding corrected loads at iteration  $i = 1$ . Using the corrected loading, a new trimmed rotor motion is determined, and the loop can start anew.

The following differences between the implementation by NLR and the original formulation by Potsdam can be noted [3]. Firstly, in Potsdam's implementation, prior to progressing to the next iteration, the convergence of the aerodynamic loads and the control inputs is assessed. If the loads and control inputs have converged to approximately three significant digits (i.e., on the order of  $\sim 10^{-2}$ – $10^{-3}$  degrees for control angles), the simulation is terminated. The implementation by NLR lacks a convergence check, and instead the number of iterations is chosen by the user a priori. A check is added prior to each CFD simulation to assess whether the final iteration  $i_{max}$  has been reached, in order to stop the simulation.

The second difference concerns the aerodynamic force and moment components that are included in the delta-airloads correction. In the implementation by Potsdam et al., the correction includes the normal force, chord force, and sectional pitching moment, while the axial force, rolling moment, and yawing moment are omitted. This is consistent with the focus of the method on the dominant sectional airloads that drive the rotor blade response. The NLR implementation includes the axial, chord, and normal force components and the sectional pitching moment in the delta-airloads correction. The rolling and yawing moments are still omitted, because their contribution to the blade motion is assumed to be small compared to the force components and the sectional pitching moment. Additionally, both the pressure and viscous contributions are retained in the transferred forces and moments.

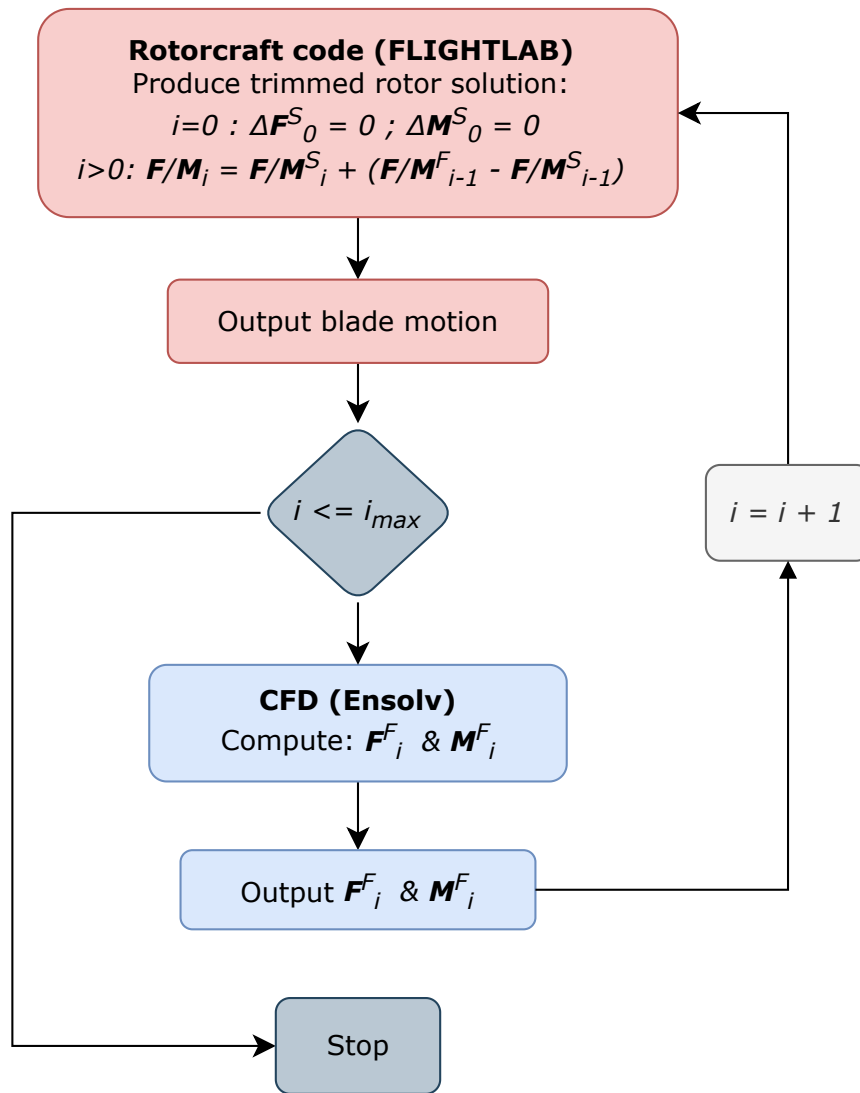


Figure 3.1: Flowchart describing NLR's implementation of the loose coupling based on the delta-airloads method.

### 3.1.3. Global Timing Scheme

The delta-airloads workflow can be rearranged to visualize the organization of the timing scheme of the global fluid-structure problem. The following serial-explicit staggered approach represents the global loose-coupling workflow, in which the structural and aerodynamic solvers are executed sequentially and exchange information once per global coupling iteration. This workflow is illustrated in Figure 3.2.

The timing and order of each global coupling iteration can be reduced to four distinct operations. First, a new iteration is started when, at a given iteration  $i - 1$ , the comprehensive analysis solver determines a new trimmed rotor motion for the given load case. At  $i = 0$ , this is achieved using solely the internal low-fidelity aerodynamic model of the comprehensive analysis. For subsequent iterations, the trim procedure incorporates the delta-airloads based on the CFD solution from the previous iteration.

Second, the determined rotor motion is transferred to the aerodynamic solver. This requires mapping the structural deformation to the aerodynamic mesh and updating the CFD setup to simulate the correct




Figure 3.2: Loose coupling staggered scheme.

rotor geometry and flight condition. Mesh deformation methods are utilised to ensure that the structural motion is correctly mapped to the blade surface and volume mesh, while maintaining mesh quality.

Third, the CFD simulation is performed using the updated geometry. The CFD simulation yields the high-fidelity aerodynamic forces and moments acting on the rotor blades, capturing the influence of complex flow phenomena that are otherwise not resolved by the comprehensive analysis.

Finally, the aerodynamic loads are extracted from the CFD and prepared to be used by the structure solver. This requires mapping the CFD loads from the aerodynamic mesh to the structural discretization of the rotor blades in a conservative manner, ensuring that the total forces and moments are conserved. The transformed loads are then shared with the comprehensive analysis solver to be used in the next iteration.

For iterations  $i > 0$ , the aerodynamic solution from the CFD simulation of the previous iteration is used to correct the structural solution. This results in an iterative process intended to approach a consistent aeroelastic response. Overall, this sequence defines one complete coupling loop.

## 3.2. Solver Methods

This section describes the two primary solvers employed within the loosely coupled framework: the comprehensive analysis code FLIGHTLAB, which provides the structural dynamics, baseline aerodynamic solution, and trim capability, and the CFD solver ENSOLV, which provides the high-fidelity aerodynamic solution. First, the formulation and modelling capabilities of FLIGHTLAB are presented, including its system equations, structural and aerodynamic model, and trim methodology. This is followed by a description of the governing equations, discretization techniques, and solution procedure of ENSOLV.

### 3.2.1. Comprehensive Analysis Method: FLIGHTLAB

The comprehensive analysis software, FLIGHTLAB, is a simulation framework used to analyse rotorcraft aeromechanics by coupling structural mechanics, aerodynamics, and flight dynamics. In the context of the partitioned ENSOLV-FLIGHTLAB framework, FLIGHTLAB's primary function is to act as the computational structural dynamics solver by providing the blade motion and baseline aerodynamic loading, which can be corrected by the CFD. Moreover, FLIGHTLAB is able to determine a consistent trim solution, which is essential to ensure that the force and moment equilibrium of the rotor is satisfied. In this section, the general formulation, structural model, aerodynamic formulation, and the trim methodology of FLIGHTLAB are described, emphasising their role for rotorcraft and coupled analysis applications.

#### System Formulation

A general rotorcraft is formulated in FLIGHTLAB as a coupled nonlinear dynamic system in which all the physical subsystems, including structure, aerodynamics, and control, contribute to the common set

of governing equations. The system can be written in generalised form as:

$$Q = f(\ddot{x}, \dot{x}, x, u) = 0, \quad y = g(\ddot{x}, \dot{x}, x, u) \quad (3.1)$$

where  $x$  represents the system states,  $u$  the inputs,  $y$  the outputs, and  $Q$  the residual forces or the force imbalance. To solve the system, the residual  $Q$  is driven to zero, meaning that the imbalances in the solution are reduced and all subsystems and physical aspects come to a consistent solution [31].

The nonlinear character of the system comes from both the individual subsystem models and their coupling. For example, the structural response depends on the instantaneous blade deformation, while the aerodynamic loads depend on the local flow state, blade motion and control inputs. These quantities are mutually dependent, since blade deformation changes the aerodynamic loading and the aerodynamic loading changes the blade response. FLIGHTLAB therefore solves the coupled nonlinear system iteratively by linearising the governing equations about the current state and updating the state vector using a Newton-Raphson procedure until convergence is reached [32].

### Structural Dynamics Model

Structurally, FLIGHTLAB models rotor blades as flexible structures using a nonlinear beam formulation. The equations of motion for the structural system can be expressed as:

$$M(x)\ddot{x} + C(x)\dot{x} + K(x)x = F_{aero} + F_{inertial} \quad (3.2)$$

where  $M$ ,  $C$ , and  $K$  denote the mass, damping, and stiffness matrices. Since these matrices depend on the instantaneous blade deformation, the structural dynamics are inherently nonlinear. The aerodynamic loads  $F_{aero}$  are directly coupled to the blade motion and act as external forcing. This formulation captures the main deformation modes that are relevant for rotor blades: flapwise bending, lead-lag motion, torsion, and axial deformation. In addition, a key physical effect that is included in the model is centrifugal stiffening, which arises from the rotor rotation and increases the effective stiffness of the blade and, thus, its resistance to bending [31].

### Aerodynamic Model

The aerodynamic model used in FLIGHTLAB computes the aerodynamic loads using a nonlinear quasi-unsteady sectional airload formulation. The model is based on blade-element theory and combines two-dimensional airfoil data with corrections for effects such as yawed flow, blade motion, and rotational three-dimensional effects. It therefore provides a computationally efficient representation of the nonlinear sectional aerodynamics required for rotorcraft flight analysis [32].

For each blade segment, the local relative airflow is determined from the structural motion, rotor rotation, freestream velocity, turbulence, and induced velocity. From this local velocity vector, the effective angle of attack, yaw angle, and local Mach number are computed. The lift, drag, and pitching moment coefficients are then obtained from airfoil tables as functions of angle of attack and Mach number. These tables provide sampled aerodynamic coefficient data; linear interpolation is used to evaluate the coefficients at the current operating condition. Using these coefficients, the sectional lift and drag are computed as:

$$dL = \frac{1}{2}\rho V_{rel}^2 cC_L(\alpha, M)dr, \quad dD = \frac{1}{2}\rho V_{rel}^2 cC_D(\alpha, M)dr \quad (3.3)$$

Additional corrections are applied to account for yawed flow and rotational three-dimensional effects, such as stall delay. The computed sectional loads are calculated in the local blade frame and are separated into circulatory and profile components. The circulatory loads are used to compute the induced velocities as part of the wake development, while the profile drag contributes only to the total aerodynamic force.

The wake is typically represented using prescribed or free-wake vortex models, where the induced velocity is determined based on the shed vortex system. The induced velocity from the wake influences the local inflow conditions and closes the aerodynamic loop.

The term "quasi-unsteady" refers to the treatment of the stall behaviour. The model treats stall as instantaneous with respect to the current effective angle of attack and does not explicitly include

time-dependent dynamic stall or non-circulatory effects. As a result, the model captures the primary nonlinear aerodynamic behaviour and motion-induced effects without the added complexity of a fully unsteady aerodynamic model [31].

### Trim Loop

After initialization of the rotorcraft model, flight condition, and load case, FLIGHTLAB determines the control inputs and rotor state required to achieve force and moment equilibrium while maintaining the prescribed flight condition. This is achieved through the trim procedure, whose outcome is a combination of control inputs, rotorcraft attitude, and rotor performance that together define a stable operating condition, referred to as the trimmed state. In this work, an azimuthal resolution of  $1^\circ$  is used for both aerodynamic and structural calculations.

The trim problem is formulated as a force and moment balance problem. Aerodynamic loads generated by the rotor are obtained by integrating sectional forces along the blade span and are transferred to the rotorcraft equations of motion, resulting in translational and rotational accelerations. A trimmed state is achieved by iteratively adjusting a set of independent trim variables until predefined trim targets converge within specified tolerances. The trim problem is therefore defined by selecting an equal number of trim variables and trim targets.

The trim variables represent the controllable degrees of freedom. For the isolated rotor considered in this work, these include the collective pitch  $\theta_0$ , longitudinal cyclic pitch  $\theta_{1s}$ , and lateral cyclic pitch  $\theta_{1c}$ . Additional variables may include rotorcraft attitude angles or flight path parameters. The trim targets define the required equilibrium conditions; in this work, these are the rotor thrust  $T$ , roll moment  $M_x$ , and pitch moment  $M_y$ . More generally, trim targets may include body accelerations, inertial accelerations, or power requirements, depending on the analysis.

The trim solution is obtained using iterative algorithms such as Newton-based or Hooke–Jeeves methods. Starting from an initial estimate of the trim variables, the nonlinear rotor equations are solved until a steady or periodic steady state is reached, typically by averaging over multiple rotor revolutions. The resulting forces and moments are then compared to the trim targets. Any residual error is used to update the trim variables, and the process is repeated until convergence criteria are satisfied or iteration limits are reached [33].

Once convergence is achieved, FLIGHTLAB provides the trimmed control inputs, rotor motion, aerodynamic loading, and time histories of the relevant quantities. In this work, the trim procedure is applied within a loosely coupled framework between ENSOLV and FLIGHTLAB and is referred to as the trim operation.

### 3.2.2. Aerodynamic Solver: ENSOLV

ENSOLV is NLR's in-house computational fluid dynamics (CFD) solver and is part of the ENFLOW simulation system. It is designed for the simulation of three-dimensional compressible flows across subsonic, transonic, and supersonic regimes, including the modelling of complex flow phenomena such as shock waves, boundary layers, and wakes [34], [35]. Within the loose coupling framework, ENSOLV provides the high-fidelity aerodynamic solutions that are used to correct the baseline aerodynamic loads obtained by FLIGHTLAB.

#### Governing Equations and Flow Model

In ENSOLV, the governing compressible Reynolds-averaged Navier-Stokes (RANS) equations, which describe the conservation of mass, momentum, and energy in the flow field, are formulated in conservative form as

$$\frac{\partial \vec{U}}{\partial t} + \nabla \cdot \vec{F}_c(\vec{U}) = \nabla \cdot \vec{F}_d(\vec{U}, \nabla \vec{U}) + \vec{S}, \quad (3.4)$$

where  $\vec{U}$  are the conserved flow variables,  $\vec{F}_c$  the convective fluxes,  $\vec{F}_d$  the diffusive fluxes, and  $\vec{S}$  source terms. The flow-state vector is defined as:

$$\vec{U} = \begin{pmatrix} \rho \\ \rho \vec{u} \\ \rho E \end{pmatrix} \quad (3.5)$$

where  $\rho$  is the density,  $\vec{u}$  the velocity vector, and  $E$  the total energy per unit mass.

In the RANS formulation, the flow variables are decomposed into mean and fluctuating components and averaged in time. This introduces the Reynolds stress tensor, which represents the effects of turbulence on the mean flow and requires modelling to close the system of equations.

In ENSOLV, turbulence is modelled using an eddy-viscosity approach, meaning that the Reynolds stresses are related to the mean velocity gradients through a turbulent viscosity. This turbulent viscosity is obtained from turbulence models, which range from algebraic formulations to additional transport equations.

### Spatial Discretization

The governing equations are discretized in ENSOLV using a cell-centered finite volume method on structured multi-block grids. This discretization is conservative by construction, since the flux balance is evaluated over each control volume. The use of structured multi-block grids allows the rotor geometry and surrounding flow domain to be divided into separate structured blocks, while retaining a finite volume formulation throughout the computational domain.

Integration of the governing equations over a control volume gives a balance of fluxes across the cell faces:

$$\frac{d}{dt} (\vec{U}_i V_i) + \sum_{f \in \partial V_i} (\vec{F}_{c,f} - \vec{F}_{d,f}) \cdot \vec{n}_f A_f = 0, \quad (3.6)$$

where  $V_i$  is the cell volume,  $A_f$  the face area, and  $\vec{n}_f$  the outward normal vector.

In ENSOLV, the convective fluxes are evaluated using a second-order central differencing scheme. Central differencing introduces little inherent numerical dissipation, meaning that resolved flow gradients and vortical structures are not strongly damped by the discretization itself. Because this scheme does not provide sufficient numerical damping for stable compressible-flow simulations, artificial diffusion is added to suppress spurious numerical oscillations. ENSOLV uses a Jameson-type artificial diffusion formulation, based on a combination of second-order and fourth-order diffusion terms. The second-order contribution provides additional damping in regions with strong gradients, while the fourth-order contribution damps high-frequency numerical oscillations in smoother parts of the flow field.

### Solution Method

The spatial discretization produces a system of nonlinear algebraic equations, which can be written in residual form as

$$\vec{R}(\vec{U}) = 0, \quad (3.7)$$

where  $\vec{R}$  represents the imbalance of fluxes in each control volume.

ENSOLV solves this system using a multi-grid method. The multi-grid approach accelerates the convergence of the solution by solving the equations on a hierarchy of grids. Coarse grids are used to reduce the global error modes, whereas fine grids are used to resolve the detailed flow features.

On each grid level, a relaxation method is applied based on an explicit multi-stage Runge-Kutta scheme:

$$\vec{U}^{(q+1)} = \vec{U}^{(q)} - \alpha_q \Delta \tau \vec{R}^{(q)}, \quad (3.8)$$

where  $\alpha_q$  are the Runge-Kutta coefficients and  $\Delta \tau$  is the pseudo-time step. Additional acceleration techniques can be employed to further enhance convergence, such as local time stepping and implicit residual averaging.

### Time Integration: Steady and Unsteady Simulations

In terms of time integration, a clear distinction is made between steady and unsteady simulations.

For steady simulations, a pseudo-time derivative is introduced:

$$\frac{\partial \vec{U}}{\partial \tau} + \vec{R}(\vec{U}) = 0, \quad (3.9)$$

where  $\tau$  is pseudo-time. For steady simulations, the pseudo-time step determines how large the update is in the artificial time domain, whereas the Runge-Kutta scheme defines how this update is applied using multiple intermediate stages. In this case, the time integration is used as a numerical tool to iteratively reduce the residual to zero.

For unsteady simulations, the physical time derivative is retained:

$$\frac{\partial \vec{U}}{\partial t} + \vec{R}(\vec{U}) = 0. \quad (3.10)$$

The time step  $\Delta t$  is selected based on accuracy and stability requirements, and the solution is advanced in physical time. Compared to steady simulations, this time-integration represents the physical evolution of the flow field rather than a numerical convergence mechanism.

### 3.3. Reference Frames and Solver Consistency

This section defines the reference frames used in this work and their relationships. Consistent definitions of the reference frames are required to ensure compatibility between the rotor motion generated in FLIGHTLAB and the aerodynamic analysis performed in ENSOLV. As these rotorcraft simulations concern rotating and deforming structures, multiple coordinate systems are defined that describe the motion, loads, and flow in a physically consistent manner.

The following reference frames are considered: inertial, body, hub, blade, and blade-section frames. The coordinate convention used to define the reference frames is based on FLIGHTLAB. For the exchange of motion and loading between FLIGHTLAB and ENSOLV, the conditions of ENSOLV are adapted to ensure consistency with the rotor system in FLIGHTLAB.

#### 3.3.1. Reference Frame Definitions: FLIGHTLAB

In FLIGHTLAB, the inertial frame is used as the global reference frame in which all rotor motion is defined. The coordinate convention in FLIGHTLAB is illustrated in Figure 3.3. In the inertial frame,  $X_i$  points upstream,  $Z_i$  points downward, and  $Y_i$  completes a right-handed system. The free-stream velocity is defined along the negative inertial  $x$ -axis, and a positive angle of attack corresponds to flow directed toward the  $(-x, -z)$  plane.

The direction of rotation of a rotor is positive counter-clockwise about the negative  $z$ -axis and is defined using the azimuth angle  $\psi$ . The azimuth angle  $\psi$  is defined such that  $\psi = 0^\circ$  when a rotor blade is aligned with the negative  $x$ -axis and increases in the direction of rotation. For a given rotor, the rotor blades are spaced uniformly in azimuth and numbered as shown.

The body-fixed frame is attached to the rotorcraft, with its origin at the rotor hub. Here,  $X_\beta$  points downstream,  $Y_\beta$  points to starboard, and  $Z_\beta$  points upward. The attitude of the rotorcraft can be controlled through pitch, where positive rotorcraft pitch means the rotor shaft is tilted aft. Additionally, the tilt of the rotor shaft can be controlled independently with respect to the rotorcraft; in this case, positive rotor shaft tilt opposes the aft pitching of the rotorcraft.

The non-rotating rotor hub frame is defined at the rotor hub and defines the rotor orientation. In this frame,  $Z_h$  is aligned with the rotor shaft,  $X_h$  points downstream, and  $Y_h$  points to starboard; thus, the  $x$ - $y$  plane defines the rotor disk. The hub frame also defines the azimuth angle, and is used as the intermediate frame between the rotorcraft and the rotating blades. When the rotor shaft is aligned with the negative  $z$ -axis, the body and hub frames are aligned.

Each blade also has a blade-fixed frame  $(X_{\beta l}, Y_{\beta l}, Z_{\beta l})$  that can be obtained from the hub frame through the azimuth rotation  $\psi$ . This frame follows the rigid-body motion of the blade. Additionally, each blade section has a local blade-section frame that is used for local blade motion and aerodynamic calculations. In this frame,  $X_s$  points radially outward,  $Y_s$  is aligned with the local chord toward the leading edge, and  $Z_s$  completes the right-handed system.

When exporting the motion of the rotor from FLIGHTLAB, all motion is expressed in the inertial frame. Due to the definition of the reference frames, this means that positive vertical displacement of the rotor occurs in the direction of the inertial negative  $z$ -axis. The rotation is given about the negative  $z$ -axis.




The aerodynamic force coefficients  $C_L$ ,  $C_D$ , and  $C_S$  are defined in the wind-axis system and are normalized by the dynamic pressure  $q_\infty = \frac{1}{2}\rho_\infty u_\infty^2$  and reference area  $S_A$ . The roll, pitch, and yaw moment coefficients are defined in the flight-mechanics body axis system as the negative  $x$ -, positive  $y$ -, and negative  $z$ -components of the aerodynamic moment vector, respectively.

### 3.3.3. Solver Consistency

To ensure the physical consistency between the simulations of FLIGHTLAB and ENSOLV, the orientation of the rotor geometry and the flight condition must correspond. This requires that the definitions of the free-stream velocity, angle of attack, and rotation axis be mutually consistent.

The rotor geometry is loaded in ENSOLV using the motion files generated by FLIGHTLAB. As a result, the rotor position and movement during a full revolution are identical in the inertial, right-handed coordinate system used in both solvers. Further guaranteeing that the motion of the rotor in ENSOLV is consistent with FLIGHTLAB requires the explicit definition of the axis of rotation. In accordance with FLIGHTLAB, the axis of rotation in ENSOLV is defined as the negative inertial  $z$ -axis.

The definitions of the free-stream velocity and angle of attack are shown in Figure 3.5 to illustrate their differences. In FLIGHTLAB, a positive free-stream velocity is defined as acting in the direction of the negative inertial  $x$ -axis, whereas in ENSOLV a positive free-stream velocity acts in the direction of the positive inertial  $x$ -axis. In addition, the convention for the free-stream angle of attack differs between the two solvers. In FLIGHTLAB, a positive angle of attack corresponds to a free-stream vector that points toward the  $(-x, -z)$  plane, while in ENSOLV a positive angle of attack corresponds to a free-stream vector that points toward the  $(+x, -z)$  plane. In other words, the conventions of the free-stream and angle of attack of FLIGHTLAB and ENSOLV are mirrored with respect to the inertial  $z$ -axis.



**Figure 3.5:** Comparison of free-stream and angle of attack conventions between FLIGHTLAB and ENSOLV. Rotor blades shown in grey.

