Non-linear Dynamics of the Flared Folding Wingtip Concept

Development and Application of a Non-linear Aeroelastic Framework for Gust-Release Dynamics

Rik Lambrichs

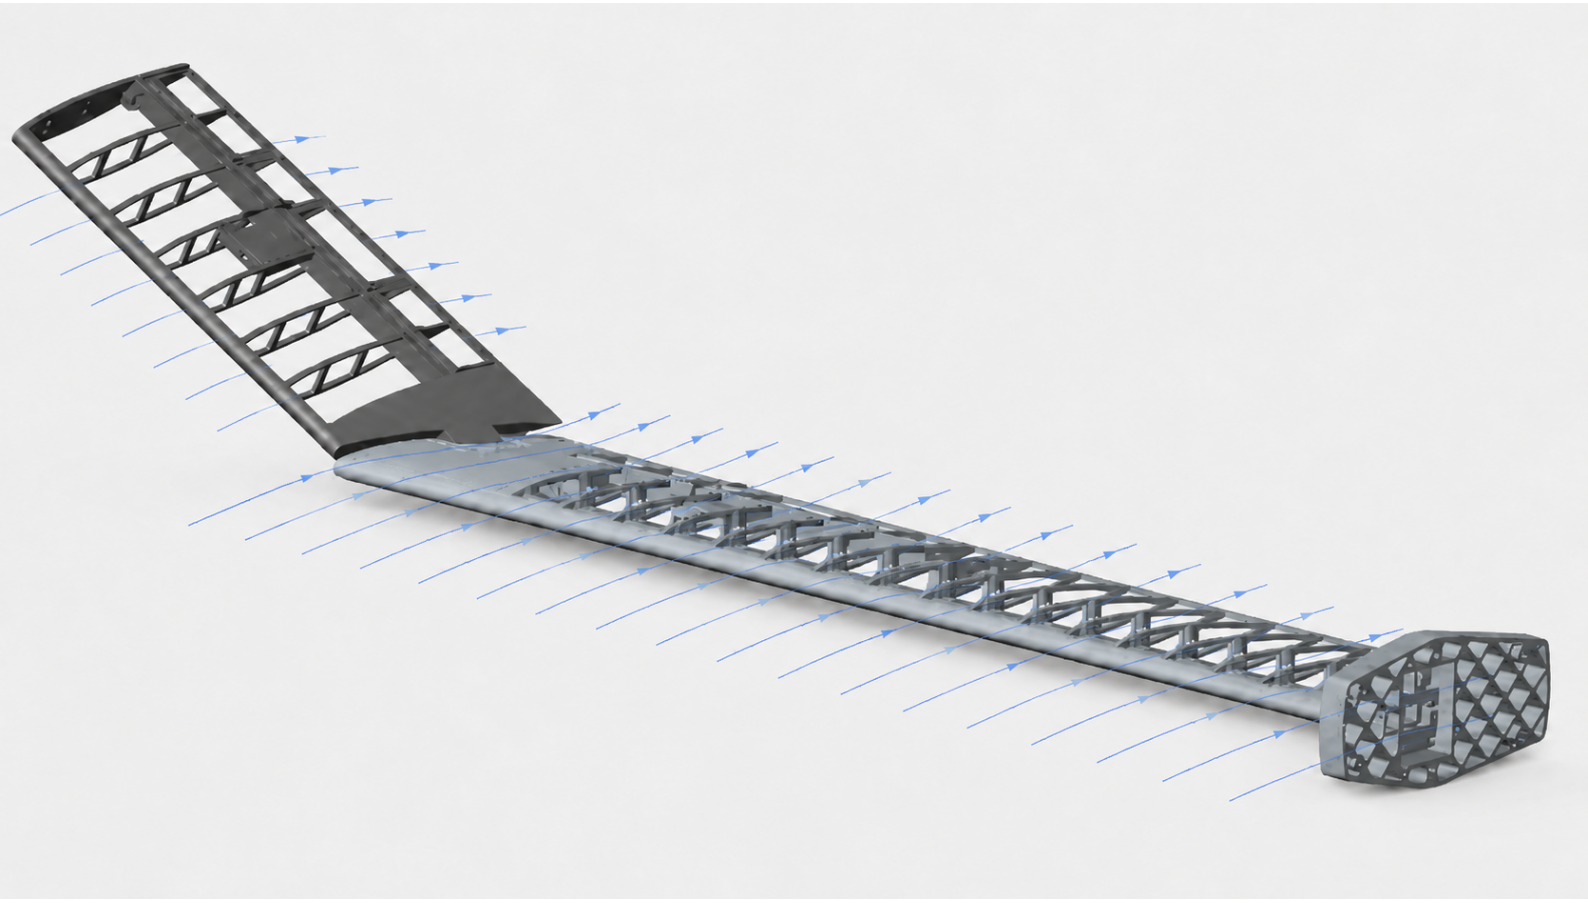

TU Delft

Non-linear Dynamics of the Flared Folding Wingtip Concept

Development and Application of a Non-linear Aeroelastic Framework for Gust-Release Dynamics

by

Rik Lambrichs

to obtain the degree of Master of Science

at the Delft University of Technology,

to be defended publicly on Friday June 5, 2026 at 09:30 AM.

Thesis committee:

Dr. ir. M.F.M. Hoogreef,	TU Delft, chairperson
Dr. C. Varriale,	TU Delft, supervisor
Dr. J. Sodja,	TU Delft, supervisor
Dr. X. Wang,	TU Delft, external examiner

Place:	TU Delft Faculty of Aerospace Engineering, Delft
Project duration:	September, 2025 – June, 2026
Student number:	5645069

Cover:	Visualisation of the created aeroelastic FFWT model and imposed gust field
Style:	TU Delft Report Style, modified by Carmine Varriale

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

TU Delft

Preface

This thesis marks the completion of my Master's degree in Aerospace Engineering at Delft University of Technology. The work presented in this report was carried out as part of the Flight Performance and Propulsion track and focuses on the non-linear dynamics of the Flared Folding Wingtip concept for gust-load alleviation.

The project has been both challenging and rewarding. It brought together several aspects of aerospace engineering that I have become increasingly interested in during my studies: aeroelasticity, structural dynamics, aerodynamic modelling and numerical simulation. Developing a coupled Simscape-UVLM framework required me to work across different modelling environments, to make many practical implementation choices, and to critically assess the reliability of the results obtained from them. Although not every step was straightforward, the process has taught me a great deal about building, validating and interpreting complex aeroelastic models.

I would like to thank my supervisors, Carmine and Jurij, for their guidance, feedback and critical questions throughout the project. Their input helped shape the direction of the work and encouraged me to keep improving both the model and the interpretation of the results. I am also grateful to the researchers and students whose previous work on the Flared Folding Wingtip formed an important foundation for this thesis.

A special thanks goes to my friend Tom, with whom I spent many hours working together at the HSL. Having someone alongside me during the long study days made the process not only more productive (at times), but also more enjoyable. I would also like to thank my girlfriend, Lonne, for her patience, support and understanding throughout this period, especially for accepting the many late-night study sessions that came with finishing this thesis and the masters as a whole.

Finally, I would like to thank my family, friends and fellow students for their support during my studies and throughout the thesis process. Their encouragement, patience and occasional distractions were all valuable in their own way.

Rik Lambrichs
Nieuwegein, May 2026

Summary

High-aspect-ratio wings offer an effective route toward improved aerodynamic efficiency, since increasing wingspan reduces induced drag. This benefit is however accompanied by increased structural loads, stronger aeroelastic coupling and practical limitations related to airport infrastructure. The Flared Folding Wingtip (FFWT) concept addresses these competing requirements by combining a folding wingtip with a hinge axis that is flared relative to the freestream. When released during a gust encounter, upward folding introduces a local nose-down change in wingtip incidence. This can unload the outboard wing and reduce the wing-root bending moment (WRBM), while allowing the wingtip to remain locked during nominal flight.

The objective of this thesis is to investigate the non-linear gust-response behaviour of an FFWT model and assess how frequency-scaled release trends, hinge properties and hinge-moment-based release thresholds affect its load-alleviation performance. A non-linear time-domain aeroelastic numerical framework is therefore developed. The framework couples a Simscape Multibody model of a flexible main wing and rigid folding wingtip to a UVLM-based aerodynamic solver. The flexible main wing is represented using a Craig-Bampton reduced-order model, while the wingtip is connected through a revolute joint aligned with the flared hinge axis. The aerodynamic model is evaluated on the instantaneous lifting-surface geometry, so that large wingtip rotations directly affect the aerodynamic lattice, panel normals, wake shedding and resulting loads. This enables the framework to capture the large-rotation, release-triggered, non-linear aeroelastic response that is central to the FFWT concept.

The framework is verified through structural, aerodynamic and coupled aeroelastic checks, and validated at trend level using previous FFWT gust-release experiments performed at TU Delft. The numerical model is not intended to reproduce the experimental setup, but captures the governing release-timing hierarchy: early release provides load alleviation, delayed release reduces the benefit, and release near the locked peak-load instant can become detrimental. This establishes the framework as a suitable tool for comparative non-linear FFWT release studies.

The simulation campaign is organised into three stages. Stage I performs a frequency-scaled release-trend assessment, examining whether the release-timing hierarchy observed in previous FFWT experiments remains visible when the gust cases are scaled to the relative-frequency range of the present structural model. The results show that release effectiveness is governed mainly by the phase of the wingtip motion relative to the locked WRBM build-up. Pre-release gives peak WRBM reductions of approximately 29–38%, while release at gust onset remains beneficial but becomes less effective with increased gust frequency. Release at the locked peak-load instant is consistently detrimental. A complementary reduced-frequency interpretation shows that large local angle-of-attack relief is not sufficient on its own, the unloading must also occur early enough relative to the structural load build-up.

Stage II fixes gust input and release instant using a Stage I reference case and varies only hinge stiffness and damping. The results show a clear trade-off between load alleviation and hinge-motion demand. Low hinge stiffness gives the largest peak WRBM reduction, but requires large hinge-angle excursions. Increasing hinge stiffness restricts wingtip motion and moves the response toward the locked-hinge response, governing peak WRBM reduction. Hinge damping mainly affects the post-release response and hinge-angle decay. Based on this trade-off, aggressive, balanced and conservative hinge-property candidates are selected for the final release-threshold study.

Stage III evaluates hinge-moment-based release thresholds for these candidates. A low, non-zero threshold can improve the response relative to simple gust-onset release, provided that release occurs during the favourable load-build-up phase. The aggressive candidate gives the largest peak WRBM reduction of 23.8%, but requires a peak hinge angle of approximately 47°. The balanced candidate reaches a similar WRBM reduction of 21.2%, while reducing the peak hinge angle to approximately 29°. The conservative candidate limits the peak hinge angle to about 14°, but its best WRBM reduction is limited to 10.9%.

Overall, the FFWT should be interpreted as a phase-sensitive aeroelastic load-alleviation mechanism rather than as a simple load-shedding device. Effective release requires the wingtip motion to develop before the critical WRBM peak is formed. Among the tested cases, the balanced hinge-property candidate with a low non-zero hinge-moment threshold provides the most favourable compromise, retaining most of the load-alleviation benefit of the aggressive setting while requiring substantially less wingtip motion.

Table of contents

List of Figures

List of Tables

List of Symbols

Roman letters

A	AIC matrix, -
AR	aspect ratio, -
b	span, m
b	UVLM RHS, m/s
C	aero. coefficient, -
c	chord, m
$\mathbf{C}_{\mathrm{CB}}$	CB damping matrix, consistent units
$C_{D,i}$	induced-drag coefficient, -
C_L	lift coefficient, -
$\mathbf{C}$	damping matrix, consistent units
C_θ	hinge damping, N·m·s/rad
D	drag, N
e	Oswald factor, -
F	force vector, N
$\mathbf{f}_a$	aero. load vector, N
f_g	gust frequency, Hz
$\hat{f}_g$	relative gust frequency, -
$\mathbf{F}_k$	frame force, N
f_s	structural frequency, Hz
$\mathbf{f}_s$	structural load vector, N
g	gravity, m/s^2
H	gust-gradient distance, m
I	identity matrix, -
I_{hinge}	hinge-axis inertia, kg·m^2
$\mathbf{K}_{\mathrm{CB}}$	CB stiffness matrix, consistent units
k_g	gust reduced frequency, -
K	induced-drag factor, -
$\mathbf{K}$	stiffness matrix, consistent units
K_θ	hinge stiffness, N·m/rad
L	lift, N
L_i	turbulence length scale, m
L_w	wake length, m
M	moment, N·m
$\mathbf{M}_{\mathrm{CB}}$	CB mass matrix, consistent units
M_h	hinge moment, N·m
$M_{h,3}$	third hinge-moment component, N·m
$M_{h,\mathrm{thr}}$	hinge release threshold, N·m
$\mathbf{M}_k$	frame moment, N·m
$\mathbf{M}$	mass matrix, consistent units
M_{min}	critical WRBM peak, N·m
N_b	bound panels, -
n_b	boundary DOFs, -
N_c	chordwise panels, -
N_f	coupling frames, -
n_m	retained modes, -
N_s	spanwise panels, -
N_{surf}	lifting surfaces, -
N_w	wake columns, -
P_{int}	interface power, W
$\mathbf{Q}_{\mathrm{CB}}$	CB force vector, consistent units
$\mathbf{q}_{\mathrm{CB}}$	CB coordinate vector, consistent units
R	range, m
$\mathbf{R}_k$	frame rotation, -
$\mathbf{r}_k$	frame position, m
$\mathbf{r}$	residual vector, consistent units
S	reference area, m^2
s	gust penetration distance, m
$\mathcal{S}$	structural state, -
t	time, s
$\mathbf{T}_{\mathrm{AS}}$	aero-structural transfer, -
$\mathbf{T}_{\mathrm{CB}}$	CB transformation, -
t_{end}	end time, s
t_g	gust-arrival time, s
$t_{M,\mathrm{locked}}$	locked-peak time, s
T_{post}	post-gust window, s
t_{rel}	release time, s
U	design gust velocity, m/s
$\mathbf{u}$	displacement vector, m and rad
$\mathbf{u}_b$	boundary displacement, m and rad
$\mathbf{u}_i$	interior displacement, m and rad
$u(s)$	gust velocity profile, m/s
V	airspeed, m/s
V_c	wake convection velocity, m/s
$\mathbf{V}_{\mathrm{ind}}$	induced velocity, m/s
V_∞	freestream velocity, m/s
$\mathbf{V}_{\mathrm{motion}}$	surface-motion velocity, m/s

V_w	wake transport velocity, m/s
WRBM	wing-root bending moment, N·m
$\mathbf{x}_i$	lattice vertex position, m
$\mathbf{x}_{i,k}^0$	reference vertex coordinates, m
$\mathcal{X}^0$	reference lattice vertices, m

Greek letters

α	angle of attack, deg
β	sideslip angle, deg
δ	deflection / perturbation, deg
$\Delta\alpha$	AoA change, deg
$\Delta\alpha_{\text{tip,kin}}$	kinematic tip AoA relief, deg
$\Delta\alpha_{\text{tip,rel}}$	relative tip AoA relief, deg
δ_g	gust amplitude, deg
ΔM_h	hinge-moment increment, N·m
$\Delta\phi$	wake potential jump, m^2/s
Δt	UVLM time step, s
ΔWRBM	WRBM reduction, %
Δx_w	wake spacing, m
η	threshold fraction, -
$\boldsymbol{\eta}$	fixed-interface coordinates, -
$\boldsymbol{\Gamma}_b$	bound circulation, m^2/s
$\boldsymbol{\Gamma}_w$	wake circulation, m^2/s
Λ	hinge flare angle, deg
Ω	spatial angular frequency, rad/m
$\boldsymbol{\omega}_k$	frame angular velocity, rad/s
$\boldsymbol{\Phi}_n$	fixed-interface modes, -
$\phi(\Omega)$	turbulence PSD, m^3/s^2
$\boldsymbol{\Psi}_c$	constraint modes, -
ρ	air density, kg/m^3
σ_i	turbulence intensity, m/s
τ	time constant, s
θ	fold angle, deg
$\dot{\theta}$	fold rate, rad/s
$\theta_{\max}$	peak fold angle, deg
ζ_h	hinge damping ratio, -

List of Acronyms

AR	Aspect Ratio
ASET	Analysis Set
CAD	Computer-Aided Design
CB	Craig-Bampton
CFD	Computational Fluid Dynamics
CMS	Component Mode Synthesis
CS-25	Certification Specifications for Large Aeroplanes
CT	Continuous Time
DLM	Doublet Lattice Method
DMIG	Direct Matrix Input at Grid
DOF	Degree of Freedom
DOFs	Degrees of Freedom
EASA	European Union Aviation Safety Agency
EU	European Union
FE	Finite Element
FEM	Finite-Element Model
FFWT	Flared Folding Wingtip
FOM	Full-Order Model
GHG	Greenhouse Gas
GLA	Gust Load Alleviation
GVT	Ground Vibration Test
IPCC	Intergovernmental Panel on Climate Change
MATLAB	Matrix Laboratory
MLA	Manoeuvre Load Alleviation
ODE	Ordinary Differential Equation
OP2	Nastran binary output file
OP4	Nastran Output4 matrix file
PCH	Nastran punch file
PROTEUS	TU Delft aeroelastic simulation framework
PSD	Power Spectral Density
PTWT	Passive Twist Wingtip
RANS	Reynolds-Averaged Navier–Stokes
RBE2	Rigid Body Element type 2
RBE3	Rigid Body Element type 3
RMS	Root Mean Square
ROFS	Reduced-Order Flexible Solid
ROM	Reduced-Order Model
SAF	Sustainable Aviation Fuel
SAH	Semi-Aeroelastic Hinge
SOL101	Nastran linear static solution sequence
SOL103	Nastran normal modes solution sequence
SOL144	Nastran static aeroelastic solution sequence
TSFC	Thrust-Specific Fuel Consumption
UVLM	Unsteady Vortex Lattice Method
VLM	Vortex Lattice Method
WRBM	Wing-Root Bending Moment
XFLR5	XFLR5 aerodynamic analysis tool

Part I

Research Definition

1

Introduction

Global warming has already warmed the planet by about $1.1°C$ on average and is amplifying extreme weather events and increasing risks like heat and drought, flooding and water scarcity, according to the latest synthesis assessment of the Intergovernmental Panel on Climate Change (IPCC) [1]. In response to global warming, the 2015 Paris Agreement set an international long-term goal to try to hold the rise in average temperature well below $2°C$ and even try to limit it to 1.5°. In addition to this, requirements were set for regularly strengthened national plans regarding finance aligned with low-emissions and development towards climate resilience [2]. In the Netherlands, these aims are translated into law and policy: the European Climate Law as devised by the European Union (EU) binds the Netherlands to climate neutrality by 2050, with at least a 55% net emissions cut by the year 2030, compared to the emissions in 1990. The Dutch government committed to this through the Dutch Climate Act and brought it into operation with the 2019 National Climate Agreement (the *Klimaatakkoord*) and its successive policies [3]. Recent assessments by the Netherlands Environmental Assessment Agency (PBL) indicate that projected 2030 reductions are in the range of $\sim$ 50% [4], which is close but still short of the legal 55% target, highlighting the need for additional action in the current decade.

A major driver of global warming is the emission of greenhouse gases (GHGs), to which the transport sector is a significant contributor. Aviation specifically accounts for roughly $2 - 2.5\%$ [5] of global energy-related fossil-fuel CO_2 emissions, but its total warming impact is even higher once non-CO_2 effects like contrails and NO_x emissions are included. Best-case estimates find that these non-CO_2 effects add warming on the same order as, or even greater than, CO_2 emissions alone [6]. As air traffic is rising back up after the COVID pandemic, the CO_2 levels in the aviation sector rose to a yearly total of approximately 950 Mt in 2023, a trend that underscores a widening gap with the Paris temperature goals if not mitigated [5]. In response, EU policies now target the aviation sector directly, where the ReFuelEU Aviation regulation mandates increased use of Sustainable Aviation Fuels (SAFs), as a measure to decrease emissions and reduce non-CO_2 impacts by a rising SAF blend of 2% in 2025 up to 70% by 2050 [7]. These SAFs are low-carbon alternatives to conventional jet fuel, produced from sustainable sources. Another class of these fuels are electricity-based SAFs (e-SAFs), which are synthesized using renewable electricity, green hydrogen and captured CO_2. Using these, the Netherlands is positioning itself within the EU regulations through an e-SAF strategy and measures at airports, for example at Schiphol Airport, where the government has proposed a cap on the annual movements around 478–480 thousand movements per year under the EU "Balanced Approach" [8].

This focus on fuel switching and GHG reduction aligns with developments on the propulsion side of aviation, where current and upcoming engines improve thermal and propulsive efficiency through geared fans, ultra-high bypass ratios and new materials. Parallel to this, alternatives regarding energy carriers, such as hydrogen and hybrid-electric options are explored, like Airbus's ZEROe program which uses hydrogen combustion and fuel-cells [9]. The Breguet range equation for jet aircraft, as presented in Equation 1.1 [10], provides insight into why the propulsion plays a large role.

$$R = \frac{V}{c}\frac{L}{D}\ln\left(\frac{W_i}{W_f}\right) \tag{1.1}$$

This equation relates the aircraft range R to the ratio of initial to final weight via $\ln\left(\frac{W_i}{W_f}\right)$, which is intuitive: for a given aircraft and flight condition, carrying a larger fuel fraction, allows a larger distance that can be covered. However, the equation also identifies two major levers for improving aircraft efficiency, namely the propulsion through the Thrust Specific Fuel Consumption c (TSFC) and the aerodynamic efficiency $\frac{L}{D}$. In this equation, V is the freestream velocity.

The aerodynamic efficiency is determined by two factors, and can thus be improved by focusing on two objectives: improving the lift, while reducing the drag. One of the core limiters of the aerodynamic efficiency when generating lift is induced drag, which is a penalty caused by the downwash associated with the trailing vortex sheet shed by a finite wing and is characterized by the standard relation presented in Equation 1.2 [11].

$$C_{D,i} = KC_L^2, \qquad K = \frac{1}{\pi e AR} \tag{1.2}$$

This shows that the induced drag is defined by a constant times the lift coefficient squared and is inversely proportional to the product $\pi e AR$. The span-efficiency factor e, or Oswald efficiency factor, accounts for how closely the actual spanwise lift distribution approaches the ideal elliptical distribution, where $e = 1$. The other factor, the Aspect Ratio (AR), is defined as the wingspan squared over the wing planform area, as presented in Equation 1.3.

$$AR = \frac{b^2}{S} \tag{1.3}$$

Increasing either e or the AR lowers the induced drag and with that improves the aerodynamic efficiency [12]. Longer, more slender wings improve the AR, but in practical design this means finding a balance between gains in wingspan against higher structural bending loads, aeroelastic constraints and gate limits at airports. This is why the industry often boosts the effective span and modifies the flow around the tip of the wing using wingtip devices like for example Whitcomb winglets, depicted in Figure 1.1.

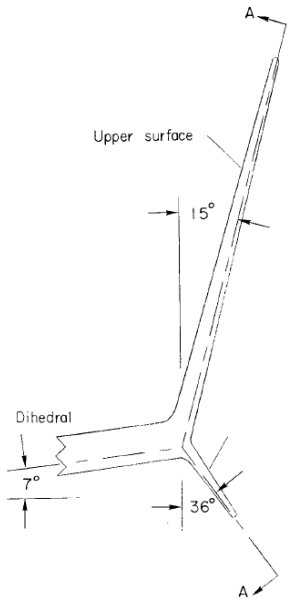

Figure 1.1: Diagram of Whitcomb's winglet, retrieved from [13].

These devices cut induced drag substantially by moving the lift distribution closer to elliptic, thus increasing e in Equation 1.2 and thereby deliver fleet fuel-burn reductions on the order of a few percent [14]. Another approach is to simply add span, by either high-AR designs or truss-braced wings like in NASA's Transonic Truss-Braced Wing study, in which wind-tunnel tests indicate a fuel-burn improvement of $\sim$ 8-10% compared to conventional wings [15]. In order to extend the wingspan, but avoid issues with gate limits, Boeing designed its newest model, the 777X, with wingtips that are locked in deployed position during flight and folded only on the ground [16].

Another approach more focused on aeroelasticity is the Flared Folding Wingtip (FFWT) concept, which combines the span extension with an additional Degree of Freedom (DOF) in flight. An FFWT is a folding wingtip in which the hinge line is rotated by a flare angle, so that it is no longer parallel to the freestream. When the hinge is released, loads induced by gusts and manoeuvres cause the tip to rotate and shed lift. In this way, the FFWT builds on the operational idea of folding wingtips, but is explicitly designed to move in flight as a passive or actively controlled load-alleviation device.