To ensure that both solvers utilise the same free-stream direction, the angle of attack defined in FLIGHTLAB needs to be transformed before being applied in ENSOLV. The following conditional transformation can account for the mirrored sign convention.

$$\alpha_{\text{EN}} = \begin{cases} 180^\circ - \alpha_{\text{FL}}, & \alpha_{\text{FL}} > 0, \\ -180^\circ - \alpha_{\text{FL}}, & \alpha_{\text{FL}} \leq 0. \end{cases} \quad (3.11)$$

Here,  $\alpha_{\text{FL}}$  denotes the angle of attack defined in FLIGHTLAB and  $\alpha_{\text{EN}}$  the corresponding angle applied in ENSOLV. This transformation ensures that the free-stream velocity has the same direction in the inertial frame in both solvers. When this convention is applied consistently, and the load case parameters and the rotor motion are transferred correctly, FLIGHTLAB and ENSOLV analyse the same flight condition despite the internal convention differences.

### 3.4. Grid Interface


In fluid-structure interaction problems, it is essential to correctly interface and couple the structural and aerodynamic discretizations of the rotor geometry. Since FLIGHTLAB and ENSOLV utilise different discretizations of the rotor blade, a dedicated mapping procedure is required to transfer the motion and aerodynamic loading between the solvers.

The data mapping must satisfy two requirements. First, the transfer of the structural deformation to the aerodynamic grid must be consistent, meaning that the deformed CFD mesh accurately represents the structural motion. Second, the transfer of the aerodynamic loads from the CFD to the structure must be conservative, meaning that the total forces and moments are preserved across the interface. Ensuring that both these requirements are met is essential for the physical consistency and numerical stability of the coupled simulation.

#### 3.4.1. Grid Interface Coupling

The structural and aerodynamic domains use fundamentally different discretizations of the same rotor blade geometry. In FLIGHTLAB, the rotor blade is modelled as a one-dimensional beam that is discretized into a set of structural elements. Additional aerodynamic control points (ACPs) are defined along the span to facilitate the application of CFD loading. ENSOLV, on the other hand, employs a three-dimensional surface mesh that is based on the vertices of the grid cell faces, accurately representing the outer geometry of the rotor blade.

Although a direct 1D to 3D mapping is possible, it is not preferred. Mapping the deformation of a one-dimensional beam onto a three-dimensional surface would result in the loss of local twist deformation, which plays a crucial role in rotor response. To preserve the local torsional deformation, the representation of the geometry in FLIGHTLAB is extended. In addition to the ACPs, points corresponding to the leading edge and trailing edge of each blade section are defined at the same spanwise locations. This effectively extends the structural representation from a one-dimensional line to a two-dimensional sectional description. As a result, the coupling interface is transformed from a 1D-3D to a 2D-3D mapping. This allows the local torsional deformation of the blade sections to be consistently represented and mapped between the structural and aerodynamic discretizations.



**Figure 3.6:** Geometric coupling interface between FLIGHTLAB and ENSOLV. (ENSOLV surface grid in grey and FLIGHTLAB points in red)

The aerodynamic grid used in FLIGHTLAB, with ACPs and the added leading edge and trailing edge points, and the surface mesh utilised in ENSOLV are shown in Figure 3.6.

### 3.4.2. Mapping Method: Volume Spline Interpolation

To transfer data across the interface, NLR's in-house aeroelastic splining utility SPLMOD is used, which constructs a smooth global mapping using a volume spline interpolation approach [36].

As inputs, this approach requires two discretizations of the same geometry. First is the structural mesh, consisting of the ACPs and the leading- and trailing-edge points defined in FLIGHTLAB. Second is the aerodynamic surface grid, consisting of the vertices of the CFD surface mesh in ENSOLV. The structural grid acts as the support points for the interpolation, while the aerodynamic surface mesh defines the locations where the transformed quantities are desired.

The volume spline interpolation is defined for support points  $\vec{x}_i \in \mathbb{R}^3$ ,  $i = 1, \dots, N$ , with function values  $f_i \in \mathbb{R}$ ,  $i = 1, \dots, N$ . The interpolation function is defined as:

$$f(\vec{x}) = \alpha_0 + \sum_{i=1}^N \alpha_i R_i(\vec{x}), \quad (3.12)$$

where the radial basis function  $R_i$  is given by  $R_i(\vec{x}) = \|\vec{x} - \vec{x}_i\|$ . The coefficients  $\alpha_i$ ,  $i = 1, \dots, N$ , are determined from the interpolation conditions  $f(\vec{x}_i) = f_i$  together with the side condition  $\sum_{i=1}^N \alpha_i = 0$ . This formulation results in a smooth interpolation, and a system of equations that depends only on the position of the support points.

For a vector field, such as forces, the same interpolation can be applied component-wise. Since the interpolation depends only on the geometry, the resulting system matrix is identical for each component and needs to be computed only once.

Evaluating the interpolation at the aerodynamic mesh points results in a linear mapping between the structural and aerodynamic representations of the geometry:

$$\vec{d}_f = \mathbf{C}_{as} \vec{d}_s, \quad (3.13)$$

where  $\vec{d}_s$  and  $\vec{d}_f$  represent the structural and aerodynamic displacements, respectively, and  $\mathbf{C}_{as}$  is the connectivity matrix obtained from the volume spline interpolation. The connectivity matrix  $\mathbf{C}_{as}$  is only dependent on the relative positions of the structural and aerodynamic discretizations and is constructed once using the undeformed geometry of the rotor.

During the coupling iterations, the rotor geometry deforms. Instead of recomputing the spline mapping, the deformation is expressed relative to the initial, undeformed rotor, and the connectivity matrix is reused to transfer the updated structural displacement to the aerodynamic mesh. As such, the deformations of the CFD surface grid remain consistent with the structural motion. The deformation is propagated into the interior grid points using transfinite interpolation [36].

The same approach is used to transfer the aerodynamic loads via a conservative mapping. The aerodynamic loads computed on the CFD surface are transferred to the structural grid as

$$\vec{F}_s = \mathbf{C}_{sa} \vec{F}_a. \quad (3.14)$$

To ensure the conservation of energy across the interface, the mapping must satisfy the condition of virtual work:

$$\vec{F}_a^T \vec{d}_f = \vec{F}_s^T \vec{d}_s. \quad (3.15)$$

Based on the mapping of the deformation, this expression leads to the relation

$$\mathbf{C}_{sa} = \mathbf{C}_{as}^T. \quad (3.16)$$

Hence, the deformation transfer is consistent and the load transfer is conservative.

### 3.4.3. Transfer of Grid Motion

In FLIGHTLAB, the trimmed rotor motion is defined in the inertial frame and stored as a description of the structural grid over a full revolution. During every global iteration of the loose coupling, this rotor

motion is shared with ENSOLV. Using the precomputed volume spline interpolation, the structural motion is transferred to the aerodynamic mesh.

For each time step, the structural displacements are mapped onto the CFD surface mesh. The rotor rotation is applied as rigid grid deformation, by rotating the body frame in ENSOLV with respect to the inertial frame. In addition, the combination of pitch (collective and cyclic), flapping, and lead-lag motion is applied as flexible grid deformation. After the updated surface mesh has been obtained, transfinite interpolation is used to propagate these displacements into the interior of the mesh. This guarantees the smooth deformation of the volume grid while the mesh quality is maintained.

For unsteady simulations, this procedure leads to the creation of a sequence of deformed CFD meshes that correspond to the motion of the rotor over one full revolution. For each iteration, the correct mesh is loaded by ENSOLV and used to compute the new flow field. For steady simulations, only a single instance is required. In this case, the initial rotor position is used, corresponding to  $\psi = 0$ .

The deformation of the CFD grid is illustrated in Figures 3.7 and 3.8 using a slice of the blade section near the blade tip ( $r/R = 0.87$  [-]). The figures show that the displacement is smoothly transferred to the aerodynamic mesh and the motion is incorporated into the surrounding grid, with the deformation gradually decaying away from the surface.




Figure 3.7: Initial, undeformed grid slice near the blade tip.




Figure 3.8: Deformed grid slice near the blade tip.

#### 3.4.4. Transfer of Aerodynamic Loading

To transfer the aerodynamic loading computed by ENSOLV to FLIGHTLAB, several processing steps are required to ensure that the loading is conserved regardless of the discretization and representation of the rotor blade employed by the solvers. The complete procedure consists of two main steps: the conservative mapping of CFD surface loading to the structural grid, and the conversion of the loading to the aerodynamic representation used in FLIGHTLAB. An additional step is required to adapt the loading azimuthally in case of steady CFD simulations. Pre-processing steps required to enable the use of SPLMOD on the CFD solution are described in Appendix B.

##### Spline-based Load Mapping

During each global coupling iteration, the converged CFD solution is exported as a data set describing the aerodynamic loading on the rotor surfaces over a full revolution. Using the volume spline interpolation that was previously described, the aerodynamic loading from the CFD is transferred to the structural mesh. The mapping transforms the loading on the vertices of the CFD mesh into a discrete representation on points of the structural mesh while conserving the total forces and moments. A comparison of the loading representations in ENSOLV and FLIGHTLAB is illustrated in Figure 3.9.

##### Azimuthal Reconstruction for Steady CFD Simulations

In case of steady CFD simulations, the aerodynamic solution describes only a single azimuthal position. However, FLIGHTLAB requires the aerodynamic loading to be defined over a full rotor revolution.


To obtain a full revolution of aerodynamic data, the steady aerodynamic force vector is rotated with the rotor. This is achieved using the rotation matrix about the axis of rotation of the rotor:

$$\mathbf{R}(\psi) = \begin{bmatrix} \cos(-\psi) & -\sin(-\psi) & 0 \\ \sin(-\psi) & \cos(-\psi) & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (3.17)$$

where  $\psi$  is the rotor azimuth. The aerodynamic force at a given azimuth is then computed as:

$$\mathbf{F}(\psi) = \mathbf{R}(\psi)\mathbf{F}_0, \quad (3.18)$$

This transformation ensures that the aerodynamic loading is consistent with the rotation of the rotor in the inertial frame. This transformation is applied separately to obtain a full revolution of loading for each blade.



**Figure 3.9:** Aerodynamic loading distribution on the rotor blade. Force distribution from ENSOLV in red. Force representation in FLIGHTLAB in blue.

### Conversion to FLIGHTLAB Representation

The aerodynamic loading in FLIGHTLAB is represented exclusively at the aerodynamic control point (ACP) locations. Therefore, for each radial section  $j$ , the forces originally acting at the leading edge, trailing edge, and ACP are combined into a single resultant force applied at the ACP. The equivalent force at the ACP is given by

$$\vec{F}_j^{\text{ACP}} = \vec{F}_{\text{LE},j} + \vec{F}_{\text{TE},j} + \vec{F}_{\text{ACP},j}. \quad (3.19)$$

To preserve the moment of the original force system, an equivalent moment about the ACP is computed for each radial section. Defining  $\vec{r}_{\text{LE}/\text{ACP},j}$  and  $\vec{r}_{\text{TE}/\text{ACP},j}$  as the position vectors from the ACP to the leading-edge and trailing-edge force application points, respectively, the equivalent moment about the ACP is

$$\vec{M}_j^{\text{ACP}} = \vec{r}_{\text{LE}/\text{ACP},j} \times \vec{F}_{\text{LE},j} + \vec{r}_{\text{TE}/\text{ACP},j} \times \vec{F}_{\text{TE},j}. \quad (3.20)$$

The force originally acting at the ACP does not contribute to the moment about the ACP, since its moment arm is zero. In this way, the original CFD force system is conservatively transferred to the ACP location while preserving both the resultant force and the equivalent moment about the ACP, consistent with the FLIGHTLAB aerodynamic representation.

## 3.5. Implementation of the Delta-Airloads in FLIGHTLAB

This section describes how the aerodynamic correction based on the CFD solution is implemented within FLIGHTLAB, and how it is incorporated into the trim process in a consistent manner. To enable this, FLIGHTLAB is extended with an additional aerodynamic load component, and the trim procedure is modified to facilitate the application of the correction.

### 3.5.1. Introduction of the AERODELTA Component

To enable the application of the aerodynamic correction within FLIGHTLAB, an additional force component is introduced, referred to as AERODELTA. This component serves as the interface through which the CFD airloads are supplied. Physically, this component is applied at the aerodynamic control points (ACPs) that are part of the discrete representation of the rotor geometry, which are elaborated on in Section 4.2.1.

The function of AERODELTA is twofold. First, it provides a dedicated interface for inputting the CFD airloads. Second, it enables the consistent computation of the difference between the CFD solution and the baseline airloads in FLIGHTLAB.

Once the aerodynamic loading from CFD has been exported to FLIGHTLAB and reconstructed at the ACP location, it can be used together with the baseline FLIGHTLAB airloads to determine the

delta-airloads correction.

### 3.5.2. Computation of Delta-Airloads

The delta-airloads represent the aerodynamic corrections that are applied to FLIGHTLAB and are defined as the difference between the aerodynamic loading from the CFD and the baseline airloads predicted internally in FLIGHTLAB. The computation of this correction is performed in the local blade frame, which is the reference frame in which FLIGHTLAB evaluates and applies the aerodynamic loading. This requires the transformation of the CFD loading from the inertial frame to the local blade frame, an operation that is handled by FLIGHTLAB internally. The reference frames are discussed in Section 3.3.

For a given blade element, the delta-airloads correction at an inner trim loop iteration  $k$  is defined as

$$\delta\vec{F}_{k-1} = \left( \vec{F}_{k-1}^F - \vec{F}_{k-1}^S \right), \quad (3.21)$$

where  $\vec{F}_{k-1}^F$  denotes the aerodynamic loading from CFD and  $\vec{F}_{k-1}^S$  the baseline airloads from FLIGHTLAB at the previous inner-loop iteration. Both quantities are expressed in the local blade frame. While the total aerodynamic loading per global iteration remains constant in the inertial frame, the aerodynamic loading  $\vec{F}_{k-1}^F$  is expressed each inner iteration in the updated local blade frame.

The delta-airloads are then applied as a correction to the aerodynamic loads in FLIGHTLAB through

$$\vec{F}_k = \vec{F}_k^S + r_k \delta\vec{F}_{k-1}, \quad (3.22)$$

where  $r_k$  is the relaxation factor applied in the  $k$ -th inner trim loop iteration,  $\vec{F}_k^S$  is the aerodynamic loading that FLIGHTLAB computes internally in the  $k$ -th inner loop, and  $\vec{F}_k$  is the resulting aerodynamic loading, including the correction, used by FLIGHTLAB to re-trim the rotor. At  $k = 1$ , the inner loop is initialized and  $\vec{F}_{k-1}^S$  equals the trimmed rotor solution from the previous global coupling iteration.

### 3.5.3. Modified Trim Procedure

Once CFD airloads have been used to determine the correction to the baseline airloads in the local blade frame, the intent is to produce an updated trimmed rotor in the next global iteration. The correction and subsequent trim are performed through an iterative loop that is illustrated in Figure 3.10.

At the global iteration  $i + 1$ , the inner trim loop is initialized using the rotor motion and CFD loading ( $\vec{F}_i^F, \vec{M}_i^F$ ) from the previous iteration  $i$ . Prior to starting the loop, the maximum number of iterations is specified; the convergence tolerance is set to ensure consistency with the chosen trim targets; and relaxation parameters are applied. To promote progressive trim convergence, a relaxation factor is applied to the delta-airloads correction during a specified number of initial iterations to allow for the gradual introduction of the delta-airloads. It is customary to start with a relatively low relaxation factor and gradually increase it to apply the full delta-airloads in subsequent iterations. This approach avoids abrupt increases in aerodynamic loading that could destabilize the trim. A minimum number of iterations is also specified to ensure that the final trim solution has applied the full correction.

At the start of each iteration of the loop, the CFD airloads are transformed into the local blade coordinate system. Subsequently, the delta-airloads are determined by computing the difference between the CFD and baseline forces and moments at each ACP. These differential airloads are saved and applied to the rotor through the AERODELTA component.

With the updated loading applied to the rotor, FLIGHTLAB proceeds with the trim procedure to determine the updated blade motion and control states of the rotor. After each trim iteration, the resulting rotor forces and moments are evaluated against the chosen trim targets. If the results lie within the prescribed tolerances, the loop is exited and the rotor motion and control states are considered to be converged. The updated rotor motion is then exported and shared as input for the next CFD simulation. If convergence is not achieved, the loop continues until the maximum number of iterations is reached.

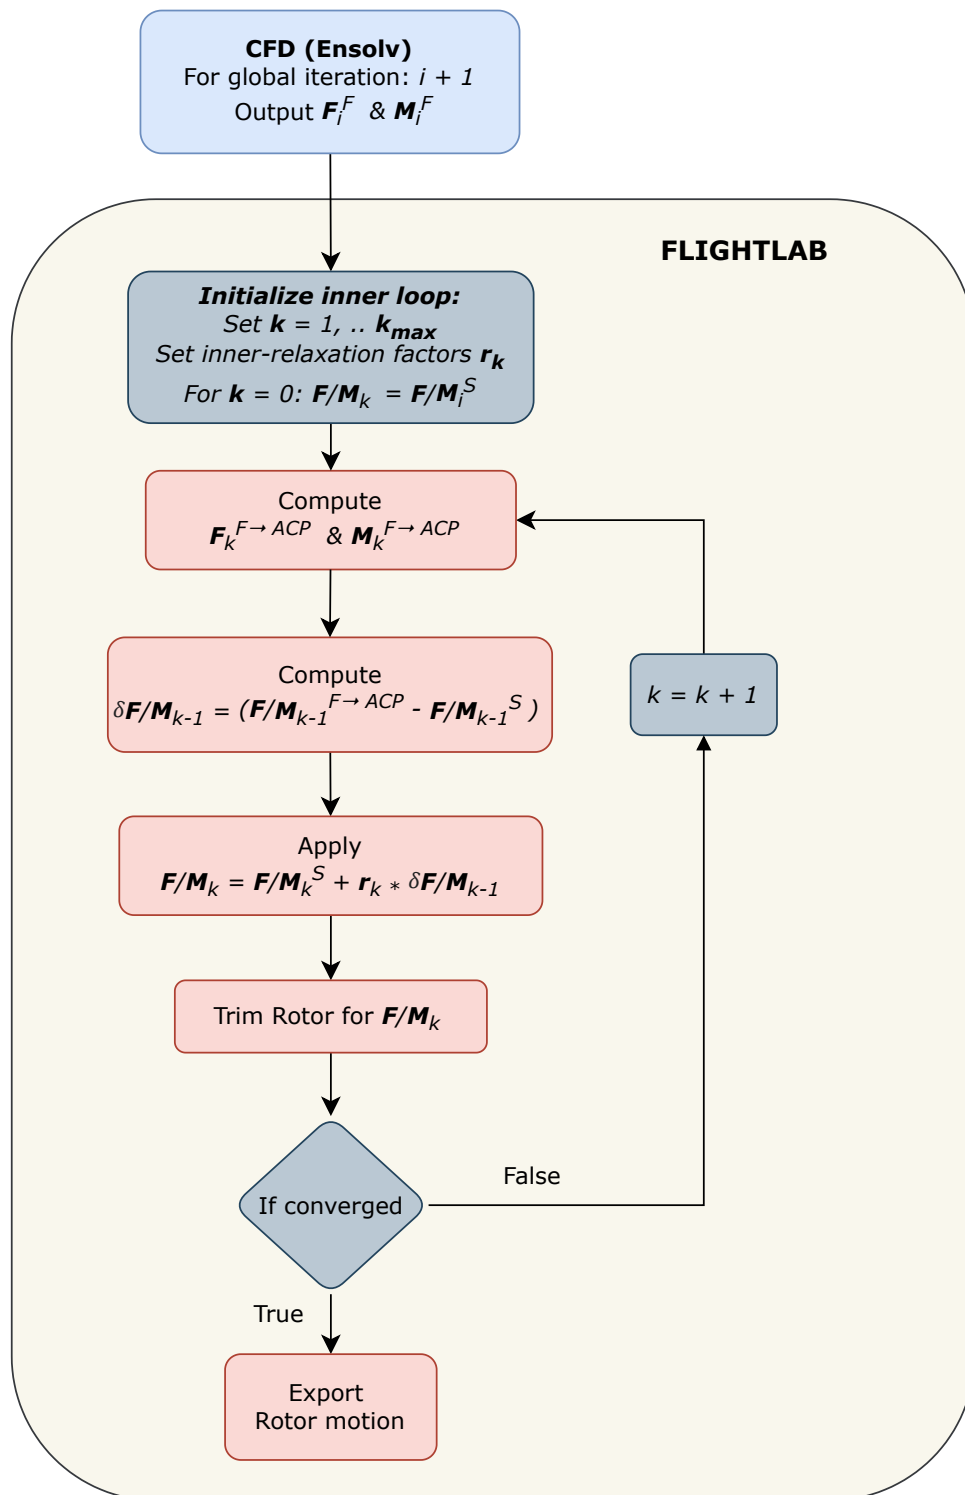


Figure 3.10: Inner loop: Flowchart of the computation and application of the delta-airloads in FLIGHTLAB.

### 3.5.4. Relaxation Strategy

The relaxation factor is introduced to improve the numerical stability of the inner trim loop. Rather than applying the full delta-airloads in a single step, the choice is made to progressively apply the delta-airloads over successive trim iterations. Given that large differences in the aerodynamic loading may be observed, this gradual mitigation decreases the risk of divergence and improves the robustness of the iterative trim process.

In FLIGHTLAB, the relaxation factor is defined as a linear ramp from the initial value  $r_0$  to unity over the first  $N_{\text{relax}}$  inner-loop iterations,

$$\mathbf{r}_k = \begin{cases} r_0 + (1 - r_0) \frac{k - 1}{N_{\text{relax}} - 1}, & k \leq N_{\text{relax}}, \\ 1, & k > N_{\text{relax}}. \end{cases} \quad (3.23)$$

The initial relaxation value  $r_0$  and the number of iterations  $N_{\text{relax}}$  over which the relaxed corrections are applied are set by the user a priori. Relaxation may be disabled entirely if preferred. The trim loop is configured to perform at least  $N_{\text{relax}} + 1$  iterations to ensure that the full corrections are applied. An additional outer-loop relaxation factor may be employed to scale the total correction at the global coupling iteration level.

# 4

## Models

This chapter describes the modelling approach of the partitioned coupling framework. It introduces the rotor geometry, the structural and aerodynamic modelling in FLIGHTLAB, and the setup of the CFD solver ENSOLV.


### 4.1. Rotor geometry

The rotor geometry analysed in this work is based on the internal NLR representation of the Mach-scaled BO105 main rotor model developed within the HeliNOVI programme [37]. The model is geometrically and dynamically scaled such that the key aeroelastic properties are preserved and accurately representative of the aeroelastic behaviour at model scale. Figure 4.1a shows the wind tunnel model that is the basis of the rotor geometry that is used in this research.

The rotor consists of four hingeless blades with a radius of  $2\text{ m}$  and a constant chord length of  $0.121\text{ m}$ , which corresponds to a solidity of  $0.077$ . The blade planform is rectangular with a constant chord along the span. A root cut-out of  $0.35\text{ m}$  is applied to the model to represent the non-lifting inboard section of the blade. The blades have a linear twist rate of  $-8^\circ/R$  and a root twist of  $5.6^\circ$ , resulting in a blade twist distribution as shown in Figure 4.1. The pre-cone angle of the rotor blade, the initial upward inclination, is  $2.5^\circ$  relative to the plane of rotation.

The blade cross-section is defined by a NACA 23012 airfoil whose trailing edge had been modified by the addition of a  $5\text{ mm}$  long tab in order to match the geometry of the full scale rotor. The rotor rotates counterclockwise when viewed from above. The nominal rotational speeds of  $1050\text{ RPM}$  is used in the simulations, which is consistent with the scaled operating conditions. Lastly, the rotor shaft tilt is set to  $0^\circ$  in order to align the axis of rotation with the vertical axis of the inertial reference frame.

The rotor geometry is analysed in isolation, and a fuselage, hub or wind tunnel support are not included. This allows the aerodynamic and aeroelastic behaviour of the rotor to be studied independently. It also allows the overall complexity of the problem at hand to be reduced and for the results to reflect the performance of the rotor system itself. This is consistent with the objectives of the present work, where the focus lies on the verification of the CFD–CSD coupling rather than full configuration aerodynamics.



(a) BO105 model main rotor geometry.[37]




Fig. 3: Blade twist distribution

(b) BO105 model blade twist distribution

Figure 4.1: BO105 model main rotor geometry and blade twist distribution.[37]

The following table presents an overview of key main rotor data.