In recent years, substantial research has focused on folding wingtips, and specifically on the FFWT, as a means of capturing high-AR benefits in cruise while alleviating loads in off-design conditions. Previous work at TU Delft has shown that releasing the flared tip during gust encounters can reduce WRBM, but has also highlighted the importance of hinge design, release timing and post-gust dynamic behaviour [17–19]. These findings motivate the present thesis, which investigates the extent to which the benefits of the FFWT can be realised in a dynamically robust manner. Before this concept and its state of the art are treated in Chapter 2, several subjects that are essential for understanding the key aspects of the FFWT are addressed in the next section. The chapter is concluded by an outline of the thesis.

1.1. Background Concepts

To prepare for the detailed discussion of the FFWT concept and its state of the art in Chapter 2, this section introduces the key physical and regulatory notions on which the remainder of the thesis builds. First, the basics of aeroelasticity are outlined, emphasizing how aerodynamic, elastic and inertial effects interact and how this framework relates to gust response. Next, gusts and their standard certification models are described, followed by a brief overview of gust load alleviation concepts, both passive and active. Together, these concepts provide the background needed to interpret the FFWT results and their impact on structural loads.

1.1.1. Aeroelasticity

Aeroelasticity studies the mutual interaction between a structure and the surrounding flow and becomes relevant whenever the interaction is significant. A distinction is made between static and dynamic aeroelasticity, where the first is a combination of purely aerodynamic and elastic forces, the latter includes inertial forces as well. Examples of static aeroelastic problems include divergence and control reversal, and dynamic phenomena include flutter and gust response. A classic tool for classification of aeroelastic problems at hand is Collar's triangle, shown in extended form in Figure 1.2 [20]. In this diagram, the three cornerstones are aerodynamic forces arising due to the flow surrounding the structure, inertial forces arising due to mass and acceleration of the structure and elastic forces due to structural properties of the aircraft.

The extended diagram shows various phenomena and where each lives in the coupling framework. Divergence and control reversal lie close to the aerodynamic-elastic edge, which aligns with the fact that they are static aeroelastic problems. Moving inward on the diagram toward flutter increases the degree of three-way coupling: unsteady aerodynamics, elasticity and inertia interact on comparable time scales. The diagram shows gusts (G) inside the triangle, because a transient flow input excites structural motion that reshapes the aerodynamics in return. The diagram helps separate static versus dynamic problems, motivates methods for predicting stability margins and transient loads and highlights how control systems can add damping. For the FFWT specifically, the gust-induced motion of the hinged wingtip is likewise a dynamic aeroelastic response: aerodynamic loads trigger tip rotation, this rotation modifies the local aerodynamic loading, and the resulting structural and inertial response determines the overall load-alleviation behaviour.

Other phenomena indicated in the diagram, such as divergence (D), flutter (F), buffeting (B), and stall and control problems (S), are classical aeroelastic instabilities and response types that are well covered in aeroelasticity literature. Since the focus of this thesis is on gust-induced load alleviation, these phenomena are not treated in detail here.

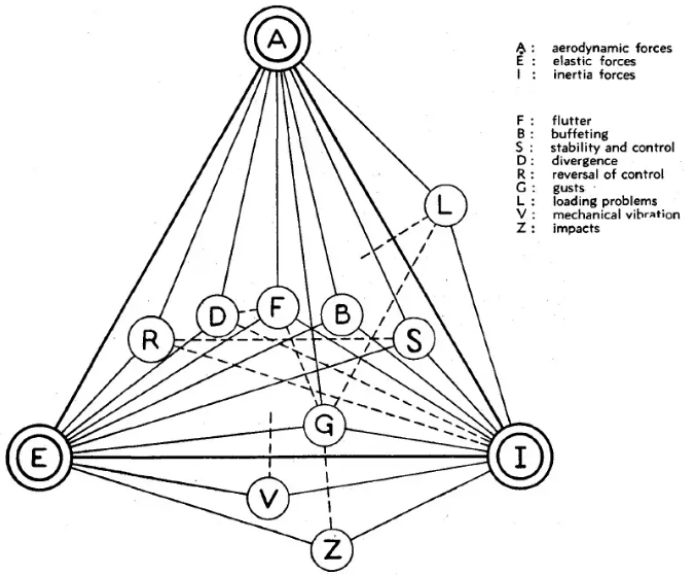

Figure 1.2: Collar's triangle of aeroelastic interactions, retrieved and edited from [20].

1.1.2. Gusts

A gust is a transient perturbation of the free-stream velocity vector that the aircraft effectively flies into. In the context of the FFWT, gusts are not only a certification load case, but also the transient excitation that can trigger wingtip motion, modify the aerodynamic loading and determine whether load alleviation is achieved. To numerically model this phenomenon, there are a number of different approaches, with different velocity profiles and directions. Three principal gust directions are considered, as depicted in Figure 1.3, namely vertical, lateral and head-on gusts.

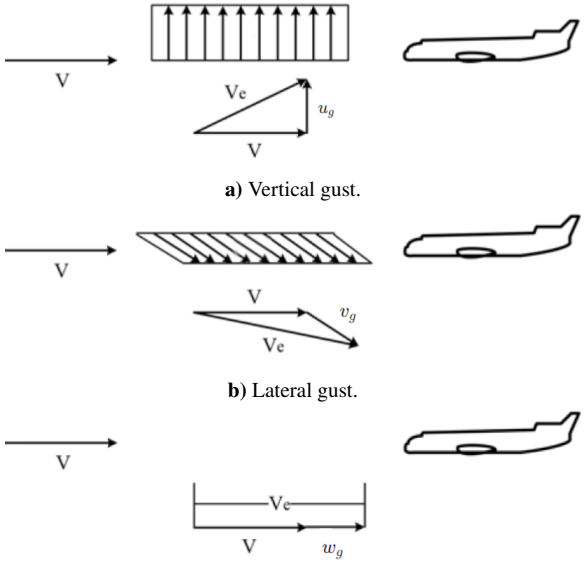

a) Vertical gust.

b) Lateral gust.

c) Head-on gust

Figure 1.3: Gust directions, retrieved and edited from [21].

In the figure [21], the relation between the effective total velocity V_e and the freestream velocity V and gust perturbation u for each direction is given. Each direction excites different primary responses:

- **Vertical (Figure 1.3a)**
  Vertical gusts, perpendicular to the wing plane, result in a sudden jump in angle of attack. In turn this leads to spikes in lift and associated aerodynamic and structural loads on the wing.
- **Lateral (Figure 1.3b)**
  Lateral gusts, parallel to the wing, produce an effective sideslip, resulting in yawing and rolling moments via the vertical tail and dihedral effects.
- **Head-on (Figure 1.3c)**
  Head-on, or longitudinal, gusts change the dynamic pressure and airspeed, affecting the thrust setting and drag of the aircraft.

For the present thesis, vertical gusts are of primary interest, because they directly induce a transient change in effective angle of attack and therefore strongly affect lift, WRBM and the gust-driven response of the FFWT. Lateral and head-on gusts are introduced for completeness and to place the selected gust cases in the broader aeroelastic context.

Since gusts affect the aerodynamic loads, they inherently affect the aircraft structure and handling qualities. For an FFWT, this is especially important because the transient load build-up during a gust determines when and how the wingtip tends to rotate, how much outboard lift is shed, and whether the resulting response reduces or amplifies the structural loads. Accordingly, the European Union Aviation Safety Agency (EASA) CS-25.341 "Gust and turbulence loads" [22] mandates certification for two types of gust, namely discrete and continuous gust.

Discrete

A discrete gust is a single, isolated turbulence event applied normal to the flight path, so vertical or lateral. As prescribed in the CS-25 regulations, the reference discrete-gust model is a 1-cosine velocity profile as shown in Figure 1.4 [23], with peak magnitude tuned by altitude and a gust-length parameter. The standard profile is described as Equation 1.4 [24].

$$\begin{cases} \dfrac{u(s)}{U} = \dfrac{1}{2}\left(1 - \cos\dfrac{\pi s}{H}\right) = \sin^2\left(\dfrac{\pi s}{2H}\right) & \text{for } 0 < s < 2H, \\ \dfrac{u(s)}{U} = 0 & \text{for } s < 0 \text{ and } s \geq 2H. \end{cases} \tag{1.4}$$

Here, $u(s)$ is the gust velocity at penetration distance s, U is the design gust velocity and H is the gust gradient distance which is the distance for the gust to reach its peak strength. Both s and H have units of length, while $u(s)$ and U have units of velocity. CS-25 regulations additionally specify a range of gradient values to sweep over to find peak responses.

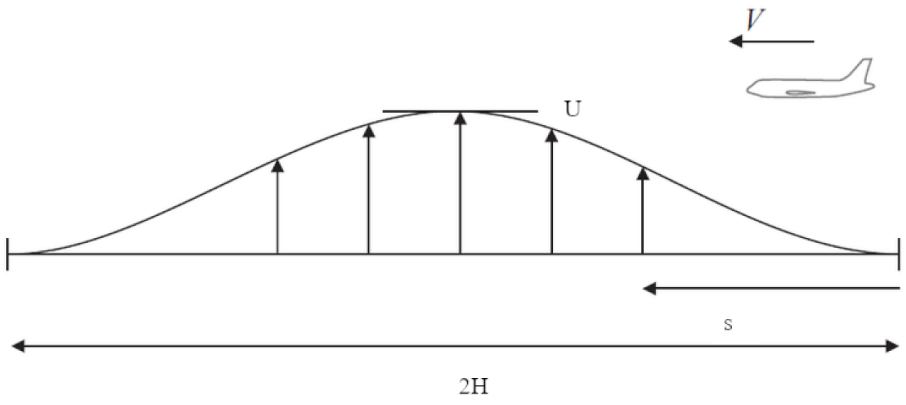

Figure 1.4: Schematic of a 1-cosine gust, retrieved and edited from [23].

The 1-cosine gust is particularly relevant in this thesis because it provides a controlled and certification-relevant transient input with which the sensitivity of the FFWT response to hinge properties, release timing and post-gust oscillations can be assessed systematically.

While the 1-cosine model is the certification default, there are many different discrete gust models. Alongside the 1-cosine model, two other gust profiles are considered as standard. The first of which is the sharp-edge gust, represented as Equation 1.5 [21], which describes an instantaneous, uniform jump or step response in gust velocity. This model is still widely used for code verification and control testing [25].

$$\frac{u(s)}{U} = 1 \tag{1.5}$$

The second gust model that is considered as standard is the linear-ramp gust, where the gust increases linearly to a peak over the gust gradient distance H. The velocity profile is represented as Equation 1.6 [21].

$$\begin{cases} \dfrac{u(s)}{U} = \dfrac{s}{H} & \text{for } 0 < s < H, \\ \dfrac{u(s)}{U} = 0 & \text{for } s < 0 \text{ and } s > H \end{cases} \tag{1.6}$$

Continuous

Although the main emphasis of this thesis is on prescribed discrete-gust transients, continuous-gust models are also relevant for assessing how the FFWT behaves under more realistic turbulence-like excitation and for positioning the results within the broader certification context. Discrete gusts remain the regulatory design case for limit loads, since these are "worst-case" transients that can define limit loads. Continuous gust models on the other hand describe the background turbulence environment and are essential to size for fatigue, tune controllers and validate handling qualities in turbulence [26]. In this approach, the atmosphere is modeled as a stationary field of turbulent velocity perturbations, where the three gust components in longitudinal direction u_g, lateral direction v_g and vertical direction w_g are characterized by Power Spectral Densities (PSDs). These PSDs contain two key parameters per axis, namely the intensity σ_i and length scale L_i. The two main stochastic turbulence spectra used in certification are the Dryden [27] and the von Kármán [26] spectra, which are presented in Equation 1.7 and Equation 1.8, respectively, where their PSDs ϕ are expressed in terms of spatial angular frequency Ω.

$$\phi(\Omega) = \sigma_i^2 \frac{L}{\pi} \frac{1 + 3(\Omega L_i)^2}{\left[1 + (\Omega L_i)^2\right]^2} \tag{1.7}$$

$$\phi(\Omega) = \sigma_i^2 \frac{L}{\pi} \frac{1 + \frac{8}{3}(1.339\,\Omega L_i)^2}{\left[1 + (1.339\,\Omega L_i)^2\right]^{11/6}} \tag{1.8}$$

These PSDs of the turbulence velocity components indicate how much gust “energy” there is at each spatial frequency and both of these spectra assume a Gaussian distribution of the gust components and a stationary turbulence field. EASA uses the von Kármán model as the reference spectrum for defining continuous turbulence [22]. In practice, these PSDs are converted to time-domain gusts by passing sequences containing white noise through shaping filters, with stochastic $u_g(t)$, $v_g(t)$ and $w_g(t)$ time histories as a result. These can then be used in simulations on loads and handling qualities.

1.1.3. Load Alleviation

Gust load alleviation (GLA) refers to passive and active measures that reduce structural and comfort-related penalties associated with gust encounters, typically by reducing critical loads such as WRBM without compromising safety or handling qualities. Closely related to this is Manoeuvre Load Alleviation (MLA), in which control or structural response is used to reduce loads generated during commanded flight manoeuvres.

Passive

Passive forms of load alleviation do not require control input, but instead reduce gust- and manoeuvre loads by allowing the wing to respond via its geometry or material such that lift and moments shed when the loads increase. The alleviation comes from built-in aeroelastic couplings and compliant mechanisms. Prime examples of this are the modern wing designs of the Airbus A350 XWB and the Boeing 787 that use composite layups with optimized stiffness directions, which enables a passive bend-twist washout resulting in load alleviation [28].

Next to these examples tailoring the entire wing, local passive devices have also been proposed that are based on a similar mechanism at the tip. Guo et al. [29] introduce the Passive Twist Wingtip (PTWT), as illustrated in Figure 1.5, in which a slender outboard section is connected to the main wing by a torsional spring about a spanwise axis. Under an upward gust, the increased bending and aerodynamic moment drive a nose-down twist of the PTWT, reducing the local angle of attack and shifting lift inboard. These concepts are related to the FFWT, since they illustrate the same general principle: load alleviation is obtained by allowing the structure or a local outboard device to respond passively to increased aerodynamic loading, thereby unloading the wingtip region.

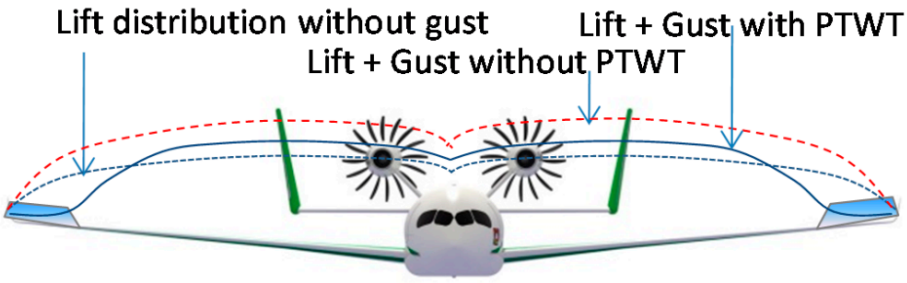

Figure 1.5: Effect of a PTWT on spanwise lift distribution, retrieved from [29].

Active

Where passive concepts rely on built-in aeroelastic couplings, active GLA uses control surfaces driven by feedback or preview logic to attenuate the loads and accelerations produced by gusts. Typically, structural or motion-related signals are fed into a controller that commands ailerons, flaperons, spoilers or elevators to redistribute lift and keep loads within limits [21, 24]. Current transport aircraft already apply this principle in MLA and GLA functions, reducing outboard lift and compensating with a more inboard load distribution to reduce WRBM [30]. Recent research extends these ideas to very flexible wings and morphing devices. Examples of this include active trailing-edge flaps on high-AR demonstrators like FLEXOP, where measured bending and twist are fed back for GLA and flutter delay [31].

The FFWT can be positioned at the interface between passive and active load alleviation. When its response is governed primarily by structural and aerodynamic characteristics, it behaves as a passive alleviation mechanism. When hinge release or hinge torque is triggered or shaped by a prescribed logic, it moves toward semi-active or active load alleviation.

1.2. Chapter Summary and Thesis Outline

This first chapter has framed the central problem in this thesis, and provided the fundamentals needed for the introduction of the FFWT concept and what it can achieve regarding load alleviation. Increasing wing span is attractive from an aerodynamic-efficiency perspective, but it also increases structural loads, aeroelastic sensitivity and operational constraints at airports. The FFWT is a promising concept, due to the fact that it may meet these competing requirements by combining span extension with load-responsive wingtip motion.

Chapter 2 will therefore review the FFWT concept and its working principles, and place it within the broader literature on gust-load alleviation, hinged wingtips and aeroelastic modelling. This review identifies the main design variables governing FFWT behaviour, such as hinge flare angle, hinge stiffness, damping, tip inertia and release timing, and highlights the limitations of existing approaches. These limitations form the basis for the knowledge gap, research objective and research questions formulated in Chapter 3. Chapter 5 then presents the non-linear Simscape–UVLM framework developed to address these questions, including the structural model, aerodynamic model, aeroelastic coupling, verification and validation procedure, and staged simulation campaign. The results of this campaign are presented in Chapter 6, where the effects of release timing, hinge properties and hinge-moment-based release thresholds on WRBM reduction and hinge-angle demand are analysed. Finally, Chapter 7 answers the research questions and summarises the main findings, while Chapter 8 discusses the principal limitations of the present work and proposes directions for future research.

2

State-of-the-art review

Chapter 1 established the central design framework motivating this thesis: increasing wingspan improves aerodynamic efficiency, but also aggravates structural loads, aeroelastic sensitivity and airport-gate constraints. The FFWT has emerged as a promising concept because it may meet these competing requirements by combining span extension with load-responsive wingtip motion. This chapter therefore reviews the state of the art on using the concept for GLA and MLA, how it has developed in literature, what has been learned from numerical and experimental studies, and which limitations remain unresolved.

To do so, the historical origin of the concept and its kinematics are first introduced in Section 2.1 and Section 2.2. The subsequent review in Section 2.3 covers parametric FFWT studies, roll-performance investigations, wind-tunnel and scaled-flight experiments, dedicated work on hinge moments and release logic, and the modelling frameworks and gust inputs adopted across the literature. The chapter concludes by identifying the knowledge gap that motivates the present thesis in Section 2.4.

2.1. Origin

The origin of the FFWT concept traces back to early hinge-based load alleviation and space-saving mechanisms. In 1922, Juan de la Cierva [32] introduced flapping hinges at the root of the rotor of his autogyro, with the intention to balance the lift of the advancing and retreating blades. This is a first use of hinging to mitigate asymmetric loads that would later inspire wing-borne concepts like the FFWT. In the Second World War, folding wings became a standard on aircraft carriers, to save space on and below deck. These concepts further developed post-war into variable-geometry programs. One example is the Bell X-5, which was the first aircraft to vary wing sweep to adapt to flight conditions [33]. The performance optimization and space-saving come together in the design of the Grumman F-14, which can change the sweep angle of its wings in flight, and includes an "oversweep" setting for compact storing.

Patents throughout this period form the basis of the FFWT concept, starting with the patent on a "wingtip device" by Irving and Davies [34], followed by the patent by Wilson et al. [35] on a "movable wingtip" specifically for the purpose of load alleviation. Related concepts like Semi-Aeroelastic Hinges (SAHs) and free-wing outboard panels reinforced the core insight of using a controlled rotational Degree-Of-Freedom (DOF) to decouple load paths under transient loads like in manoeuvres and during gusts [36, 37]. In civil transport, the concept of folding the wingtip resurfaced to meet airport gate requirements and the need for higher Aspect Ratio (AR) wings. The Boeing 777X includes a ground-only folding wingtip to adhere to gate requirements on airports, but still allow an increase in AR in flight [38]. In parallel, Airbus explored folding wingtips for civil airliners in the mid 2000s to increase the AR without breaching airport gate limits. This eventually led to the development of the bio-inspired folding demonstrator AlbatrossONE [39] and the full-aircraft demonstrator Extra Performance Wing [40], allowing maturing of the technology into full flight instead of ground-only solutions.

Much of the modern research on hinged wingtips then converged on the “flared” hinge geometry, in which the hinge axis is tilted by a flare angle (Λ) relative to the freestream so that upward rotation reduces the local angle of attack. This transformed the folding wingtip from a purely operational device into a load-responsive aeroelastic concept, with a passive restoring tendency in motion and a reduced transmission of bending moment to the wing root. From this point onward, the FFWT literature developed along several main directions that are particularly relevant to the TU Delft research line followed in this thesis: early parametric studies mapped the effects of flare angle, tip inertia

and hinge stiffness [41]; wind-tunnel campaigns validated the alleviation mechanism and exposed the limits of linear prediction tools at large rotations [42, 43]; and controlled-release studies highlighted the importance of release timing, damping and post-gust dynamics [17]. These directions form the basis of the detailed review presented in the following sections.

2.2. The Flared Folding Wingtip Concept

High-AR wings promise a lower induced drag and improved cruise efficiency, but they also face two practical limitations: airport gate constraints and a higher sensitivity to gust and manoeuvre loads. The discussed conventional folding tips only address the first of these by folding on the ground to pass gate requirements, while flying at full span. The FFWT could serve as a solution to both issues in one device by hinging the outboard panel of the wing on a flared axis relative to the freestream, as presented in Figure 2.1.

Geometrically, the flare angle (Λ) is measured between the hinge line and the freestream direction. If the angle is zero, the system reduces to a conventional, streamwise hinge with no angle-of-attack relief, whereas increasing the angle strengthens the coupling between upward rotation and nose-down change in local angle of attack. The fold angle θ describes the rotation about the hinge axis.

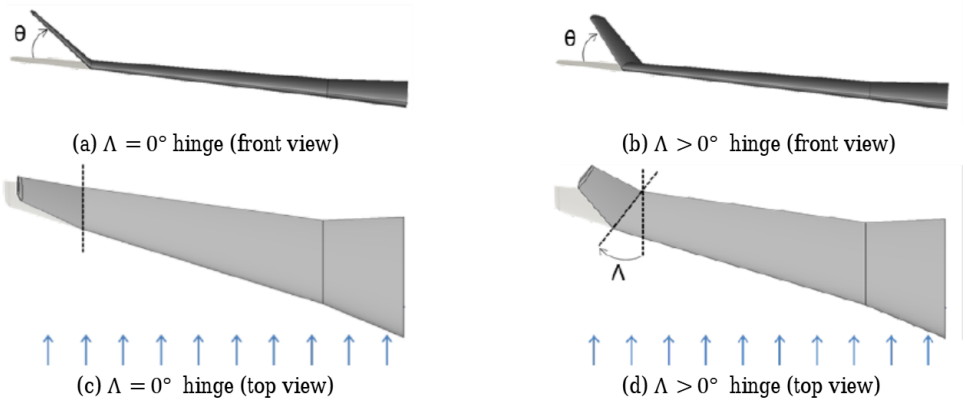

Figure 2.1: FFWT concept hinge orientation, retrieved from [41].

When released, the tip's upward rotation thus reduces its local angle of attack due to the fact that the upward fold projects partly into a nose-down rotation relative to the incoming flow [41]. The tip can thus remain locked for span efficiency in steady flight, and rotate upward through a load-triggered release so that its effective local angle of attack is reduced. Subsequently, this allows for load alleviation by unloading the outer wing and reducing WRBM when needed while preserving the high AR benefits the remainder of the flight. Assuming small fold angle deflections and a positive flare angle, this leads to the relation in Equation 2.1.

$$\Delta\alpha \;=\; -\arctan\big(\tan\theta\,\sin\Lambda\big) \approx -\theta\sin\Lambda \qquad (\text{for } |\theta| << 1) \tag{2.1}$$

This relation is useful for explaining the sign and first-order scaling of the FFWT mechanism, the large-angle kinematic angle of attack definition used in this research is defined later in Subsection 5.6.1. Structurally, the FFWT introduces a single rotational degree of freedom at the outboard panel, so that, once unlocked, part of the outer-span load can be shed by rotation rather than being transmitted as bending moment to the wing root.

2.3. Review of FFWT Research

Once the FFWT was formulated as a load-responsive flared-hinge concept, the literature developed along several main lines. Early work focused on parametric studies to map the influence of flare angle, tip inertia and hinge properties. This was followed by investigations into roll performance, wind-tunnel campaigns to validate the load-alleviation mechanism, scaled-flight demonstrations, and dedicated studies on hinge moments, release logic and modelling strategy. The following sections review these themes in turn.

2.3.1. Parametric Studies

Early parametric work on the FFWT, that forms the baseline for the subsequent research, used low-fidelity aeroelastic models that were fast enough to sweep design variables and expose first-order trends. The preliminary study of Castrichini et al. [41] used a representative model of a transport aircraft, presented in Figure 2.2, to show that hinging the outboard tip about a flared axis enables shedding of outer-span lift and reduction of the WRBM during 1-cosine gust cases. For the structural model, a stick model with lumped masses is used and for the aerodynamic model a Double Lattice Method [44] is used. Both models consist of a wingtip connected to the wing, increasing the wingspan of the aircraft by 25%. The results [41] show that increasing the flare angle and reducing the tip inertia generally boosts gust-load alleviation, but these gains trade against a reduced flutter speed. This makes stability an important parameter for sizing as well. Building on this, Castrichini et al.[45, 46] first introduced non-linear hinges with a piecewise stiffness and later a negative-stiffness spring to hold the tip in cruise and add the ability to unlock rapidly under load.

The results highlight that, although performance is mainly dominated by the stiffness of the hinge, the flare angle and the tip mass, the release threshold and hinge damping play an additional but critical role in the performance. In parallel, high-fidelity aeroelastic co-simulation studies by Pattinson et al. [47] demonstrate that coupling a flexible wing with a folding tip refines the sensitivities regarding the hinge orientation, stiffness and wingtip mass. These early FFWT parametric studies are invaluable for mapping design variables, but because they rely on linear unsteady aerodynamics and small-deflection structure, they risk over-predicting load relief once tip rotations, separation, or release dynamics dominate.

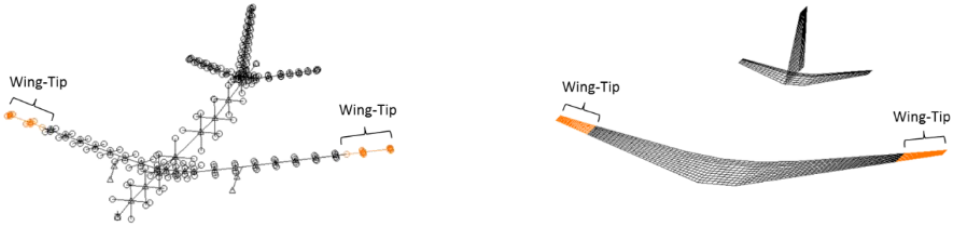

(a) Structural model (b) Aerodynamic model

Figure 2.2: Representative model of a transport aircraft, retrieved from [41].

These early parametric studies established the main FFWT design levers and demonstrated that significant alleviation is possible in principle. However, because they rely largely on linear aerodynamics and small-deflection structural assumptions, their predictive value becomes limited once large tip rotations, non-linear release behaviour or wake distortion become important.

2.3.2. Roll Performance

Beyond gust alleviation, hinged and flared tips affect lateral-directional dynamics and stability. The effective local angle of attack drops as the panel starts to fold, and this reduces the aerodynamic roll damping C_{ℓ_p} and delays tip stall. At the same time, the tip motion can modify the dihedral effect C_{ℓ_β}, as well as the adverse-yaw derivatives C_{n_p} and C_{n_β} [48]. The governing factors are the hinge stiffness or damping, the release logic, tip inertia and the wing's out-of-plane stiffness.

Dussart et al. [49] and Healy et al. [50] examine the roll manoeuvres using relatively stiff wings and tuned hinges, and report that unlocking the tip generally reduces aerodynamic roll damping and can increase roll rate for a given aileron input, because the outboard unloading lowers induced-drag growth and delays tip stall during the transient. This allows for more of the command to feed the angular acceleration. Dussart et al. rely on reduced-order models and Healy et al. use a sub-scale rig, and both extrapolate trends with largely linear and inviscid aerodynamics, which make the results strongest for moderate angles and attached flow.

In contrast, two independent studies by Ajaj et al. [48] and Sanghi et al. [51] demonstrate that this benefit is conditional rather than universal through analyzing very flexible, high AR configurations. In this case, the hinge motion can increase the dihedral effect and adverse-yaw contributions, but it can damp the roll authority or even excite Dutch roll tendencies, especially at higher dynamic pressure. Gains concentrate therefore at lower values for dynamic pressure and as this parameter and the flexibility rise, improvements shrink or can even reverse. These results are obtained through computational Vortex-Lattice Methods, both steady (VLM) and Unsteady (UVLM), with parameterized hinge dynamics.

Regarding the governing parameters that determine whether the device helps or hurts, the four studies agree, namely the hinge stiffness, release logic, tip inertia and the structural stiffness. The design implication that follows from this is to treat the FFWT as part of the lateral-control system, with a scheduled release with dynamic pressure and roll-rate targets. Tuning the FFWT together with the aileron could realize the benefits without eroding the directional stability. The roll-performance literature therefore suggests that the FFWT cannot be treated purely as a passive load-alleviation device, because it also affects lateral-directional stability and control. Whether the concept improves or degrades roll performance depends strongly on structural flexibility, hinge dynamics and operating condition, which implies that GLA and handling qualities must be assessed together.

2.3.3. Experimental Studies

Moving to the experimental studies, Cheung et al. establish the first experimental backbone for FFWTs across three progressively more realistic test articles. The first [42] consists of a proof-of-concept wing with a hinged outboard tip, as shown in Figure 2.3, in order to isolate the basic mechanism under discrete gusts and manoeuvres.

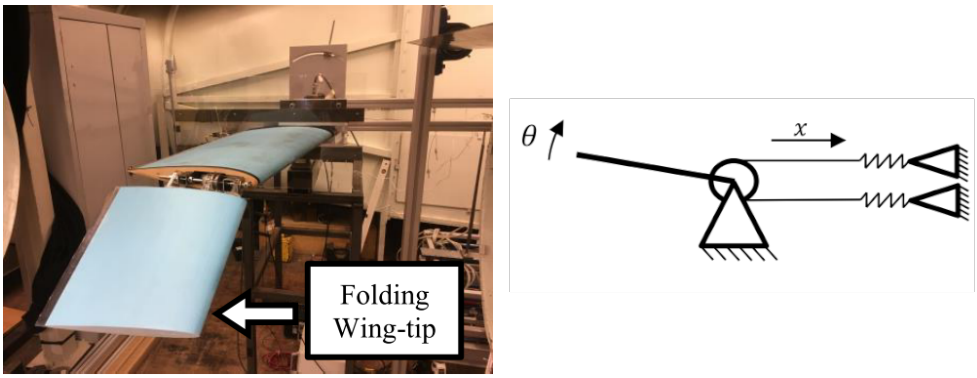

Figure 2.3: First wind-tunnel setup by Cheung et al., retrieved from [42].

In this setup, discrete gusts and simple manoeuvre inputs are used to verify the mechanism anticipated in preliminary parametric studies [46]. When the tip is unlocked about a flared axis, the outboard section sheds load and the measured WRBM decreases relative to the locked baseline, confirming the FFWT GLA capability. Comparison with a linear DLM model shows good trend-level agreement in small-perturbation cases, but as tip rotations become larger and hinge dynamics dominate, the discrepancies grow. The second [43] test article, as presented in Figure 2.4, is a flexible, high AR wing model with a folding tip to capture structural coupling and spanwise load redistribution.

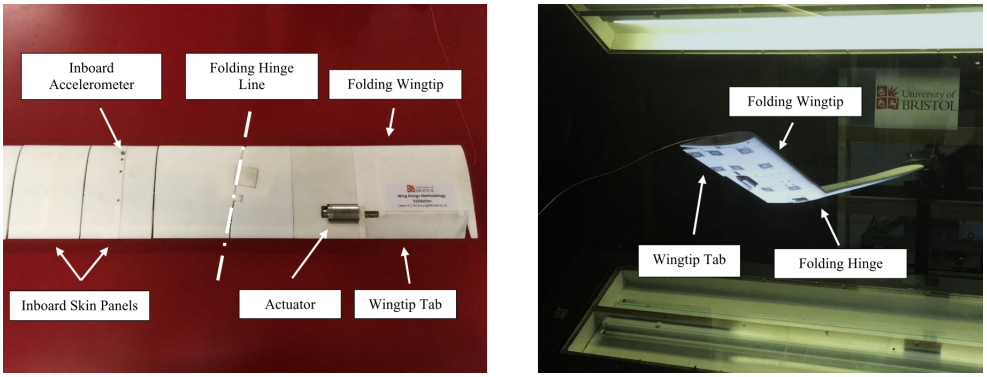

Figure 2.4: Second wind-tunnel setup by Cheung et al., retrieved from [43].

Results show that under discrete gusts and simple manoeuvre inputs, the free, hinged configuration reduces the measured WRBM relative to the locked baseline. The alleviation level in this case is now strongly dependent on hinge stiffness, tip inertia and the flare angle. Additionally, the flexible article now reveals clearer timing sensitivity between the release phase and the gust peak, as well as post-gust oscillations in the hinge angle. Compared to the first setup, the trend-wise agreement with the linear DLM model remains, but only for small motions. Phase lag and amplitude are increasingly poorly estimated once rotations grow and flexibility dominate.

The third [50] test article in Figure 2.5 is a variant that combines the tip with a trailing edge control surface, to probe timing and control integration. The setup uses small tip preset and flap deflection phased to the gust, and with that produces larger WRBM reductions than tip-only cases, while also helping damp post-gust hinge oscillations. The alleviation level proves highly sensitive to timing: wrongly timed release or flap commands erode or even reverse the benefit. Like the other two setups, the trend compared to the linear DLM model remains acceptable for small motions, but large rotations increase the error in estimation of the phase and amplitude, showing the need for time-domain, non-linear aeroelastic tools and phase-aware control logic.

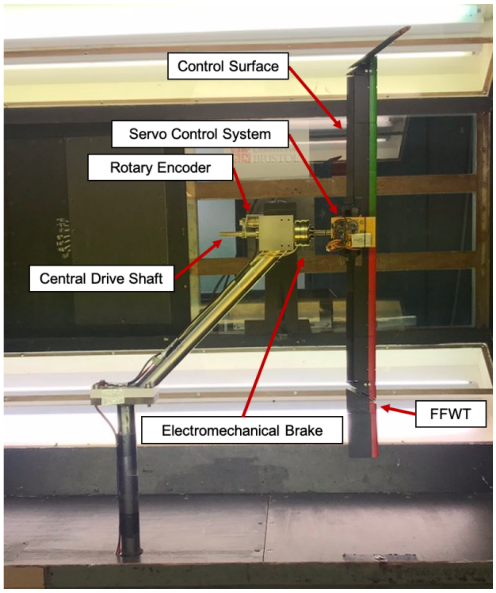

Figure 2.5: Third wind-tunnel setup by Cheung et al., retrieved from [50].

Córcoles et al.[52] report a wind-tunnel program on a flexible semi-span wing with a flared, hinged tip that isolates when the tip should be allowed to move. Using 1-cosine gusts and a hinge configured as locked, free or threshold-released with adjustable damping, the study shows that early release consistently yields the largest WRBM reduction. Near-peak or late release can erode or even reverse the benefit relative to the locked case, thereby confirming the results obtained by Cheung et al. If the threshold is too high, the motion is delayed and the alleviation level is reduced. A threshold that is too low leads to nuisance motion. Adding hinge damping suppresses post-gust oscillations, but excessive damping blunts the response. The results turn "release timing" and "hinge damping/threshold" into primary design variables.

The underlying TU Delft MSc thesis by Córcoles [17] documents this semi-span campaign in greater detail, including the full test matrix, measured WRBM time histories and hinge-angle traces for the locked, free and threshold-released configurations. These data form a key reference for validation of numerical FFWT models.

Taken together, the experimental studies validate the physical load-alleviation mechanism of the FFWT, but also show that its performance is highly sensitive to timing, damping and structural flexibility. They further highlight the limitations of linear aeroelastic prediction tools once large-amplitude tip motion becomes important.

Scaled Flight Demonstrations

Scaled flight sits between wind-tunnel experiments and full-scale use, and therefore allows the mechanism to be tested in real atmospheric disturbances and for integration of the hinge hardware, sensing, actuation and control scheduling to be validated. Also, possible additional aeroelastic couplings and failure modes could be exposed. Within this space, Airbus's AlbatrossONE [39] program is the most clearly documented public flight effort directly focused on the FFWT implementation, as illustrated in Figure 2.6. AlbatrossONE is best described as a semi-aeroelastic hinged-wingtip demonstrator closely related to the FFWT concept, while the exact public documentation of its flare geometry is limited.

Figure 2.6: AlbatrossONE folding wingtips in GLA position, screenshot retrieved from [53].

The sub-scale demonstrator flies with semi-aeroelastic hinged outboard panels to explore load-responsive tip motion, release logic and damping in free flight. These form the key ingredients an FFWT needs when the hinge axis is flared so that the upward rotation reduces local angle of attack. The AlbatrossONE provides a platform to test when the tip should be allowed to move, how fast it should move and how it is coordinated with conventional control surfaces, thereby bridging the gap from tunnel-verified physics to flight-worthy forms of the FFWT concept. Although Airbus has not released WRBM time-histories, phase maps or hinge-angle traces in the public domain, preventing quantitative comparisons, the public reports state that freely flapping tips can alleviate wing loads and avoid tip stall. This aligns with the wind-tunnel takeaways, where a free tip means outboard unloading, dependent on the configuration.

In parallel, Airbus UpNext's Extra Performance Wing [40] demonstrator pushes the hinged outboard devices more toward a flight-relevant architecture on a scaled business jet platform. The program integrates folding wingtips with phase-aware control through forward gust sensing, and aims to combine span-efficiency with gust and manoeuvre load management. The demonstrator entails all required integration factors for an FFWT implementation, namely robust hinge hardware, real-time release scheduling tied to dynamic pressure and measured disturbances, as well as coordination with conventional control surfaces.

Complementary to these experimental efforts, Hoyng [18] developed and ground-tested a highly flexible FFWT demonstrator at TU Delft. The work combines a slender, high AR wing with a flared folding tip and trailing-edge control surfaces, and focuses on structural design, Ground Vibration Testing (GVT) and mode-shape identification. The resulting aeroelastic model and measured mode set form a practical bridge between idealized FFWT concepts and hardware-realistic implementations.

2.3.4. Hinge Moments

In addition to WRBM alleviation, the hinge moment and its use in defining a release strategy have been explored. Castrichini et al.[45] first introduced a passive non-linear hinge spring to hold the tip during cruise and allow it to deflect only above a load threshold. This can be represented by a piecewise linear stiffness law, as presented in Equation 2.2, where the hinge behaves as effectively rigid up to a prescribed threshold moment M_{max}, and then switches to a very low stiffness once this threshold is exceeded.

$$M_{\text{spring}} = -K_\theta \theta, \qquad K_\theta = \begin{cases} 1.0 \times 10^{12}\ \text{Nm/rad} & \text{if } K_\theta\theta \le M_{\max} \text{ and } 0 < t < t_{\text{release}}, \\ 1.0 \times 10^{0}\ \text{Nm/rad} & \text{if } K_\theta\theta > M_{\max} \text{ and } t \ge t_{\text{release}}. \end{cases} \tag{2.2}$$

This way, the hinge behaves as effectively rigid before the threshold is exceeded and behaves as almost free once it is released. Parametric studies using this model show that increasing M_{max} delays the release, allowing loads to build up before the wingtip starts to rotate, which consequently degrades the gust-load alleviation performance. The release instant, defined by the combination of M_{max} and the actuation delay is critical for achieving the beneficial FFWT behaviour. The spring moment and stiffness are defined following the schematic in Figure 2.7.

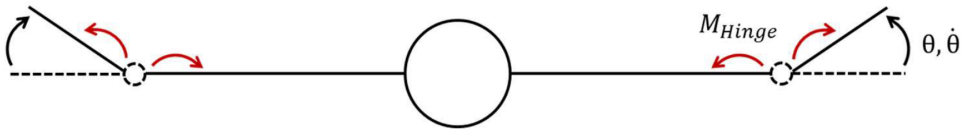

Figure 2.7: Schematic of the applied hinge moments and tip rotation, retrieved from [54].

Building on this, a non-linear negative-stiffness wingtip spring device was proposed by Castrichini et al.[46], involving a hinge that is very stiff around the planar position in static conditions, but gets much softer during dynamic motion. The objective is to keep the wingtip stiff during cruise, while allowing a fast, large-amplitude rotation during a gust encounter and a passive return to the planar configuration afterwards. Simulations of 1-cosine gusts show that this mechanism can provide load alleviation, as well as the ability to restore planar configuration. Its effectiveness however remains sensitive to how quickly the wingtip reacts to the gust and to the chosen stiffness parameters of the hinge. These two hinge concepts illustrate that hinge moment, release threshold and actuation latency are primary design drivers for FFWT implementation.

Subsequent work translates and extends these concepts. Experimental SAH and FFWT campaigns study dynamic release timing and report hinge-moment time histories and vary the hinge stiffness to study how the load path and release behaviour change, confirming that softer hinges reduce peak WRBM but increase hinge motion and sensitivity to timing [42, 54, 55]. More recent SAH scaling and flight-dynamics studies treat the hinge as a floating or semi-aeroelastic joint that releases under manoeuvre or gust loads, and use hinge moment as one of the key indicators in sizing and control-law design [36, 56]. In parallel, actively actuated flared folding wingtips are being explored, in which the hinge torque is provided by a controller rather than a purely passive spring. This effectively generalizes the threshold for the moment and the negative stiffness concept into a fully programmable release logic [57].

2.3.5. Modelling Framework

After assessing the parametric and experimental studies, it is important to review the modelling frameworks that underpin the FFWT research, with an emphasis on how aeroelastic solvers, aerodynamic solvers and gust inputs have evolved alongside the concept. The following subsections summarize this progression and extract relevant pointers to justify the modelling choices made in this thesis.

Aeroelastic Solvers

The aeroelastic modelling of flared and hinged wingtips follows a clear chronological progression. As mentioned, the first systematic FFWT and SAH studies by Castrichini et al.[41] use representative transport-aircraft models, built in the commercial multibody dynamics code LMS-Virtual Lab Motion (VLMotion) [58], in which the wing is modelled as a stick with lumped masses and a small set of generalized bending and torsion coordinates. The wingtip then enters as an additional rotational DOF about a flared hinge line. The LMS-VLMotion is coupled to unsteady aerodynamic loads obtained from a DLM formulation, so that flutter margins, worst-case gust loads and the effect of non-linear hinge springs on WRBM can be assessed. These multibody DLM models are fast enough for large parametric sweeps in hinge stiffness, tip inertia and flare angle, but still rely on small-deflection assumptions in structural description and linearized aerodynamics.

As the state of the concept moves to wind-tunnel testing, the aeroelastic solvers shift toward finite-element based models in MSC Nastran [59]. For the rigid semi-span hinged tip demonstrator, Cheung et al.[42] build a detailed MSC Nastran finite-element model of the wing and hinge hardware, to compute aeroelastic predictions using the built-in DLM module in Nastran's aeroelastic solver. The folding tip is represented by nodes, connected via elements aligned with the hinge axis. This way, different spring and damping laws can be imposed. Following on this [43], a flexible high AR wing is tested, where the same Nastran-DLM combination is used as framework. The aluminium spar and 3D-printed sections are modelled with beam elements and concentrated masses. The hinged tip is then introduced by adding a rotational spring at the outboard station, while the aerodynamic loads again come from a Nastran DLM panel model. Subsequent hinged-wingtip studies [55], essentially adopt the same Nastran-DLM aeroelastic architecture and then validate it experimentally. Which in some way set MSC Nastran with DLM panels and hinge elements as the industry standard for low- to mid-fidelity FFWT load-alleviation studies.

Complementary to these studies focused on wind-tunnel experiments, several FFWT studies that are more flight-dynamics oriented like the research by Ajaj et al.[48] use modal-superposition models derived from MSC Nastran as the structural backbone. A small set of bending, torsion and hinge modes are then retained and embedded in the flight-dynamics equations. These approaches still start from the finite-element models in Nastran, but export reduced-order aeroelastic state-space models for use in the time-domain simulations of trim, roll manoeuvres and handling qualities.

Recent work by Conti et al.[60] moves beyond Nastran, towards tailored multibody aeroelastic formulations for SAH and FFWT systems. In the quasi-steady analysis of the SAH, a multibody formulation is introduced with a linearly flexible main wing and rigid folding wingtips attached through revolute hinges. This explicitly accounts for finite rotations and geometric non-linearities in the tip motion, while modelling the main wing via a modal description. Mastracci et al.[61] extend this framework to an unsteady aeroelastic multibody formulation, including local geometric non-linearities and aerodynamic forces applied in a follower sense, i.e. with their direction tied to the instantaneous orientation of the folding wingtip, time marching flutter and gust analyses. In both of these studies, the solver is an in-house multibody aeroelastic code tailored to systems consisting of a flexible wing and rigid tip, with large tip rotations. Earlier multibody concepts also appear in broader SAH and articulated-wing flight dynamics studies, which model the aircraft as rigid bodies connected by elastic hinges and solve the resulting non-linear multibody equations coupled to quasi-steady aerodynamics [62].

In parallel to these developments, Ghosh [19] implements the FFWT concept in the in-house PROTEUS framework as a multibody aeroelastic system with a flared hinge degree of freedom and linearized unsteady aerodynamics. The PROTEUS model captures hinge motion, release dynamics and load alleviation within a flexible 3D aircraft representation, and highlights solver-robustness issues that arise when large rotations and strongly coupled hinge dynamics are present.

Aerodynamic Solvers

Regarding the aerodynamic solvers, the development has largely been described in the previous section, but the FFWT and SAH studies follow the hierarchy used in gust and aeroelastic analysis. The DLM has a central role in the research as early studies use DLM panel models to provide unsteady loads. These are coupled to low-order structural or multibody representations to predict flutter margins and discrete 1-cosine gust responses [46]. As the concept moves to wind-tunnel testing, the same approach is embedded as Cheung et al. [42, 43] use MSC Nastran's built-in DLM module to compute WRBM alleviation for rigid and flexible folding wingtip test articles. In parallel, lower-order strip theory and steady or quasi-steady vortex-lattice models are used in research more oriented towards flight dynamics [48], in order to generate stability and control derivatives for wings fitted with flared or folding wingtips. Thereby trading some unsteady fidelity for computational speed.

A shift toward time-domain vortex methods appears once large tip rotations and non-linear hinge laws become central. In the dedicated modelling study by Cheung et al.[63], the DLM implementation is augmented with an UVLM. As tip deflections grow and wake deformation becomes more important, the results between the two diverge, which has led to recent non-linear FFWT and SAH frameworks to adopt UVLM as the preferred unsteady aerodynamic solver when large rotations and non-planar wakes must be resolved.

Gust Input

Looking at the gust models that are used in the FFWT and SAH studies, they form a narrow subset of the variants introduced in Subsection 1.1.2. Almost all numerical and experimental FFWT work focuses on vertical, discrete gusts, with the standard 1-cosine profile as standard. Sharp-edge gusts and linear-ramp profiles appear mainly as auxiliary cases for code verification and response characterization, rather than as primary design drivers.

Continuous gust models, such as Dryden or von Kármán spectra, are far less common in the FFWT literature, being occasionally used in broader handling-qualities and loads studies to characterize background turbulence. Current FFWT and SAH work remains however predominantly framed around deterministic discrete gusts that expose worst-case transient loads and timing effects. These standard profiles will be reused and extended in this thesis as building blocks for a wider gust library, including variations in directionality, shape and phasing that are more representative of realistic atmospheric disturbances.

The modelling literature therefore reveals a persistent trade-off between computational efficiency and physical fidelity. Lower-order frameworks are well suited to broad parametric studies, while higher-fidelity time-domain methods are needed to resolve large rotations, non-linear release behaviour and extended gust scenarios in a more realistic way.

2.4. Knowledge Gap

Overall, the existing FFWT literature establishes that a flared hinge with an outboard folding panel can alleviate gust and manoeuvre loads, and that the key design parameters are flare angle, hinge stiffness and damping, tip inertia and release logic. Parametric multibody studies and wind-tunnel campaigns consistently show that appropriately tuned hinges can reduce peak WRBM under discrete gusts and simple manoeuvres. However, the literature also reveals several limitations that remain relevant for the present thesis.