Table 4.1: Mach-scaled BO105 main rotor key data

Property	Main rotor
Rotor type	hingeless
Number of blades	4
Nominal rotor speed	1050 rpm (220 m/s)
Direction of rotation (view from top)	Anti-clockwise
Rotor radius	2 m
Radius scale factor	2.455
Tip Mach number	0.64
Chord	0.121 m
Root cut-out	0.35 m
Solidity	0.077
Pre-cone angle	2.5°
Linear twist rate	-8°/R
Airfoil	NACA 23012
Shaft tilt (aft)	0°

## 4.2. FLIGHTLAB Model Set-up


This section describes the modelling choices for the implementation of the loose coupling in FLIGHTLAB. The intent is to provide a physically representative structural dynamics solution that can be coupled to CFD. The rotor modelled in FLIGHTLAB is the isolated rotor system described in the previous section.

### 4.2.1. Structural discretisation model

In FLIGHTLAB, the rotor structure is discretized using 18 non-linear beam elements (finite elements). Near the root and the tip, the spacing between the beam elements is decreased in order to more accurately capture blade structural behaviour. The structural properties of the blade are based on the reference Bo105 model from the HeliNOVI programme. This includes the representative mass and stiffness distributions for the scaled rotor system.

Each beam element represents the three translational degrees of freedom (axial, lead-lag, flap), and, in the case of a flexible rotor blade, the corresponding three rotational degrees of freedom. The element also accounts for the elastic and inertial loading that arise from the translational and angular accelerations of the rotor, and account for the resulting centrifugal stiffening and Coriolis effects and the subsequent structural response. The full blade span, including the root cut-out region, is structurally discretised.

The distribution of beam elements along the span of the rotor blade is illustrated in figure 4.2.



**Figure 4.2:** Definition of the structural and aerodynamic elements on the rotor blade in FLIGHTLAB.

The Bo105 rotor is a hingeless configuration, meaning that the rotor blades are attached to the hub without discrete flapping or lead-lag hinges and instead via a rigid connection at the root of each rotor blade. As a result, the primary blade degrees of freedom arise from structural elasticity, with flapwise bending, lead-lag bending, and torsional deformation governed by the blade structural properties. This differs from articulated rotor systems, where hinges allow independent flapping, lead-lag, and feathering motions. In the case of a rigid rotor model in FLIGHTLAB, flapping and lead-lag motion cannot occur and only prescribed pitching (feathering) motion is possible.

During simulations the choice can be made to structurally represent the rotor as a flexible or rigid system. In the rigid configuration, flapping and lead-lag motion are suppressed and only prescribed pitching (feathering) motion is retained. In the flexible configuration, both rigid and elastic deformation are allowed.

### 4.2.2. Aerodynamic Model

The aerodynamic model in FLIGHTLAB is based on a lifting-line blade element formulations, in which the sectional airloads are taken from airfoil lookup tables. These lookup tables are derived from the HeliNOVI database, and are based on experimental measurements using the NACA 23012 airfoil for varying Mach Number and angle of attack.

A quasi-unsteady formulation is employed, but no additional stall delay or dynamic stall model is used. Along the span, the airfoil characteristics are assumed to be uniform, and no three-dimensional rotational corrects are applied. The baseline aerodynamic model therefore utilises a simplified two-dimensional sectional approach, and spanwise effects are only introduced through the inflow model.

The aerodynamic forces in FLIGHTLAB are determined at discrete aerodynamic computation points (ACP), also referred to as blade collocation points. These point define where the induced velocities and local flow contritions are computed. This rotor model is discretized into 11 ACPs. They are located on the lifting section of the rotor blade, not on the root cut-out, and their spanwise locations along the blade correspond to the endpoints of the structural elements. Along the blade section, the ACPs for this rotor model are located at the quarter-chord position. The locations of the ACPs along the span are shown in figure 4.2.

### 4.2.3. Numerical discretisation

Modelling the rotor in FLIGHTLAB is done by performing multiple rotor revolutions until a converged, periodic steady-state trim solution is reached, meaning that the aerodynamic loads and blade motion repeat from one revolution to the next. This periodic solution represents the trimmed operating condition of the rotor. These revolution are performed using an azimuthal resolution of  $1^\circ$ , corresponding to 360 time steps per rotor revolution. In terms of convergence criteria, the trim procedure is performed until the trim target

### 4.3. ENSOLV

This section describe the modelling choices in ENSOLV. This includes the flow and turbulence modelling, the temporal and spatial discretisation, the computational domain and mesh. These choices are made to ensure the description and analysis of the rotor is consistent with FLIGHTLAB while maintaining a computationally efficient setup.

#### 4.3.1. Flow and turbulence modelling

The flow is modelled using three-dimensional compressible Reynolds-averaged Navier-Stokes (RANS) equations and adapted for both steady and unsteady simulations as discussed in Section 3.2.2. Steady simulations are performed using pseudo time stepping, whereas the unsteady simulations are performed in an unsteady URANS framework using dual-time stepping, allowing the time-accurate resolution of the rotor flow field.

The spatial discretization of governing equations is based on second-order central differencing with artificial dissipation for numerical stability. In addition, a fourth-order finite-volume scheme is utilised to the spatial accuracy. This high-order method prevents unphysical dissipation of (smaller) turbulent structures inside massively separated flows. As such, the small-scale vortical structures are preserved, which is essential for capturing rotor wake dynamics and blade-vortex interactions [38].

Turbulence is modelled using the three-dimensional TNT Explicit Algebraic Reynolds Stress Model (EARSM). This model accounts for the anisotropy of the Reynolds stresses and is computationally efficient within a RANS framework..

#### 4.3.2. Temporal discretisation

For steady simulation, convergence of the solution is assessed based on the reduction of the residuals in the computational domain. The residuals were found to no longer decrease significantly after a minimum of 3500 pseudo-time iterations, discussed in Appendix A. The number of pseudo-iterations for the steady simulations is therefore set to between 3500 – 4000.

For unsteady simulations, the physical time-step is chosen such that the CFD solution is consistent with the azimuthal discretisation of FLIGHTLAB. In this context, an azimuthal resolution of  $1^\circ$  is used, which corresponds to 360 physical time-steps per rotor revolution. The rotor speed of 1050, rpm corresponds to

$$f = \frac{1050}{60} = 17.5, \text{ rev/s}, \quad (4.1)$$

so that the period of one rotor revolution is

$$T = \frac{1}{f} = \frac{1}{17.5} \approx 5.714 \times 10^{-2}, \text{ s}. \quad (4.2)$$

The physical time-step is therefore

$$\Delta t = \frac{T}{360} \approx 1.59 \times 10^{-4}, \text{ s}. \quad (4.3)$$

Therefore, in the present coupling strategy, one CFD solution is obtained per FLIGHTLAB azimuthal step. This is not a strict requirement for the CFD solver, since smaller azimuthal increments can also be used to support the wake development, but it avoids an additional interpolation or resampling operation when transferring the aerodynamic loads to FLIGHTLAB.

Within each physical time-step, the dual-time stepping procedure performs 40 inner pseudo-time iterations. Relaxation within pseudo-time is carried out using a 5-stage Runge-Kutta scheme with a CFL number of 4.5. This ensures that the solution is iteratively converged at each azimuthal position before advancing to the next time-step.

In addition, for the unsteady simulations the number of revolutions is set to include multiple rotor revolutions. This allows the flow field to develop and adjust to the rotor motion, as well as reduce the residuals in the computational domain. However, the analysis of the residuals in the unsteady CFD simulations in Appendix A showed that the residuals fail to reduce in the computational domain when

the number of revolutions is increased. This is a reported issue in the computational domain that was used to perform these CFD simulations, and calls for further analysis. This is outside of the scope of this work.

### 4.3.3. Domain

The CFD models a full four-bladed rotor rather than a single-blade sector with periodic boundary conditions. This allows the complete rotor wake to be represented and allows for the assessment of asymmetric flight conditions. This rotor is modelled in the computational domain that consists of a spherical far-field region with a radius of  $R = 50$  m measured from the origin, as shown in Figure 4.3. The large domain size is chosen to minimize artificial reflection or blockage effects, and allow for the proper development of the wake. To match the orientation in FLIGHTLAB, the rotor geometry is oriented such that positive vertical displacement is along the negative  $z$ -axis.




Figure 4.3: Spherical computational domain with radius  $R = 50$  m.

The motion of the rotor is represented using a rotating and deforming mesh. The main rotation of the rotor is applied as rigid-body rotation about the  $-z$ -axis at the prescribed angular velocity. The blade deformation that is obtained from FLIGHTLAB is imposed on the CFD mesh as discussed in Section 3.4.3. Both the rotor geometry and the inner mesh move, while the far-field boundary remains stationary, resulting in a moving inner mesh embedded within a non-rotating outer domain. The governing equations are solved in an inertial reference frame and not in the rotating body frame.

### 4.3.4. Boundary conditions

The rotor blades are modelled as solid walls and a no-slip condition is prescribed with an adiabatic wall temperature boundary condition. At the outer boundary of the spherical domain, pressure-based far-field boundary conditions are prescribed using the free-stream conditions defined in the simulations. These boundaries are defined in the inertial reference frame, allowing the wake to develop naturally without artificial rotation of the external flow field.


### 4.3.5. Computational grid

The computational domain is discretised using a structured multi-block grid and consists of approximately 5000 structured blocks and 3 million cells. The grid is refined in regions where strong gradients occur, and includes the blade root and blade tip regions. Overall, the grid is progressively coarsened toward the outer spherical boundary, where flow gradients are small. Figure 4.4 shows a grid slice near the blade tip, visualising the grid refinement. Figure 4.5 shows a top view of the rotor and surrounding multi-block grid topology.

The grid is also used for the imposed blade motion and elastic deformation. The displacements of the blade surface grid points are prescribed from the mapped structural deformation. These surface

displacements are then propagated to the remaining grid by applying the ENFLOW grid-deformation method: first, the displacements of block vertices and edges outside the solid geometry are computed using a volume-spline method, after which the displacement of the interior grid points is obtained using transfinite interpolation. In this way, the blade deformation is imposed on the CFD surface mesh and extended into the volume mesh.

For this research, a relatively coarse grid level was chosen. This is a deliberate modelling choice, since the primary objective is to verify the implementation of the loose coupling framework rather than to obtain a fully grid-converged aerodynamic solution. No grid-refinement study was performed for the present coupled simulations, and therefore no quantitative discretisation error estimate is provided. The resulting aerodynamic loads should therefore be interpreted as suitable for assessing the coupling procedure, while further grid refinement, particularly near the wake and tip-vortex regions, is required before drawing conclusions on grid-converged aerodynamic accuracy.



**Figure 4.4:** Local grid resolution near the blade tip region.

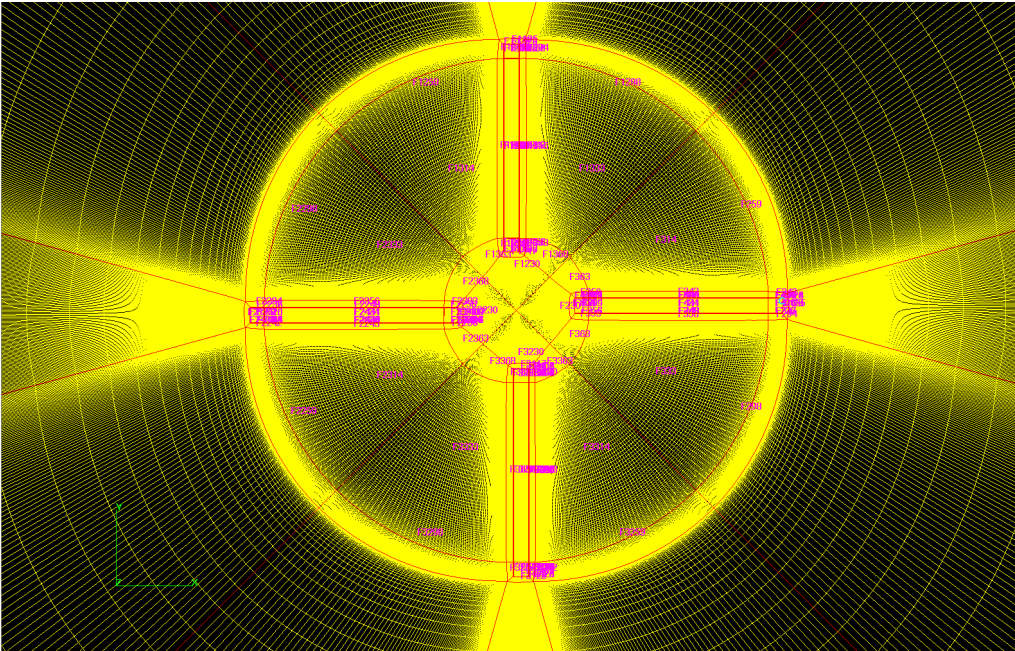


Figure 4.5: Top view of the rotor and structured multi-block grid topology.

# 5

## Results

This chapter presents the verification and initial validation of the loose CFD-CA coupling framework analysed in this work. The intent of this chapter is to assess whether the utilised delta-airloads method is capable of producing numerically stable, physically consistent and practically useful coupled rotor solutions.

The results are ordered based on increasing complexity. First, the chapter describes the design of experiment and the associated test matrix, followed by verification cases for simplified hover flight. This is extended to more demanding cases to assess the influence of modelling choices, such as steady versus unsteady CFD and rigid versus flexible blade behaviour. Lastly, the coupled results are compared with the HART II experimental database to provide an initial assessment of the model behaviour.

### 5.1. Design of Experiment

The set of simulations is designed to assess the partitioned CFD-CA framework in a structured and progressive manner. The primary goal is to verify that the transfer of blade deformations from FLIGHTLAB to ENSOLV, and the transfer of CFD load corrections back to FLIGHTLAB, are implemented correctly within the delta-airloads approach. This is assessed by monitoring the convergence of the coupled solution, the numerical stability of the coupling iterations, and the resulting changes in rotor loads and blade response. In addition, the simulations provide an initial comparison with experimental data for a representative rotorcraft flight condition.

The test matrix is therefore organised in increasing order of physical and numerical complexity. The simulations begin with simplified hover cases used for verification, followed by controlled extensions in modelling fidelity, and end with validation against the baseline descending-flight from the HART II database. The progression ensures that the fidelity of the implementation is established before the coupled results are interpreted physically.

Throughout this chapter, the coupled framework is assessed using the following four criteria:

1. verification of the transfer of aerodynamic information between the solvers,
2. numerical stability and convergence of the global coupling iterations,
3. assessment of the coupled trim state and aerodynamic response, and
4. comparison with experimental data for the validation case.

The comparison with the HART-II experiment must be interpreted within the limitations of the present CFD set-up. As discussed, a relatively coarse mesh was chosen to decrease the computational costs and simulation time. The reduced resolution of the CFD mesh will affect the extent to which the numerical results can be directly compared to the experimental measurements. The comparison is used therefore an initial validation-oriented assessment, and is included to evaluate whether the coupled framework produces physically reasonable load distributions, captures the main trends of the experimental data, and remains numerically stable for a representative rotorcraft flight condition.

### 5.1.1. Analysis Strategy

The intent of the design of experiment is to verify and validate the loose coupling framework based on the delta-airloads method. Initial attempts by NLR to apply the coupling to an isolated flexible rotor in forward flight using unsteady CFD resulted in unstable behaviour and unphysical responses. These early results indicated that the complete problem was too complex to directly identify the capabilities and limitations of the implementation. This motivated the use of an incremental verification and validation strategy.

The simulations are therefore organised by progressively increasing the complexity of the problem. First, simplified hover cases are used for verification. The cases are then extended to time-resolved aerodynamic simulations, forward flight and structural flexibility. Finally, the predictive capabilities of the loose coupling are assessed by comparing the results with experimental data. This allows implementation errors to be identified before the remaining discrepancies are considered as numerical or modelling limitations.

This strategy is supported by findings from literature, which indicate that aerodynamic modelling is often a dominant source of error in loosely coupled CFD-CA methods. Therefore, convergence of the loose coupling alone does not demonstrate the physical correctness of the results. A coupled solution may converge smoothly while still being inaccurate due to limitations in the CFD solution, for example related to turbulence modelling, wake resolution and sectional pitching moment accuracy.

To support this analysis strategy, a set of quantities is collected to assess convergence, physical consistency and the solver-to-solver exchange of information. The analysis focuses primarily on aerodynamic loads, delta-airloads and trim behaviour, since these quantities are central to the coupling methodology. From ENSOLV, the aerodynamic forces and moments on the rotor surface are extracted together with supporting data such as the solution convergence and wake structure. From FLIGHTLAB, the baseline aerodynamic loads, the applied delta-airloads and the trim variables are recorded. The structural motion is considered only in terms of blade pitch and flapping motion, since a detailed analysis of the structural response is outside the scope of this research.

### 5.1.2. Verification Test Cases

The verification phase focuses on assessing whether the delta-airloads method is implemented correctly between FLIGHTLAB and ENSOLV. The baseline flight condition that is selected for verification is hover flight. These hover test cases are used to establish confidence in the numerical behaviour and implementation of the coupling framework.

In hover, the rotor has no velocity relative to the air, either vertical or horizontal. Hence, the rotor motion and aerodynamic loading are periodic, and the rotor flow field is time-invariant in the rotating frame. As a result, the dynamics and aerodynamics of the rotor are greatly simplified. Hover therefore is the ideal load case to start assessing the correctness of the coupling methodology.

A further advantage of hover is that it is inherently steady and allows for the use of steady CFD simulations resulting in reduced computational costs and faster coupling iterations. The use of steady CFD simulations does require the adaptation of transfer of aerodynamic loads between ENSOLV and FLIGHTLAB in the current implementation. Since a steady CFD solution provides a single instance of the aerodynamic loads that is azimuthally invariant, as opposed to an unsteady CFD solution that contains multiple time-steps of aerodynamic data over a rotor revolution.

Simulating a true hover flight in CFD is not practical, because in the absence of inflow, the wake is not convected away from the rotor. Therefore, a small vertical inflow of 1  $m/s$  is prescribed, which corresponds to a slow vertical climb. This alteration allows the wake to convect enough in order to form a steady state solution that corresponds to hover flight. This inflow conditions is added to ENSOLV as well as FLIGHTLAB to ensure consistency.

To further reduce the complexity of the problem, the first verification cases consider a rigid rotor. By only allowing rigid deformation of the rotor blades, the influence of the structural model is removed and with it the coupled aeroelastic effects. For this reason, the verification can focus on the implementation and numerical behaviour of the methodology.

Following the successful execution of the rigid rotor cases the flexible deformation of the structure is

reintroduced in a limited manner. Rigid and flexible deformation are enabled in FLIGHTLAB, while in ENSOLV the steady simulation is performed using a mean deformed rotor shape. This approach allows the behaviour of the aeroelastic effects to be studied in FLIGHTLAB, without sacrificing the use of steady CFD. As a result, the complexity of the problem is extended with the addition of a flexible structure while maintaining low computational costs.

Finally, to assess the influence of the operating condition additional hover cases are analysed for which the rotor thrust is varied. These cases highlight the behaviour of the coupling for different loading conditions and the sensitivity to increased aerodynamic loading. The selected thrust levels are based on cases used in the HeliNovi project [37] to ensure representative operating conditions.

### 5.1.3. Extension to Unsteady CFD

Following the verification of the coupling framework in steady hover flight, the simulations proceed by increasing the complexity of the framework and problem in a controlled manner. These extensions intended to assess whether the loose coupling remains stable and physically consistent when the modelling assumptions are extended brought closer to the intended application.

The first extension to the loose coupling implementation is the transition from steady to unsteady CFD simulations. This approach provides FLIGHTLAB with a full revolution of aerodynamic data rather than a single steady-state solution. When applied to the rigid rotor hover flight scenarios simulated during verification the physical solution is expected to remain unchanged due to the steady-state of hover. Instead, unsteady CFD is used in order to isolate the influence of the time resolution of the aerodynamic solution, any differences in the behaviour of the CFD simulations or any coupling-related effects. For a flexible rotor, the transition to unsteady CFD is expected to have a more direct physical effect. In that case, the aerodynamic loading over a full revolution is able to interact with the periodic blade motion, leading to additional aeroelastic effects.

After the addition of unsteady CFD, the experimental campaign is further extended to forward flight conditions. While maintaining periodicity over a full revolution, forward flight introduces azimuthal asymmetry in the aerodynamic loading and structural response of the rotor. These conditions therefore evaluate the ability of the coupling to handle an azimuthally changing delta-airloads correction and asymmetric flow field. These forward flight cases are therefore used as the first assessment of the loose coupling under time-varying aerodynamic condition. In line with the previous experiments, these forward flight simulations are performed for using both the rigid and flexible rotor so that the effect of aeroelasticity can be assessed separately from the effect of the unsteady flowfield.

### 5.1.4. Validation Case Selection

The final step of the experimental campaign, the validation, is performed after the verification has demonstrated that the coupling is stable and consistently able to simulate the intended flight scenario. Validation focusses on assessing the predictive capability of the framework and requires a comparison with experimental data that is representative of the coupled rotor dynamics and aerodynamics. The results of validation should indicate whether the framework suits the needs of NLR and where the predictive capabilities could be improved.

For this purpose, the HART II (Higher Harmonic Control Aeroacoustic Rotor Test) experimental campaign is selected as the main validation case. HART II was conducted in the German-Dutch Wind Tunnel (DNW) open-jet facility and was specifically designed to provide an extensive, multi-disciplinary data set for the validation of aeroelastic rotor simulations in descending forward flight. This flight condition is characterised by strong blade-vortex interactions, which cause pronounced unsteady aerodynamic effects and significant coupling between structure and flow.[22]

The baseline descending flight condition of HART II is selected as the reference validation case. The strong blade-vortex interactions that occur during this flight condition make this data set particularly suitable for the validation of coupled simulations. The database includes synchronised measurements of sectional airloads at  $r/R = 0.87$  and elastic blade deformations, enabling direct comparison with the coupled CFD-CA simulations. The choice of the HART II is further motivated by the fact that the experimental data set has also been extensively used as a validation benchmark to assess coupled CFD-CA approaches. An additional advantage of the HART-II campaign is that it utilises the same

rotor geometry that is used by NLR.

### 5.1.5. Summary of Test Matrix

The complete set of test cases analysed as part of this research is summarised in Table 5.1. The matrix describes the flow conditions, rotor configuration, trim targets, CFD set-up and the intended objective of each case. The test cases are ordered to reflect the progressive increase in complexity from verification to validation. This ordering highlights the importance of each case to the overall methodology.

**Table 5.1:** Test matrix for verification and validation cases used in the CFD-CA coupling study.

Case ID	$V_\infty$ [m/s]	AOA [°]	Rotor pitch [°]	Hub $F_z$ [N]	Hub $M_x$ [Nm]	Hub $M_y$ [Nm]	D.o.F.	CFD	Flight Type
C100-RGD-STD	1	-90	0	3000	0	0	Rigid	steady	Hover
C400-FLX-STD	1	-90	0	3000	0	0	Flex	steady	Hover
C100-RGD-UNS	1	-90	0	3000	0	0	Rigid	unsteady	Hover
C400-FLX-UNS	1	-90	0	3000	0	0	Flex	unsteady	Hover
C102-RGD-STD	1	-90	0	6000	0	0	Rigid	steady	Hover
C402-FLX-STD	1	-90	0	6000	0	0	Flex	steady	Hover
C102-RGD-UNS	1	-90	0	6000	0	0	Rigid	unsteady	Hover
C402-FLX-UNS	1	-90	0	6000	0	0	Flex	unsteady	Hover
C851-RGD-UNS	12.3	0	0	3774.7	-33.1	132.4	Rigid	unsteady	Forward flight
C851-FLX-UNS	12.3	0	0	3774.7	-33.1	132.4	Flex	unsteady	Forward flight
C929-RGD-UNS	32.9	0	4.5	3300	-20	-20	Rigid	unsteady	Descending flight
C929-FLX-UNS	32.9	0	4.5	3300	-20	-20	Flex	unsteady	Descending flight

The aerodynamic inflow conditions are determined by the free-stream velocity, angle of attack and rotor pitch. The free-stream velocity, angle of attack and the rotor pitch follow the sign convention used by FLIGHTLAB. The free-stream velocity  $V_\infty$  is defined as positive along the negative  $x$ -axis and the angle of attack is positive toward the negative  $x$ , negative  $z$ -plane. The rotor pitch, the pitch angle of the rotor hub, is defined as positive aft in FLIGHTLAB.