First, many parametric studies, and much of the associated experimental interpretation, rely on linearised aeroelastic frameworks such as LMS-VLMotion multibody models or MSC Nastran coupled with DLM aerodynamics. These tools are efficient and useful for trend identification, but they are based on small-deflection assumptions and simplified hinge representations. Consequently, large wingtip rotations, non-linear hinge behaviour, release transients and hysteresis-like effects are only partially represented.

Second, more recent multibody formulations can represent geometric non-linearity more realistically, but they are often implemented as dedicated solvers and are not always coupled in a transparent and reusable way to time-domain unsteady aerodynamics. This limits their applicability for systematic studies of FFWT release timing, hinge-property variation and post-release dynamic behaviour.

Third, the way hinge moments and release logic are treated. While several studies introduce piecewise stiffness laws, negative-stiffness concepts and threshold-based release, the hinge moment is often treated mainly as a response quantity rather than as a primary release-design variable. There is still a need for a time-domain framework that directly links predicted hinge moment, release thresholds, release timing, hinge-property selection, WRBM reduction and post-gust motion.

Consequently, there is still no openly documented, time-domain non-linear framework that combines a flexible-wing/rigid-tip multibody model, configurable hinge non-linearity, unsteady aerodynamics and hinge-moment-based release logic for systematic assessment of FFWT gust-load alleviation.

3

Research Proposal

Chapter 2 showed that the FFWT concept can alleviate gust-induced loads, but also highlighted several unresolved issues. Existing studies remain dominated by linearized aeroelastic frameworks, simplified hinge representations and a narrow set of prescribed gust inputs. As a result, the combined effects of non-linear hinge behaviour and time-domain release logic on load alleviation performance and post-gust dynamics are not understood within a single openly documented framework.

The present thesis addresses this gap by building on the recent TU Delft FFWT research line. Córcoles's wind-tunnel campaign [17] provides semi-span WRBM time histories and hinge-angle traces for locked, free and threshold-released configurations. Hoyng's flexible demonstrator [18] supplies a validated structural layout and mode set for a high-AR wing with a flared folding tip. Ghosh's PROTEUS implementation [19] demonstrates a multibody FFWT model within a linear aeroelastic framework and highlights the solver-robustness issues that arise once hinge motion and release dynamics become strongly coupled.

Against this background, this chapter formulates the research objective and research questions that guide the remainder of the thesis.

3.1. Research Objective

In response to the knowledge gap identified in Section 2.4, this thesis develops a non-linear Simscape-UVLM framework that combines a multibody representation of a flexible main wing and rigid folding tip with configurable hinge non-linearity, time-domain release logic and high-fidelity unsteady aerodynamics. This leads to the following research objective:

The objective of this thesis is to develop and validate a non-linear, time-domain aeroelastic framework that couples an FFWT Simscape multibody model to a UVLM-based unsteady aerodynamic solver, and to use this framework to design and assess hinge-moment-based release strategies for a flared folding wingtip, quantifying their load-alleviation performance.

3.2. Research Questions

To quantify this objective, it is translated into one main research question and a set of supporting sub-research questions addressing validation, release-strategy formulation and performance assessment. Starting with the main research question:

Which combinations of FFWT design parameters and hinge-release settings reduce wing-root bending moment relative to a locked-hinge baseline, while keeping post-gust motion and aeroelastic behaviour within acceptable limits?

This research question is supported by the following sub-research questions.

1. How accurately can the framework reproduce existing FFWT reference cases in terms of wing-root bending moment, hinge-angle dynamics and aeroelastic modes?
2. How can hinge-moment-based release strategies be formulated and implemented within the framework?
3. How do key FFWT design parameters affect dynamic behaviour, load alleviation performance and post-gust tip motion for standard vertical discrete gusts?

The non-linear Simscape framework that is developed as one of the objectives of this thesis follows this latest generation of aeroelastic solvers. The main wing is modelled as a flexible body, the flared folding wingtip as a rigid body, and the connection as an explicit flared hinge degree of freedom with configurable parameters. This accommodates large tip rotations, release dynamics and contact effects, while remaining compatible with high-fidelity unsteady aerodynamic and gust-input models.

The aerodynamic modelling in this thesis is positioned within this hierarchy: a UVLM-based formulation is selected to remain comparable to established DLM results, yet compatible with the large-rotation, non-linear hinge dynamics captured by the Simscape aeroelastic framework.

Part II

Scientific Article

Non-linear Simscape–UVLM Assessment of Release Strategies based on Hinge Moment for Gust-Load Alleviation with a Flared Folding Wingtip

Author: **R. Lambrichs** *Delft University of Technology, Delft, NL*

Research Supervisors: **C. Varriale, J. Sodja** *Delft University of Technology, Delft, NL*

Abstract: High-aspect-ratio wings offer a direct method to improved aerodynamic efficiency through reduced induced drag, but their increased span also increases structural loads, aeroelastic sensitivity and creates airport-gate compatibility challenges. The Flared Folding Wingtip (FFWT) concept addresses these competing requirements by combining an outboard folding panel with a flared hinge axis. When released during a gust encounter, the folding motion can reduce the local wingtip angle of attack and unload the outboard wing, thereby reducing the Wing-Root Bending Moment (WRBM). Previous numerical and experimental studies have demonstrated the potential of the concept, but also show that its performance depends strongly on hinge dynamics, release timing and post-gust oscillatory behaviour. This paper presents the development and application of a non-linear, time-domain aeroelastic framework for analysing FFWT release dynamics. The framework couples a Simscape Multibody representation of a flexible main wing and rigid folding tip to an aerodynamic solver based on an Unsteady Vortex Lattice Method (UVLM). The model is used to assess prescribed release strategies as well as based on hinge moment, under discrete vertical gust excitation. The results show that the FFWT response is governed primarily by release phase: early release reduces the critical WRBM peak, whereas release at the locked peak-load instant consistently increases the critical WRBM response. A subsequent hinge-parameter study shows that low hinge stiffness improves load alleviation but increases demands on the hinge angle, while damping mainly affects post-gust dynamic quality. For the simulated configuration, the best compromise is obtained with a low-to-moderate post-release stiffness, sufficient damping, and a low hinge-moment threshold, retaining most of the peak-load reduction of the most compliant setting while substantially reducing hinge-angle demand.

Keywords: Flared folding wingtip, gust load alleviation, aeroelasticity, hinge-moment release, unsteady vortex lattice method, non-linear multibody dynamics.

Nomenclature

A^n	Aerodynamic interface wrench set at time step n
c	Reference chord length
C_θ	Hinge damping coefficient
C_{D_i}	Induced-drag coefficient
C_L	Lift coefficient
e	Oswald span-efficiency factor
f_g	Gust frequency
$\hat{f}_g$	Relative gust frequency, f_g/f_s
f_s	Reference structural frequency
I_{hinge}	Effective hinge inertia
k_g	Aerodynamic reduced gust frequency
K_θ	Hinge stiffness
L_w	Wake length
$M_{\min}$	Critical WRBM peak in the gust-load direction
M_h	Hinge-moment component used for release logic
$M_{h,\text{thr}}$	Hinge-moment release threshold
q_{CB}	Craig-Bampton reduced coordinate vector
R_k	Rotation matrix of interface frame k
r_k	Position vector of interface frame k
S^n	Structural interface state at time step n
t_{rel}	Release time
T_{CB}	Craig-Bampton transformation matrix
u	Physical displacement vector
V_∞	Freestream velocity
$\Delta\alpha_{\text{tip,kin}}$	Kinematic local wingtip angle-of-attack relief
ΔM_h	Gust-induced hinge-moment increment
ΔWRBM	WRBM reduction relative to the locked baseline
η	Non-dimensional hinge-moment threshold fraction
θ	Fold angle about the flared hinge axis
$\dot{\theta}$	Fold rate
Λ	Hinge flare angle
ζ_h	Hinge damping ratio

1 Introduction

Increasing aerodynamic efficiency remains a central objective in the development of future transport aircraft. While improvements in propulsion systems, hydrogen-based concepts and sustainable aviation fuels are important routes toward reduced emissions [1, 2], aerodynamic and structural efficiency remain equally important. For a fixed aircraft weight and flight condition, induced drag is one of the main aerodynamic penalties associated with lift generation and is inversely related to aspect ratio [3, 4]. Increasing wing span is therefore an attractive way of reducing induced drag and improving cruise efficiency. However, larger span also increases the Wing-Root Bending Moment (WRBM), raises sensitivity to gust and manoeuvre loads, strengthens aeroelastic coupling and introduces airport-gate compatibility constraints. Conventional folding wingtips, such as those implemented on the Boeing 777X, address the gate-compatibility problem by folding on the ground, but do not directly provide in-flight load alleviation [5]. This motivates wingtip concepts that combine the aerodynamic benefit of increased span with a mechanism for reducing critical in-flight loads. The Flared Folding Wingtip (FFWT) is one such concept. It consists of an outboard wingtip connected to the main wing through a hinge whose axis is rotated relative to the freestream direction, as indicated as V_{inf} in Figure 1. This hinge flare couples upward folding motion to a nose-down change in local wingtip incidence. For conceptual interpretation, the geometric angle-of-attack relief produced by a fold angle θ and flare angle Λ is often written in the literature as:

$$\Delta\alpha_{\text{tip,geom}} = -\arctan\left(\tan\theta\sin\Lambda\right) \approx -\theta\sin\Lambda \quad (1)$$

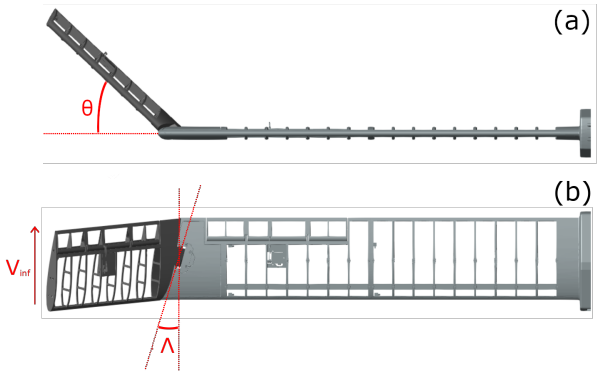

Figure 1: Front (a) and top (b) view of flared folding wingtip concept.

This relation illustrates the basic FFWT mechanism: positive upward folding of a flared tip produces a nose-down change in local incidence. However, it is used here only as a conceptual small-deflection interpretation. The post-processing metric used in the present work therefore retains the finite-angle geometric effect of the flared hinge, rather than relying on the small-angle approximation.

Thus, when the wingtip is released during a gust encounter, upward folding can reduce the local angle of attack of the outboard panel, shed lift near the wingtip and reduce the contribution of the outer span to WRBM. In the locked configuration, the wing retains the aerodynamic benefit of increased span during nominal flight. In the released configuration, part of the outboard aerodynamic load can be absorbed through wingtip motion rather than being transmitted directly to the root.

The FFWT therefore sits between passive and active load alleviation: the folding response is governed by aerodynamic, inertial and structural effects, while the decision to release or control the hinge may be imposed through a release law or actuator.

Previous numerical and experimental studies have demonstrated the load-alleviation potential of the FFWT. Early parametric studies showed that a flared hinge can reduce WRBM during gust encounters and identified flare angle, tip inertia, hinge stiffness and hinge damping as key design parameters [6–8]. Wind-tunnel experiments subsequently confirmed that unlocked or released flared wingtips can reduce WRBM relative to a locked configuration under discrete gust excitation [9–11]. These studies also showed that the benefit is strongly dependent on release timing. Early release can reduce the critical gust-load peak, whereas release close to the locked peak-load instant may erode the benefit or even increase the load response [11, 12]. The FFWT should therefore not be interpreted as a simple load-shedding device, but as a phase-sensitive aeroelastic mechanism in which the timing of local wingtip unloading relative to the global WRBM build-up is critical.

Despite this progress, several modelling and design challenges remain. Much of the existing FFWT literature relies on linearised aeroelastic tools, simplified hinge representations or frequency-domain aerodynamic methods. These approaches are efficient and useful for identifying first-order trends, but they become restrictive when large wingtip rotations, non-linear hinge behaviour, release transients and post-gust oscillations are important [13]. More recent multibody formulations can represent geometric non-linearity more directly, but are often developed as dedicated solvers and are not always coupled to time-domain unsteady aerodynamic methods suitable for large-rotation responses [14, 15]. In addition, the hinge moment is often treated primarily as a response quantity rather than as a central release-design variable. There is therefore a need for a unified time-domain framework that links hinge-moment evolution, threshold crossing, release timing, WRBM alleviation and post-gust dynamic behaviour.

The objective of this research is to assess whether release strategies based on the hinge moment can improve gust-load alleviation for a FFWT under discrete vertical gust excitation. To this end, a non-linear time-domain aeroelastic framework is developed by coupling a Simscape Multibody representation of a flexible main wing and rigid flared folding tip to a UVLM-based unsteady aerodynamic solver evaluated on the instantaneous lifting-surface configuration. This paper makes four main contributions. It first develops and verifies a coupled Simscape-UVLM framework suitable for large FFWT rotations. The framework is then used to assess prescribed release timing over a range of gust time scales. Next, the hinge-stiffness and hinge-damping trade-offs are quantified in terms of WRBM reduction, hinge-angle demand and post-gust response. Finally, release thresholds based on hinge moment are evaluated for selected hinge-property candidates.

The paper is organised as follows. Section 2 describes the methodology used for the structural model, multibody implementation, aerodynamic formulation and aeroelastic coupling. Section 3 subsequently presents the verification and validation of the numerical model. Section 4 presents the release-timing, hinge-parameter and hinge-moment-threshold results. Section 5 discusses the physical interpretation, design implications and limitations of the findings. Finally, Section 6 summarises the main conclusions and identifies directions for future work.

2 Methodology

The methodology is designed to resolve the coupled non-linear dynamics of a flexible main wing with a rigid flared folding wingtip undergoing large hinge rotations. The framework combines a reduced-order structural model exported from Simcenter/Nastran, a Simulink/Simscape Multibody implementation of the flexible–rigid multibody system, and a UVLM-based aerodynamic solver evaluated on an updated lifting-surface lattice. The methodology is organised into the following two steps: construction of the numerical model and the verification and validation of that model. An overview of the coupled aeroelastic toolchain and the exchanged variables is shown in Figure 2.

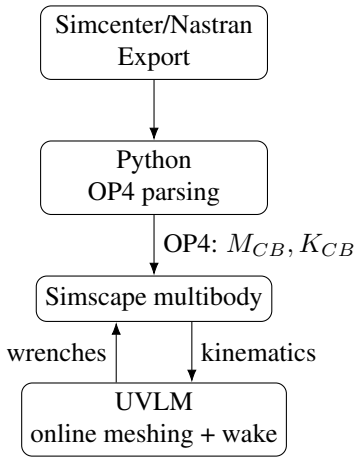

Figure 2: Overview of the coupled aeroelastic toolchain.

Structural Model and Modal Reduction

The FFWT model is represented by a flexible main wing and a rigid folding wingtip. The flexible main wing extends from the root to the hinge and is based on the full-order finite-element model developed by Hoyng [16]. The wingtip is treated as a rigid body, which is an assumption that is consistent with previous FFWT modelling studies and experimental observations that local tip deformation is small compared with the main-wing deformation under comparable loading [6, 10, 12, 17]. This partition preserves the dominant structural flexibility while keeping the model efficient enough for repeated time-domain simulation.

The full-order structural dynamics for the main wing are defined as:

$$\mathbf{M}\ddot{\mathbf{u}} + \mathbf{C}\dot{\mathbf{u}} + \mathbf{K}\mathbf{u} = \mathbf{F} \tag{2}$$

Here, $\mathbf{u}$ is the displacement vector and $\mathbf{M}$, $\mathbf{C}$ and $\mathbf{K}$ are the mass, damping and stiffness matrices, respectively. The flexible main wing is reduced using Craig-Bampton Component Mode Synthesis [18, 19]. The retained boundary coordinates are kept physical, while the interior dynamics are represented by fixed-interface modes. The Craig-Bampton transformation is written as:

$$\begin{bmatrix} \mathbf{u}_b \\ \mathbf{u}_i \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{\Psi}_c & \mathbf{\Phi}_n \end{bmatrix}}_{\mathbf{T}_{CB}} \begin{bmatrix} \mathbf{u}_b \\ \boldsymbol{\eta} \end{bmatrix} = \mathbf{T}_{CB}\mathbf{q}_{CB} \tag{3}$$

Galerkin projection subsequently gives the reduced system:

$$\mathbf{M}_{CB}\ddot{\mathbf{q}}_{CB} + \mathbf{C}_{CB}\dot{\mathbf{q}}_{CB} + \mathbf{K}_{CB}\mathbf{q}_{CB} = \mathbf{Q}_{CB} \tag{4}$$

The retained boundary set consists of the root, the hinge and nine spanwise aerodynamic coupling nodes along the flexible main wing that coincide with the location of spanwise ribs. Each retained node contains six physical degrees of freedom, giving $n_b = 66$ retained boundary coordinates. The resulting Craig-Bampton mass and stiffness matrices are exported from Simcenter/Nastran in OP4 format, parsed using the `pyNastran` toolbox [20], and converted to MATLAB-compatible files for import into Simscape.

Simscape Multibody Implementation

The reduced flexible main wing is imported into Simscape Multibody using a Reduced-Order Flexible Solid (ROFS) block [21]. A custom unit system of mm/kg/s is enforced to remain consistent with the Simcenter/Nastran matrix export. The rigid wingtip is implemented using a File Solid block with CAD-derived inertial properties [16, 22]. Additional frames are defined on the wingtip for the hinge connection and aerodynamic load transfer.

The main wing and wingtip are connected through a Revolute Joint block [23]. Rigid transform blocks align the default revolute axis with the physical flared hinge axis. The hinge-axis orientation follows the geometric definition of the demonstrator developed by Hoyng [16], for which the wingtip hinge line is set at a flare angle of $\Lambda = 15^\circ$ with respect to the freestream direction. The joint provides the fold angle $\theta(t)$ and fold rate $\dot{\theta}(t)$. Depending on the simulation case, the hinge is kept locked, released as a free hinge, or assigned finite post-release stiffness K_θ and damping C_θ. After release, the hinge motion is not prescribed; it follows from the coupled aerodynamic, inertial and structural response.

Aerodynamic loads are applied as equivalent wrenches at the structural coupling frames using External Force and Torque blocks [24]. To preserve work consistency, the same physical frames are used for motion extraction and load application. This ensures that load directions and moment arms evolve consistently with the large wingtip rotations.

Aerodynamic Model

The aerodynamic model is based on the UVLM formulation used in the PROTEUS framework developed at TU Delft [17, 25–27]. UVLM is selected because it operates in the time domain, represents unsteady lift build-up through bound and shed vorticity, and can accommodate non-planar lifting-surface motion at moderate computational cost [28, 29]. This makes it more suitable for the present large-rotation FFWT problem than the frequency-domain Doublet Lattice Method (DLM), which is mainly appropriate for small perturbations.

The main wing and folding tip are represented as two lifting surfaces. The main wing uses 10×39 panels and the wingtip uses 10×15 panels, giving 540 bound panels in total. Spanwise clustering is applied near the hinge and tip, while chordwise spacing is kept uniform. The wake is convected with a prescribed convection velocity approximately equal to the freestream velocity. The wake spacing is $\Delta x_w/c = 0.1$, the target wake length is $L_w/c = 20$, and the effective UVLM time step is 2.2×10^{-3} s.

At each aerodynamic evaluation, the no-penetration condition is imposed at the panel collocation points:

$$\mathbf{n}_j^n \cdot \left(\mathbf{V}_\infty + \mathbf{V}_{\text{motion},j}^n + \mathbf{V}_{\text{ind},j}^n\right) = 0 \qquad j = 1, \ldots, N_b \tag{5}$$

The resulting linear system is solved for the bound circulation vector. Aerodynamic forces and moments are then evaluated using an unsteady formulation based on the Kutta-Joukowski condition and condensed into equivalent interface wrenches.

A prescribed-wake model is used rather than a free-wake model. In the continuous-time UVLM formulation, wake transport can be expressed using the potential jump $\Delta\phi$ as:

$$\frac{\partial \Delta\phi}{\partial t} + \mathbf{V}_w \cdot \nabla \Delta\phi = 0 \tag{6}$$

Although induced-velocity wake roll-up is neglected, new wake rows are shed from the instantaneous trailing-edge geometry. Consequently, large wingtip rotations still produce a prescribed but non-planar wake. The aerodynamic model is therefore suitable for low-speed, attached-flow and comparative aeroelastic analysis, but it does not capture viscous separation, dynamic stall, skin-friction drag or compressibility effects.

Aeroelastic Coupling

The coupled model uses a partitioned time-marching scheme. At each coupling instant, Simscape provides the pose and rates of the structural interface frames. The aerodynamic solver reconstructs the lattice, evaluates the UVLM loads and returns equivalent interface wrenches to the structural model. The structural state passed to the aerodynamic solver is:

$$\mathcal{S}^n = \{\mathbf{r}_k^n, \mathbf{R}_k^n, \dot{\mathbf{r}}_k^n, \boldsymbol{\omega}_k^n\}_{k=1}^{N_f} \tag{7}$$

The aerodynamic output is the corresponding set of interface wrenches:

$$\mathcal{A}^n = \{\mathbf{F}_k^n, \mathbf{M}_k^n\}_{k=1}^{N_f} \tag{8}$$

Online mesh updating is used to reconstruct the instantaneous aerodynamic lattice. For a lattice vertex associated with interface frame k:

$$\mathbf{x}_i^n = \mathbf{r}_k^n + \mathbf{R}_k^n \mathbf{x}_{i,k}^0 \tag{9}$$

The corresponding rigid-body contribution to the local surface velocity is then:

$$\mathbf{V}_{\text{motion},j}^n = \dot{\mathbf{r}}_k^n + \boldsymbol{\omega}_k^n \times \left(\mathbf{x}_j^n - \mathbf{r}_k^n\right) \tag{10}$$

Thus, the flared wingtip rotation enters the aerodynamic simulation through updated panel positions, normals, trailing-edge geometry and motion boundary conditions. The baseline coupling is weakly coupled, with one aerodynamic evaluation per structural coupling step. Aerodynamic loads are held constant over the interval $[t_n, t_{n+1}]$ to avoid an algebraic loop between structural motion and aerodynamic loading.

3 Verification and Validation

The framework is verified and validated using a ladder of structural, aerodynamic, coupled static and dynamic trend comparisons. This section establishes the model credibility required for the comparative release-strategy studies presented in Section 4.

Table 1: Verification and validation ladder used for the coupled Simscape–UVLM framework.

Level	Purpose
Structural model	Verify ROM properties, static compliance and imported flexible-body dynamics.
Aerodynamic model	Validate locked-wing UVLM loads, lift-curve behaviour and coordinate consistency.
Coupled static response	Verify aero-to-structural load transfer and static aeroelastic equilibrium.
Dynamic FFWT response	Validate the release-timing hierarchy and phase-sensitive load-alleviation behaviour against Córcoles.

Structural Verification

At structural level, the Craig-Bampton matrices are checked to determine whether they preserve the relevant properties of the full-order finite-element model. The retained boundary set consists of the root, hinge and aerodynamic coupling frames, giving $n_b = 66$ physical boundary coordinates. A modal convergence study is used to select $n_m = 80$ fixed-interface modes, resulting in a 146×146 reduced model. Within the validation bandwidth $f_{\text{val}} = 130$ Hz, the selected ROM gives a mean absolute eigenfrequency error of 1.02% and a maximum error of 4.05% relative to the full-order model. The matrices are also checked for symmetry, dimensional consistency, positive definiteness and numerical conditioning.

After confirming that the matrices are suitable for the Simscape flexible body, the imported model is verified against Simcenter static and dynamic reference responses. Static SOL101 load cases are used to check the dominant hinge-compliance channels under representative force and moment loading. The dominant channels agree closely, while near-zero spanwise compliance is treated as a secondary sign-convention check because very small absolute displacements can produce misleadingly large percentage errors. This verifies that the low-frequency stiffness behaviour of the imported flexible body is consistent with the finite-element reference.

Dynamic impulse-response comparisons are then used to verify that the imported model also preserves the transient bending- and torsion-dominated dynamics. The results are summarised in Table 2.

The bending-dominated $u_z(t)$ response gives matching dominant frequency in Simcenter and Simscape, while the torsion-dominated $\theta_y(t)$ response differs by less than 1% in dominant frequency. The high correlation coefficients and low Normalised Root-Mean-Square Error (NRMSE) values confirm that the imported flexible-body dynamics are sufficiently accurate for the subsequent aeroelastic simulations.

Table 2: Summary of structural dynamic validation against Simcenter impulse responses.

Response	f_{SC} [Hz]	f_{SS} [Hz]	corr. [-]	NRMSE [-]
$u_z(t)$	3.03	3.03	0.998	0.065
$\theta_y(t)$	32.42	32.72	0.966	0.072

Aerodynamic Validation

At aerodynamic level, the UVLM implementation is validated against XFLR5 Vortex-Lattice-Method (VLM) and lifting-line estimates for the locked rigid-wing configuration [3, 30], before enabling aeroelastic feedback. This isolates the UVLM implementation from structural deformation and checks whether the aerodynamic lattice, incidence definition, reference area, force extraction and coordinate transformations are implemented consistently. After these corrections, the PROTEUS UVLM lift curve agrees with XFLR5 VLM to within approximately 1% over the investigated angle-of-attack range.

The induced-drag coefficient is used as an additional consistency check. The UVLM-induced drag shows a systematic 4–5% difference relative to XFLR5. This discrepancy is considered acceptable for the present purpose because $C_{D_i} \propto C_L^2$ amplifies small lift differences, and because the UVLM induced-drag estimate is reconstructed from the spanwise circulation distribution.

The aerodynamic validation therefore indicates that the locked-wing lifting response is captured with sufficient fidelity for comparative low-speed, attached-flow aeroelastic simulations.

Coupled Static Verification

After separate checks for the structural and aerodynamic model, the coupled numerical model is first verified at static level. This step is necessary since agreement of isolated subsystems does not guarantee that aerodynamic wrenches are applied at the correct structural locations, in the correct frame convention, or with consistent moment arms. The coupled checks are therefore focused on consistency in load transfer and static aeroelastic equilibrium.

First, a frozen-load comparison is performed. Aerodynamic wrenches are computed with the UVLM solver for a locked-wing condition and then applied statically in both Simscape and a SOL101 reference model. This isolates the aero-to-structural load transfer from the unsteady feedback loop. The resulting vertical deflections agree with a mean percentage error below 1%, while the wingtip displacement differs by 1.597%. This confirms that the aerodynamic loads are transferred into the structural model consistently.

Second, a low-speed, low-incidence locked-hinge SOL144 comparison is used to assess the coupled static aeroelastic equilibrium. The comparison is deliberately restricted to $V_\infty = 20$ m/s and $\alpha = 1^\circ$ so that the response remains within the range where the linear Simcenter reference solution is meaningful. Most local deflection differences remain close to 2%, with a wingtip difference of -3.1%.

The coupled framework therefore reproduces both the magnitude and the spanwise trend of the static aeroelastic deformation in the low-speed locked-hinge regime.

Dynamic Trend Validation

Finally, a dynamic trend validation is performed against the gust-response experiments of Córcoles [11, 12]. This final validation step assesses whether the complete coupled framework is able to reproduce the dynamic FFWT behaviour that was observed experimentally.

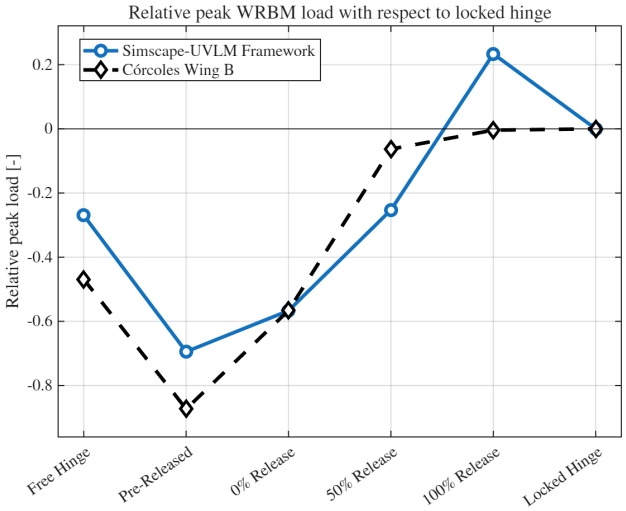

Figure 3: Dynamic trend validation against the 5 Hz gust-release experiment of Córcoles for Wing B [12].

The comparison uses the $\alpha = 5^\circ$, $\delta_g = 2.5^\circ$, $f_g = 5$ Hz case for Córcoles' Wing B.

The numerical model created in Simscape is not geometrically or structurally identical to this Wing B, but is an enlarged version of it. Therefore, the comparison is not interpreted as a pointwise validation of absolute WRBM magnitudes, but instead is used as a trend-level validation of the release-timing hierarchy. This is the relevant validation target for the present work, because the subsequent simulations use the framework primarily for comparative assessment of release strategies rather than for load prediction at certification level. The relative peak-load values from the reference case are compared with the recreated Simscape–UVLM results in Figure 3.

In this figure, negative values indicate alleviation relative to the locked hinge, whereas positive values indicate a larger peak load than the locked reference. The comparison recreates the relative peak-load alleviation metric reported by Córcoles, and is therefore used to assess whether the Simscape–UVLM framework reproduces the experimentally observed release-timing hierarchy. The framework reproduces the governing trend, where free and pre-released cases provide substantial alleviation, gust-onset release remains beneficial, delayed release reduces the benefit, and release near the locked-reference peak becomes ineffective or detrimental. The comparison also clarifies the interpretation of release timing. Nominally identical release labels do not necessarily correspond to identical aeroelastic phase relations, because the experiment includes finite release-mechanism response time whereas the numerical model applies prescribed release instants. This distinction proves important: in the validation comparison, the objective is to reproduce the relative release-timing trend observed experimentally. In the subsequent parametric studies, however, the primary metric is the critical WRBM peak in the gust-load direction, $M_{\min}$, rather than the absolute maximum of the WRBM signal.

This avoids interpreting phase-shifted upward rebound in early-release cases as the critical gust-load response. The validation is therefore interpreted in terms of effective release phase relative to the gust-induced WRBM build-up. On this basis, the framework captures the key FFWT mechanism required for the present study: early wingtip motion can unload the outboard wing before the critical root-load peak forms, while late release can mainly excite post-gust motion without providing useful peak-load alleviation.
Taken together, these steps in the verification and validation ladder establish the framework as sufficiently credible for the comparative release-strategy studies that follow. The structural, aerodynamic and coupled static checks verify the numerical implementation of the model components and their interfaces, while the dynamic comparison with Córcoles confirms that the framework reproduces the release-timing hierarchy that governs FFWT gust-load alleviation.

4 Results

The results are organised around the three design questions addressed in this work. First, the effect of release timing is studied by varying the gust time scale and comparing locked, free and prescribed-release cases. Second, the post-release hinge stiffness and damping are varied to quantify the trade-off between load alleviation, hinge-angle demand and post-gust response. Finally, selected hinge-property candidates are evaluated using a threshold release law based on hinge moment.

Evaluation Metrics

The validated numerical model is assessed using structural, kinematic and release-related response quantities. The primary load metric is the critical WRBM peak in the gust-load direction. In the adopted sign convention, the locked gust response produces a negative WRBM peak. The critical peak for each case is therefore defined as:

$$M_{\min} = \min_{t_g \le t \le t_{\text{end}}} \text{WRBM}(t) \tag{11}$$

t_g is the gust onset time in this equation and the corresponding peak-load reduction relative to the locked baseline is:

$$\Delta_{\text{WRBM}} = \left(1 - \frac{|M_{\min,\text{case}}|}{|M_{\min,\text{locked}}|}\right) 100\% \tag{12}$$

Positive values of Δ_{WRBM} indicate alleviation of the critical downward WRBM peak, whereas negative values indicate a larger critical peak than the locked reference. This definition focuses on the gust-load direction and avoids interpreting phase-shifted rebound in the opposite direction as the critical load-alleviation metric.

Additional response quantities are used to interpret the load metric, including the peak hinge angle, hinge moment, release time, post-gust dynamic behaviour and local wingtip angle-of-attack relief. The latter is computed from the exact local kinematics of the flared folding tip rather than from the small-angle relation in Equation 1.

The local unloading mechanism is interpreted using a geometric wingtip incidence-relief metric. This metric isolates the incidence change caused by the flared hinge orientation and fold angle. The additional apparent-angle contribution from hinge-rate motion is included in the coupled UVLM simulation through the moving-surface boundary condition, but is not included in this scalar post-processing metric.

The relevant tip-frame velocity components are defined as:

$$u_T = V_\infty \cos\alpha \cos\Lambda \tag{13}$$

$$w_T = V_\infty \left(\sin\alpha \cos\theta - \cos\alpha \sin\theta \sin\Lambda\right) \tag{14}$$

The local flow-model angle of attack is then defined as:

$$\alpha_{\text{tip}}(t) = \text{atan2}(w_T, u_T) \tag{15}$$

The spanwise velocity component is treated as sideslip and is therefore not included in the chord-normal incidence angle.

The local AoA relief is evaluated relative to the locked local tip angle:

$$\Delta\alpha_{\text{tip,rel}}(t) = \alpha_{\text{tip,locked}} - \alpha_{\text{tip}}(t) \tag{16}$$

Positive values therefore correspond to nose-down local AoA relief.

Since only unloading generated before the locked-reference WRBM peak can reduce the critical peak load, the useful relief is reported as:

$$\Delta\alpha_{\text{tip}} = \max_{t_g \le t \le t_{M,\text{locked}}} \Delta\alpha_{\text{tip,rel}}(t) \tag{17}$$

where

$$t_{M,\text{locked}} = \arg\min_{t_g \le t \le t_{\text{end}}} \text{WRBM}_{\text{locked}}(t) \tag{18}$$

Frequency-Scaled Characterisation

The first analysis examines whether the gust-load alleviation behaviour of the FFWT can be characterised using scaled gust-frequency parameters. This is motivated by the fact that the effectiveness of a released wingtip depends not only on the magnitude of its motion, but also on whether the resulting local unloading occurs on a time scale that is relevant to the global WRBM build-up. Two non-dimensional measures are therefore considered: a structural relative frequency and an aerodynamic reduced frequency.

The structural relative frequency is defined as:

$$\hat{f}_g = \frac{f_g}{f_s} \tag{19}$$

Here, f_g is the gust frequency and f_s is the selected reference structural frequency of the model. The gust frequencies are selected by matching the relative-frequency ratios of the Córcoles reference cases [12] to the modal frequencies of the present configuration, presented in Table 3.

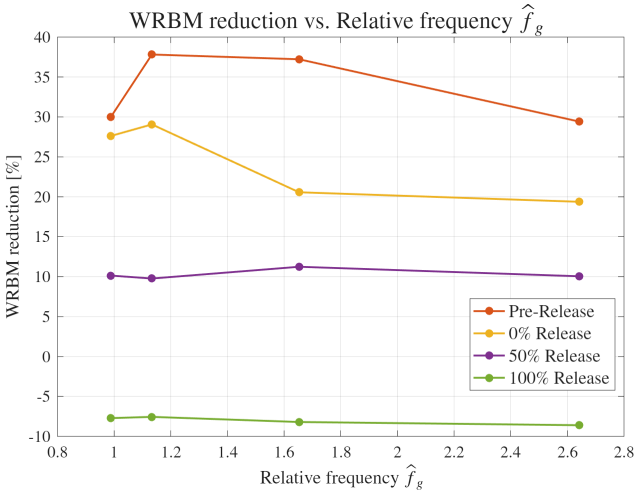

(a) Peak WRBM reduction as a function of relative gust frequency $\hat{f}_g$.

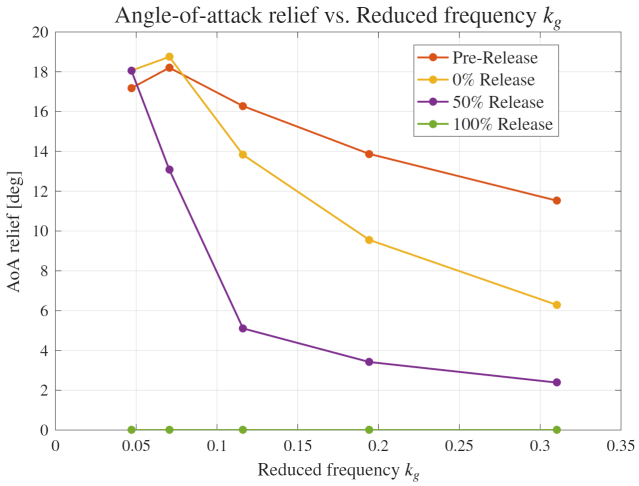

(b) Kinematic wingtip angle-of-attack relief available before the locked WRBM peak as a function of reduced gust frequency k_g.

Figure 4: Frequency-scaled interpretation of FFWT gust-load alleviation using structural relative frequency and aerodynamic reduced frequency.

Table 3: Structural reference frequencies used for the relative-frequency scaling.

Hinge condition	**Reference mode**	f_s [Hz]
Free hinge	Flapping mode	0.9
Locked hinge	First bending mode	1.7

This gives four scaled gust cases presented in Table 4.

Table 4: Gust cases selected by structural relative-frequency scaling.

Case	**Scaling target**	f_g [Hz]	$\hat{f}_g = f_g/f_s$ [-]
1	Flapping-scaled	1.02	1.13
2	Bending-scaled	1.68	0.99
3	Intermediate	2.81	1.65
4	High-frequency	4.49	2.64

For each gust case, a locked reference simulation is first used to define the prescribed release instants. The response is then evaluated for locked, free, pre-released, gust-onset release, mid-build-up release and peak-load release conditions. The last three correspond to release at 0%, 50% and 100% of the locked-reference WRBM build-up.

The same cases are also interpreted using the aerodynamic reduced gust frequency,

$$k_g = \frac{\omega_g c}{2V_\infty} = \frac{\pi f_g c}{V_\infty} \tag{20}$$

In this definition, c is the reference chord and V_∞ is the freestream velocity. An additional low-frequency anchor to clarify the low-frequency trend is included at $f_g = 0.68$ Hz, corresponding to $\hat{f}_g = 0.40$ and $k_g = 0.047$.

The resulting peak WRBM reductions for the scaled gust based on relative frequency are shown in Figure 4a, while the corresponding reduced-frequency interpretation using the available kinematic tip angle-of-attack relief is shown in Figure 4b.

Early release provides the largest reduction of the critical downward WRBM peak, while release at the locked peak-load instant is consistently detrimental. The corresponding numerical values are summarised in Table 5, together with the relative and reduced frequencies. The pre-release cases reduce the critical WRBM peak by approximately 29-38%, while release at gust arrival remains beneficial but loses effectiveness as the gust excitation becomes faster. The 50% release case gives only limited alleviation, close to 10% in all cases. The 100% release case increases the critical downward WRBM peak by approximately 8% relative to the locked baseline.

Table 5: Peak WRBM reduction in the gust-load direction for the structurally scaled gust cases.

	Case 1	**Case 2**	**Case 3**	**Case 4**
f_g [Hz]	1.02	1.68	2.81	4.49
$\hat{f}_g$ [-]	1.13	0.99	1.65	2.64
k_g [-]	0.070	0.116	0.194	0.310
Pre-release [%]	37.8	30.0	37.2	29.4
0% release [%]	29.1	27.6	20.6	19.4
50% release [%]	9.8	10.1	11.2	10.0
100% release [%]	-7.6	-7.7	-8.2	-8.6

The timing of the locked WRBM peak explains this behaviour. Increasing the gust frequency reduces the time available between gust arrival and the locked critical peak. The time from gust onset to the locked peak decreases from 1.039 s for $f_g = 1.02$ Hz to 0.357 s for $f_g = 4.49$ Hz. The time between the 50% release and the locked peak similarly decreases from 0.508 s to 0.131 s. The useful release window therefore becomes progressively narrower as the gust excitation becomes faster. As a result, delayed release initiates wingtip motion after much of the critical load build-up has already occurred.

The hinge-angle responses show that peak hinge deflection alone does not determine the load reduction. Several delayed-release cases reach large hinge angles but provide only modest WRBM alleviation because the motion is poorly phased relative to the locked WRBM build-up. Conversely, early-release cases generate local unloading while the gust load is still developing, which reduces the critical root load even when the peak hinge angle is not the largest. This confirms that the FFWT release mechanism is primarily phase-sensitive: the wingtip must rotate early enough for the local aerodynamic unloading to affect the initial WRBM excursion.

The reduced-frequency interpretation adds the aerodynamic time-scale perspective. The angle-of-attack relief in Figure 4b is evaluated only up to the locked-reference WRBM peak, so it represents the local geometric unloading available during the critical load-build-up phase. This distinction is important: the peak-load release cases show almost no useful relief because the hinge is released when the relevant unloading window has effectively closed. Conversely, the low-k_g cases allow delayed release to remain more effective because the gust-induced loading develops more slowly relative to the convective time scale. The global WRBM benefit therefore depends not only on the magnitude of the local tip unloading, but also on whether that unloading is generated before the critical root-load peak forms.

The scaled-frequency interpretation therefore organises the response trends, but does not collapse the FFWT behaviour onto a single universal parameter. The structural relative frequency is useful for comparing gust time scales against the modal response of the model, while the reduced frequency helps interpret whether local wingtip unloading occurs early enough in aerodynamic time to affect the WRBM build-up. The load-alleviation performance is therefore characterised by frequency scaling only in combination with release phase.

These results support the selection of the gust-onset release case at $f_g = 2.81$ Hz as the reference condition for the hinge-property study. This case is a practical triggered-release condition rather than an idealised free or pre-released configuration. It gives clear peak WRBM alleviation, remains outside the near-quasi-steady low-k_g range, and avoids the consistently detrimental behaviour of peak-load-triggered release.

Hinge-Stiffness and Damping Sensitivity

The hinge-property sensitivity study is performed at the reference condition selected from the frequency-scaled analysis, summarised in Table 6.

This condition is used because it gives clear peak load alleviation while remaining sensitive to release timing. The hinge-property sweep is then defined using the effective wingtip inertia about the flared hinge axis.

Table 6: Reference condition used for the hinge-property sensitivity study.

Parameter	Value
Angle of attack	5 [°]
Freestream velocity	10 [m/s]
Gust amplitude	2.5 [°]
Gust frequency	2.81 [Hz]
Relative gust frequency	1.65 [−]
Reduced gust frequency	0.194 [−]
Release time	5.2 [s]
Release condition	Gust onset

From the Simscape rigid-body properties, this inertia is estimated as:

$$I_{\text{hinge}} = \mathbf{e}_h^T \mathbf{I} \mathbf{e}_h, \qquad \mathbf{e}_h = \begin{bmatrix} \cos\Lambda & \sin\Lambda & 0 \end{bmatrix}^T \tag{21}$$

This gives $I_{\text{hinge}} \approx 0.0232\ \text{kg m}^2$.

The stiffness values are selected on a logarithmic scale to resolve the transition from free-like to locked-like post-release behaviour:

$$K_\theta \in \{0.3,\ 1,\ 3,\ 10\}\ \text{N m/rad}. \tag{22}$$

The damping sweep is defined using damping ratios:

$$\zeta_h \in \{0,\ 0.02,\ 0.10,\ 0.20\}, \tag{23}$$

The damping ratios are selected to range from an undamped reference to a strongly damped response, allowing the influence of damping on post-release oscillation decay to be isolated from the stiffness-driven change in peak hinge motion. The corresponding damping coefficient for each stiffness level and damping ratio is computed using:

$$C_\theta = 2\zeta_h \sqrt{K_\theta I_{\text{hinge}}}. \tag{24}$$

The purpose of the sweep is not to identify a unique optimum, but to map how the response moves between the free-hinge and locked-hinge limits as stiffness and damping are varied and to identify representative hinge-property regimes for the subsequent threshold-release analysis.

The results of the hinge-property sensitivity study are summarised in Figure 5. The representative high-damping cases selected from this sensitivity study are listed in Table 7, and are used to structure the following discussion.

Table 7: Representative hinge-property trade-off cases at $\zeta_h = 0.20$.

Candidate	K_θ [Nm/rad]	Δ_{WRBM} [%]	$\theta_{\max}$ [°]
Aggressive	0.3	17.0	51
Balanced	1.0	11.0	30
Conservative	3.0	5.6	14

The stiffness sweep at fixed damping ratio $\zeta_h = 0.20$ shows that hinge stiffness is the dominant parameter governing the load-alleviation/motion-demand trade-off. The most compliant case, $K_\theta = 0.3\ \text{Nm/rad}$, gives the largest WRBM reduction, approximately 17%, but also demands the largest peak hinge motion, approximately 51°.

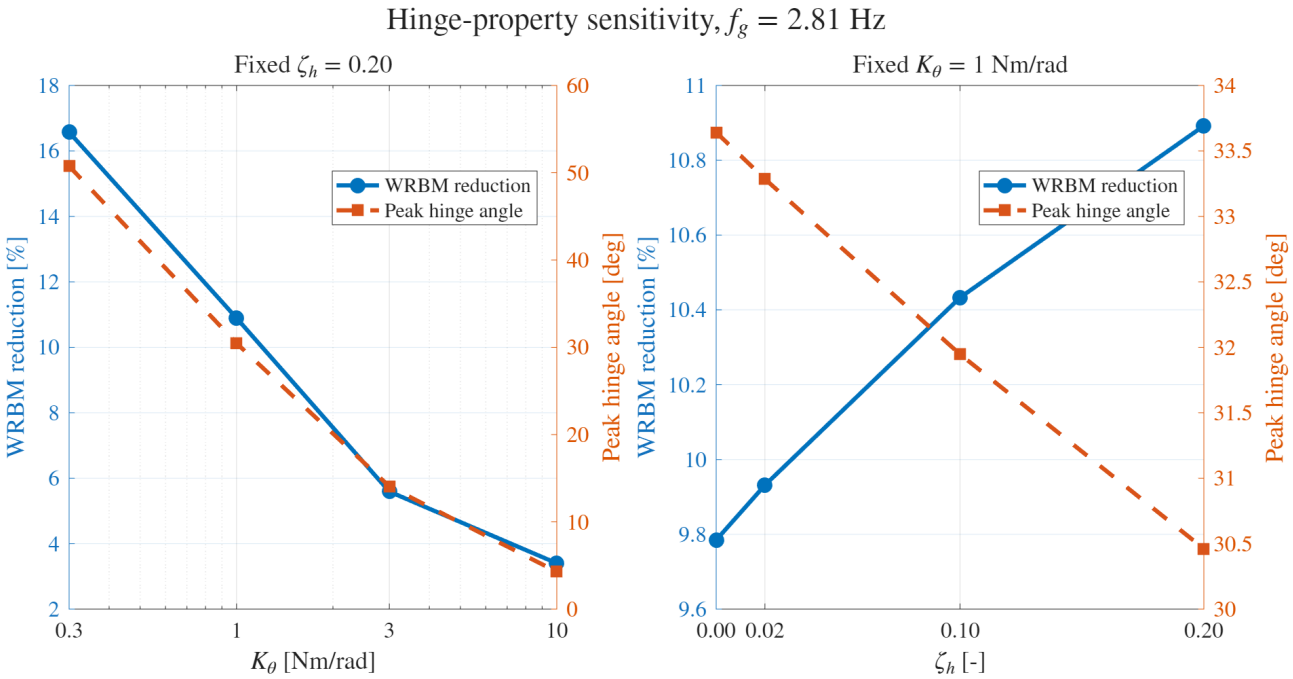

Figure 5: Representative hinge-property sensitivity for the gust-onset release case at $f_g = 2.81$ Hz. The left panel shows a stiffness sweep at fixed damping ratio $\zeta_h = 0.20$, while the right panel shows a damping-ratio sweep at fixed stiffness $K_\theta = 1$ Nm/rad.

Increasing the stiffness to $K_\theta = 1$ Nm/rad reduces the WRBM benefit to approximately 11%, but also lowers the peak hinge angle to about 30°. At $K_\theta = 3$ Nm/rad, the peak-load reduction decreases further to approximately 5–6%, while the peak hinge angle is reduced to about 14°. The stiffest case, $K_\theta = 10$ Nm/rad, approaches the locked-hinge limit and provides only limited load alleviation, around 3–4%, with a peak hinge angle below 5°.

The damping sweep at fixed stiffness $K_\theta = 1$ Nm/rad shows that damping has a weaker influence on the first peak response than stiffness. Increasing ζ_h from 0 to 0.20 produces only a modest increase in peak WRBM reduction, from approximately 10% to 11%, while reducing the peak hinge angle from roughly 34° to 30°. Thus, damping is useful for moderating hinge motion and improving the post-release dynamic response, but it does not change the primary stiffness-controlled trade-off between load alleviation and hinge-angle demand.