For the hover and forward flight cases, the rotor pitch is set to zero and the inflow condition is imposed via the angle of attack. This formulation ensure that inertial and hub reference frame align, and that the gravitational forces are correctly accounted for. In the descending flight case of the validation effort, the wind tunnel conditions are replicated by pitching the rotor by  $4.5^\circ$  while the angle of attack is set to 0.

The trim targets are specified in terms of the hub force and moments ( $F_x$ ,  $M_x$ ,  $M_y$ ) and correspond to the quantities that FLIGHTLAB uses to determine the equilibrium state of the rotor.

As a whole, this test matrix describes a structured and systematic assessment of the loose coupling implementation, that begins with controlled verification test and ends with the validations against experimental data.

## 5.2. Verification

Before assessing the predictive capabilities of the coupled framework, the numerical consistency of the coupling interface and the implementation of the delta-airloads are verified. This verification is divided into three parts. First, transfer of the aerodynamic loading from ENSOLV to FLIGHTLAB is assessed, including the conservation of the forces and moments. Second, the additional transformation that is required for the use of steady CFD simulations is verified. Third, the section is concluded with the verification of the computation and application of the delta-airloads in FLIGHTLAB.

The verification of the transfer of the rotor motion to the ENSOLV grid was qualitatively discussed in Section 3.4.3 and is not repeated here. The mapping of the rotor motion to the ENSOLV surface mesh was visually inspected to be consistent with the rotor motion exported by FLIGHTLAB. For future analysis, a dedicated utility that is able to quantify the mapping error is recommended.

### 5.2.1. Transfer of aerodynamic loading from CFD to CA

The transfer of the aerodynamic loading from ENSOLV to FLIGHTLAB for the computation of the delta-airloads is part of the interface between both solvers and therefore directly responsible for the consistency and accuracy of the coupling. During this transfer, it is crucial that the forces and moments are conserved. Any loss or artificial generation of loading will lead to inconsistencies between ENSOLV and FLIGHTLAB, and will comprise the physical validity of the coupled results.


The verification of this transformation is particularly important because in the original implementation a discrepancy was identified after transferring the loads to FLIGHTLAB. To quantify the impact of this error and to correct the implementation, a comparison of the aerodynamic loading before, during and after the transformation is presented in Table 5.2. The table compares the original CFD solution represented by total force  $F_z$ , intermediate solutions produced as part of the transformation, referred to as Centre to Vertex and SPLMOD, and the final loading at the ACPs in FLIGHTLAB. The details behind the operations of Centre to Vertex and SPLMOD are discussed in Appendix B and Section 3.4.2, respectively.

**Table 5.2:** Comparison of aerodynamic loading, represented by total force  $F_z$ , between CFD and FLIGHTLAB for C100 case.  $F_x$  and  $F_y$  are approximately 0 and thus omitted.

CFD estimate [N]	Centre to Vertex [N]	SPLMOD [N]	FLIGHTLAB [N]	Diff. [N]	Diff. [%]
2348	2430	2430	2430	82	+3.5
5241	5465	5465	5465	224	+4.3

The discrepancy in the data was identified to be caused by the operation Centre to Vertex, and was shown to cause an artificial increase in loading. The issue was found to be caused by the incorrect implementation of the transfer of the aerodynamic loading from the cell face centres to the cell vertices of the aerodynamic surface mesh, performed by the Centre to Vertex operation. This transformation was corrected for the present implementation. The verification of this operation are also described in Appendix B. The other operations were shown to conserve the total loading.

Based on the force total in Table 5.2, the transfer of the aerodynamic loading from the ENSOLV mesh to the FLIGHTLAB using the spline mapping was observed to be conservative. This can be further supported by the results of the transfer for an unsteady CFD simulation. Figure 5.1 compares the aerodynamic force  $F_z$  as a function of azimuth on the CFD and CA discretisations for Case C851.



**Figure 5.1:** Comparison of the aerodynamic solution from CFD and the spline-mapped aerodynamic loading on the CA discretization. Case C851: Forward flight & rigid rotor.

Figure 5.1 shows that the aerodynamic force  $F_z$  on the FLIGHTLAB grid closely follows the CFD solution over the full revolution. The two curves closely overlap, indicating that the integrated force is preserved by the spline mapping at each azimuthal position. This verifies the conservative transfer of the total aerodynamic loading.

The resulting distribution on the leading edge, trailing edge and ACP locations in FLIGHTLAB is then summed to the ACPs as shown in Figure 5.2. This representation shows how the aerodynamic loading from the CFD is expressed in accordance with FLIGHTLAB.

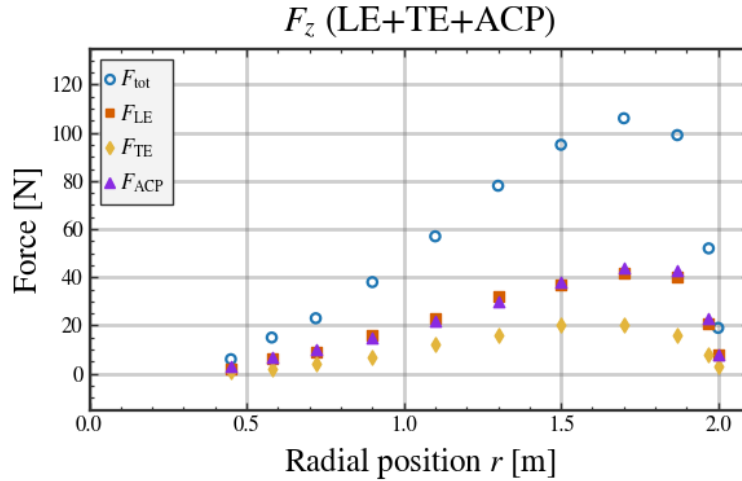


Figure 5.2: Discrete force  $F_z$  per ACP, determined by CFD and transformed to the representation used by FLIGHTLAB.

Taken together, these results show that the transfer of the aerodynamic loading from ENSOLV to FLIGHTLAB is conservative and consistent.


### 5.2.2. Alteration for Steady CFD

In the case of a steady CFD simulations, the computed airloads corresponds to a single instance in time and in space, and thus a single azimuth. However, FLIGHTLAB requires the aerodynamic loading to be defined for a full rotor revolution, and in terms of the set number of azimuthal positions. To obtain a full-revolution of aerodynamic data, the steady aerodynamic force vector is rotated with the rotor. The definition of the rotation matrix was described in Section 3.4.4.

The rotation only affects the in-plane force components ( $x, y$ ) while the  $z$ -component of the force remains unchanged. Once rotated, the  $x$ - and  $y$ -components of the force vary sinusoidally with the azimuth, whereas the  $z$ -component remains constant. This behaviour is consistent with the expected physics and visualised in figure 5.3.

Here, the azimuth on the horizontal axis follows the definition of FLIGHTLAB defined in Figure 3.3, and is zero when a blade is aligned with the negative  $x$ -axis. The four blades are shown to be identical in hover in the case of a rigid rotor. This behaviour confirms that the transformation of the steady CFD airloads into the rotational representation for FLIGHTLAB is implemented correctly.

**Case: 102 | loc: ac | r/R = 0.75 [-]**



**Figure 5.3:** Azimuthal variation of the aerodynamic force components obtained by rotating the steady CFD solution. Case C102, chosen location: ACP and at 75% span.

### 5.2.3. Verification of Delta-Airloads Implementation

The implementation of the delta-airloads correction was verified by assessing the evolution of the thrust produced by the four rotor blades. The correction is based on the difference between the CFD airloads and the FLIGHTLAB lifting-line airloads. The verification was performed for a hover case, where the CFD solution was steady and the FLIGHTLAB trim target was set to 3 kN. Including the rotor weight of 98 N, the expected total rotor thrust is therefore approximately  $3100 \pm 50$  N, based on the trim tolerances used in FLIGHTLAB. The thrust forces obtained during two successive global coupling iterations are given in Table 5.3, using the notation introduced in Section 3.5.

**Table 5.3:** Verification of delta-airloads for global coupling iterations. Total rotor thrust in z for a steady hover case with trim target of 3 kN.

Iteration	State	F [N]	$F_i^{LL}$ [N]	$F_i^{CFD}$ [N]	$\Delta F$ [N]
Iteration 1	Initial	-2349	-3058	-2349	709
	After trim	-3072	-3781	-2349	709
Iteration 2	Initial	-3023	-3781	-3023	758
	After trim	-3079	-3837	-3023	758

The table shows that the delta-airloads correction is computed correctly from the difference between the internal FLIGHTLAB loads and the CFD loads. In the first global iteration, the initial thrust difference of 709 N indicates that FLIGHTLAB overestimates the aerodynamic loading compared to ENSOLV. By applying the delta-airloads correction, the aerodynamic loading used by FLIGHTLAB is reduced from  $F_i^{LL}$  to  $F_i$ , to match the CFD value. After the correction is applied, FLIGHTLAB retrim the rotor to recover the required thrust. This can be seen from the increase in  $F_i^{LL}$  from  $-3058$  N to  $-3781$  N during the first iteration. The corrected thrust  $F_i$  then increases from  $-2349$  N to  $-3072$  N, which brings the rotor back to the intended thrust level.

The second global iteration confirms that the retrimmed rotor motion also corresponds to an increase in loading in the CFD solution. The initial CFD-based thrust increases from  $-2349$  N in the first iteration to  $-3023$  N in the second iteration. The new delta-airloads correction of 758 N is then applied, after which FLIGHTLAB again retrim the rotor and increases the aerodynamic loading toward the required thrust level.

These results verify the intended behaviour of the delta-airloads implementation which is to iteratively replace the FLIGHTLAB airloads with CFD airloads and use the trim procedure to recover the required thrust. As. At convergence, global coupling iterations are performed the trim constraints are satisfied

and the correction is expected to converge as the CFD aerodynamic forces become unchanged between iterations. At convergence, the corrected airloads used in FLIGHTLAB are equal to the CFD airloads, and no internal aerodynamic contribution is retained in the final coupled solution.

### 5.3. Rigid Rotor Simulations

The intended behaviour of the delta-airloads correction is to replace the internally computed aerodynamic loads of FLIGHTLAB with those obtained from the CFD. By repeating this process in successive iterations, the expected result is that the rotor motion output by FLIGHTLAB, when simulated in CFD, should reproduce the same airloads that FLIGHTLAB internally computes.

The successful verification of all the operations of the partitioned coupling should lead to the successful execution of the coupling implementation. This section discusses the global coupling behaviour of the partitioned CFD-CA coupling, specifically for the rigid rotor. Due to an error in loading the aerodynamic solution into FLIGHTLAB, the result of the C102 (6kN) case using unsteady CFD was corrupted. A lack of time meant the problem could not be addressed and this case is therefore omitted.

#### 5.3.1. C100: Hover (3 kN), Steady CFD, Rigid Rotor

This C100 steady rigid-hover case served as verification case of the delta-airloads coupling strategy. Figure 5.4 shows the global iteration history of the total force  $F_z$ . No outer relaxation factor was used to apply the CFD airloads. The force shown represents the instantaneous total rotor hub force. Due to the azimuthal invariance of steady hover, the instantaneous values are equal to the mean force over a full revolution. The contributions in  $x$  and  $y$  are omitted as they are approximately 0.




Figure 5.4: Global iteration history of the total vertical force  $F_z$  for the C100 rigid hover case (3 kN).

Figure 5.4 shows the evolution of four variables related to the hub  $F_z$  force across the global coupling iterations. The blue curve represents the thrust  $F^{CFD}$  computed by ENSOLV based on the rotor motion that was determined in the previous global iteration. The red curve indicates the trim target set in FLIGHTLAB, which for this case is set to  $3000 \pm 50$  N. The yellow curve represents the total aerodynamic thrust defined as the trim target plus the rotor weight, and is equal to  $F$ , the FLIGHTLAB loads corrected by CFD. The purple curve is the effective aerodynamic loading  $F^{LL}$  in FLIGHTLAB without the deduction by the delta-airloads. The effective aerodynamic loading should always equal to the trimmed aerodynamic thrust plus the delta-airloads correction. It must be noted that the blue curves

represents the CFD input at the given global coupling iteration, whereas the remaining curves represent either outputs or internally used quantities within FLIGHTLAB.

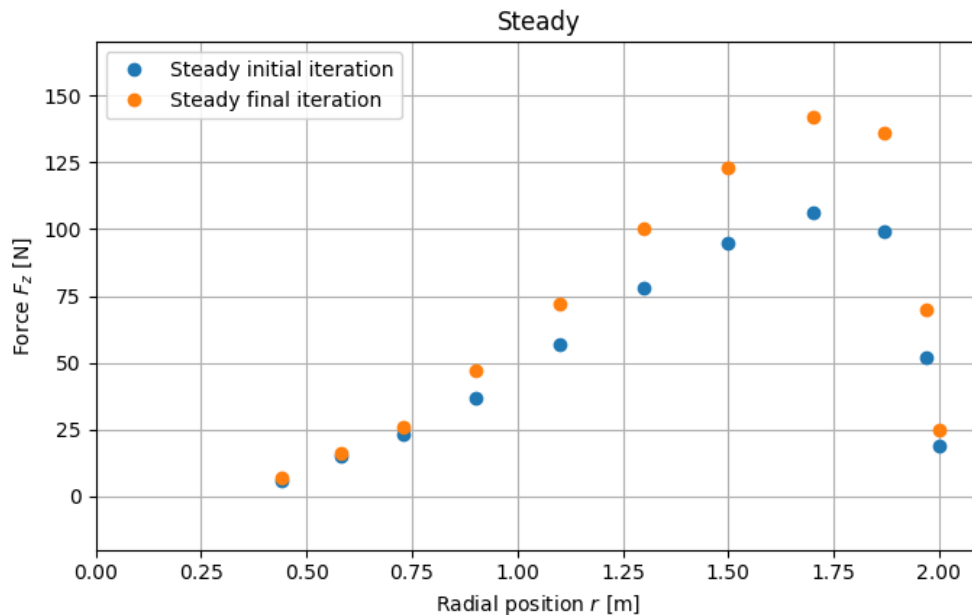
In the first global iteration, the thrust computed by CFD based on the initial rotor motion greatly undershoots the aerodynamic thrust predicted by FLIGHTLAB. The CFD force is approximately 2300 N, while FLIGHTLAB required a thrust of 3100 N. As a result, a correction of approximately 700 N is computed. This correction is applied to the aerodynamic loading in FLIGHTLAB, effectively reducing it by 700 N. Since the trim and thrust targets remain fixed at 3000 and 3100 N, FLIGHTLAB has to compensate by increasing the internally computed aerodynamic loading by the same amount in order to achieve a trimmed rotor solution. The resulting effective aerodynamic loading therefore lies 700 N above the original solution, and equals the sum of the internally computed airloads and the delta-airloads correction. Following the correction, the difference reduces rapidly and a converged solution is reached within five global iterations.

Table 5.4 summarises the evolution of the collective pitch and the  $F_z$  correction. The cyclic pitch component are omitted and each equal to zero. The collective pitch increases from  $5.45^\circ$  to approximately  $6.6^\circ$  following the correction, and reaches convergence in the fifth iteration, illustrating FLIGHTLAB's response to the correction. The delta-airloads correction is initially large, reflecting the initial difference between CFD and FLIGHTLAB. The difference is maintained in the following iteration, but decreased from the fourth iteration indicating the convergence of the aerodynamic loading. The stabilisation of both  $\theta_0$  and  $\Delta F_z$  confirms that a consistent trim state is reached after approximately five global iterations.

**Table 5.4:** Swashplate collective pitch  $\theta_0$  and  $\Delta$ -airloads correction per global iteration (C100).

Iteration $n_i$	1	2	3	4	5	6	7	8	9	10
$\theta_0$ [ $^\circ$ ]	5.45	6.60	6.67	6.68	6.55	6.52	6.51	6.50	6.51	6.53
$\Delta F_z$ [N]	710	758	750	632	634	631	635	646	657	658

Figure 5.5 compares the spanwise distribution of the discrete vertical force  $F_z$  in the first and final global iteration. The increase in force is observed over the entire blade span and in particular at the outboard region of the blade. This redistribution is consistent with the higher collective pitch and reflects the corrected aerodynamic loading obtained through the coupled analysis.



**Figure 5.5:** Spanwise distribution of the vertical force  $F_z$  in the first and final global iteration for the C100 rigid hover case.

The increased collective and outboard loading indicate that the discrepancy between FLIGHTLAB and CFD is likely related to the modelling of the induced flow. The local aerodynamic loads are primarily

governed by the effective angle of attack and Mach number. Since the Mach number is determined by the rotor speed and radial position, the observed load difference is attributed to differences in the inflow condition of FLIGHTLAB and ENSOLV. In lifting flight, a larger induced downwash reduces the effective angle of attack. Hence, the initial overprediction of the aerodynamic loads by FLIGHTLAB is consistent with an underprediction of the induced velocity, resulting in a higher effective angle of attack and increased sectional loading.

### 5.3.2. C102: Hover (6 kN), Steady CFD, Rigid Rotor

The C102 case corresponds to a steady rigid-hover condition with a trim target of  $F_z = 6$  kN. In contrast to C100, an outer relaxation factor was applied in the coupling procedure in order to assess its effect on convergence behaviour. Figure 5.6 shows the global iteration history of the total vertical thrust. The aerodynamic loading increases progressively in the first three iterations and reaches a converged value in the fourth iteration. The effect of the relaxation factor is clear, there is no overshoot in the CFD solution and the convergence of the loading is smoother. The final trimmed force slightly undershoots the 6 kN trim target but remains within the prescribed trim tolerance in FLIGHTLAB.




Figure 5.6: Global iteration history of the total vertical force  $F_z$  for the C102 rigid hover case (6 kN) with outer relaxation.

Table 5.5 summarises the evolution of the collective pitch and the  $\Delta$ -airloads correction. Due to the applied relaxation, both variable increase gradually over the first four iterations before stabilising. However, this table does not reflect that the initial discrepancy between CFD and FLIGHTLAB is approximately 770 N, which is similar to the difference observed in the previous C100 case. Moreover, the magnitude of the converged value of the correction is also nearly identical, suggesting that the difference between the aerodynamic models does not scale with thrust.

Table 5.5: Collective pitch  $\theta_0$  and delta-airloads correction per global iteration (C102).

Iteration $n_i$	1	2	3	4	5	6	7
$\theta_0$ [°]	10.00	10.60	10.70	10.98	10.97	10.96	10.95
$\Delta F_z$ [N]	430	510	678	675	669	662	656

Figure 5.7 shows the spanwise  $F_z$  distribution in the first and final iterations. The increase in loading is again observed across the full blade span. In particular in the outboard region, which is consistent with the higher collective pitch in the converged state.




Figure 5.7: Spanwise distribution of the vertical force  $F_z$  in the first and final global iteration for the C102 rigid hover case.

### 5.3.3. C100: Hover (3 kN), Unsteady CFD, Rigid Rotor

This C100 case corresponds to a rigid rotor in steady hover with a trim target of ( $F_z = 3$  kN). In contrast to the previous cases, an unsteady CFD simulation was employed. Since the simulated operating condition is a steady, azimuthally invariant hover case, the time-averaged outcome is expected to be consistent with the steady CFD result. An outer relaxation factor was applied during the first three global coupling iterations in order to progressively introduce the aerodynamic loading correction in FLIGHTLAB. Although this relaxation was not strictly necessary for the rigid-hover case, since comparable steady-CFD rigid-rotor simulations converged without it, it was retained here as a conservative measure to introduce the CFD-based aerodynamic correction gradually; the factor was initialized at 0.3 and increased to 1.0 over four iterations, as described in Section 3.5.4.

The convergence behaviour of the coupled solution is illustrated in Figure 5.8, which shows the global iteration history of the mean total vertical hub force. The overall convergence behaviour is similar to that observed in the earlier rigid rotor cases that utilised steady CFD simulations. The application of the relaxation factor gradually increases the aerodynamic loading correction during the first few iterations and ensures a smooth convergence of the solution by approximately the fourth iteration. The final force levels are also consistent with the magnitudes obtained in the previously presented steady CFD cases.

Table 5.6 summarises the evolution of the collective pitch and the mean total  $F_z$  correction during the global coupling iterations. The collective pitch increases from  $5.45^\circ$  to approximately  $6.5^\circ$  as the aerodynamic correction is progressively applied across the iterations. Owing to the applied relaxation factor, the correction initially remains relatively small and increases gradually. The final converged values of both the collective pitch and the aerodynamic correction are consistent with those obtained earlier using steady CFD simulations.

Table 5.6: Collective pitch  $\theta_0$  and delta-airloads correction per global iteration (C100-UNS-RGD), with final steady CFD values for comparison.

Iteration $n_i$	1	2	3	4	5	6	7	Steady CFD final
$\theta_0$ [°]	5.45	5.70	6.18	6.57	6.53	6.53	6.52	6.53
$\Delta F_z$ [N]	145	460	700	670	668	666	665	658

While the mean values are useful for analysing the convergence behaviour of the coupled solution, they do not fully describe the aerodynamic load distribution. The azimuthal variation of the thrust force




Figure 5.8: Global iteration history of the mean total vertical hub force for Case C100-UNS-RGD.

during the coupling process is therefore examined in Figure 5.9 and 5.10, which show the hub force distribution as a function of azimuth for the first and sixth iterations, respectively. The hub forces are observed to be essentially invariant with respect to azimuth. This behaviour was verified previously and is expected for the present configuration, as it results from the rigid blade assumption and the steady hover loading condition.





Figure 5.9: Hub force  $F_z$  as a function of azimuth at iteration 1. Figure 5.10: Hub force  $F_z$  as a function of azimuth at iteration 6.

The spanwise loading distribution is shown in Figure 5.11, where the  $F_z$  distribution along the blade is compared between the first and final iterations. An increase in loading is observed across the entire blade span, with the most pronounced change occurring in the outboard region of the blade. The resulting distributions are consistent with those reported previously for the steady CFD case.




Figure 5.11: Spanwise distribution of  $F_z$  during the first and final coupling iterations for Case C100. Comparison of steady and unsteady CFD.

### 5.3.4. C851: Forward Flight, Unsteady CFD, and Rigid Rotor

This C851 case represents a forward flight condition with trim targets of  $F_z = 3774.7$  N,  $M_x = -33.1$  Nm, and  $M_y = 132.4$  Nm. The freestream velocity is aligned with the rotor hub plane. An outer relaxation factor is applied during the first four coupling iterations. The factor was initialized at 0.3 and increased to 1.0 over four iterations, as described in Section 3.5.4.

Figure 5.12 shows the evolution of the mean total hub forces and moments. For the hub forces, global convergence is reached after approximately four iterations. For the hub moments, an error occurred during the export of the altered CA aerodynamic loading, which therefore appears as zero. Nevertheless, the CFD solution and trim targets still converge and agree with each other, indicating that the coupled solution remains valid.

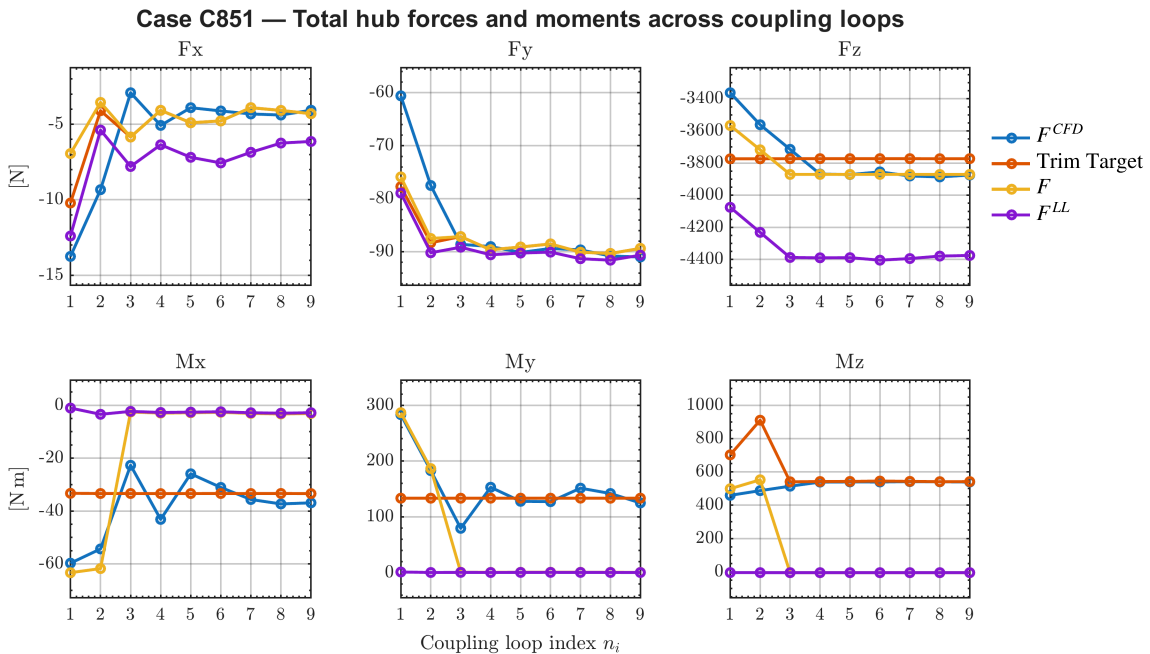


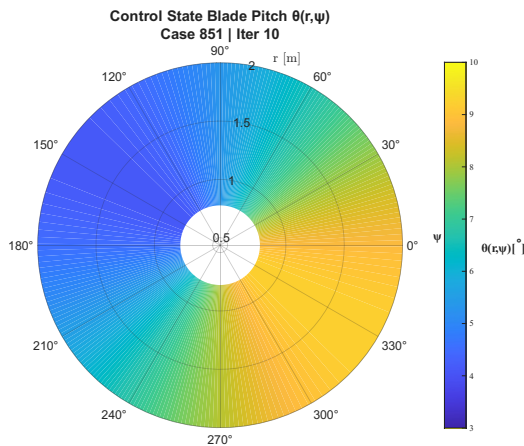
Figure 5.12: Evolution of the mean hub forces and moments during the coupling procedure for Case C851.