These trends show why WRBM reduction alone is insufficient for selecting a hinge configuration. The most compliant hinge gives the strongest peak-load reduction, but at the cost of large wingtip motion. Increasing stiffness reduces the hinge-angle demand, but progressively suppresses the folding mechanism responsible for outboard unloading. The selected candidates therefore represent three points along this trade-off: an aggressive low-stiffness case, a balanced intermediate-stiffness case, and a conservative higher-stiffness case. The $K_\theta = 10$ Nm/rad case is not retained because it behaves too close to the locked baseline and provides limited alleviation.

Based on this trade-off, three candidates are carried forward into the hinge-moment threshold-release analysis. The aggressive candidate uses $K_\theta = 0.3\text{N m/rad}$ and retains the strongest load-alleviation behaviour. The balanced candidate uses $K_\theta = 1.0\text{N m/rad}$ and provides a less compliant but still effective response. The conservative candidate uses $K_\theta = 3.0\text{N m/rad}$ and reduces hinge-angle demand more strongly.

Hinge-Moment Threshold Release Strategy

The final analysis evaluates whether the prescribed-release trends can be reproduced using a local hinge-moment trigger. Three representative hinge-property candidates are selected from the stiffness and damping study: an aggressive, a balanced and a conservative configuration. These candidates represent increasing levels of post-release hinge stiffness and decreasing hinge-angle demand.

The release variable is the gust-induced increment of the hinge-moment component used for release logic, where $M_{h,0}$ is the pre-gust locked-reference hinge moment:

$$\Delta M_h(t) = M_h(t) - M_{h,0}, \tag{25}$$

The threshold is defined from the locked-reference response as:

$$M_{h,\text{thr}} = \eta \max_{t \geq t_g} \left| \Delta M_h^{\text{locked}}(t) \right| \tag{26}$$

In this equation, η is the non-dimensional threshold fraction. The hinge is released when

$$|\Delta M_h(t)| \geq M_{h,\text{thr}}, \qquad t \geq t_g \tag{27}$$

The tested threshold fractions are

$$\eta \in \{0,\ 0.05,\ 0.10,\ 0.15,\ 0.20,\ 0.25\}. \tag{28}$$

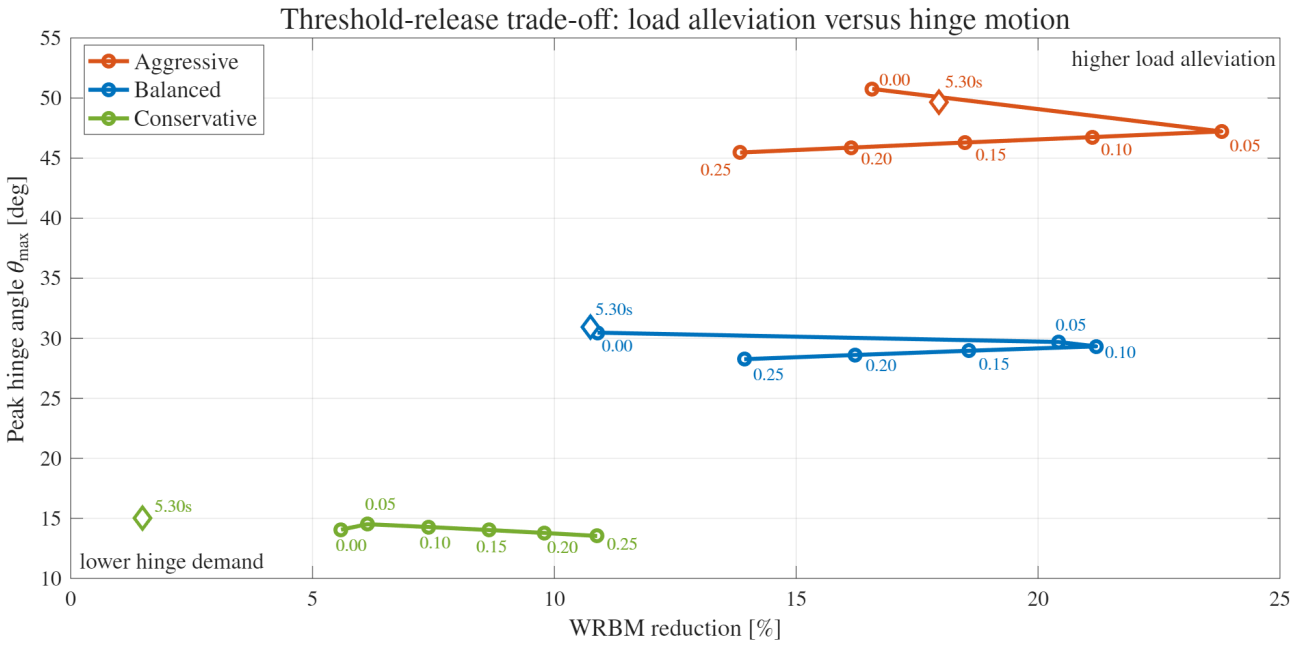

Figure 6: Peak WRBM reduction versus peak hinge angle for the aggressive, balanced and conservative hinge-property candidates. Circular markers represent threshold-release cases, with labels indicating the threshold parameter η. Diamond markers indicate the prescribed-release reference case at $t_{\mathrm{rel}} = 5.30$ s, included as an intermediate timing reference for each candidate.

A prescribed-release case at $t_{\mathrm{rel}} = 5.30$ s is also included for each candidate as an intermediate timing reference.

The extracted release times show that the threshold sweep does not produce a broad timing range. The $\eta = 0$ case releases at gust onset, $t = 5.200$ s, while the non-zero thresholds cases release between 5.422 s and 5.481 s. The spread of the non-zero threshold releases is therefore only 0.059 s. This clustering occurs because the selected threshold levels are crossed on the same steep rising branch of the locked-reference hinge-moment response. The threshold sweep is therefore interpreted as a local phase-sensitivity study around the release window when triggered by hinge moment, rather than as a full release-time sweep.

The resulting trade-off between load alleviation and hinge motion is shown in Figure 6, and the representative best cases are summarised in Table 8. The aggressive and balanced candidates both show a favourable low non-zero threshold. For the aggressive candidate, the maximum WRBM reduction is 23.8% at $\eta = 0.05$. For the balanced candidate, the maximum reduction is 21.2% at $\eta = 0.10$. In both cases, increasing the threshold beyond this favourable range progressively reduces the peak-load benefit, because the wingtip motion is initiated too late to fully affect the WRBM build-up. The conservative candidate behaves differently: its WRBM reduction increases over the tested threshold range, reaching 10.9% at $\eta = 0.25$, but the absolute load-alleviation level remains well below the aggressive and balanced candidates.

Table 8: Representative threshold-release trade-off cases.

	Aggressive	**Balanced**	**Conservative**
Selected threshold	$\eta = 0.05$	$\eta = 0.10$	$\eta = 0.25$
t_{rel} [s]	5.422	5.440	5.481
Δ_{WRBM} [%]	23.8	21.2	10.9
θ_{max} [°]	47.2	29.3	13.5

The hinge-angle demand is governed primarily by the selected hinge-property candidate. The aggressive candidate produces the largest hinge excursions, with peak angles above 45° for all threshold cases. The balanced candidate remains near 28–31°, while the conservative candidate stays near 14–15°. Increasing η generally reduces the peak hinge angle for the aggressive and balanced candidates because the wingtip has less time to rotate before and during the main load transient.

From a pure peak-load perspective, the aggressive $\eta = 0.05$ case is the best-performing configuration. However, the balanced $\eta = 0.10$ case provides the most attractive compromise: it retains most of the WRBM reduction of the aggressive case while reducing the peak hinge angle from approximately 47° to 29°. The conservative candidate successfully limits hinge motion, but suppresses too much of the folding mechanism to be competitive in peak-load alleviation under the present gust condition.

Overall, the threshold-release results show that release based on hinge moment can improve the response relative to simple gust-onset release, but only when the threshold triggers during the favourable load-build-up phase. The balanced hinge-property candidate with a low non-zero hinge-moment threshold is therefore the most promising release-strategy combination tested in this study.

5 Discussion

The results show that the FFWT response is governed by the timing of the folding motion relative to the gust-induced WRBM build-up. Across the release-timing cases, the most effective cases are those in which the wingtip is already free, or is released early enough for the folding motion to develop before the locked-hinge WRBM peak is formed. Conversely, release at the locked peak-load instant is consistently detrimental. This confirms that the FFWT does not behave as a purely load-triggered relief device where larger load automatically implies a better release opportunity. Instead, it is a phase-sensitive aeroelastic mechanism: the local aerodynamic unloading must occur during the formation of the critical root load.

This phase sensitivity also explains why peak hinge angle alone is not a reliable indicator of load alleviation. Several delayed-release cases generate large hinge deflections but produce only limited WRBM reduction, because the corresponding angle-of-attack relief occurs after the critical WRBM excursion has already developed. The reduced-frequency interpretation supports the same conclusion.

At low reduced frequency, the gust develops more slowly and delayed release can remain effective because the wingtip still has time to produce local unloading. At higher reduced frequency, the available release window narrows and the same delayed release logic becomes less effective. The relevant design quantity is therefore not only the magnitude of the local angle-of-attack relief, but its timing relative to the locked WRBM peak.

The hinge-parameter study shows that hinge stiffness is the dominant parameter controlling the trade-off between load alleviation and wingtip motion. A compliant hinge allows the wingtip to rotate sufficiently to unload the outboard wing, giving the strongest WRBM reduction.
However, this benefit comes with large hinge-angle excursions. Increasing stiffness restricts the wingtip motion and reduces the hinge-angle demand, but also suppresses the folding mechanism that produces the aerodynamic unloading. The response therefore transitions from free-like behaviour at low stiffness to locked-like behaviour at high stiffness.

Hinge damping mainly affects the post-release dynamic quality rather than the first WRBM peak. Increasing damping generally reduces the persistence of hinge motion and can lower the post-gust oscillatory response. However, damping does not compensate for excessive stiffness: once the hinge becomes too stiff, the load-alleviation mechanism is already weakened. Damping must therefore be selected together with stiffness, not treated as an independent stabilising correction. The hinge-property sweep also shows that intermediate stiffness is not automatically optimal, since it can combine limited load alleviation with appreciable residual hinge–wing oscillation.

The release analysis based on hinge-moment provides a more physically meaningful trigger than prescribed release time, because it uses a local load quantity at the hinge. However, the results show that the threshold cannot be selected solely by magnitude. The non-zero threshold cases are clustered in time because the locked-reference hinge moment crosses the selected thresholds on a steep rising branch. As a result, small changes in threshold correspond to small but dynamically relevant changes in release phase. The threshold should therefore be tuned using the locked-response phase and the expected gust time scale. A hinge-moment trigger is useful only if it releases the hinge during the favourable load-build-up phase.

The comparison between the aggressive, balanced and conservative candidates highlights that the preferred design is not the case with maximum peak-load reduction alone. The aggressive setting provides the strongest unloading but requires large hinge excursions, while the conservative setting limits motion at the cost of much weaker WRBM alleviation. The balanced candidate therefore represents the most useful compromise for the present configuration, because it retains most of the load-reduction benefit while substantially reducing the wingtip-motion demand.

These findings imply that FFWT release design should not be based on peak WRBM reduction alone. A favourable strategy must reduce the critical WRBM peak, limit excessive wingtip motion, and avoid large post-gust oscillations.

For the present model and gust condition, the balanced hinge-property candidate with a low non-zero hinge-moment threshold provides the best overall compromise. This conclusion is specific to the investigated configuration and gust case, but the underlying mechanism is more general: effective FFWT load alleviation requires a release law that is phase-aware, not merely load-aware.

The main limitation of the present interpretation is that the aerodynamic model is based on potential-flow UVLM assumptions. The framework captures unsteady lifting-surface effects and large geometric changes through online meshing, but it does not model viscous separation, dynamic stall or skin-friction drag. The results should therefore be interpreted as comparative aeroelastic trends under attached-flow conditions.

In addition, the final release-law assessment is based on a selected discrete gust condition. Further work should extend the threshold logic to a broader set of gust amplitudes, frequencies and directions, and should investigate whether the same balanced release strategy remains favourable under continuous turbulence and manoeuvre loading.

6 Conclusions

This paper presented a non-linear time-domain aeroelastic framework for assessing release strategies for a flared folding wingtip under discrete vertical gust excitation. The framework combines a reduced-order flexible main-wing model and rigid folding tip in Simscape Multibody with a UVLM aerodynamic solver evaluated on the instantaneous lifting-surface configuration.

After verification and trend-level validation against experimental release-timing trends, the model was used to assess prescribed release timing, hinge-property effects and hinge-moment-based release thresholds.

The main finding is that FFWT load alleviation is governed by timing rather than release alone. A release is beneficial only if the wingtip motion develops early enough to unload the outer wing before the critical WRBM peak forms; releasing near the locked peak-load instant is already too late and can worsen the response. The hinge stiffness and release threshold therefore act mainly as timing and motion-shaping parameters: lower stiffness improves load relief but demands larger hinge motion, while moderate stiffness combined with a low non-zero threshold gives a more balanced compromise.

These conclusions remain specific to the investigated demonstrator geometry, discrete gust cases, hinge-property range and attached-flow UVLM assumptions. Future work should extend the assessment to certification-relevant 1-cosine gust gradients, positive and negative gust amplitudes, selected sharp-edge or ramp inputs and continuous turbulence. The idealised instantaneous release should also be replaced by a finite-latency release model including sensing, triggering and lock-disengagement delays, followed by robustness assessment under manoeuvre loading, structural limits and full-aircraft integration effects.

References

[1] Airbus. First ZEROe Engine Fuel Cell Successfully Powers On. Airbus web page, 2024. URL https://www.airbus.com/en/newsroom/stories/2024-01-first-zeroe-engine-fuel-cell-successfully-powers-on. Accessed: 2026-04-29.

[2] European Commission. ReFuelEU Aviation. European Commission web page, 2024. URL https://transport.ec.europa.eu/transport-modes/air/environment/refueleu-aviation_en. Accessed: 2026-04-29.

[3] John D. Anderson. *Fundamentals of Aerodynamics*. McGraw Hill, 7 edition, 2024. ISBN 9781264151929.

[4] NASA Glenn Research Center. Induced Drag Coefficient. NASA Glenn Beginner's Guide to Aeronautics, 2023. URL https://www1.grc.nasa.gov/beginners-guide-to-aeronautics/induced-drag-coefficient/. Accessed: 2026-04-29.

[5] Boeing Commercial Airplanes. 777-9 Airplane Characteristics for Airport Planning. Technical report, Boeing, 2020.

[6] A. Castrichini, V. Hodigere Siddaramaiah, D. E. Calderon, J. E. Cooper, T. Wilson, and Y. Lemmens. Preliminary Investigation of Use of Flexible Folding Wing-Tips for Static and Dynamic Loads Alleviation. *The Aeronautical Journal*, 121:73–94, 2017. doi: 10.1017/aer.2016.108.

[7] A. Castrichini, V. Hodigere Siddaramaiah, D. E. Calderon, J. E. Cooper, T. Wilson, and Y. Lemmens. Nonlinear Folding Wing Tips for Gust Loads Alleviation. *Journal of Aircraft*, 53(5):1391–1399, 2016. doi: 10.2514/1.C033474.

[8] A. Castrichini, J. E. Cooper, T. Wilson, A. Carrella, and Y. Lemmens. Nonlinear Negative Stiffness Wingtip Spring Device for Gust Loads Alleviation. *Journal of Aircraft*, 54(2):627–641, 2017. doi: 10.2514/1.C033887.

[9] Ronald C. M. Cheung, Djamel Rezgui, Jonathan E. Cooper, and Thomas Wilson. Testing of a Hinged Wing-Tip Device for Gust Loads Alleviation. *Journal of Aircraft*, 55(5):2050–2067, 2018. doi: 10.2514/1.C034811.

[10] Ronald C. M. Cheung, Djamel Rezgui, Jonathan E. Cooper, and Thomas Wilson. Testing of Folding Wingtip for Gust Load Alleviation of a Flexible High Aspect Ratio Wing. *Journal of Aircraft*, 2020. doi: 10.2514/1.C035732.

[11] Xavier Carrillo Córcoles, Christoph Mertens, Andrea Sciacchitano, Bas W. van Oudheusden, Roeland De Breuker, and Jurij Sodja. Effect of Wing Stiffness and Folding Wingtip Release Threshold on Gust Loads. *Journal of Aircraft*, 60(6):1917–1936, 2023. doi: 10.2514/1.C037108.

[12] Xavier Carrillo Córcoles. An Experimental Study on Flared Folding Wingtips: Effects of Wing Stiffness, Aeroelastic Tailoring and Hinge Release Threshold on Gust Load Alleviation Performance. Master's thesis, Delft University of Technology, 2022. URL http://repository.tudelft.nl/.

[13] R. C. M. Cheung, C. Wales, D. Rezgui, J. E. Cooper, and T. Wilson. Modelling of Folding Wing-Tip Devices for Gust Loads Alleviation. In *AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*. American Institute of Aeronautics and Astronautics, 2018. doi: 10.2514/6.2018-0462.

[14] Claudio Conti, Francesco Saltari, Franco Mastroddi, Thomas Wilson, and Andrea Castrichini. Quasi-Steady Aeroelastic Analysis of the Semi-Aeroelastic Hinge Including Geometric Nonlinearities. *Journal of Aircraft*, 58(5):1168–1178, 2021. doi: 10.2514/1.C036115.

[15] Paolo Mastracci, Francesco Saltari, Franco Mastroddi, Thomas Wilson, and Andrea Castrichini. Unsteady Aeroelastic Analysis of the Semi Aeroelastic Hinge Including Local Geometric Nonlinearities. *AIAA Journal*, 60(5):3147–3165, 2022. doi: 10.2514/1.J061108.

[16] J. P. Q. Hoyng. Design and Manufacturing of a Demonstrator for the Flared Folding Wingtip Concept. Master's thesis, Delft University of Technology, 2025. URL http://repository.tudelft.nl/.

[17] Aakash Ghosh. Development of an Aeroelastic Model for a Flared Folding Wing Tip: An Exploration into the Multibody Framework of PROTEUS. Master's thesis, Delft University of Technology, 2025.

[18] Roy R. Craig and Mervyn C. C. Bampton. Coupling of Substructures for Dynamic Analyses. *AIAA Journal*, 6(7):1313–1319, 1968. doi: 10.2514/3.4741.

[19] Roy R. Craig and Andrew J. Kurdila. *Fundamentals of Structural Dynamics*. John Wiley & Sons, 2 edition, 2006. ISBN 9780471430445.

[20] Steve Doyle. pyNastran: Nastran BDF/F06/OP2/OP4 File Reader/Writer/Viewer. Software package, 2024.

[21] David Balbuena. Modeling Flexible Bodies Using the Reduced Order Flexible Solid Block. Guy on Simulink, 2021. URL https://blogs.mathworks.com/simulink/2021/03/22/modeling-flexible-bodies-using-the-reduced-order-flexible-solid-block/. Accessed: 2026-04-29.

[22] MathWorks. File Solid Block. Simscape Multibody Documentation, 2024. URL https://nl.mathworks.com/help/sm/ref/filesolid.html. Accessed: 2026-04-29.

[23] MathWorks. Revolute Joint Block. Simscape Multibody Documentation, 2024. URL https://nl.mathworks.com/help/sm/ref/revolutejoint.html. Accessed: 2026-04-29.

[24] MathWorks. External Force and Torque Block. Simscape Multibody Documentation, 2024. URL https://nl.mathworks.com/help/sm/ref/externalforceandtorque.html. Accessed: 2026-04-29.

[25] R. De Breuker, M. M. Abdalla, and Z. Gürdal. Energy-Based Aeroelastic Optimisation of a Morphing Wing. Technical report, Delft University of Technology, 2007.

[26] Noud Werter. *Aeroelastic Modelling and Design of Aeroelastically Tailored and Morphing Wings*. PhD thesis, Delft University of Technology, 2017. URL https://research.tudelft.nl/en/publications/aeroelastic-modelling-and-design-of-aeroelastically-tailored-and-/.

[27] Noud P. M. Werter, Roeland De Breuker, and Mostafa M. Abdalla. Continuous-Time State-Space Unsteady Aerodynamic Modeling for Efficient Loads Analysis. *Journal of Aircraft*, 56(3):905–916, 2017. doi: 10.2514/1.J056068.

[28] Joseba Murua, Rafael Palacios, and J. Michael R. Graham. Applications of the Unsteady Vortex-Lattice Method in Aircraft Aeroelasticity and Flight Dynamics. *Progress in Aerospace Sciences*, 55:46–72, 2012. doi: 10.1016/j.paerosci.2012.06.001.

[29] Rafael Palacios, Joseba Murua, and Robert Cook. Structural and Aerodynamic Models in Nonlinear Flight Dynamics of Very Flexible Aircraft. *AIAA Journal*, 48(11):2648–2659, 2010. doi: 10.2514/1.J050513.

[30] XFLR5. XFLR5. Software web page. URL https://www.flow5.tech/xflr5/xflr5.html. Accessed: 2026-04-29.

Part III

Additional material

5

Methodology

The research proposal in Chapter 3 defined the central question of this thesis as identifying which combinations of FFWT design parameters and hinge-release settings provide meaningful WRBM alleviation while maintaining acceptable post-gust motion and aeroelastic behaviour. This chapter therefore presents the methodology used to implement a coupled aeroelastic framework set up to simulate a flexible main wing with a rigid, flared folding wingtip undergoing large rotations. The modelling objective is a time-domain, geometrically non-linear multibody formulation that remains compatible with unsteady lifting-surface aerodynamics and gust excitation. The resulting toolchain couples a reduced-order structural model exported from Simcenter/Nastran, a Simulink/Simscape Multibody implementation of the flexible–rigid multibody system, and a Unsteady Vortex Lattice Method (UVLM) aerodynamic solver evaluated on the instantaneous configuration via online meshing.

This methodology is structured to address the research questions directly. First, the structural and aerodynamic subsystems and their coupling are formulated. Next, the resulting Simscape-UVLM framework is verified and validated against structural reference solutions, aerodynamic benchmarks and published FFWT response trends. Finally, the validated framework is used in a staged simulation campaign to assess the influence of hinge parameters and hinge-moment-based release strategies on load alleviation performance and post-gust dynamics.

Figure 5.1 provides a system-level overview of the numerical workflow and the exchanged variables. The structural backbone originates from a high-fidelity Full Order finite-element Model (FOM) in Simcenter/Nastran. A Craig-Bampton reduction is performed and the resulting reduced matrices are exported in OP4 format. These matrices are converted through a scriptable Python pipeline into MATLAB-compatible MAT files, after which the reduced model is instantiated in Simscape via a Reduced-Order Flexible Solid (ROFS) block. Aerodynamic loads are computed by a UVLM solver evaluated on an aerodynamic lattice that is reconstructed at every coupling step from the structural interface kinematics.

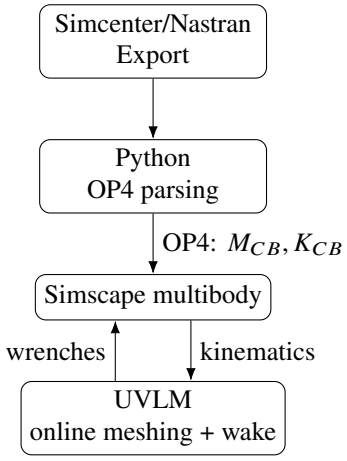

Figure 5.1: Overview of the coupled aeroelastic toolchain.

5.1. Structural Model

The structural model used in this thesis decomposes the FFWT demonstrator into two domains: a flexible main wing and a rigid wingtip. This division reflects the physical configuration of the demonstrator and provides a practical balance between structural fidelity and computational efficiency within the coupled aeroelastic framework.

The main wing, which extends from root to hinge, is modelled as a flexible structure using the FOM developed by Hoyng [18]. This representation captures the distributed mass, stiffness, and inertial properties required to accurately resolve the wing's dynamic response under aerodynamic loading.

The outboard wingtip is treated as a rigid body which is an assumption that is consistent with established practice in FFWT modelling, where the wingtip high stiffness relative to the flexible main wing justifies neglecting local deformation [19, 41, 43]. Experimental work by Córcoles [17] demonstrated that wingtip deflections are negligible when compared to main wing displacements of 15–20% semispan under identical aerodynamic loading, supporting the rigid-body assumption. Similarly, Hoyng [18] validated this modelling choice through ground vibration testing, achieving modal frequency within 10% of experimental results across both locked and free-hinge configurations using a rigid wingtip representation. The rigid-body wingtip is therefore coupled to the flexible wing through a multibody formulation in Simscape.