The evolution of the control states is given in Table 5.7. Due to the asymmetry of forward flight, non-zero cyclic pitch inputs are required. The collective pitch increases gradually from approximately  $6^\circ$  to  $6.7^\circ$  during the coupling process. The lateral cyclic pitch, which tilts the rotor disc sideways and therefore redistributes lift between the advancing and retreating sides of the rotor, changes from approximately  $-1.6^\circ$  to  $-2.1^\circ$ . The longitudinal cyclic pitch, responsible for forward and rearward rotor tilt, increases from about  $1^\circ$  to nearly  $1.4^\circ$ .


**Table 5.7:** Evolution of the control states for Case C851 over the global coupling iterations. .

Control state	1	2	3	4	5	6	7	8	9	10
Collective $\theta_0$ [ $^\circ$ ]	5.98	6.25	6.48	6.74	6.74	6.74	6.77	6.74	6.71	6.71
Lateral cyclic $\theta_{1c}$ [ $^\circ$ ]	-1.61	-1.98	-2.21	-2.09	-2.13	-2.12	-2.11	-2.14	-2.15	-2.13
Longitudinal cyclic $\theta_{1s}$ [ $^\circ$ ]	1.01	1.19	1.36	1.35	1.38	1.36	1.36	1.37	1.38	1.39

The converged control states lead to the blade pitch distributions shown in Figures 5.13 and 5.14. The first figure presents the blade pitch distribution derived from the control states, while the second figure shows the total blade pitch after including the blade twist distribution. With the freestream velocity directed toward the  $0^\circ$  azimuth, the control inputs reduce the pitch on the advancing side and increase it on the retreating side of the rotor disk. This redistribution is required to balance the dissymmetry of lift that arises in forward flight. When the blade twist distribution is included, the resulting total pitch distribution exhibits a maximum near the blade root between approximately  $270^\circ$  and  $30^\circ$ .



**Figure 5.13:** Blade pitch distribution derived from the control states.



**Figure 5.14:** Total blade pitch distribution including blade twist.

Figures 5.15 and 5.16 show the azimuthal variation of the hub force components during the first and final iterations of the coupling procedure. In the initial iteration, clear discrepancies exist between the CFD solution and the trim target. The CFD solution exhibits larger oscillation amplitudes, with an overshoot in  $F_x$  and underpredictions in  $F_y$  and  $F_z$ .

After convergence of the coupling, shown in Figure 5.16, the CFD and trim target solutions show much closer agreement. The hub forces  $F_x$  and  $F_y$  clearly exhibit the dominant  $4/\text{rev}$  harmonic associated with the pitch variations of each of the four blades in forward flight. The difference in amplitude for  $F_x$  and  $F_y$  is noticeable between the CFD solution and the baseline airloads, and shows where FLIGHTLAB is in need of a correction. The corrected FLIGHTLAB loads closely follow the CFD solution, demonstrating that the coupled method successfully reproduces the aerodynamic loading.

The azimuthal distributions of the hub moments during the first and final iterations are shown in Figures 5.17 and 5.18. In the first iteration, the agreement between CFD and FLIGHTLAB is already good for the rolling moment  $M_x$ . Larger discrepancies occur for  $M_y$ , where FLIGHTLAB underpredicts both the mean value and oscillation amplitude. The largest difference occurs for  $M_z$ , which exhibits a significant offset due to the initial mismatch in thrust.

After convergence, excellent agreement is obtained for all moment components, as shown in Figure




Figure 5.15: Azimuthal distribution of hub forces for Case 851 during the first coupling iteration.




Figure 5.16: Azimuthal distribution of hub forces for Case 851 after convergence (iteration 10).




Figure 5.17: Azimuthal distribution of hub moments during the first coupling iteration.

5.18. The CFD solution closely matches the trim targets in both mean value and azimuthal variation, indicating that the aerodynamic and structural solutions represent the same trimmed rotor state.




Figure 5.18: Azimuthal distribution of hub moments in the tenth coupling iteration.

To further examine the aerodynamic behaviour, the spanwise and azimuthal distribution of the vertical force  $F_z$  is compared between the CFD and the baseline FLIGHTLAB solutions in Figures 5.19 and 5.20. In FLIGHTLAB, the distribution is nearly azimuthally invariant, whereas the CFD solution exhibits a characteristic sickle-shaped pattern, indicating local reductions in lift between approximately  $40^\circ$ – $110^\circ$  and  $250^\circ$ – $340^\circ$ .

Since no discontinuity exists in the blade pitch distribution, these variations in the CFD solution are likely related to rapid variations in the inflow conditions and local blade–vortex interactions, phenomena that do not seem to be captured by the aerodynamic model in FLIGHTLAB. This interpretation is supported by the flow visualisation in Figure 5.21, which shows the Q-criterion iso-surfaces used to identify vortical structures in the wake. Multiple blade-wake interactions are visible and both root and tip vortices can be seen to interact with the following blades. In the regions where the reductions in lift occurred, the blades dissect through a large wake structure, which could alter the local inflow conditions. Moreover, several tip vortices are visible from the different blades and multiple revolutions appear to be present. However, the different wake regions merge together into large vortical structures.




Figure 5.19: Spanwise and azimuthal distribution of  $F_z$  from the CFD solution. Represented as discrete forces at the ACP's in FLIGHTLAB.




Figure 5.20: Spanwise and azimuthal distribution of baseline  $F_z$  from FLIGHTLAB. Represented as discrete forces at the ACP's in FLIGHTLAB.




Figure 5.21: Q-criterion visualisation of the rotor wake for Case C851.

### 5.3.5. Summary of the Rigid Rotor

These test cases using the rigid rotor show that the loose coupling based on the delta-airloads correction is stable and converges rapidly in both hover and forward flight. The main findings are summarised below.

#### Hover Flight and Steady CFD

Across the investigated steady rigid-hover cases (3 and 6 kN), the coupling framework demonstrates rapid convergence. Independent of thrust level and the use of outer relaxation, the solution converges within approximately four to five global iterations. The relaxation applied in the 6 kN case, initially 0.3 and increased to 1.0 over four iterations, primarily improves the smoothness of convergence by preventing overshoot, without significantly affecting the number of iterations required.

A consistent trend is observed in the  $\Delta$ -airloads correction. In both cases, the initial CFD-FLIGHTLAB discrepancy corresponds to a correction of approximately 750–780 N. As the coupling progresses, this value reduces and stabilises around 630–660 N. Notably, the magnitude of the converged correction remains of similar order despite the doubling of thrust from 3 to 6 kN.

This indicates that the difference between CFD and the internal aerodynamic model of FLIGHTLAB does not scale strongly with collective pitch or total thrust within the investigated operating range. The discrepancy therefore appears to originate from a systematic modelling difference, most likely associated with the predicted induced flow and resulting effective angle of attack, rather than a thrust-dependent nonlinear aerodynamic effect.

Overall, the delta-airloads strategy yields a stable and repeatable coupled solution for steady rigid-hover conditions, with predictable convergence characteristics and a consistent correction magnitude.

#### Hover Flight and Unsteady CFD

For the rigid rotor case utilising the unsteady CFD solution (C100, 3 kN), the loose CFD-CA coupling again demonstrates stable and rapid convergence. The global solution converges within approximately four iterations, and the outer relaxation factor ensures a smooth addition of the aerodynamic correction during the initial coupling loops. The final collective pitch and total thrust closely match the values that were found for the corresponding steady CFD case, indicating that the use of steady or unsteady aerodynamics does not significantly alter the converged trim state for a rigid rotor in hover. This is

expected, since the rigid hover case represents a nominally steady aerodynamic problem in the rotating frame, for which both steady and unsteady CFD should converge to the same mean flow state.

The azimuthal distribution of the hub forces remains essentially invariant, which was expected for a rigid rotor operating in steady hover. As a consequence, the delta-airloads correction remains nearly uniform in azimuth and the corrected spanwise loading is consistent with the increased collective pitch. Overall, the results confirm that the coupling approach remains stable when unsteady CFD is employed and produces results consistent with those obtained using steady aerodynamics.

### **Forward Flight and Unsteady CFD**

For the rigid rotor forward flight case (C851), the loose coupling remains stable and converges within four global iterations, even given the asymmetric unsteady loading. The use of an outer relaxation factor ensures a smooth application of the aerodynamic correct and evolution of the coupled solution during the initial iterations.

Unlike the hover cases that require solely a non-zero collective pitch, the forward-flight requires the lateral and longitudinal cyclic pitch components to be non-zero to reach a trim state. Adjusting these control states during the coupling iterations redistributes the lift over the rotor disk and compensates for the dissymmetry of lift between the advancing and retreating sides of the rotor. The converged control states reflect that the aerodynamic corrections introduced by the CFD solution successfully adjust the trim state determined by FLIGHTLAB.

After convergence, the mean and azimuthal variations of the hub forces and moments are consistent between ENSOLV and FLIGHTLAB. In particular, the correct FLIGHTLAB solution more accurately captures the periodic 4/rev variation that dominates the loading. While this is an expected result, given that the CFD loading is intended to replace the baseline loading in FLIGHTLAB, it demonstrates that the delta-airloads is also able to correct both the mean and unsteady component of the aerodynamic loading.

The comparison of the baseline FLIGHTLAB and CFD load distributions over a full revolution shows that the aerodynamics of forward flight are strongly influenced by the flowfield, and highlights the capabilities and limitations of the solvers. In particular, the presence of blade-vortex interactions leads to localised decreases in lift in the CFD solution, a phenomenon that FLIGHTLAB is not able to resolve and take into account. This shows the added benefit of the increased fidelity of the CFD solution and confirms that the loose coupling approach is capable of a trimmed rotor for unsteady aerodynamic conditions, such as forward flight.

## **5.4. Flexible Rotor Simulations**

Following the assessment of the rigid rotor, the behaviour of the loose coupling was also assessed for a flexible rotor. This meant that the number of degrees of freedom of each structural blade element was increased to allow for in- and out-of-plane bending, and torsional deformation of the rotor blade. The structural modelling is based on the Bo105 blade model that was used by NLR as part of the HeliNOVI project. This structural model was constructed based on the reported structural properties of the wind tunnel model of the rotor [39].

Due to an error in the data-handling procedure, the results of the C400 (3 kN) case using unsteady CFD and the C402 (6 kN) case using steady CFD were corrupted and failed to converge. The error was caused by an incorrect rotor-motion file being loaded during the coupling procedure. This issue was identified and corrected, but the simulations could not be rerun within the available time. They are therefore omitted from the present analysis. The remaining simulations provide sufficient information to assess the ability of the implementation to handle a flexible rotor.

### **5.4.1. C400: Hover (3 kN), Steady CFD, Flexible Rotor**

The C400 case models a flexible rotor in steady hover with a trim target of  $F_z = 3$  kN, where blade flexibility is enabled in FLIGHTLAB. The CFD solution is obtained from a steady simulation using the mean rotor shape over one revolution as input. During the first four coupling iterations, an outer relaxation factor is applied, increasing from 0.3 to 1.0 as described in Section 3.5.4.

**Table 5.8:** Collective pitch  $\theta_0$  for the first five global iterations (flexible C400), with steady C100 results for comparison.

Iteration $n_i$	1	2	3	4	5	Steady C100 final
$\theta_0$ [°]	5.65	6.15	6.52	6.85	6.73	6.53

### Trim convergence and control evolution

Figure 5.22 shows the global iteration history of the total vertical thrust. The convergence behaviour is similar to the rigid blade hover cases. During the first four iterations, the aerodynamic loading increases progressively before converging in the fifth iteration. The relaxation factor ensures a smooth progression of the loading and only a minor overshoot is observed in the fourth loop. Overall, the global convergence is comparable to the rigid rotor case.


**Figure 5.22:** Global iteration history of the total vertical force  $F_z$  for the flexible hover case C400 (3 kN).

Table 5.8 contains the evolution of the collective pitch. Compared to the rigid case at the same thrust target, the collective is slightly higher during each iteration. This behaviour is consistent with the torsional deformation of the blade that was observed to reduce the spanwise blade pitch distribution. Due to the decreased blade pitch, a higher collective pitch is required to achieve the same thrust target.


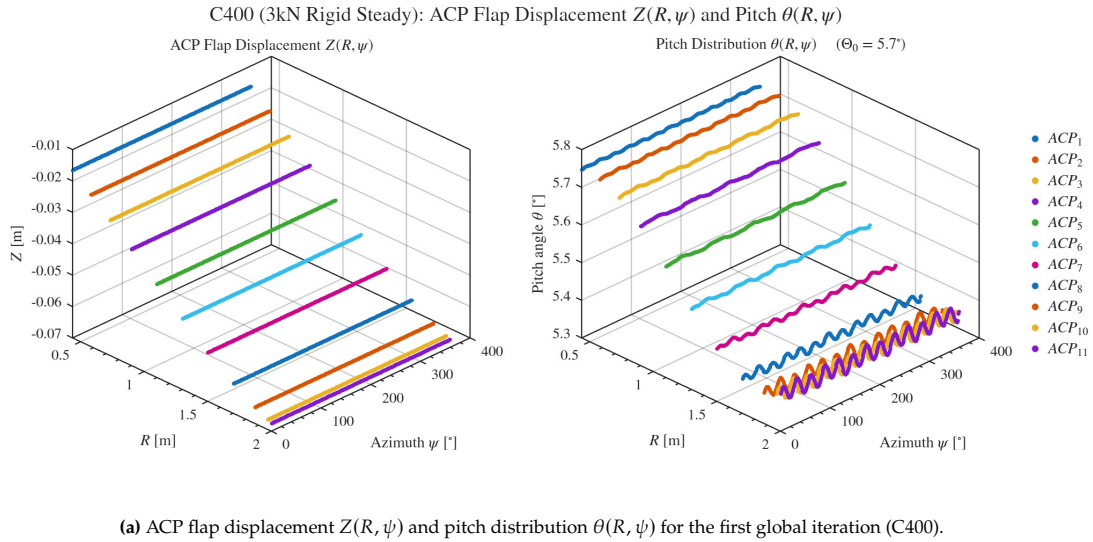
### Blade motion

Before assessing the azimuthal distribution of the force, the flexible blade motion in the initial and final iteration are illustrated in Figures 5.23a and 5.23b. These figures show the flap displacement  $Z(R, \psi)$  and pitch distribution  $\theta(R, \psi)$  along the blade span during a full revolution. Here, the pitch distribution is the relative pitch distribution, which is obtained by subtracting the blade twist distribution from the instantaneous rotor orientation. This should roughly equal to the control state pitch of the rotor blade.


The initial blade motion is characterised by constant downward blade flapping, as discussed previously in Appendix A, and this is accompanied by both spanwise decreasing blade pitch and local blade pitch oscillations due to torsional blade deformation ( $13/rev$ ). Compared to the collective pitch of  $5.7^\circ$ , the blade pitch at the tip has been significantly reduced, which explains the increased collective pitch compared to the rigid blade case. Following the addition of the correction, the blade flapping is periodic, but is destabilized, and now shows increased blade flapping towards the tip and azimuthal oscillations. Similarly, in the converged solution the addition of the aerodynamic correction has greatly increased the amplitude of the blade pitching oscillations.

### Hub loads

To explain the destabilization of the blade motion, the hub forces as a function of azimuth in the first iteration are shown Figure 5.24a. While the steady CFD solution is observed to be azimuthally invariant,



**Figure 5.23:** Comparison of the flexible blade motion for Case C400 in the initial and converged coupling iterations.



(a) Hub force  $F_z(\psi)$  as a function of azimuth in the first global iteration (C400). (b) Delta-airloads correction  $\Delta F_z(\psi)$  as a function of azimuth in the first global iteration (C400).

**Figure 5.24:** Initial azimuthal loading and correction for the flexible hover case C400.

the variations in the flexible rotor motion determined by FLIGHTLAB lead to oscillating periodic hub loads. Combined they lead to the computed delta-airload correction, shown in Figure 5.24b, which is also azimuth dependent and oscillates at the same frequency as the blade pitch oscillation (13/rev).

The compounded effect of applying the azimuthally varying correction is seen in the loads of the

converged solution, shown in Figures 5.25a and 5.25b. While the mean loading in FLIGHTLAB is equivalent to that of ENSOLV, both the amplitude and the frequency of the oscillations have been greatly altered and the resulting loading is no longer approximately azimuthally invariant. The application of the non-uniform correction in combination with FLIGHTLAB attempting to achieve a periodic motion that agrees with the prescribed non-uniform aerodynamic correction leads to a periodic, but non-representative solution for a hovering rotor. The trim procedure in FLIGHTLAB is observed to be sensitive to changes in the mean aerodynamic loading, and the amplitude and the frequency of loading oscillations, and its interactions with the blade motion.

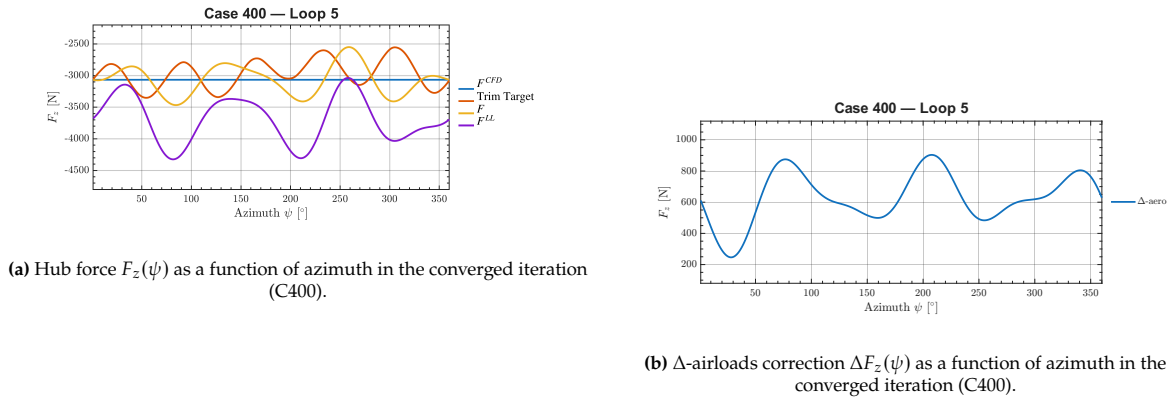


Figure 5.25: Converged azimuthal loading and correction for the flexible hover case C400.

Overall, this coupling converges in terms of hub thrust and provides a computationally efficient way to assess the interaction between ENSOLV and FLIGHTLAB. However, combining the steady CFD solution with a fully flexible rotor introduces an inherent modelling inconsistency. The steady aerodynamic solution is azimuthally invariant in hover, whereas the flexible blade response determined by FLIGHTLAB contains azimuthal variations. As a result, the delta-airloads correction becomes azimuth-dependent and feeds these variations back into the structural solution. Although the coupled solution satisfies the FLIGHTLAB trim criteria, the resulting blade motion no longer corresponds to a steady hover solution. Instead, the correction amplifies the existing oscillatory components in the blade flapping and pitching motion, as FLIGHTLAB attempts to obtain a periodic response that is consistent with the imposed non-uniform aerodynamic loads. Therefore, the issue is not a lack of numerical convergence of the coupling procedure, but a physically non-representative aeroelastic response caused by the mismatch between a steady CFD correction and a time-dependent flexible rotor motion.

Although a converged trimmed rotor state is reached, the final solution no longer accurately represents the original stable hover flight scenario, since the locations of the azimuthal and spanwise variations of the correction, and therefore the blade response, change between each iteration. An alternative solution would have been to apply only an azimuthally invariant correction, which would adjust the mean aerodynamic loading while leaving the oscillatory part of the blade response unaffected. This strategy would introduce a uniform increase in aerodynamic loading and could potentially result in a more stabilised rotor motion while still correcting the rotor thrust. However, this was not investigated due to a lack of time.

#### 5.4.2. C402: Hover (4 kN), Unsteady CFD, Flexible Rotor

This C402 case models a flexible rotor in steady hover using unsteady CFD-simulations. Due to an initialisation mistake, this load case has a trim target of  $F_z = 4$  kN, this however should not affect the behaviour of the results and coupling. An outer relaxation factor was applied during the first seven global coupling iterations.

##### Trim convergence and control evolution

Figure 5.26 shows the global iteration history of the mean total force in  $z$ . The aerodynamic loading increases progressively and converges in the eighth iteration. The relaxation factor ensures a smooth application of the loading. In terms of magnitude, the initial difference between the CFD solution and




Figure 5.26: Global iteration history of the mean total vertical hub force for Case C402.

Table 5.9: Coupling loop convergence history (Case C402).

Iteration	1	2	3	4	5	6	7	8
Collective $\theta_0$ [deg]	7.08	7.33	7.42	7.53	7.63	7.69	7.74	7.75


the trim target in FLIGHTLAB is very similar to that found in the other cases despite the difference in trim target.

Table 5.9 presents the evolution of the collective pitch. The collective pitch increases from  $7.08^\circ$  in the first iteration to  $7.75^\circ$  in the eighth iteration. The adjustments between iterations are initially large and gradually decrease as the trim solution approaches convergence. By the final iteration, the collective pitch stabilizes, suggesting that a consistent aerodynamic loading and trim have been reached.


Despite the convergence of the global coupling behaviour, previous coupling results showed that in the case of a flexible rotor, elastic blade deformation could introduce azimuthal and spanwise variations in the solutions, as well as between global iterations. The flexible blade motion and the azimuthal distribution of the force in the initial iteration are illustrated in Figures 5.27a and 5.27b.

The initial blade motion is characterised by increased upward blade flapping and is accompanied by both spanwise decreasing blade pitch and local blade pitch oscillations due to torsional blade deformation. The influence of these azimuthal and spanwise variations on the azimuthal distribution of force lead to a highly oscillatory CFD solution that matches the torsional blade deformation. The aperiodic, increasing trend that is visible in the thrust force  $F_z$  in the initial iteration is most likely caused by the flow field that is forming in the CFD solution, as this trend is not observed in the flexible blade motion. In addition, a discontinuity in the force distributions in the CFD solution between the beginning and end of the first and second revolution was observed. While FLIGHTLAB determines the rotor motion as a continuous motion when trimming, the rotor motion that is exported to ENSOLV is ensured to be periodic as shown in Appendix A. The discontinuity is therefore attributed to how ENSOLV performs the unsteady CFD simulations. While a clear cause was not found, the impact on the result is limited.

The final globally converged solution in terms of the flexible blade motion and the azimuthal distribution of the force is shown in Figures 5.28 and 5.27c. Comparing the azimuthal force distributions, the initial amplitude variations are on the order of  $\approx 10$  N, whereas the difference between the minima and maxima in the final solution is approximately  $\approx 1000$  N. The frequency of the oscillations has also been affected due to the correction, while simultaneously introducing large peaks in the aerodynamic loads.



(a) Flexible blade motion during the first global iteration (C402).



(b) Azimuthal thrust distribution during the first global iteration (C402).

(c) Azimuthal thrust distribution in the converged iteration (C402).