Based on this structural partition, the flexible-wing full-order model is reduced through Component Mode Synthesis (CMS) to obtain a computationally efficient reduced-order model (ROM) suitable for repeated time-domain aeroelastic simulation. The remainder of this section first describes the full-order finite-element model and then presents the CMS-based reduction procedure used to construct the structural ROM.

5.1.1. Full-order Finite Element Model

The FOM used as baseline for the modelling in this thesis is adopted directly from the demonstrator wing design and FEM model developed by Hoyng [18]. In Hoyng's workflow, the CAD assembly is discretized in Simcenter/Nastran FX into a finite-element representation that captures the primary load-carrying components and key attachments with sufficient fidelity for modal analysis and aeroelastic simulation. This FOM serves as the validated structural baseline for subsequent ROM construction, which is described in Subsection 5.1.2. The structural dynamics of the discretised system are governed by a system of coupled, second-order differential equations, presented in Equation 5.1.

$$\mathbf{M}\ddot{\mathbf{u}} + \mathbf{C}\dot{\mathbf{u}} + \mathbf{K}\mathbf{u} = \mathbf{F} \tag{5.1}$$

Here, $\mathbf{u} \in \mathbb{R}^n$ is the displacement vector, $\mathbf{M}$, $\mathbf{C}$, and $\mathbf{K}$ denote the mass, damping, and stiffness matrices respectively, and $\mathbf{F}$ is the external force vector. The structural discretization makes use of a selection of element types from the Nastran library listed below, each selected to efficiently capture the stiffness and inertial properties [64]:

- **0D elements**: Lumped masses that represent localized properties and discrete connections between specific degrees of freedom.
- **1D elements**: Beam-type elements that discretise slender load-carrying members via line geometry, where sectional properties are derived from CAD profiles. This approach reduces the element count for the main spar and leading/trailing edges while maintaining fidelity for modal and stress analysis.
- **2D elements**: Shell- or plate-type elements that model thin-walled components through mid-surface representation. Thickness and material properties are introduced as element attributes.

In addition to the various types of these continuum elements representing the wing structure, the FEM assembly uses connection elements to represent how elements and components are connected with each other, to enforce kinematic relations between components and define interface points for coupling. Three distinct element types serve complementary roles [64]:

- **CBUSH elements**: Used to provide generalized spring-damper connections with user-defined stiffness and damping. In the FOM, CBUSH is applied to the hinge mechanism to enable tuning between locked and free states using high and low rotational stiffness, facilitating investigation of load alleviation across the operational envelope.
- **RBE2 elements**: Rigid Body Elements that enforce rigid kinematic constraints by linking a primary node to multiple dependent nodes for specified degrees of freedom. This produces very high constraint stiffness,

making RBE2 effective for stiff load-transfer paths such as rib-to-spar connections. However, RBE2 must be deployed cautiously to avoid artificial stiffening of the assembly.

- **RBE3 elements**: Distribute forces and moments across multiple nodes via weighted averaging, without imposing rigid kinematic constraints. This flexibility makes RBE3 ideal for defining interface nodes required for aeroelastic coupling and reduced-order model construction, where loads must be transmitted without introducing artificial stiffness.

The specific materials for each structural component, including wall thicknesses, laminate properties, and element parametrisations, follow established practices for high-AR flexible wings and are inherited from Hoyng's FEM model [18]. Table 5.1 summarizes the key model parameters for each component, including the hinge stiffness formulation. Notable choices include the main spar thickness as a parametric tuning variable for global bending stiffness. The Oracover skin density set to zero, thereby reflecting negligible mass contribution and behaving as just an aerodynamic surface. Variable rotational stiffness of the hinge mechanism is used to enable investigation of both locked and free-hinge load alleviation modes.

Table 5.1: Full-order model component characteristics, following Hoyng [18].

Component	Element Type	Key Parameters	Purpose
Leading edge	1D Beam	CAD-derived section	Bending stiffness
Trailing edge	1D Beam	CAD-derived section	Bending stiffness
Ribs (front)	2D Shell	5 mm thickness	Torsional rigidity
Ribs (aft)	2D Shell	5 mm thickness	Torsional rigidity
Main spar	2D Shell	Variable thickness	Global bending, tuning
Oracover skin	2D Shell	25 μm, $\rho = 0$	Aerodynamic surface
Hinge	CBUSH	K_{rot} variable[I]	Locked/free states

[I] Rotational stiffness tuned by Hoyng [18] for locked/free hinge behaviour.

Although the FOM provides the high-fidelity structural baseline essential for capturing distributed inertia and stiffness, its large number of DOFs makes it computationally ineffective to try and implement a repeated time-domain simulation within the coupled aeroelastic framework. A ROM is therefore created through Component Mode Synthesis (CMS) that retains selected interface DOFs, while condensing interior structural dynamics onto a small set of fixed-interface normal modes.

5.1.2. Component Mode Synthesis

To enable efficient time-domain simulation within the coupled aeroelastic framework, the FOM is condensed using Craig-Bampton CMS [65]. This reduction technique divides the structural degrees of freedom into boundary DOFs that are retained for coupling to aerodynamics and rigid wingtip and interior DOFs that are condensed onto a modal basis. This yields a ROM with computational efficiency more suitable for parametric studies while preserving the dominant dynamic characteristics of the wing.

The Craig-Bampton reduction method is selected over alternatives for multiple reasons. Firstly, it preserves selected interface DOFs exactly, enabling direct coupling to external subsystems without loss of spatial fidelity at the boundaries. Secondly, it captures dynamic response beyond the first few modes through a compact set of fixed-interface normal modes, unlike static condensation methods, like Guyan reduction, that neglect higher-frequency contributions. Lastly, it is widely established in aerospace structural dynamics [66] and natively supported in Simcenter/Nastran's superelement workflow.

Craig-Bampton

Starting from the FOM equations of motion Equation 5.1, the displacement vector is partitioned into boundary denoted with subscript b and interior DOFs denoted with subscript i, as shown in Equation 5.2.

$$\mathbf{u} = \begin{bmatrix} \mathbf{u}_b \\ \mathbf{u}_i \end{bmatrix} \tag{5.2}$$

The retained boundary degrees of freedom are defined explicitly using an Analysis Set (ASET) in the Simcenter/-Nastran input deck. Then, $\mathbf{u}_b \in \mathbb{R}^{n_b}$ contains the displacements at the selected interface nodes that are defined

by ASET, and $\mathbf{u}_i \in \mathbb{R}^{n_i}$ contains all remaining interior DOFs. Craig-Bampton CMS constructs a reduction basis that combines constraint, or static, modes representing the interior deformation due to unit boundary displacements, and fixed-interface normal modes obtained with boundary DOFs clamped and truncated to n_m retained modes up to a prescribed cut-off frequency. Using these basis vectors, the Craig–Bampton transformation can be written as Equation 5.3.

$$\begin{bmatrix} \mathbf{u}_b \\ \mathbf{u}_i \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{\Psi}_c & \mathbf{\Phi}_n \end{bmatrix}}_{\mathbf{T}_{CB}} \begin{bmatrix} \mathbf{u}_b \\ \boldsymbol{\eta} \end{bmatrix} = \mathbf{T}_{CB}\, \mathbf{q}_{CB}, \qquad \mathbf{q}_{CB} = \begin{bmatrix} \mathbf{u}_b \\ \boldsymbol{\eta} \end{bmatrix} \tag{5.3}$$

Here, $\boldsymbol{\eta} \in \mathbb{R}^{n_m}$ are the modal coordinates corresponding to the retained fixed-interface modes and $\mathbf{q}_{CB}$ are the Craig-Bampton generalised coordinates, while I is the identity matrix. Subsequently substituting Equation 5.3 into Equation 5.1, the full-order equations are not satisfied exactly in general, leaving a residual $\mathbf{r}(\mathbf{q}_{CB})$ defined in Equation 5.4.

$$\mathbf{r}(\mathbf{q}_{CB}) = \mathbf{M}\mathbf{T}_{CB}\ddot{\mathbf{q}}_{CB} + \mathbf{C}\mathbf{T}_{CB}\dot{\mathbf{q}}_{CB} + \mathbf{K}\mathbf{T}_{CB}\mathbf{q}_{CB} - \mathbf{F} \tag{5.4}$$

Therefore, a Galerkin projection is used, requiring the residual to be orthogonal to the trial subspace spanned by $\mathbf{T}_{CB}$, which results in the reduced operators presented in Equation 5.5

$$\mathbf{M}_{CB} = \mathbf{T}_{CB}^T\mathbf{M}\mathbf{T}_{CB}, \qquad \mathbf{C}_{CB} = \mathbf{T}_{CB}^T\mathbf{C}\mathbf{T}_{CB}, \qquad \mathbf{K}_{CB} = \mathbf{T}_{CB}^T\mathbf{K}\mathbf{T}_{CB}, \qquad \mathbf{Q}_{CB} = \mathbf{T}_{CB}^T\mathbf{F} \tag{5.5}$$

This corresponds to the standard Galerkin framework commonly used in finite elements and projection-based model reduction [67, 68]. After applying this projection, the reduced system as presented in Equation 5.6 is obtained.

$$\mathbf{M}_{CB}\,\ddot{\mathbf{q}}_{CB} + \mathbf{C}_{CB}\,\dot{\mathbf{q}}_{CB} + \mathbf{K}_{CB}\,\mathbf{q}_{CB} = \mathbf{Q}_{CB} \tag{5.6}$$

A key property of this formulation is that the boundary DOFs, $\mathbf{u}_b$, remain physical coordinates in the reduced model, which enables direct application of interface constraints and aerodynamic loads at the selected interface nodes.

Interface Nodes

As mentioned, the ASET collects all six DOFs at each selected interface node and constitutes the boundary set (B-set) for the Craig-Bampton reduction and all remaining structural DOFs belong to the interior set. This ensures that the constraint modes are constructed exactly for the physically relevant coupling interfaces. For the application in this thesis, the nodes are divided in retained boundary nodes and nodes meant for intermediate coupling. The boundary nodes consist of the root and the hinge interface nodes defined as `RBE2` spiders, as presented in Figure 5.2, providing a rigid kinematic link between the master interface grid and the surrounding structural grids.

The interface nodes are added for the extraction of physical displacements and application of resulting aerodynamic loads outputted by the aerodynamic solver. The interface nodes are realised using `RBE3` spider elements to ensure physically meaningful load transfer and robust coupling to the structure, while avoiding artificial stiffening of the flexible wing when adding intermediate coupling points for load application and motion extraction. The number of these nodes is selected to provide sufficient spatial resolution for the transfer of three-dimensional aerodynamic loads while keeping the boundary DOF count small enough to preserve ROM compactness. As a precedent, Ghosh [19] models the flexible main wing using a relatively coarse discretisation that is nevertheless reported to be adequate for FFWT aeroelastic behaviour: the main wing is represented with ten structural beam elements. Therefore, nine aerodynamic interface nodes are distributed spanwise along the flexible main wing in addition to the root and hinge interface nodes, resulting in $n = 11$ boundary nodes and $n_b = 66$ retained boundary DOFs.

The selection of retained interface DOFs and fixed-interface modes is justified in the structural verification and validation section in Subsection 5.5.1. The retained ASET/interface nodes are summarised in Table 5.2.

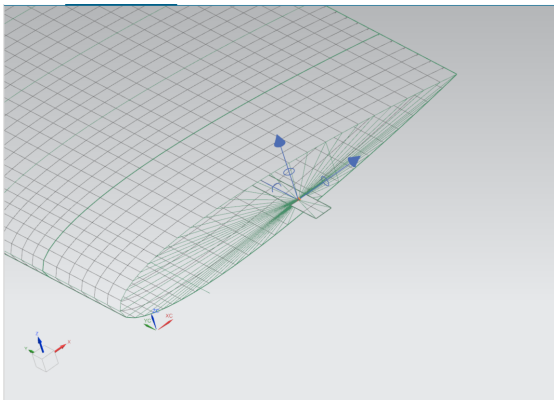

a) Root interface spider (RBE2).

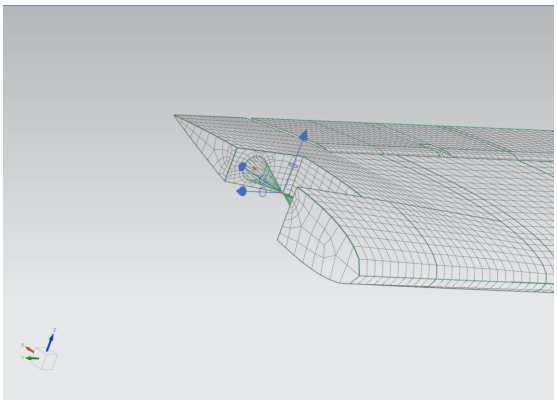

b) Hinge interface spider (RBE2).

Figure 5.2: RBE2 spider elements at main wing root and hinge location.

Table 5.2: Retained interface nodes (ASET) used for Craig-Bampton reduction.

Interface frame	Node ID	Role	Spider	x[mm]	y[mm]	z[mm]
F_1	203692	Interface node 1	RBE3	115.00	869.50	0.00
F_2	203680	Interface node 2	RBE3	115.00	759.50	0.00
F_3	203662	Interface node 3	RBE3	115.00	649.50	0.00
F_4	203638	Interface node 4	RBE3	115.00	530.50	0.00
F_5	203620	Interface node 5	RBE3	115.00	420.50	0.00
F_6	203608	Interface node 6	RBE3	115.00	310.50	0.00
F_7	203600	Interface node 7	RBE3	115.00	200.50	0.00
F_8	132269	Interface node 8	RBE3	115.00	93.00	0.00
F_9	204553	Root	RBE2	115.00	31.00	0.00
F_{10}	204437	Interface node 9	RBE3	115.00	977.00	0.00
F_{11}	204562	Hinge	RBE2	115.60	1107.80	0.00

5.1.3. Simcenter ROM Export

In the sense of the envisioned coupled aeroelastic simulation environment in MATLAB/Simscape, the Craig-Bampton ROM generated in Simcenter/Nastran should be exported in a compatible format. Simcenter provides several mechanisms for exporting reduced structural models, that are each optimized for different downstream applications:

- **Superelement matrices**: Exports the Craig-Bampton matrices ($\mathbf{M}_{CB}$, $\mathbf{K}_{CB}$, $\mathbf{T}_{CB}$) in native Nastran format, preserving the full block structure including boundary-modal coupling terms [69, 70].
- **Flexible body**: Modal-neutral file for direct import into commercial multibody dynamics solvers. Stores mass-normalized modal coordinates with unit modal mass matrix. [71, 72].
- **DMIG bulk data**: Direct Matrix Input at Grid (DMIG) points, enables manual assembly of reduced matrices within subsequent Nastran models. Primarily used for multi-level substructuring [70].
- **External superelement**: Binary or formatted files for hierarchical assembly in large-scale models, like aircraft fuselage and wing sections combined. [69].

For the application in this thesis, the export format must satisfy two main requirements that follow directly from the Craig-Bampton reduction described in Subsection 5.1.2. First, the reduced model must retain the physical boundary DOFs and their associated mass and stiffness contributions, since these coordinates are used directly for aeroelastic coupling. Second, the exported format has to be compatible with a custom state-space formulation, so provide explicit access to the mass and stiffness matrices in order to construct the state-space representation in Simscape.

DMIG and external superelement formats are designed for multi-level sub-structuring within Nastran and do not provide the direct matrix access required for coupling to external solvers [70]. The flexible body format is optimized

for monolithic import into commercial multibody solvers, like for instance Simcenter Motion or Adams, and does not expose the reduced mass and stiffness matrices required for custom state-space formulation. The superelement matrix export is actually a perfect fit to the requirements: The physical mass distribution is preserved and the mass and stiffness matrices are directly accessible, and is therefore selected as the export method for this thesis.

Matrix Extraction

The Nastran superelement data can be exported from Simcenter in several formats. Binary `.op2` files are compact and efficient, but are platform-dependent and typically require dedicated readers. ASCII punch `.pch` files can be read as is, but become inefficient for large matrices because of their verbose formatting. The `.op4` format, or formatted Output4, provides a structured ASCII representation specifically intended for matrix exchange between solvers. Although this format also requires a dedicated reader, it can be accessed readily through open-source Python libraries that support native `.op4` parsing [73].

For this reason, the reduced model is exported in `.op4` format. In Simcenter, this is configured by selecting the SOL103 superelement solution and enabling the `MATOP4` option. Simcenter then modifies the `EXTSEOUT` case-control command in the Nastran input deck accordingly [64], instructing Nastran to write the reduced Craig-Bampton matrices in Output4 format. In the resulting file, the primary matrices of interest are the reduced stiffness matrix $\mathbf{K}_{\mathrm{CB}}$ (`KXX`) and the reduced mass matrix $\mathbf{M}_{\mathrm{CB}}$ (`MXX`). Nastran also exports additional matrices, including the constraint transformation matrix `BXX1` and the modal eigenvector transformation $\mathbf{T}_{\mathrm{CB}}$ `VX`, which contains the retained boundary DOFs and fixed-interface mode shapes. These are not used directly in the present aeroelastic coupling, but are retained in the export for traceability and reference.

The `.op4` file is parsed in Python using the `pyNastran` library [73]. The extraction script reads the required matrices, retrieves `MXX` and `KXX`, and writes them to a MATLAB-compatible `.mat` file using `scipy.io.savemat` for subsequent import into the Simscape multibody model. The complete extraction scripts are provided in Appendix A. This two-step workflow from Nastran to Python and from Python to MATLAB avoids MATLAB's limited native `.op4` support while enabling automated validation checks before integration.

For reproducibility, it is noted that the matrices are stored in the `.op4` file by name and with a fixed DOF ordering. In the present Craig-Bampton export, the generalized coordinates follow the standard convention of retained boundary DOFs first, followed by the retained modal coordinates.

5.2. Simscape Multibody

The coupled aeroelastic model is assembled in MATLAB's Simulink as a block-diagram representation of the overall simulation workflow, including the aerodynamic module and the data exchange logic. The structural dynamics, however, are realised using Simscape Multibody, which represents mechanical systems through an energy-based physical network formulation. This is done by using the blocks that are part of the Simscape Multibody library. During simulation, Simulink coordinates the global execution and time stepping, while the Simscape engine formulates and solves the multibody dynamics as a constrained system subject to joint and kinematic constraints. Signals exchanged between the algorithmic Simulink domain and the physical Simscape domain are passed via physical-signal interfaces, which in turn enables aerodynamic forces and moments to be applied as wrenches at the defined interface nodes and the resulting structural kinematics to be extracted for coupling.

The Craig-Bampton ROM exported from Simcenter/Nastran FX is imported into MATLAB and instantiated inside a Simscape Multibody model. As mentioned, only the flexible main wing is represented using the ROM in combination with the ROFS block [74]. The relevant settings of this block are added in Table 5.3.

A custom unit system (mm/kg/s) is enforced to match the unit convention used in the exported reduced matrices from Nastran. This avoids stiffness/inertia scaling errors that would otherwise manifest as incorrect eigenfrequencies and static compliance. The number of interface frames is set to ten, which is consistent with the ROM boundary set ($n = 11$ nodes, $n_b = 66$ DOFs), and defines the locations where interface wrenches are applied and structural kinematics are extracted for the aeroelastic coupling. The coordinates of these frames are manually added in the ROFS block as well. Interface frame orientations are kept as identity, so no local rotations, to maintain a consistent axis convention between the FE model, Simscape, and the aerodynamic solver. The stiffness and mass matrix names (K and M) reference the imported Craig-Bampton operators in the MATLAB workspace, maintaining a traceable export-import pipeline. The ROM is evaluated about linearised mean axes conditions, which is consistent with the underlying linear elastic assumptions of the Craig-Bampton formulation. The imported geometry from the STEP file is used for visualisation only, facilitating verification of frame placement and deformation directions.

Table 5.3: ROFS block configuration.

Setting group	Value
Unit system	Custom
Length / mass / time units	mm / kg / s
Number of interface frames	11 (root, hinge, 9 aerodynamic coupling frames)
Interface frame orientations	Identity (no local rotation applied)
Reduced-order stiffness matrix name	`K`
Reduced-order mass matrix name	`M`
Reference conditions	Linearized mean axes
Additional reduction inside ROFS	None
Visualization geometry	Partitioned geometry; visual properties from STEP file

The outboard wingtip is modelled as a rigid body using a File Solid block [75], which is defined as a solid with properties derived from CAD file. It imports the wingtip geometry for visualisation and defines the rigid body's reference frame. The inertial properties of the rigid wingtip like the mass, centre of mass and inertia tensor are assigned in the File Solid block based on the wingtip design data, using CAD-derived values to ensure dynamic consistency with the demonstrator configuration [18]. The standard File Solid reference frame is not coincident with the physical hinge line and therefore, two additional frames are defined on the wingtip body via the File Solid interface:

- **Hinge frame W1:** located on the hinge axis, which is subsequently used to connect the wingtip to the flexible main wing imported through the ROFS block and to measure hinge kinematics and reaction moment.
- **Wingtip aerodynamic interface frame W2:** used as the aerodynamic coupling point for the wingtip lattice and for the application of condensed aerodynamic wrenches consisting of force and moment.

These frames are used consistently in both the Simscape and the aerodynamic model. Transform Sensor blocks are used to measure the pose and rates of the interface frames relative to the world frame, providing $\{\mathbf{r}_k, \mathbf{R}_k, \dot{\mathbf{r}}_k, \boldsymbol{\omega}_k\}$ directly for the coupling with the UVLM solver.

The relevant parameters of the File Solid block and the additional wingtip frames are summarised in Table 5.4.

Table 5.4: File Solid block configuration.

Setting	Value
File Name	`[Tip_Clean.stp]`
Unit Type	Custom
Unit	mm
Type	Calculate from Geometry
Based On	Custom Mass (*mass specified by user*)
Frames defined	Frame1 = W1 (hinge frame), Frame2 = W2 (aero interface frame)

The unit settings (custom, mm) ensure geometric consistency with the ROM and the overall model unit system, while the "Calculate from Geometry" option preserves the CAD-based visual representation. The "Custom Mass" option enforces the experimentally relevant inertial properties independent of CAD material assumptions.

The connection between the flexible main wing and the rigid wingtip is realised using a Revolute Joint block [76]. The joint is defined between the main-wing hinge interface frame ROFS block and the wingtip hinge frame File Solid block, and constrains the relative motion to a single rotational DOF about the prescribed hinge axis. The hinge axis is aligned with the physical hinge line of the demonstrator through Rigid Transform blocks on either side of the Revolute Joint. Not only does this apply the flare angle, but it also aligns the revolute axis of the block, which is the Z-axis by default, by applying a rotation sequence about base axes Y and Z of 90° and 15° respectively. The Rigid Transform block on the main wing side of the hinge introduces this sequence, and the Rigid Transform block on the wingtip side of the hinge counters it. The orientation of the wingtip and correct working of the hinge is subsequently confirmed through a straight-forward gravity check, starting the simulation and setting the Revolute Joint to 'disengaged'.

The output of the Revolute Joint block provides the fold angle $\theta(t)$ and fold rate $\dot{\theta}(t)$, which are used for monitoring and validation of the FFWT motion and also as an additional consistency check against the interface-frame angular velocity $\boldsymbol{\omega}$ obtained from the Transform Sensors. Depending on the intended hinge representation, joint internal mechanics can include rotational stiffness and damping to represent hinge friction or restoring elements. In the baseline configuration used for aeroelastic simulations in this thesis, the hinge includes stiffness and damping, with relevant parameters listed in Table 5.5. Joint limits are applied as required to prevent non-physical fold angles outside the demonstrator motion envelope.

Table 5.5: Revolute Joint block configuration.

Setting	Value
Equilibrium Position	0 deg (Compile-time)
Spring Stiffness	100 N·m/rad (Compile-time)
Damping Coefficient	3 N·m·s/rad (Compile-time)
Specify Lower Limit	−90°
Specify Upper Limit	90°
Mode Configuration	*[default]*
Direction	Follower on Base
Resolution Frame	Base

The actuation and motion input settings are disabled ("Automatically computed") such that the hinge motion is a result of the coupled dynamics and applied loads rather than being prescribed. The direction and resolution-frame options are chosen to ensure follower behaviour and an unambiguous sign convention for $\theta(t)$ under large rotations. Constraint/total force and torque outputs are disabled in the baseline model since hinge loads are reconstructed from the measured reaction moments at the hinge frames, avoiding redundant signals unless required for debugging.