**Figure 5.27:** Initial blade motion and azimuthal thrust distributions for Case C402.

The frequency of these large peaks match the pitch displacement of the rotor blade highlighting the cause of the loading variations. The large magnitude of the low frequency peaks also means that the initial high frequency variations have become unobservable in the final solution.

Previously, the azimuthal variance of the correction appeared to destabilize the blade motion when FLIGHTLAB attempted to retrim the rotor. However, for unsteady CFD simulations, each azimuthal instance should be representative of the rotor blade at that particular instant and should therefore not necessarily cause the blade motion to diverge. Regardless, the blade motion of the converged solution does worsen, as for every outer coupling the blade flapping and pitching increase and become more chaotic in nature. In this case, the increase in blade pitching amplitude is smaller than in the final steady CFD solution ( $\sim 2.5^\circ$  versus  $\sim 8^\circ$ ), but it is still significantly greater than the initial blade pitch variation ( $\sim 0.05^\circ$ ).

While the global convergence behaviour looks satisfactory, the converged blade motion is clearly not representative of hover due to the azimuthal variations in blade motion and loading. The results from ENSOLV are not the cause despite their coarseness and the discontinuity. Based on the current results, the increase in flap and pitch can be attributed to the influence of the sectional pitching moment on the structural flexibility. With the rotor blade free to torsionally deform the pitching moment forces the blade to twist and bend, and combined with the azimuthal variations that the correction causes the blade motion is worsened.

As a part of the verification, the exchange of the sectional pitching moment to FLIGHTLAB was confirmed to correspond to that in ENSOLV. However, the results of these simulations show that the cause is a mismatch between the sectional pitching moment distributions between FLIGHTLAB and ENSOLV, and that this mismatch causes excessive blade pitch. Due to limited access to the data,




Figure 5.28: Flexible blade motion of the converged solution (C402).

this analysis did not analyse the sectional pitching moment any further and an exact cause could be determined. Given the destabilization of the blade flapping, it can be proposed that assessing the pitching moment would result in improving the overall agreement and convergence of the blade motion and loading. Other influences such as the aerodynamic modelling and the structural modelling should also be considered.

### 5.4.3. C851: Forward Flight, Unsteady CFD, and Flexible Rotor

Troubled by the same problems as the previous cases, the forward-flight case C851 failed to successfully produce a trimmed rotor in the 8th coupling iteration due to the worsening of the blade motion. Regardless the results of the first seven iterations is shown here. For brevity, the detailed discussion of the destabilisation of the blade motion is omitted here. Instead, this section focuses on the evolution of the control states and on the hub forces and moments.

The evolution of the mean hub force history is shown in Fig 5.29. While, initially the agreement between FLIGHTLAB and ENSOLV seems to improve as the corrected aerodynamic loading approaches the CFD total, a clear disagreement is observed in the 7th iteration for  $F_x$  and  $F_y$ . However, the accuracy of the coupling solution cannot be assessed solely on the convergence mean total hub-force.

Table 5.10 presents the evolution of the trim controls over the first seven coupling iterations. All three control state continue to increase in magnitude as the coupling progresses and show no sign of slowing





Figure 5.29: Evolution of the mean total hub forces for Case C851 over the coupling iterations.

down. The trim solution can be observed to drift rather than settle on stable fixed values.


**Table 5.10:** Approximate control-state evolution for Case C851 over the first seven coupling iterations.

Quantity	1	2	3	4	5	6	7
Collective $\theta_0$ [°]	6.12	6.36	6.49	6.62	6.72	6.79	6.94
Lateral cyclic $\theta_{1c}$ [°]	-1.42	-1.53	-1.61	-1.77	-1.94	-2.14	-2.22
Longitudinal cyclic $\theta_{1s}$ [°]	0.17	0.23	0.28	0.35	0.46	0.58	0.72

The divergence of the coupled solution can be illustrated by how the hub loads change during the coupling process. Figures 5.30a and 5.30b compare the hub-force over the rotor azimuth for the first and seventh iterations, respectively. Initially, the solutions from FLIGHTLAB and ENSOLV are comparable to the results for the rigid rotor and the initial distributions of both solver are quite closely and similar, with the exception of the phase of  $F_x$  and  $F_y$ , and the trim target. However, in the seventh iteration the  $F_x$  and  $F_y$  components of the trim solution in particular have exploded in magnitude. The earlier observed increase in amplitude variation and phase disagreement in the  $F_z$  components are also observed.



(a) Azimuthal distribution of hub forces for Case 851 during the first coupling iteration (flexible rotor blade).




(b) Azimuthal distribution of hub forces for Case 851 during the 7th coupling iteration (flexible rotor blade).


**Figure 5.30:** Evolution of the azimuthal hub-force distributions for Case C851.

Similar behaviour is observed in the hub moments, shown in Figures 5.31a and 5.31b. Large differences between both solvers are observed for the  $M_x$  and  $M_y$  components initially, which are increased significantly in the seventh iteration to unrealistic values as compared to the rigid simulations. The  $M_z$

components on the other hand do not differ significantly, but the difference in magnitude, amplitude and oscillations does increase between iterations.



(a) Azimuthal distribution of hub moments for Case 851 during the first coupling iteration (flexible rotor blade).

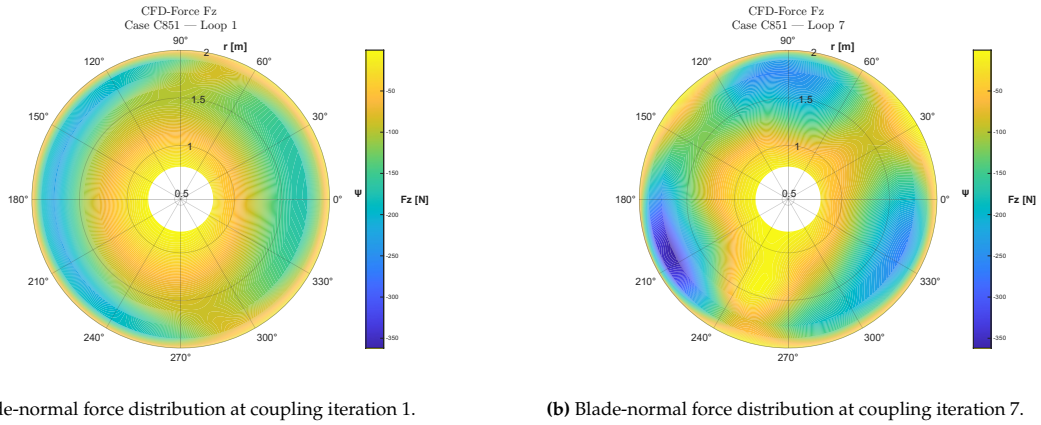


(b) Azimuthal distribution of hub moments for Case 851 during the 7th coupling iteration (flexible rotor blade).

**Figure 5.31:** Evolution of the azimuthal hub-moment distributions for Case C851.

Finally, the distribution of the aerodynamic thrust force  $F_z$  can be shown to further illustrate the effect of the divergence of the loose coupling. Figures 5.32a and 5.32b compare the  $F_z$  distribution versus the blade radius and the azimuth. Initially, the sickle-shape can be observed that was also found for the rigid rotor in forward flight and is likely caused by the interaction between rotor blade and wake which reduces the lift locally. In the seventh iteration, the distribution has become asymmetric and random, which can be attributed to the destabilized blade motion.

Overall, the control state drift and the disagreement between solvers in terms of total and distributed loads indicate that the loose coupling approach does not converge when structural flexibility is included for the Case C851.



**Figure 5.32:** Comparison of the CFD blade-normal force distribution for Case C851 at the first and seventh coupling iterations. Represented on the aerodynamic FLIGHTLAB grid.

#### 5.4.4. Summary of the Flexible Rotor

The flexible-rotor cases demonstrate that the current loose coupling is not successful once structural flexibility is included. In contrast to rigid-rotor cases, the addition of blade elasticity introduces a complex aeroelastic response making blade deformation, aerodynamic loading and trim strongly interdependent. As a result, the coupled solution becomes much more sensitive to the aerodynamic correction and the consistency between both solvers.

The results indicate that the present implementation is unable to handle the increased complexity. Despite the apparent global convergence of some cases, the resulting blade motion and loading are not representative of the simulated cases unlike the rigid rotor simulations. This suggests that there is an unresolved mismatch in the interaction of the aerodynamic loading with the structural response and the trim process. In this research, the exact source of this mismatch has not been identified due to the limited time available.

Overall, the flexible-rotor simulations have shown to be unsuccessful in their current form, but have critically been useful to demonstrate that the present implementation is unable to accurately capture the aeroelastic behaviour of a flexible rotor.

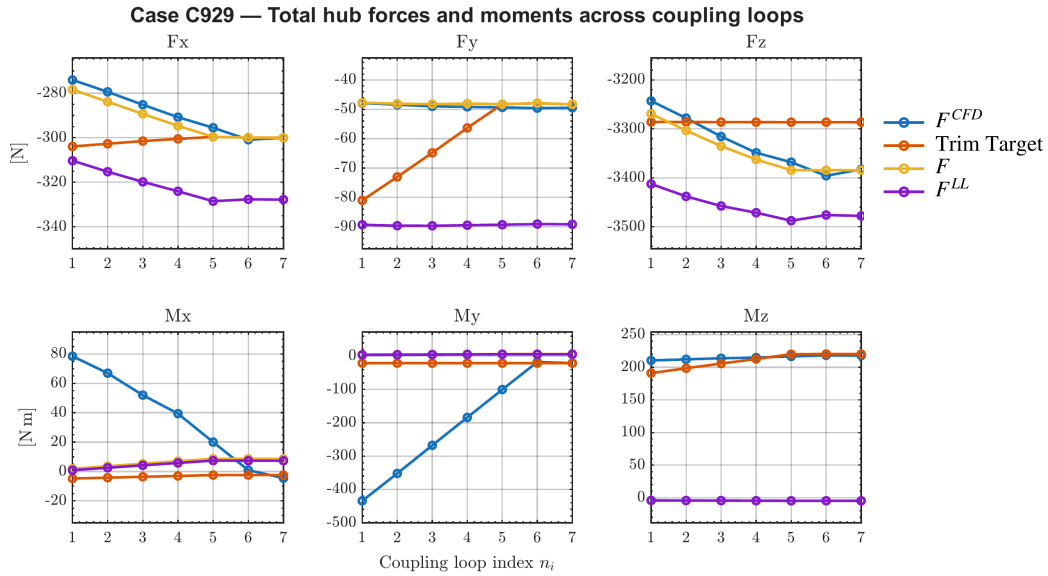
## 5.5. Validation

This section describes the validation of the loose coupling against the HART-II baseline case that was described in the test matrix. The HART-II campaign has been used extensively as a benchmark because it provides a well-documented database containing rotor motion, sectional airloads and wake trajectories, and contains strong BVI-induced effects. The baseline case was performed using a 40% scaled Bo105 rotor in the DNW open-jet facility, and corresponds to a descending flight. Due to wind-tunnel interference, the geometric shaft angle of  $5.3^\circ$  corresponds to an effective disk angle of attack of approximately  $4.5^\circ$ . Under these conditions, the BVI events occur near azimuth angles  $\psi \approx 50^\circ$  and  $\psi \approx 300^\circ$  in the first and fourth quadrants of the rotor disk. Here, the vortex axis and the blade leading edge are approximately parallel [30]. The control angles used to set the rotor motion of the HART-II experiment are given in Table 5.11.

**Table 5.11:** Experimental trim control angles of the HART-II experiment

	Collective	Lateral cyclic	Longitudinal cyclic
Symbol	$\theta_0$	$\theta_{1c}$	$\theta_{1s}$
Value [°]	3.8	1.92	-1.34

The HART-II campaign and the associated workshop showed that this baseline case is particularly challenging because accurate prediction depends on more than global trim convergence. They found that the correct preservation of the wake geometry, vortex strength, elastic blade response and the timing



**Figure 5.33:** Evolution of the mean hub forces and moments during the coupling procedure for Case C929. (Rigid rotor)

of the BVI on both the advancing and retreating sides of the disk greatly affected the quality of the solution. The present validation therefore considers both the global convergence of the coupled solution and compares its sectional airloads to the experimental results of the baseline case. When interpreting these comparisons, it should be noted that the conventions used in the HART-II experiments differ from those of FLIGHTLAB, which explains the sign of the hub rolling moment  $M_x$ .

### 5.5.1. C929: Descending Flight, Unsteady CFD, and Rigid Rotor

This C929 case represents the baseline descending flight condition that was part of the HART-II campaign and was simulated using an rigid rotor. An outer relaxation factor is applied during the first five coupling iterations.


#### Trim convergence and control evolution

Figure 5.33 shows the evolution of the mean total hub forces and moments. Global convergence of all hub force and moment components is observed and the relaxation factor ensure smooth convergence. The initial difference between the mean total forces is relatively small in comparison to the differences that were noted for the hover cases, indicating the CFD solution and the FLIGHTLAB solution more closely approaches each other. The mean  $M_x$  and  $M_y$  values do differ more significantly initially, which can be investigated based on the azimuthal distribution of these components.

Table 5.12 presents the evolution of the trim controls. As the relaxation factor and the applied aerodynamic correction increase, control states monotonically rise or fall before reaching a plateau in the sixth iteration. Compared to the control states of the HART-II experiment reported in Table 5.11, the signs of the cyclic pitches are flipped. This can be attributed to the difference in definition of the control state as determined by the collective and cyclic pitches between FLIGHTLAB and HART-II. They do however represent the same blade motion with respect to the rotor orientation and flight condition.

**Table 5.12:** Evolution of the control states for Case C929 over the global coupling iterations for the simulated rigid rotor.

Control state	1	2	3	4	5	6	7
Collective $\theta_0$ [°]	2.65	2.72	2.78	2.84	2.88	2.93	2.91
Lateral cyclic $\theta_{1c}$ [°]	-2.29	-2.19	-2.09	-1.99	-1.88	-1.79	-1.79
Longitudinal cyclic $\theta_{1s}$ [°]	1.84	1.81	1.77	1.72	1.66	1.60	1.58



**Figure 5.34:** Azimuthal distribution of hub forces for Case 929 during the first coupling iteration (rigid rotor blade).

The result of the loose coupling underestimates the collective pitch by nearly  $\sim 1^\circ$ , whereas the cyclic pitch are within  $\sim 0.25^\circ$  of the experimental values. Contributing factors may be the lack of the fuselage or the lack of structural flexibility, which both can lead to altered inflow conditions or decreased blade pitching thereby requiring a higher collective pitch when included.

In the HART-II campaign, the participating parties reported generally similar control angles that fell within  $0.5^\circ$  of the experimental values. In general, the CFD-CA methodologies were reported to the same tendency of over- or under-predicting a particular control angle for the same case. Moreover, little difference was found between the trim angles when refining the surface and wake grids, although larger differences were noted when the same methods were run with different grids or different CA methods [30].

### Hub loads

The azimuthal distribution of the hub forces in loop 1 and loop 7, shown in Figures 5.34 and 5.35, illustrate the convergence of the hub forces. In the first loop, the results from ENSOLV already agrees with the FLIGHTLAB loading reasonably well for all three components, indicating that the difference between the aerodynamic loading in FLIGHTLAB and ENSOLV is relatively small. The difference with the trim target in case of the  $F_x$  and  $F_y$  component however is clear, indicating that the retrimmed rotor motion and CFD solution do not yet agree. By Loop 7, the agreement is substantially improved, and the solutions have nearly collapsed on to each other containing only minor differences at the minima and maxima. All components are expected to collapse the same solution, except for the trim target in  $F_z$  as it equals the aerodynamic loading minus the weight of the rotor. Overall, this demonstrates that the coupling has produced a shared solution.

In addition, the azimuthal distributions of the hub moments in loop 1 and loop 7 are shown in Figures 5.36 and 5.37. In the first loop, the differences are large as ENSOLV predicts substantially oscillations in the moments  $M_x$  and  $M_y$  that are not capture by the trim target. The  $M_z$  component on the other is much closer in mean level and phase but still shows a difference in amplitude. By loop 7, the agreement is significantly improved, and ENSOLV and FLIGHTLAB almost fully overlap for all three components. The altered CA solution, however, remains close to zero and does not reproduce the converged moments, indicating that an error has occurred during post-processing with the creation of this data as this variable is expected to agree with the CFD solution. Overall, the progression indicates that the hub moments show complete agreement between ENSOLV and the target.

### Blade loading and wake structure

The final converged aerodynamic thrust distribution ( $F_z$ ) from ENSOLV is illustrated in Figure 5.38. The load distribution is dominated by the radial distribution of lift, with the outer section carrying the largest load. The distribution is axisymmetric, as expected from forward descending flight, with the




Figure 5.35: Azimuthal distribution of hub forces for Case 929 during the 7th coupling iteration (rigid rotor blade).




Figure 5.36: Azimuthal distribution of hub moments for Case 929 during the first coupling iteration (rigid rotor blade).




Figure 5.37: Azimuthal distribution of hub moments for Case 929 during the seventh coupling iteration (rigid rotor blade)




Figure 5.38: Blade-normal force distribution from CFD.

advancing and retreating sides of the rotor disk producing roughly the same distribution. A notable observation is the relatively limited visible influence of BVI on the aerodynamic loading. The effects of BVI are the localized reductions in loading near azimuth angles  $50^\circ$  and  $300^\circ$ , where at the outboard blade distinct low-force regions are present. These locations agree with the locations of BVI events that were reported in the experimental campaign. However, the high-frequency oscillations that were associated with BVI events in the experimental data are not clearly resolved here. This suggests that the CFD solution captures the dominant unsteady loading pattern and the asymmetry of the wake, but that the wake resolution is insufficient to fully preserve the finer BVI-induced variations in blade loading.

A qualitative assessment of the wake is shown in Figure 5.39 where the vortical structures are visualised using iso-surfaces of the Q-criterion coloured by Mach number. The wake is characterised by root and tip vortices that persist over several blade passages and additional wake regions along the span of the blades. The interaction of the passing blades with the wake structure is clear in the regions near azimuth angles  $50^\circ$  and  $300^\circ$  where the presence of BVI events was reported by the HART-II campaign, and extend to other regions of the rotor disk. The preservation of the wake structures over several revolutions is limited. HART-II comparisons showed that preserving the vortical structures is essential for capturing the unsteady airloads and the aeroelastic response, and that finer wake resolution leads to an improved preservation of the wake and the prediction of the BVI-induced loading. Although the present visualisation is qualitative, it does indicate that the coupled solution partially resolves the vortex structures that are expected for the HART-II descending-flight condition.




Figure 5.39: Q-criterion iso-surface coloured by Mach number for Case C929 with a rigid rotor.

### 5.5.2. C929: Descending Flight, Unsteady CFD, and Flexible Rotor

The descending flight case was also simulated using flexible rotor blades. The earlier cases showed that the simulations that included structural flexibility were unsuccessful due to the complex aeroelastic response that is introduced. A similar result is expected for this case, which is discussed in this section.

#### Trim convergence and control evolution

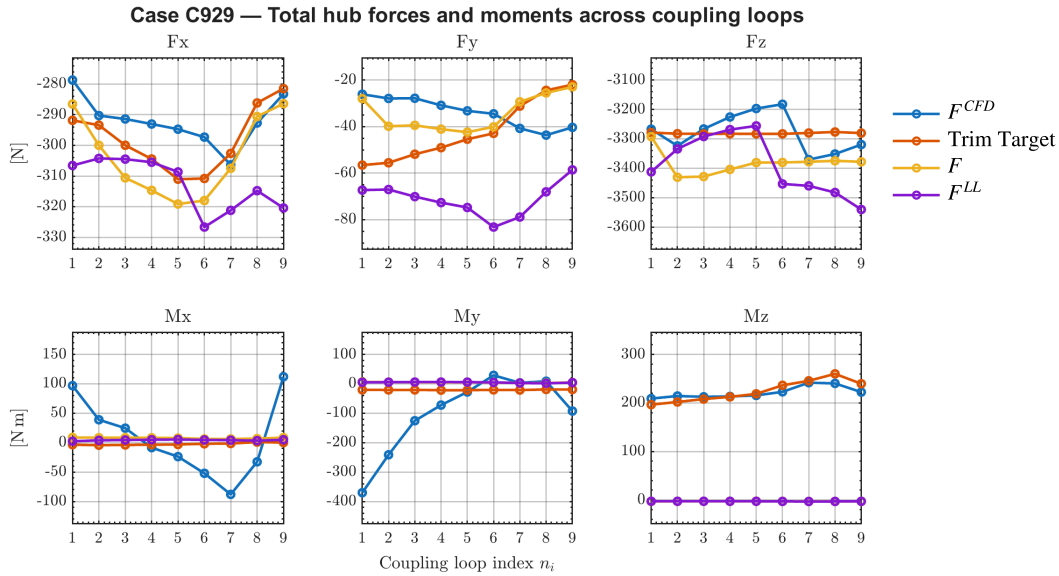
The evolution of the mean total hub forces and moments is illustrated in Figure 5.40. In the initial iteration, the agreement of the hub forces and  $M_z$  between FLIGHTLAB and ENSOLV is relatively close, whereas the  $M_x$  and  $M_y$  components differ significantly. As the coupling progresses, the behaviour of all components becomes chaotic and is no longer comparable to the global convergence behaviour that was reported earlier for successful coupling results. This is indicative of the destabilizing effect that the inclusion of the structural flexibility has when implemented unsuccessfully. Similarly, the evolution of the control states of Table 5.13 also highlights the divergence of the solution. The collective pitch initially increases steadily and reaches a peak in the seventh iteration before decreasing again instead of reaching a plateau. Moreover, the lateral and longitudinal cyclic pitches both increase and decrease with a random pattern and also do not reach a converged value. Each iteration therefore produces a different trimmed rotor motion indication that a converged global solution is not reached. Comparing these control states with the experimental values is therefore not valuable.

Table 5.13: Case C929 — Control state evolution across coupling (flexible blade)

Iteration $n_i$	1	2	3	4	5	6	7	8	9
Collective $\theta_0$ [°]	3.22	3.67	3.87	4.04	4.22	4.40	4.52	4.49	4.43
Lateral cyclic $\theta_{lc}$ [°]	-1.31	-1.61	-1.55	-1.62	-1.63	-1.51	-1.29	-1.26	-1.73
Longitudinal cyclic $\theta_{ls}$ [°]	1.06	1.10	0.91	0.75	0.58	0.39	0.68	1.19	1.81

#### Hub loads

The azimuthal distribution in loop 1 and loop 9, shown in Figures 5.41 and 5.42 illustrate the evolution of the hub forces. Initially, the aerodynamic hub forces of ENSOLV and FLIGHTLAB agree in terms of phase and mean level, but not in terms of amplitude. In particular, the amplitude of the altered solution



**Figure 5.40:** Evolution of the mean hub forces and moments versus coupling iterations for Case C929 (flexible rotor).

in FLIGHTLAB in the  $z$  direction varies by nearly  $\sim 1200\text{ N}$ , whereas for the rigid rotor the thrust varied by  $\sim 600\text{ N}$ . Moreover, the trimmed hub loading is highly oscillatory, which can be attributed to the inclusion of the structural motion when FLIGHTLAB determines the equilibrium loading based on all loading contributions. The exact interaction mechanism is not further studied in this analysis. However, this indicates that the initial rotor motion is destabilized even without the inclusion of a minimal aerodynamic correction.


Compared to the ninth loop, the  $F_x$  and  $F_y$  components of the ENSOLV and altered FLIGHTLAB loading remains relatively unchanged, whereas the magnitude of the trim target has massively increased. The addition of the correction has caused FLIGHTLAB to determine an equilibrium condition that is massively different than the original scenario, and not a solution that both solvers agree on. The same can be said for the  $F_z$  components, where the amplitude variation has been increased to  $\sim 2000\text{ N}$ . Overall, the hub forces also illustrate that the solution of the coupling changes between iterations, becomes non physical and diverges.

The azimuthal distribution in loop 1 and loop 9, shown in Figures 5.43 and 5.44, illustrate the evolution of the hub moments. In the first loop, the solutions from both solvers clearly do not agree with each other, by the solutions are nearly identical to the rigid rotor solution in the first iterations, given the added oscillations to the trim target solution. The quality of the solution in the initial iteration, before the aerodynamic correction is fully applied, is therefore still representative of the physical problem at hand.


In comparison the hub moments from the ninth iterations show once more that the quality of the solution has deteriorated. In line with the growing control states, the amplitude variation of the CFD solutions have increased significantly, and the same can be said for the trim target. Additionally, the frequency and amplitude of the oscillations in the  $z$  component of the trim target have affected. Overall, these results demonstrate once more that the inclusion of structural flexibility greatly affects the predictive capabilities of the loose coupling, highlights the sensitivity of the current implementation to a mismatch in the handling of the aeroelastic response of the rotor.

### Blade loading


Finally, the quality of the solution can be shown using the distribution of the normal force over the rotor disk, illustrated in Figure 5.45. As opposed to the rigid rotor load distribution, this distribution is not dominated by the radial distribution of lift and is not axisymmetric as is expected for this flight



**Figure 5.41:** Azimuthal distribution of hub forces for Case 929 during the first coupling iteration (flexible rotor blade)



**Figure 5.42:** Azimuthal distribution of hub forces for Case 929 during the ninth coupling iteration (flexible rotor blade)



**Figure 5.43:** Azimuthal distribution of hub moments for Case 929 during the ninth coupling iteration (flexible rotor blade)




Figure 5.44: Azimuthal distribution of hub forces for Case 929 during the ninth coupling iteration (flexible rotor blade)




Figure 5.45: Blade-normal force distribution from CFD for case C929 in loop 7 (flexible rotor blade).

condition. Instead, the distribution shows a chaotic pattern and included a peak of lift that opposes the positive vertical direction. The blade motion that is the result from this coupling is therefore non representative of a flexible rotor in descending flight and has produced a non physical result. This can once again be attributed to the aeroelastic interaction that is introduced by the structural flexibility and the inability of the current coupling implementation to handle it accurately and consistently.

### 5.5.3. Comparison with Experimental Data

This section compares the experimental airloads of the HART II with results of the loose coupling for the rigid and flexible rotors. The objective is to assess whether the coupled methodology captures the correct aerodynamic behaviour associated with descending flight and BVI. The comparison focuses on the span station  $r/R = 0.87$  that was instrumented as part of the HART-II campaign and where strongest BVI-induced loading effects were reported [30].

### Quantities of Interest

The experimental airloads that are compared as part of the validation are the sectional normal force and pitching moment. The sectional normal force coefficient is defined as

$$C_{nM2} = \frac{\frac{dL'}{dy}}{\frac{1}{2}\rho a^2 c'} \quad (5.1)$$

and the sectional pitching moment coefficient about the quarter-chord is defined as

$$C_{mM2} = \frac{\frac{dM_{c/4}}{dy}}{\frac{1}{2}\rho a^2 c^2} \quad (5.2)$$

where  $dL'/dy$  denotes the sectional lift per unit span,  $dM_{c/4}/dy$  is the sectional moment per unit span about the quarter-chord,  $\rho$  is the freestream density,  $a$  is the speed of sound, and  $c$  is the local blade chord.

In the HART-II campaign, the sectional airloads were obtained experimentally by integrating the measured surface pressure distributions at the span station. The sectional normal force per unit span was computed from the pressure difference between lower and upper surfaces as

$$\frac{dL'}{dy} = \int_0^c (p_\ell(x) - p_u(x)) dx, \quad (5.3)$$

and the quarter-chord pitching moment per unit span was determined through

$$\frac{dM_{c/4}}{dy} = \int_0^c (p_\ell(x) - p_u(x)) (x - 0.25c) dx. \quad (5.4)$$


From the CFD simulations, the sectional airloads were obtained by integrating the sectional normal force and pitching moment distributions along the span and taking the derivative in order to obtain the load distribution as a function of the span. These distribution were then evaluated at the radial station  $r/R = 0.87$  for a full revolution and normalised as given above.

### Sectional Normal Force

The comparison of the sectional normal-force coefficients over a full revolution at the radial position  $r/R = 0.87$  is shown in Figure 5.46. The experimental value is positive throughout the revolution, as expected from a lifting rotor, and varies smoothly while containing localized high-frequency oscillations near approximately  $50^\circ$  and  $300^\circ$ . These rapid fluctuations are consistent with the reported BVI that perturb the aerodynamic loading at these specific azimuths [30]. On the advancing side, the blade loading increases as the blade experiences its highest relative velocity ( $\sim 90^\circ$ ). The loading then decreases as the blade passes from the advancing to the retreating side, before rising again on the retreating side and being disturbed by BVI again.

The rigid rotor results approach the experimental mean level and overall trend reasonable well. The main lift increases and decreases are captured, including the broader lift peaks and the reductions in the BVI-affected regions. However, it does not reproduce the high-frequency oscillations that are reported by the measurements. Since these oscillations are associated with BVI-induced loading effects rather than the elastic motion such as flapping or torsion, the absence of these oscillations is not a consequence of the absence of structural flexibility. This is supported by the findings of the HART-II campaign, which found that these high-frequency oscillations are not present in the structural motion of the blade [30]. Instead, the lack of oscillations suggests that the CFD simulation does not resolve BVI or preserve the wake with sufficient accuracy or resolution.

The flexible result illustrate that adding structural flexibility strongly amplifies the variation in sectional normal force. While the sensitivity to aeroelastic phenomena is expected to increase due to the flexibility, the response is clearly exaggerated compared to the experiment. The magnitude, the irregularity of



**Figure 5.46:** Normal force predictions at 87 % radial station. Comparison between HART-II experiment, and simulated rigid and flexible rotor.

the distribution are most likely unphysical. This is consistent with the earlier findings that the loose coupling did not converge, and this results should therefore not be interpreted as a valid physical result. Overall, the rigid model result is comparable to the experimental and provides a credible approximation, while the flexible rotor shows the unstable aeroelastic response of the current implementation.


The flexible-rotor result shows the opposite tendency. Instead of providing a refinement of the rigid result, it strongly amplifies the variation in sectional normal force and introduces large peaks and valleys that are much larger than those observed experimentally. While structural flexibility is expected to alter the sectional airloads through aeroelastic coupling, the magnitude and irregularity of the predicted response are clearly non-physical. This behaviour is consistent with the earlier observation that the flexible-rotor loose coupling did not converge, and the resulting sectional normal-force distribution should therefore not be interpreted as a valid physical prediction. Overall, the rigid-rotor solution provides a credible approximation of the measured normal-force trend, whereas the flexible-rotor result reflects the unstable aeroelastic response of the current implementation.

### Sectional Pitching Moment

The sectional pitching moment, Figure 5.47, shows a poorer agreement between the experiment and the simulations than the normal force coefficient, and any direct comparison must be made with caution. The HART-II campaign reported that although the pressure instrumentation density was sufficient for integrating pressures to the normal force, the sensitivity of the moments to the sensors is much higher [30]. The reported mean values of the pitching moment may therefore be biased and less reliable than the normal force coefficient.

First, the experimental pitching moment remains negative over the full revolution, implying a nose-down tendency that opposes the direction of the torsional blade flapping. A broad minimum is visible as the blade advances into the free-stream and progresses to the retreating side. Rapid localized oscillations are also present, particularly near the regions where BVI was reported. However, the high-frequency content of the aerodynamic moment is not relevant to the blade dynamics or other physics, such as rotor noise, making it unimportant whether it is predicted. The low-frequency content on the other hand is an important source of the torsional excitation and damping and must therefore be sufficiently predicted.

The rigid rotor result only capture the mean level of the pitching moment and remains nearly constant throughout the revolution. Due to the lack of flexibility, the rigid solution does not reproduce the low-frequency variation that influences the torsional motion of the blade. Moreover, the rigid model also misses the unsteady variations, but since the high-frequency content is not the most important



**Figure 5.47:** Pitching moment predictions at 87 % radial station. Comparison between HART-II experiment, and simulated rigid and flexible rotor.

feature of the flow, this is less critical than it is for the sectional lift.

The flexible rotor result shows the opposite problem. Rather than being smooth, the azimuthal variations are strongly exaggerated. Similarly to the normal force coefficient, large peaks and valleys are present that are far greater than those of the experimental data. This indicates that the added structural flexibility has not improved the moment prediction in the present implementation, but has instead amplified an unrealistic aeroelastic interaction. This suggests that the phase and amplitude of the aeroelastic coupling are not being captured correctly, and that FLIGHTLAB and ENSOLV fail to produce an agreed-upon solution.

Overall, the moment results support the same findings drawn from the normal-force analysis. The rigid rotor provides a more credible baseline result because it preserves the correct general sign and mean level without introducing unphysical variations. The flexible rotor suffers from a worsening aeroelastic response until the rotor motion in FLIGHTLAB is unable to comply with the trim convergence criteria, and introduces non-representative variations.

#### 5.5.4. Validation Summary

Taken together, the sectional-airload comparisons show that the rigid-rotor loose-coupling result captures the dominant mean aerodynamic loading and the approximate azimuthal locations of the BVI disturbances, but does not recover the fine-scale unsteady loading associated with the experimental BVI signatures. This indicates that the present methodology produces a physically plausible baseline solution, yet still lacks sufficient wake preservation and local interaction fidelity to fully reproduce the HART-II airload response. By contrast, the flexible-rotor solution does not improve the agreement with experiment and instead amplifies non-physical variations in both the sectional normal force and pitching moment. The experimental comparison therefore confirms the earlier conclusion drawn from the integrated hub loads: the present loose-coupling implementation is capable of producing a credible rigid-rotor solution for the HART-II baseline case, but the aeroelastic coupling with structural flexibility is not yet stable or reliable.

# 6

## Conclusion

This thesis investigated the implementation, verification and assessment of a loosely coupled CFD-CA framework between ENSOLV and FLIGHTLAB based on the delta-airloads method. The framework couples the high-fidelity CFD solver ENSOLV to FLIGHTLAB, the comprehensive rotorcraft analysis environment. The objective was not only to establish the coupling framework, but also to determine whether the coupling approach represents the intended methodology, whether the numerical interface is conservative and consistent, and whether the coupled simulations produce stable, converged and physically credible rotor solutions

The results show that the successful development and use of a loosely coupled framework requires three conditions. First, the coupling workflow must represent the delta-airloads methodology and the rotor aeromechanical problem correctly. Second, the numerical interface must transfer motion, deformation, aerodynamic loads, forces, and moments consistently between ENSOLV and FLIGHTLAB. Third, the coupled solution must not only converge, but also remain physically representative and predictively reliable for the selected flight condition and rotor configuration.

Methodological consistency was achieved by extending FLIGHTLAB through the AERODELTA component, which is used to apply the aerodynamic correction. This component preserves the role division between the two solvers: FLIGHTLAB provides the trim solution, rotor motion and structural response, while ENSOLV determines the high-fidelity aerodynamic solution. The aerodynamic correction is obtained from the difference between the CFD airloads and the baseline FLIGHTLAB airloads, and is applied during the trim process of FLIGHTLAB. The corrected rotor response is then subsequently used to determine the next CFD solution. The periodic exchange of motion and loading therefore also means that the coupling procedure represents the rotor aeromechanical problem in a loosely coupled sense. The aerodynamic correction affects trim, blade motion and the structural response, while the updated motion and deformation affect the subsequent CFD solution. The framework therefore models the interaction between aerodynamics, structures and trim control that is required to model the behaviour of rotorcraft in reality. However, this representation relies on simplifications, most notably the reduction of the three-dimensional aerodynamic pitching moment to a one-dimensional input for FLIGHTLAB, that limit the extent to which the methodology is represented and, as discussed below, plausibly contribute to the instability observed for flexible rotor configurations.

Consistency of the numerical interface conditions was achieved through verification, and the verification of the reference-frame consistency, deformation transfer, aerodynamic load mapping, force and moment reconstruction, and delta-airloads computation were all confirmed within the coupling workflow. Overall, the verification established that the coupling performs the intended operations. However, this verification did not prove the predictive capability and reliability of the loose coupling and did not guarantee that every converged coupled solution represented the rotor physics accurately. While numerical convergence is a necessary condition for successful coupled simulations, it is not sufficient.

To assess the coupled solutions, a test matrix of simulations was constructed containing both rigid and flexible rotor configurations. The test matrix included hover, forward flight and descending flight, and

used both steady and unsteady CFD simulations. This allowed the coupling implementation to be assessed from simplified verification cases up to increasingly complex conditions.

The rigid rotor cases provide a clear, initial demonstration of stable and successful coupling behaviour. In hover, the coupled simulations converged smoothly and rapidly, and the delta-airloads correction evolved smoothly over the global iterations. This behaviour was supported by the use of relaxation, which limited abrupt changes in the applied correction and improved the robustness of the coupling. Across the investigated hover cases, the converged delta-airloads correction stabilised at a similar magnitude, approximately 630–660 N, regardless of whether the thrust target was 3 kN or 6 kN. This load-independence indicates that the discrepancy between the CFD and FLIGHTLAB aerodynamic models originates from a systematic modelling difference, most likely associated with the predicted induced flow and effective angle of attack, rather than from a thrust-dependent nonlinear effect. The use of unsteady CFD in hover further showed that a time-accurate aerodynamic solution did not destabilise the coupling, provided that the solution remained periodic and sufficiently regular. Beyond simplified hover, forward flight showed that the coupling could handle azimuthally varying loading, cyclic pitch inputs, and more representative operating conditions. In the descending-flight validation case, the results also reproduced the dominant mean loading and the approximate azimuthal locations of the BVI events. The rigid rotor results therefore showed that the framework is capable of predicting dominant loading trends, overall trim behaviour, and unsteady flow and loading features.

At the same time, the validation results revealed clear limitations. The comparison with the HART-II experimental data showed that the mean level predicted by the rigid rotor is comparable, but that the detailed sectional airloads are not reproduced with sufficient accuracy. In particular, the detailed BVI-events are not captured sufficiently, indicating that the wake resolution and wake preservation are not adequate to predict the detailed sectional loading fluctuations due to BVI. This finding, however, was partly expected as the choice was made to utilise a coarse CFD mesh in order to decrease the computational costs of the loose coupling.

The flexible rotor simulations on the other hand exposed the main weaknesses of the present implementation, and the addition of structural flexibility failed to successfully increase the physical accuracy of the coupled solution. The intent of the added structural flexibility was for the aerodynamic loads and moments to affect the blade deformation, which then altered the local aerodynamic state, leading to inclusion of aeroelastic feedback in the next CFD solution. However, this feedback was not controlled with sufficient consistency, and caused the rotor blades to chaotically flap and feather in the updated trim solution.

As a result, the rotor state in the completed flexible-rotor cases either became physically unrepresentative or failed to converge. Two additional cases, C400 with unsteady CFD and C402 with steady CFD, were excluded from this assessment due to a separate, identified data-handling error unrelated to the coupling methodology itself; the flexible-rotor instability reported here is therefore based specifically on the cases that ran to completion. This disrupting feedback was found to be related to the causal pathway relating the sectional pitching moment to the torsional response and section lift. Inaccuracies or inconsistencies in the aerodynamic pitching moment directly affect the elastic torsional response in FLIGHTLAB. This torsion changes the angle of attack and therefore the sectional loading distribution, which then affects the next coupling iterations and continues to compound over the next iterations as a new trim state is determined and CFD simulations are performed. The flexible validation case, demonstrated this problem clearly: the flexible solution did not improve the agreement with the experimental data, but instead produced exaggerated variations in normal force and pitching moment, and even failed to reach a converged solution. The present framework therefore is unable to provide reliable predictions for flexible rotor configurations.

Given the verification of the transfer and application of the aerodynamic loading from ENSOLV to FLIGHTLAB, the feedback appears to result from a mismatch between the value of the pitching moment correction and the structural modelling of the rotor blade. Possible causes could be related to quality of the CFD solution, the magnitude of the pitching moment, and the transformation of the pitching moment from a three-dimensional distribution to one-dimensional representation. Additionally, this could indicate that the structural model used in FLIGHTLAB to model the rotor blades does not accurately model the aeroelastic behaviour and additional structural validation is required. Due to a lack of time and limited access to data, this could not be further examined. Recommendations for future

work are included in the next chapter.

As previously mentioned, an additional limitation of the present implementation is the resolution and preservation of the rotor wake. While the coarse CFD mesh was able to capture the approximate locations of the unsteady loading features, it was unable to resolve and preserve the wake and the wake-induced loading of BVI events required for the accurate prediction. Although, the effects on the rigid rotor are limited, they become critical for flexible rotor because unsteady loads directly influence blade deformation as found in the HART-II campaign. Improved wake resolution and preservation is therefore required before the framework can be used reliably for flexible aeroelastic load prediction.

A third limitation is related to the interpretation of convergence. A stable global iteration history, a converged trim state or converged total hub loads have shown to not necessarily imply a physically representative rotor solution. This is especially important for flexible rotor simulations, where sectional loads, pitching moments, blade deformation and the phase behaviour of quantities can remain unrealistic while global quantities seem acceptable.

Overall, the ENSOLV-FLIGHTLAB framework assessed in this thesis is a verified loosely coupled CFD-CA implementation based on the delta-airloads methodology. It successfully transfers aerodynamic loads, applies the aerodynamic correction, determines an updated trimmed rotor and produced stable and converged coupled solutions for rigid rotor configurations. The current framework is therefore useful for rigid rotor analysis. However, the current framework is clearly limited in terms of its application for flexible aeroelastic rotor load prediction. The limitations are: primarily the coupled interaction between aerodynamic pitching moment and torsional response, and also the resolution and preservation of wake and wake-induced loading. Recommendations for future work are included in the next chapter.

# 7

## Recommendations

The results of this thesis show that the ENSOLV–FLIGHTLAB coupling provides a verified implementation of the delta-airloads methodology, but that further development is required before the framework can be used reliably for flexible aeroelastic rotor-load prediction. The following recommendations are proposed.

First, the prediction and transfer of the aerodynamic pitching moment should be investigated in detail. The flexible rotor cases showed strong sensitivity to the aerodynamic moment distribution, since the pitching moment directly affects elastic torsion. This torsional response alters the local angle of attack and therefore the sectional loading, which then feeds back into the next coupling iteration. Future research should therefore verify the pitching moment distribution from ENSOLV, the corresponding baseline moment in FLIGHTLAB, and the resulting applied delta-moment. This verification should be performed at the sectional level and over azimuth, and take into account how the moment distribution is preserved throughout the 3D-to-1D transformation from ENSOLV to FLIGHTLAB. Additionally, it is recommended to examine the phase relationship between the aerodynamic correction, the blade torsional response and the trim updates, since the present results suggest that the coupled aeroelastic response is not consistently represented.

Second, before further flexible simulations are performed the structural blade response in FLIGHTLAB should be verified. The present results showed that structural flexibility strongly affects the coupled solution, but the reliability of this response depends on the structural model, which for this application was assumed to remain valid on the basis of earlier validation withing the HeliNOVI project. Future research should verify the validity of the structural model in terms of the natural frequencies, mode shapes, torsional stiffness, flap response, and pitch–torsion coupling of the flexible blade model. This should be done independently of the CFD coupling, so that structural model errors can be separated from coupling errors. Two approaches are recommended to perform this verification. The first is to perform the loose, partitioned coupling using experimental airloads instead of the CFD-based correction. In this case, if the resulting blade motion and control states do not match the motion and control states reported in the experiment, this would indicate that the structural model itself is the source of the discrepancy, rather than the aerodynamic correction. The second is to construct a tight coupling, in which the aerodynamic and structural information is exchange at each time step, to assess whether the aeroelastic instability is still present when the periodic exchange interval is removed.

Third, the wake resolution and wake preservation in ENSOLV can be improved. The validation cases showed that the coupling captures the approximate locations of the unsteady loading features, but not the detailed BVI events. This affects quality of the sectional-load prediction and becomes more critical for flexible rotors. Future research should therefore investigate mesh optimization with specific attention to wake preservation over multiple rotor revolutions, beginning with a dedicated grid-sensitivity study focused on the blade tip region, the near wake, and the downstream vortex convection.

Fourth, strict criteria for CFD periodicity and residuals can be introduced to control the quality of the CFD solution before it is used in the coupling loop. The delta-airloads methods assumes that the

aerodynamic correction is a periodic rotor state. The residual histories of the unsteady simulation indicate that the flow solution does not always reduce cleanly over multiple revolutions, and if the solution has not reached a sufficiently periodic condition, the correction may contain transient effects that are then imposed in FLIGHTLAB. Future work should therefore monitor CFD apply periodicity and convergence criteria based on sectional airloads and residuals over successive iterations.

Fifth, the relaxation strategy should be further developed. Relaxation improved the stability of the rigid rotor cases, but CFD simulations are costly compared to internal FLIGHTLAB iterations, which comparatively takes only a fraction of time and computing power to run. To minimize the number of CFD simulations to reach a coupled solution, the relaxation strategy should omit the outer relaxation factor and only apply an inner relaxation in FLIGHTLAB to iteratively apply the aerodynamic correction fully during each global iteration. At each inner iteration, the relaxed correction is then applied and the rotor is trimmed until the convergence criteria are met. Each global coupling iterations applies the full delta-airloads correction and the number of CFD simulation is kept to a minimum.

Sixth, the convergence criteria of the global coupling procedure should be expanded. The present work shows that a stable global iteration history or a converged trim state does not guarantee a reliable rotor-load prediction. This is especially important for flexible rotor simulations. Future convergence assessment should include global hub loads and trim states, but also sectional normal force, pitching moment, blade deformation and phase behaviour. A coupled solution should only be considered successful when both numerical convergence and physical consistency are satisfied.


Seventh, the validation should be extended beyond the current comparison. The HART-II case provides an important reference for descending flight and blade–vortex interaction, but additional validation cases are needed to establish the general applicability of the framework. Future validation should include cases with different thrust levels, advance ratios, descent rates, and rotor configurations. Where possible, both integrated loads and sectional airloads should be compared against experimental data.

In summary, the immediate recommendation is not to broaden the scope of the framework, but to consolidate it. The present work has established a solid basis for loose CFD–CSD coupling at NLR. The next phase should focus on improving aerodynamic fidelity, introducing objective convergence criteria, and resolving the flexible aeroelastic inconsistencies that currently limit the applicability of the method.

# References

- [1] W. Johnson, "A history of rotorcraft comprehensive analyses," NASA Ames Research Center, Tech. Rep. NASA/TP-2011-215971, 2011.
- [2] A. T. Conlisk, "Modern helicopter rotor aerodynamics," *Progress in Aerospace Sciences*, vol. 37, no. 5, pp. 419–476, 2001.
- [3] M. Potsdam, H. Yeo, and W. Johnson, "Rotor airloads prediction using loose aerodynamic/structural coupling," *Journal of Aircraft*, vol. 43, no. 3, pp. 732–742, 2006. doi: 10.2514/1.14006.
- [4] R. A. Ormiston, "Revitalising advanced rotorcraft research – and the compound helicopter," *The Aeronautical Journal*, vol. 120, no. 1223, pp. 83–129, 2016.
- [5] A. Datta, "Fundamental understanding, prediction and validation of rotor vibratory loads in steady level flight," Ph.D. dissertation, University of Maryland, College Park, 2004.
- [6] A. Datta, J. Sitaraman, I. Chopra, and J. D. Baeder, "Cfd/csd prediction of rotor vibratory loads in high-speed flight," *Journal of Aircraft*, vol. 43, no. 6, pp. 1698–1709, 2006. doi: 10.2514/1.18915.
- [7] H. Bijl, A. van Zuijlen, A. de Boer, and D. J. Rixen, "Fluid-structure interaction: An introduction to numerical coupled simulation," Delft University of Technology, Faculty of Aerospace Engineering, Lecture Notes, May 2008, Course wb1417 Fluid-Structure Interaction.
- [8] C. Farhat and M. Lesoinne, "Two efficient staggered algorithms for the serial and parallel solution of three-dimensional nonlinear transient aeroelastic problems," *Computer Methods in Applied Mechanics and Engineering*, vol. 182, no. 3-4, pp. 499–515, 2000. doi: 10.1016/S0045-7825(99)00206-6.
- [9] W. Johnson, *Rotorcraft Aeromechanics* (Cambridge Aerospace Series). Cambridge, UK: Cambridge University Press, 2013, vol. 36, ISBN: 9781107028074.
- [10] P. Sarker and U. K. Chakravarty, "On the dynamic response of a hingeless helicopter rotor blade," *Aerospace Science and Technology*, vol. 115, p. 106741, 2021. doi: 10.1016/j.ast.2021.106741.
- [11] M. Smith. "Master lecture: Rotary-wing aerodynamics analysis." YouTube lecture, accessed March 15, 2026. [Online]. Available: <https://www.youtube.com/watch?v=rS01-rMSr2g>.
- [12] E. Roca León, "Aeromechanical simulations for the optimization of helicopter rotors in forward flight," Master's thesis, UNIVERSITE DE NICE-SOPHIA ANTIPOLIS - UFR Sciences, 2014.
- [13] W. J. McCroskey, "Vortex wakes of rotorcraft," NASA Ames Research Center, NASA Technical Memorandum NASA-TM-110822, 1995.
- [14] G. Li and Y. Cao, "Numerical simulation of the wake generated by a helicopter rotor in icing conditions," *Fluid Dynamics & Materials Processing*, vol. 17, no. 2, pp. 235–252, 2021. doi: 10.32604/fdmp.2021.014814.
- [15] J. G. Leishman, *Principles of Helicopter Aerodynamics*, 2nd ed. Cambridge: Cambridge University Press, 2006.
- [16] K. Mulleners and M. Raffel, "The onset of dynamic stall revisited," *Experiments in Fluids*, vol. 54, no. 6, p. 1519, 2013. doi: 10.1007/s00348-013-1519-9.
- [17] W. Johnson, "Rotorcraft comprehensive analysis for design and analysis," NASA Ames Research Center, Tech. Rep. NASA/TP-2012-216012, 2012.
- [18] A. Datta, M. Nixon, and I. Chopra, "Review of rotor loads prediction with the emergence of rotorcraft cfd," *Proceedings of the European Rotorcraft Forum*, 2005.
- [19] W. G. Bousman, "Uh-60a airloads catalog," NASA Ames Research Center, Tech. Rep. NASA/TM-2003-212265, 2003.
- [20] J. W. Lim, B. G. van der Wall, and M. Potsdam, "Airloads prediction of the hart ii rotor using loose cfd/csd coupling," *Journal of the American Helicopter Society*, 2007.

- [21] A. Datta, J. Sitaraman, and I. Chopra, "Cfd/csd prediction of rotor vibratory loads in high-speed flight," in *American Helicopter Society Annual Forum*, 2012.
- [22] B. G. van der Wall, "The hart ii test – a major step forward in rotorcraft bvi research," in *Proceedings of the European Rotorcraft Forum*, 2003.
- [23] J. S. Park, O. J. Kwon, and D. J. Lee, "Cfd/csd coupling for rotorcraft airloads prediction using dymore and kflow," *Aerospace Science and Technology*, 2013.
- [24] S. N. Jung, J. S. Park, and O. J. Kwon, "Cfd/csd coupled analysis including rotor-fuselage interaction effects," *Journal of Aircraft*, 2013.
- [25] D. You, O. J. Kwon, and J. S. Park, "Comparative study of structural modeling in cfd/csd coupling for rotorcraft," *Aerospace Science and Technology*, 2013.
- [26] M. Amiraux, A. Le Pape, and P. Beaumier, "Assessment of aerodynamic models in cfd/csd coupling for rotorcraft," in *European Rotorcraft Forum*, 2012.
- [27] B. Passe, A. Sridharan, and J. Baeder, "Computational investigation of coaxial rotor interactional aerodynamics in steady forward flight," in *33rd AIAA Applied Aerodynamics Conference*, 2015. doi: 10.2514/6.2015-2883.
- [28] J. Sitaraman, M. Potsdam, B. Jayaraman, A. Wissink, and A. Datta, "Application of the helios computational platform to rotorcraft problems," in *49th AIAA Aerospace Sciences Meeting*, 2011. doi: 10.2514/6.2011-1123.
- [29] B. Uekermann et al., "Coupling cfd and csd solvers using precice for rotorcraft applications," *Computers & Fluids*, 2016.
- [30] M. J. Smith et al., "The hart ii international workshop: An assessment of the state of the art in cfd/csd prediction," *CEAS Aeronautical Journal*, vol. 4, no. 4, pp. 345–372, 2013. doi: 10.1007/s13272-013-0078-8.
- [31] Advanced Rotorcraft Technology, Inc., *Flightlab theory manual, volume two*, Theory manual, Fremont, CA, USA, n.d.
- [32] Advanced Rotorcraft Technology, Inc., *Flightlab theory manual, volume one*, Theory manual, Sunnyvale, CA, USA, n.d.
- [33] Advanced Rotorcraft Technology, Inc., *Flightlab x-analysis user manual*, User manual, Sunnyvale, CA, USA, 2023.
- [34] J. Kok, "Numerical design of ensolv version 3.20: A flow solver for 3d euler/navier-stokes equations in arbitrary multi-block domains," Netherlands Aerospace Centre (NLR), NLR-CR-2000-620, 2000.
- [35] J. Kok, "User guide of ensolv/ensens version 9.20: A flow and adjoint solver for aerodynamic, aeroelastic, and aeroacoustic applications using 3d multi-block structured grids," Royal NLR – Netherlands Aerospace Centre, Tech. Rep. NLR-TR-2025-XXX, Jan. 2025.
- [36] S. P. Spekrijse, B. B. Prananta, and J. C. Kok, "A simple, robust and fast algorithm to compute deformations of multi-block structured grids," National Aerospace Laboratory NLR, Tech. Rep. NLR-TP-2002-105, Feb. 2002, Proceedings of the 8th International Conference on Numerical Grid Generation in Computational Field Simulations, Honolulu, Hawaii, USA, June 3–6, 2002.
- [37] O. Dieterich, "HeliNOVI report: Wind tunnel model database," HeliNOVI Consortium, European Commission GROWTH Project GRD1-2001-40113, Deliverable D2.1-1 Issue 1.0, Sep. 2002, Confidential report, 25-09-2002.
- [38] J. C. Kok, "A high-order low-dispersion symmetry-preserving finite-volume method for compressible flow on curvilinear grids," National Aerospace Laboratory NLR, Tech. Rep. NLR-TP-2008-775, Dec. 2008.
- [39] S. Voutsinas, "HeliNOVI Report: Report on Pre-Test Aeroelastic Calculations," HeliNOVI Consortium, Deliverable D 2.1-2, Jun. 2026, Issue 1.0, confidential report, EC GROWTH project GRD1-2001-40113.



# Appendix: Solver Verification

Prior to analysing the loose coupling implementation, both solvers were verified to ensure that they were correctly initialised and represented the same rotor configuration and operating condition. The details of that analysis are not repeated here. This section instead focuses on assessing the quality of the solution produces by each individual solver and verifying that the modelling assumptions are satisfied.

## A.1. FLIGHTLAB Verification

The trimmed rotor solution that is obtained from FLIGHTLAB must adhere to the modelling assumptions, namely the rigid or flexible deformation, the azimuthal invariance of the solution in hover and the periodicity of the motion. In addition, the trim targets must be satisfied and the resulting control states must be consistent with the operating condition.

### A.1.1. Periodicity of the Rotor Motion

For all flight scenarios treated in this research, the structural response must be periodic over one rotor revolution. To verify this, two revolutions of the rotor motion were exported and compared. To assess the periodicity, the spatial differences between the 33 ACP locations of the blades at the start of the revolutions were evaluated.

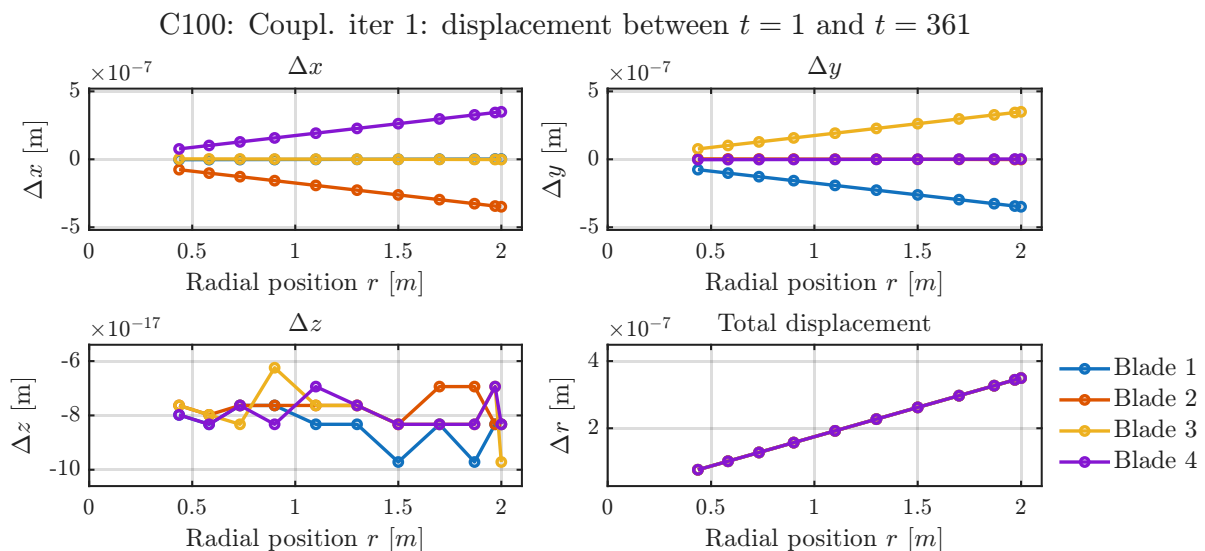


Figure A.1: Assessment of the periodicity of the rotor motion over two consecutive revolutions

Figure A.1 shows that the spatial differences in each direction and in total are negligible. The rotor motion that is output by FLIGHTLAB is confirmed to be periodic.

### A.1.2. Rigid and Flexible Behaviour

The structural rotor model used in FLIGHTLAB corresponds to the rotor geometry defined previously and two configurations are considered: a rigid rotor model and a flexible rotor model. The rotor is hingeless, thus no physical flapping or lead-lag hinges are present.

When structural flexibility is disabled the rotor behaves rigidly. As a result, rigid blade motion can only occur from prescribed pitch through the control states. Figure A.2 compares the mean vertical blade position along the span over a full revolution for the rigid rotor in hover for the initial and final iteration of the loose coupling. The z-position of the blade remains unchanged between the iterations, indicating that the blade does not deform and the pre-cone angle has been maintained.




Figure A.2: Vertical rigid blade position in hover for iterations 1 and 8. Case C100

In addition, figure A.3 shows the blade pitch along the span across a full revolution for a rigid hover scenario. This blade pitch was determined by subtracting the local spanwise twist distribution from the absolute blade pitch that is determined based on the blade motion in the inertial frame. This blade pitch should correspond to the blade pitch that is prescribed due to the control states. The magnitude of the blade pitch is not of interest here, but the behaviour of the solution is. Most importantly, the blade pitch is observed to be constant versus blade span and azimuthally invariant. This confirms the correct rigid behaviour and the azimuthally invariance of the rigid hover solution.

When structural flexibility is enabled, elastic deformation allows the rotor to flap, lead-lag and deform torsionally. Figure A.4 shows the z-position and pitch distribution as functions of the spanwise and azimuthal coordinates during the initial global iteration.

The z-displacement remains azimuthally invariant and shows that the flexible blade has flapped downward compared to the rigid rotor position. The blade pitch on the other hand gradually decreases toward the tip and torsional oscillations with a frequency of 13/rev ( $\approx 227$  Hz) appear that increase in amplitude. The presence of these oscillations indicate that the flexible degrees of freedom are active, but they do not correspond to any of the blade's reported torsional eigenfrequencies [37]. Although, these oscillate do cause the rotor motion to not be azimuthally invariant for the present work it is sufficient that these observations confirm that elastic deformation occurs, and that the resulting motion is stable and periodic. Further, detailed assessment of the structural mode shapes and their dynamic behaviour is outside the scope of this research. This structural behaviour is therefore treated as a property of the flexible rotor motor, and further analysed in the loose coupling framework.

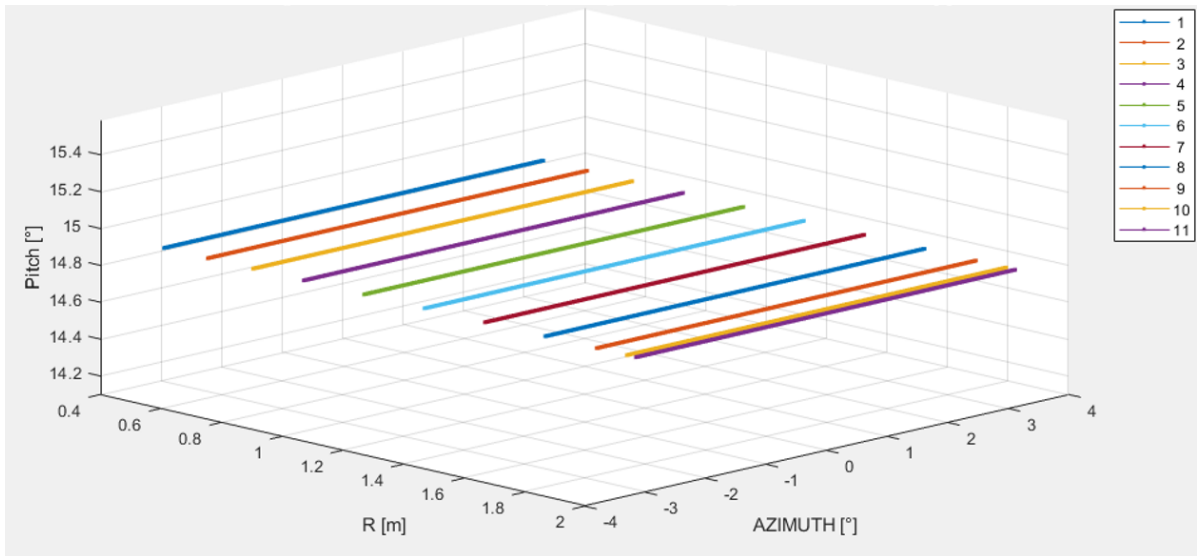


Figure A.3: Relative blade pitch deformation in the ACPs versus span and azimuth.




Figure A.4: Rotor blade flapping and pitch deformation in the ACPs versus span and azimuth. Case C400.

### A.1.3. Trim Targets and Control States

The rotor is trimmed to operating condition based on the trim targets  $F_z$ ,  $M_x$  and  $M_y$ . The convergence of the trim procedure is verified internally in FLIGHTLAB by assessing if the trim residuals are below the specified tolerance, in this case  $\approx 50N$ . Figure A.5 shows the hub loads obtained from the trimmed solution for a rigid rotor in hover. The loads satisfy the trim targets of  $F_z = 3000 N$ ,  $M_x = 0 Nm$  and  $M_y = 0 Nm$ . Moreover, the loads can be assumed to be azimuthally invariant.

The resulting collective and cyclic pitches are summarised in table A.1. For the hover case, only the collective pitch is required and the cyclic pitch inputs converge to zero. This is consistent with the expected azimuthal invariance of a rotor in steady hover.

Table A.1: Trimmed control states for Case C100 rigid hover.

Control Input	Value [°]
Collective pitch $\theta_0$	5.5
Longitudinal cyclic $\theta_{1c}$	0.0
Lateral cyclic $\theta_{1s}$	0.0




Figure A.5: Hub loads for the trimmed rigid hover solution (Rigid hover case C100).

## A.2. Verification of the CFD solution

Once the CFD has been correctly set-up, the results of the simulations need to be analysed to ensure that they are converged, stable and suitable to be transferred to FLIGHTLAB. This section assess the CFD solutions for both the steady and unsteady simulations that are performed by ENSOLV. The verification focuses on the convergence of the governing equations as indicated by the time history of the residuals, the convergence behaviour of the rotor loading, and the periodicity of the rotor loads in case of the unsteady simulations. Checking the quality of the simulations ensures that flow transients have decayed and that the aerodynamic loading is correct and consistent with the assumptions. Coincidentally, it informs the user whether the set number of time-step is insufficient, which causes the simulations to be terminated prematurely. Establishing the quality of the CFD solutions is a prerequisite for the subsequent coupling and validation steps.

### A.2.1. Residuals

Residuals are used to assess whether the governing equations are being satisfied throughout the computational domain and therefore provide a first indication of numerical convergence. To evaluate the convergence history of these CFD simulations, both Root-Mean-Square (RMS)- and maximum-norm residual measures are used for the continuity, momentum, and energy equations as a function of the iteration index. The RMS residual corresponds to the  $L_2$  norm of the residual vector  $r$ , normalised by the number of cells in the mesh, and is defined as

$$L_2 = \left( \frac{1}{N} \sum_{i=0}^N |r_i|^2 \right)^{1/2}, \quad (\text{A.1})$$

where  $r_i$  is the local residual vector in cell  $i$  and  $N$  is the number of cells in the mesh. This norm provides greater weighting to larger residual components due to the quadratic term and is therefore representative of the global convergence behaviour of the solution. In addition, the maximum-norm residual corresponds to the  $L_\infty$  norm and is defined as

$$\|r\|_\infty = \max |r_i|, \quad (\text{A.2})$$

and considers the highest residual in the domain, thus making it sensitive to localised, increased residuals that may persist even when the overall solution appears converged in an RMS sense.

For the steady simulations, convergence of the solution is expected to show as the monotonic reduction of the residuals toward a low and nearly constant level. In the case of the unsteady simulations, the residuals are not expected to decay as much since the flow solution remains time-dependent and periodic over the rotor revolution. In that case, the residual history can instead be assessed for boundedness, repeatability and the absence of divergence. Solely the residuals are not sufficient to confirm that the aerodynamic solution is converged and suitable to be transferred to FLIGHTLAB, but they are essential to verify the numerical quality of the CFD solution.

### A.2.2. Steady CFD solution behaviour

Figure A.6 shows the convergence history of the continuity, momentum, and energy equations, using the Root-Mean-Square (RMS)- and maximum-norm measures as a function of the iteration index in the initial CFD simulation of the loose coupling.




Figure A.6: RMS- and max-norm residuals of the continuity, momentum, and energy equations vs. the iteration index. Case C100-RGD-STD initial CFD simulation.

In the initial steady CFD simulation, all residuals decrease monotonically relative to their original values, indicating that initial imbalances in the cells are being resolved. After  $t \approx 3000$  iterations, the residuals reach an plateau and further iterations are unable to improve the solution. This could be attributed to local flow phenomena, the mesh discretisation or limitations of the aerodynamic model. Overall, the RMS residuals converge smoothly to low levels, indicating global convergence of the solution. The  $L_\infty$  residuals remain higher, indicating that local region of imbalance persist in the solution. The final solution is free of oscillations and shows comparable trends for the mass, momentum and energy equations. An appropriate number of time-steps for the steady CFD simulation is at least  $t = 3500$ .

### Aerodynamic loading

To connect the residual convergence to physical quantities, the convergence of the aerodynamic loading is assessed by analysing the total aerodynamic thrust  $F_z$  as a function of the simulation time-step. Figure A.7 shows the force evolutions of  $F_z$ . Initially, the force oscillates as the flowfield adjusts from the initial condition and shows noticeable transients. As the number of iterations is increased, these oscillations decay and the force converges to a stable value. The force can be observed to converge at an earlier iteration than the residual, which indicates that despite the residuals an converged aerodynamic solution has been obtained.




Figure A.7: Total lift for the steady C100 hover case.

### A.2.3. Unsteady CFD solution behaviour

In the unsteady simulations, the full rotor revolution is resolved in time and the flow-field is simulated over a set number of revolutions. Figures A.8 and A.9 show the residual histories for the first and final loose coupling iterations of rigid hover simulated using an unsteady simulation.

In contrast to the steady simulation, the residuals exhibit a repeating pattern during each revolution. A characteristic box pattern, with an increase and decrease that occur approximately at the same azimuthal positions in every revolution can be identified. The residual magnitudes do not reduce at all as the simulation progresses, and the pattern remains present in the final coupling iteration. This behaviour indicates that local imbalances persist throughout the simulation, which appear linked to spatial positions in the domain. A clear cause or solution to this phenomenon was not identified and is therefore a real issue associated with the unsteady CFD simulations. Given the overarching goal of verifying the implementation of the coupling the choice was maybe to treat this behaviour of a property of the numerical set-up. Although, the residuals do not converge they remain bounded and consistent over successive iterations.

To assess the physical convergence of the solution, the total thrust  $F_z$  is evaluated over the final two revolutions of the simulation, as shown in Figure A.10. In the initial coupling iteration, the thrust is not azimuthally invariant or periodic, and exhibits a gradual increase over the revolutions. Since hover should lead to identical azimuthal conditions, this behaviour indicates that the CFD flow-field has possibly not fully developed into a hover state. Exporting this solution to FLIGHTLAB would introduce an asymmetry and a discontinuity between the start and the end of the revolutions. In the final coupling




Figure A.8: Residual history for the unsteady simulation during coupling loop 1 (Case C100).




Figure A.9: Residual history for the unsteady simulation during coupling loop 8 (Case C100).

iteration, the thrust variation is greatly reduced and remains within  $\approx 1N$  during a full revolution. The loading has approximately become periodic and azimuthally invariant, although small variations remain. This solution is sufficiently constant for the load transfer to FLIGHTLAB.




Figure A.10: Lift force distributions for unsteady CFD (Case 100) at different coupling loops.

# B

## Appendix: Pre-Processing of CFD solution for Data Exchange

This section describes the details of the pre-processing operation that ensures the representation of the CFD solution is compatible with the representation in FLIGHTLAB and can be mapped correctly. This operation will be referred to as Centre to Vertex.

In ENSOLV, the aerodynamic solution is saved in terms of the integral forces. The integral forces are vector functions and act at the cell-face centre of each surface cell. They are defined as follows,

$$\mathbf{F}_{cell} = \frac{\mathbf{F}_{faces}}{\frac{1}{2} \rho_{\infty} u_{r,\infty}^2 S_A}, \quad (\text{B.1})$$

where  $\rho$  is free-stream air density,  $u_{r,\infty}$  the reference free-stream velocity and  $S_A$  the reference cell area. The force vector  $\mathbf{F}_{faces}$  acting on the cell faces is defined as,

$$\mathbf{F}_{faces} = - (p - p_{\infty}) \mathbf{A} - \rho \mathbf{u} (\mathbf{u} - \mathbf{v}) \cdot \mathbf{A} + \boldsymbol{\tau} \cdot \mathbf{A}, \quad (\text{B.2})$$

where  $p$  is the local pressure,  $p_{\infty}$  the freestream pressure,  $\rho$  the density,  $\mathbf{u}$  the flow velocity,  $\mathbf{v}$  the face (or grid) velocity,  $\boldsymbol{\tau}$  the viscous stress tensor, and  $\mathbf{A}$  the face area vector. The force on each cell is composed of the contributions due to pressure force, convective momentum and viscous forces. [35]


To map the aerodynamic loads from ENSOLV to FLIGHTLAB, the transformation that mapped the deformation can be reversed. This transformation however, relates the cell vertices of the structural grid to the vertices of aerodynamic grid, and not the cell-face centres. In order to reverse the mapping, the integral forces need to be represented at cell vertices and thus transformed. Figure B.1 shows how the cell face-centred grid and the cell vertex-grid relate to each other.

To transfer the integral forces from the cell face-centre to the cell vertices, each rectangular cell face is divided into four equal quarter areas. The face-centred force  $\mathbf{F}_f$ , acting on the face  $A_f$ , is redistributed to the four corner vertices based on the weight of the associated quarter area. For rectangular cells, each quarter has an area  $A_f/4$ , resulting in a force contribution

$$\mathbf{F}_{f \rightarrow v_k} = \frac{A_f/4}{A_f} \mathbf{F}_f = \frac{1}{4} \mathbf{F}_f, \quad k = 1, \dots, 4, \quad (\text{B.3})$$

where the index  $k$  denotes the four vertices of the face  $f$ . The redistributed total force  $\mathbf{F}_v$  at a vertex ( $v$ ) is obtained by summing the contributions from all adjacent faces,

$$\mathbf{F}_v = \sum_{f \in \mathcal{F}(v)} \frac{1}{4} \mathbf{F}_f. \quad (\text{B.4})$$



**Figure B.1:** Left: cell face-centred grid. Middle: cell vertex grid. Right: overlap of both.

This area-based redistribution conserves the integral forces and allows for the use of the spline mapping to the structural grid. The corrected redistribution was implemented in the coupling framework and subjected to verification. To verify the force conservation, the total force obtained from the CFD solver was compared to the total force obtained after the redistribution to the cell vertices. The results are summarized in Table B.1. The total force in  $x$  and  $y$  have been omitted as they are approximately zero in hover flight.

**Table B.1:** Verification of force conservation during face-centre to vertex redistribution.

Quantity	Total force $F_z$ [N]
CFD (face-centred)	-2530
After redistribution (Vertex-centred)	-2530

The redistributed force corresponds to the original CFD force total exactly. The verification therefore confirms that the updated redistribution conserves the integral forces during the transformation from cell face-centres to cell vertices.