The revolute hinge introduces an orientation discontinuity between the flexible main wing and rigid wingtip. By defining consistent interface frames on both sides of the joint and extracting their kinematics directly from Simscape, the aerodynamic lattice can be updated in the instantaneous configuration and the returned aerodynamic wrenches can be applied as follower loads, ensuring that both the load direction and moment arms evolve correctly with the FFWT fold angle.

5.2.1. MATLAB Import

As described in Subsection 5.1.3, the Craig-Bampton superelement matrices are exported from Simcenter/Nastran in .op4 format and subsequently imported into MATLAB through a scriptable Python-based pipeline. After file-parsing, the extracted matrices are written to a MATLAB-compatible .mat file using `scipy.io.savemat` [77]. This two-step workflow avoids MATLAB's limited native `.op4` support while preserving full traceability of matrix naming and DOF ordering.

MATLAB then loads the resulting `.mat` file and the reduced Craig-Bampton mass and stiffness matrices $\mathbf{M}_{CB}$ and $\mathbf{K}_{CB}$ are loaded into the ROFS block. The ROFS interface is defined using $n = 10$ boundary nodes, corresponding to the order the DOFs are exported in the ROM. This leads to:

- Nine spanwise interface frames: $F_1, \ldots, F_8$ and F_{10}
- Root interface frame: F_9
- Hinge interface frame: F_{11}

Each boundary node contributes six degrees of freedom in the required order presented in Equation 5.7.

$$\mathbf{U}_i = [T_{x,i},\ T_{y,i},\ T_{z,i},\ R_{x,i},\ R_{y,i},\ R_{z,i}], \qquad i = 1, \ldots, 10 \tag{5.7}$$

In this way, the retained boundary coordinate vector satisfies $\mathbf{u}_b \in \mathbb{R}^{6n} = \mathbb{R}^{66}$. This DOF requirement and ordering are enforced by construction, since the ROFS block requires each boundary node to contribute six DOFs in this exact order. Subsection 5.1.2 defines the Craig-Bampton ROM and its operators. For the MATLAB/Simscape implementation, the reduced coordinate vector is assembled as Equation 5.8

$$\mathbf{q}_{CB} = \begin{bmatrix} \mathbf{u}_b \\ \boldsymbol{\eta} \end{bmatrix}, \qquad \mathbf{u}_b \in \mathbb{R}^{66} \tag{5.8}$$

In this equation $\mathbf{u}_b$ contains the physical translations and rotations of the $n = 11$ retained interface frames, and $\boldsymbol{\eta} \in \mathbb{R}^m$ contains the retained modal coordinates.

The ROM dynamics imported into Simscape can be written in the standard second-order form as Equation 5.9.

$$\mathbf{M}_{CB}\,\ddot{\mathbf{q}}_{CB} + \mathbf{C}_{CB}\,\dot{\mathbf{q}}_{CB} + \mathbf{K}_{CB}\,\mathbf{q}_{CB} = \mathbf{Q}_{CB} \tag{5.9}$$

No damping matrix is exported as part of the Simcenter/Nastran Craig-Bampton reduction. In the present framework, structural damping of the flexible main wing is therefore introduced directly in Simscape through a uniform damping setting applied to the reduced flexible body. A constant damping ratio of $\zeta = 0.03$ is used throughout the simulations. This choice provides a simple and numerically robust representation of structural dissipation without introducing an additional fitted damping model at ROM level. Since the emphasis of this thesis is on comparative aeroelastic response and hinge-release behaviour, the damping treatment is intended as a practical reduced-order approximation.

The generalized load vector $\mathbf{Q}_{CB}$ collects the loads associated with the Craig-Bampton coordinates. In the present multibody implementation, aerodynamic interface wrenches applied at the coupling frames act directly on the retained boundary coordinates of the ROM. The interaction between the flexible main wing and the rigid wingtip is transmitted through the revolute joint and associated multibody constraints at the hinge frame, so that hinge loads enter the ROM implicitly through the coupled Simscape formulation rather than through an explicitly assembled external load vector.

The fidelity of the imported ROM in terms of modal properties and static compliance is assessed in Subsection 5.5.1.

Load Mapping and Work Conjugacy

Aerodynamic loads computed by the UVLM are transferred to the structural subsystem through the aerodynamic coupling frames of the main wing and for the rigid wingtip, through the wingtip interface frames. At each coupling frame F_i the aerodynamic solver provides an equivalent wrench as presented in Equation 5.10.

$$\mathbf{w}_i = \begin{bmatrix} \mathbf{F}_i \\ \mathbf{M}_i \end{bmatrix}, \qquad i \in \mathcal{I} \tag{5.10}$$

Then, $\mathbf{F}_i$ is the force applied at the frame origin and $\mathbf{M}_i$ is the moment about that origin. In Simscape Multibody these wrenches are applied using the External Force and Torque block [78], which applies forces at the origin of the connected frame and torques about the corresponding axes. For the flexible main wing, the coupling frames coincide with retained boundary frames of the ROFS block and therefore act directly on the physical boundary coordinates $\mathbf{u}_b$ of the Craig-Bampton model.

To avoid spurious energy gain or loss at the coupling interface, the exchanged loads and kinematics are formulated to satisfy work conjugacy. The instantaneous mechanical power transmitted through the interface is defined as Equation 5.11.

$$P_{\text{int}}(t) = \sum_{i \in \mathcal{I}} \left(\mathbf{F}_i^{\mathsf{T}} \mathbf{v}_i + \mathbf{M}_i^{\mathsf{T}} \boldsymbol{\omega}_i \right) \tag{5.11}$$

Here, $(\mathbf{v}_i, \boldsymbol{\omega}_i)$ are the translational and angular velocities of coupling frame F_i. For the Craig-Bampton structural coordinates in Equation 5.8, the corresponding generalised power is $\mathbf{Q}_{CB}^{\mathsf{T}} \dot{\mathbf{q}}_{CB}$. Work conjugacy is satisfied when the mapping between interface wrenches and generalised forces is defined such that Equation 5.12 holds.

$$\sum_{i \in \mathcal{I}} \left(\mathbf{F}_i^{\mathsf{T}} \mathbf{v}_i + \mathbf{M}_i^{\mathsf{T}} \boldsymbol{\omega}_i \right) = \mathbf{Q}_{CB}^{\mathsf{T}} \dot{\mathbf{q}}_{CB} \tag{5.12}$$

In practice, this condition is achieved by applying the aerodynamic wrenches at the same physical frames that define the retained boundary DOFs of the ROFS block, so that the boundary generalised forces correspond directly to forces and moments about the boundary frames in the same axis convention and units.

The External Force and Torque block applies wrenches in the selected resolution frame. In this thesis, aerodynamic forces and moments are applied in the world frame to ensure consistent follower-load behaviour under large FFWT rotations. When the aerodynamic solver returns loads expressed in a local interface frame, they are rotated to the world frame using the instantaneous rotation matrix of that frame in Equation 5.13.

$$\mathbf{F}_{i,\mathrm{w}} = \mathbf{R}_i \, \mathbf{F}_i, \qquad \mathbf{M}_{i,\mathrm{w}} = \mathbf{R}_i \, \mathbf{M}_i \tag{5.13}$$

The load direction and moment arm then evolve correctly with the fold angle. This is particularly important at the hinge, where the orientation discontinuity between the flexible main wing and rigid wingtip must be represented as a true large-rotation effect rather than as a small perturbation.

With structural kinematics and load application points defined, the next section introduces the aerodynamic model used to compute unsteady loads on the instantaneous configuration.

5.3. Aerodynamic Model

As discussed in Subsection 2.3.5, large rotations and non-planar wake development has caused FFWT research to shift from frequency-domain methods like the DLM toward time-domain vortex methods capable of resolving these phenomena. While DLM can provide trend-level predictions under small perturbations, wind-tunnel results reported in FFWT literature [42] indicate that it over-predicts gust load alleviation since it cannot represent time-resolved wake evolution and the coupling between elastic wing motion and wingtip pitch dynamics. For this reason, the aerodynamic model in this thesis is based on the UVLM, which operates in the time domain and captures the unsteady lift build-up through the evolution of bound and shed vorticity. UVLM operates in the time domain, accommodates non-planar wakes, and updates panel normals with structural motion, providing the necessary balance between fidelity and computational efficiency for coupled aeroelastic analysis of the FFWT concept.

To be able to capture unsteady lift build-up and wake evolution in the time domain at moderate computational cost, this thesis employs a vortex-based lifting-surface method in the UVLM family. The UVLM discretises the lifting surface into panels carrying bound vorticity and enforces a flow tangency boundary condition at panel control points, while vorticity shed from the trailing edge forms a wake that is advected downstream to satisfy Kelvin's theorem [79].

The UVLM implementation used in this thesis follows the potential-flow formulation used in the PROTEUS framework that was developed at Delft University of Technology. The framework builds on earlier morphing-wing aeroelastic work by De Breuker et al.[80] and subsequently extended by Werter et al.[81], who integrated a dynamic solver and the continuous-time (CT) UVLM aerodynamic formulation. In the original CT formulation, the wake transport is written in terms of the jump in velocity potential across the wake surface $\Delta\phi$ and the resulting aerodynamic model can be cast in a compact state-space form suitable for aeroelastic simulations. [82]

In many applications, these CT state-space UVLM models are linearised about a static equilibrium and used for small perturbation responses, like flutter or small-amplitude gust response. The key implementation choice in this thesis is to evaluate the UVLM at the instantaneous configuration through online meshing. Essentially, at every time step the aerodynamic lattice geometry consisting of panel corner points, control points and panel normals, is reconstructed from the structural kinematics and the aerodynamic influence system is reassembled. This ensures that large FFWT rotations enter the aerodynamics through the updated panel normals and trailing-edge kinematics, rather than through perturbations about a fixed reference configuration. This approach is constant with standard non-linear time-domain UVLM practice in coupled aeroelasticity [83].

The discretisation strategy, wake modelling approach, aerodynamic load extraction, and applicability limits are detailed in the following subsections.

5.3.1. Discretisation and UVLM Operator

The lifting surfaces are discretised into a structured lattice of $N_c \times N_s$ quadrilateral panels in the chordwise and spanwise directions. In the vortex-ring formulation adopted in PROTEUS, each panel is represented by a closed vortex ring consisting of four straight vortex filaments [19]. The unknown bound circulation strengths are collected

in the vector $\mathbf{\Gamma}_b \in \mathbb{R}^{N_b}$, where N_b is the total number of bound panels over all lifting surfaces. A collocation point is defined for each panel at three-quarter chord and mid-span location, and the instantaneous panel normal vector $\mathbf{n}_j^n$ is computed from the deformed panel geometry. A consistent vertex ordering is employed so that vortex filament directions and panel normals follow a right-hand convention.

To improve resolution in regions with strong spanwise gradients, the main wing uses a one-sided sine-based clustering towards the hinge, while the folding wingtip uses a cosine spacing clustered towards both the hinge and the outboard tip. Chordwise spacing is kept uniform and is selected to balance leading-edge resolution, trailing-edge wake shedding and computational cost, noting that the UVLM runtime scales strongly with the total bound-panel count.

The flared hinge geometry is represented in the aerodynamic mesh through the local chord-direction definition rather than through an explicit non-rectangular cut-out. This preserves the two-surface UVLM representation of the main wing and folding wingtip while introducing the planform slant associated with the flared hinge. The prescribed flare angle is denoted by Λ and is applied as a rotation of the local chord-direction vector in the aerodynamic x–y plane. For each spanwise mesh station, the local chord-direction vector is written as Equation 5.14.

$$\mathbf{c}_{\text{def}}(\xi) = \begin{bmatrix} \cos\phi(\xi) \\ \sin\phi(\xi) \\ 0 \end{bmatrix}, \tag{5.14}$$

In this relation, ξ is the local non-dimensional spanwise coordinate of the corresponding lifting surface. On the main wing, the chord direction is kept aligned with the global freestream direction over most of the span and is then smoothly ramped toward the prescribed hinge flare near the wingtip. Using the ramp-start location $\xi_0 = 0.75$, this is implemented as Equation 5.15.

$$\phi_{\text{main}}(\xi) = \begin{cases} 0, & \xi < \xi_0, \\ \Lambda \dfrac{\xi - \xi_0}{1 - \xi_0}, & \xi \geq \xi_0. \end{cases} \tag{5.15}$$

On the folding wingtip, the local chord direction starts with the same flare angle at the hinge and relaxes back to the straight chord direction at the outboard tip as defined in Equation 5.16.

$$\phi_{\text{tip}}(\xi) = \Lambda(1 - \xi). \tag{5.16}$$

The resulting reference aerodynamic mesh is illustrated in Figure 5.3.

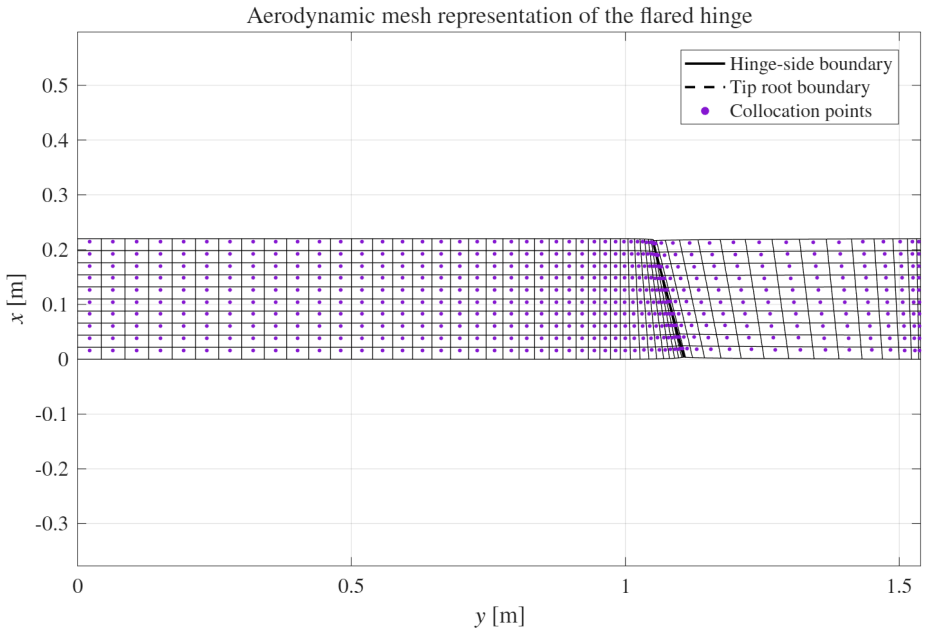

Figure 5.3: Top-view aerodynamic mesh representation of the flared hinge. The main wing and folding wingtip are retained as two structured lifting surfaces, while the flared hinge effect is introduced through a spanwise variation of the local chord-direction vector.

This treatment should be interpreted as an approximate flared-cut representation. It captures the slanted hinge-line effect in the aerodynamic lattice while retaining the existing main-wing and wingtip surface definitions. It does not explicitly construct a fully non-rectangular cut-out geometry with overlap removal at the hinge interface.

The discretisation parameters used for the UVLM implementation are reported in Table 5.6.

Table 5.6: UVLM discretisation settings used in the coupled simulations.

Parameter	Symbol	Value
Number of lifting surfaces	N_{surf}	2
Number of structural frames	N_{frame}	12
Main-wing chordwise panels	$N_{c,\text{main}}$	10
Main-wing spanwise panels	$N_{s,\text{main}}$	39
Wingtip chordwise panels	$N_{c,\text{tip}}$	10
Wingtip spanwise panels	$N_{s,\text{tip}}$	15
Total bound panels	N_b	540
Chord length	c	0.220 m
Wake spacing	Δx_w	2.2×10^{-2} m
Reduced wake spacing	$\Delta x_w / c$	0.1
Target wake length	L_w / c	20
Wake columns	N_w	201
Effective UVLM time step	Δt_{UVLM}	2.2×10^{-3} s
Main-wing flare-ramp start	ξ_0	0.75

The UVLM enforces that the normal component of the flow velocity at each collocation point vanishes. For a collocation point j the boundary condition reads as Equation 5.17.

$$\mathbf{n}_j^n \cdot \left(\mathbf{V}_\infty + \mathbf{V}_{\text{motion},j}^n + \mathbf{V}_{\text{ind},j}^n (\boldsymbol{\Gamma}_b^n, \boldsymbol{\Gamma}_w^n) \right) = 0, \qquad j = 1, \ldots, N_b \tag{5.17}$$

Here, $\mathbf{V}_\infty$ is the freestream velocity, $\mathbf{V}_{\text{motion},j}^n$ is the kinematic velocity of the moving surface at the collocation point, and $\mathbf{V}_{\text{ind},j}^n$ is the induced velocity from both bound and wake vorticity via Biot-Savart superposition. Collecting all collocation conditions yields a linear system for the bound circulations in Equation 5.18.

$$\mathbf{A}\left(\mathbf{X}^{n_r}, \mathbf{X}_w^{n_r}\right) \boldsymbol{\Gamma}_b^n = -\mathbf{b}^n \tag{5.18}$$

The index n_r denotes the most recent time step at which the geometry-dependent UVLM operator was assembled. If re-linearisation is disabled, $n_r = 0$ throughout the simulation and the operator remains based on the reference aerodynamic mesh. If re-linearisation is performed every time step, $n_r = n$ and the operator follows the instantaneous aerodynamic geometry. Intermediate options are also implemented, where the operator is updated periodically or only when prescribed deformation thresholds are exceeded.

For the simulations presented in this thesis, re-linearisation is disabled unless explicitly stated otherwise. Therefore, the UVLM operator remains fixed about the reference aerodynamic mesh, while the instantaneous aeroelastic coupling enters through the updated panel kinematics, panel normals, wake-shedding geometry and aerodynamic load mapping.

5.3.2. Wake Modelling

As mentioned, in the PROTEUS CT-UVLM formulation the wake is represented as a surface connected to the trailing edge and wake transport is expressed in terms of the potential-jump variable $\Delta\phi$ across the wake. Its transport is then governed by the advection equation presented in Equation 5.19.

$$\frac{\partial \Delta\phi}{\partial t} + \mathbf{V}_w \cdot \nabla \Delta\phi = 0 \tag{5.19}$$

In this equation, $\mathbf{V}_w$ is the wake convection velocity. Under small-disturbance assumptions in the linearised setting, the wake is convected with the free-stream velocity, so $\mathbf{V}_w \approx \mathbf{V}_\infty$, which yields the simplified relation in Equation 5.20.

$$\frac{\partial \Delta\phi}{\partial t} + V_\infty \frac{\partial \Delta\phi}{\partial x} = 0 \tag{5.20}$$

This corresponds to a prescribed-wake model, where previously shed wake vorticity is transported with a prescribed velocity rather than being convected by its own induced velocity field, which would be a free-wake roll-up. [79] In the CT formulation, the wake-advection equation is discretised in space while retaining a continuous time derivative, resulting in a compact state-space representation, where the spatial discretisation of the wake transport employs a discontinuous Galerkin method [82].

The Kutta condition then enforces smooth flow departure at the trailing edge by requiring continuity of the potential jump between the wing-body and wake along the same streamline, as shown in Equation 5.21.

$$\Delta\phi_b\big|_{TE} = \Delta\phi_w\big|_{TE}, \tag{5.21}$$

This relation couples the bound circulation on the lifting surface to the wake strength shed at the trailing edge.

In the coupled system in this thesis however, wake geometry is advanced with a prescribed convection velocity, taken as $\mathbf{V}_c = \mathbf{V}_\infty$, while new wake elements are shed from the instantaneous trailing-edge geometry at each time step. Consequently, even though the wake is prescribed, so no induced-velocity roll-up, it becomes non-planar when the trailing edge undergoes large FFWT rotations and elastic deformation, because the newly shed wake row inherits the instantaneous trailing-edge orientation.

5.3.3. Aerodynamic Loads

Once the bound circulation distribution $\mathbf{\Gamma}_b^n$ has been obtained from Equation 5.18 and the wake has been updated, aerodynamic forces and moments are evaluated on the instantaneous lattice using an unsteady Kutta-Joukowski-based formulation. The resulting panel-wise loads are integrated to obtain the net force and moment on each aerodynamic component, after which they are condensed into equivalent wrenches at the aerodynamic coupling frames for transfer to the structural model.

The aero-to-structural mapping is performed using a transfer matrix $\mathbf{T}_{AS}$ that maps the aerodynamic load vector $\mathbf{f}_a^n$ containing panel-wise forces, or equivalent sectional loads, to the structural nodal load vector $\mathbf{f}_s^n$ as Equation 5.22.

$$\mathbf{f}_s^n = \mathbf{T}_{AS}\,\mathbf{f}_a^n \tag{5.22}$$

The transfer is constructed to conserve the resultant force and moment at the coupling interface, ensuring that the structural solver receives an energetically consistent representation of the aerodynamic loading. For the FFWT configuration, this wrench condensation is performed in the instantaneous local frame of each coupling interface so that the hinge-induced orientation discontinuity is captured consistently as the wingtip rotates.

5.3.4. Model Applicability

The aerodynamic model employed in this thesis is a UVLM-type, potential-flow lifting-surface method. As such, its applicability is primarily to low-speed, attached-flow conditions in which unsteady lift is dominated by the evolution of bound and shed vorticity rather than viscous separation effects. The UVLM formulation assumes inviscid, incompressible, irrotational flow and therefore does not represent viscous phenomena such as boundary-layer growth, skin-friction drag, separation, dynamic stall, or shock effects. Consequently, predictive fidelity degrades at high angles of attack and in manoeuvres where separation occurs.

In the present framework, the wake is advanced using a prescribed convection velocity, $\mathbf{V}_c \approx \mathbf{V}_\infty$. This neglects free-wake roll-up and wake deformation driven by induced velocities and strong wake-body interaction. These effects may become relevant in highly unsteady motion or when the wake passes close to the lifting surfaces, and are not captured by the prescribed-wake model.

A key limitation of linearised UVLM implementations is that aerodynamic operators and load directions are evaluated about a fixed reference configuration, which restricts validity when large rotations change panel normals and trailing-edge kinematics. This thesis extends the applicability of the UVLM model with respect to geometric non-linearity by evaluating the aerodynamic influence system on the instantaneous configuration through online meshing. The lattice

geometry, panel normals and trailing-edge position are updated each coupling time step, so that large FFWT rotations enter the aerodynamics directly via the instantaneous surface orientation. In combination with wake shedding from the instantaneous trailing edge, this yields a prescribed but potentially non-planar wake that reflects large wingtip motion.

As with any UVLM implementation, accuracy depends on discretisation N_c, N_s, wake length and time step Δt. The wake is truncated to a finite extent, and load predictions are therefore sensitive to wake resolution and convection assumptions. These parameters are selected in this thesis to balance numerical stability, computational cost and the fidelity required for coupled aeroelastic time-domain simulations.

5.4. Aeroelastic Coupling

After defining the structural and aerodynamic model, the coupled aeroelastic framework is implemented as a time-marching scheme in which the structural solver advances the structural state and provides the instantaneous kinematics required by the aerodynamic solver. The aerodynamic solver then evaluates unsteady aerodynamic loads on the updated lattice and returns equivalent interface wrenches to the structural model. This staggered update is consistent with standard UVLM-based aeroelastic coupling strategies, and is particularly suitable here because large FFWT rotations are introduced through the instantaneous geometry via the online meshing discussed in Subsection 5.3.1 [83].

In the following, the pose of interface frame k is denoted by $\{\mathbf{r}_k^n, \mathbf{R}_k^n\}$ and its translational and angular velocity by $\{\dot{\mathbf{r}}_k^n, \boldsymbol{\omega}_k^n\}$, all evaluated at time step t_n.

5.4.1. Online Meshing

Online mesh updating reconstructs the instantaneous aerodynamic lattice from a reference grid using the structural kinematics provided at the coupling interfaces. This step is essential for the FFWT application, because large wingtip rotations alter local panel normals, trailing-edge kinematics and moment arms, which must be reflected directly in the UVLM boundary condition and load extraction.

A reference aerodynamic grid is defined in the undeformed configuration, consisting of N_v lattice vertices $\mathbf{X}^0 = \{\mathbf{x}_i^0\}_{i=1}^{N_v}$. At time step t_n, each vertex is assigned to a coupling interface frame or to an interpolation region between frames, and the instantaneous vertex positions are obtained by applying the corresponding rigid-body transformation. For a vertex associated with interface frame k this reads as Equation 5.23.

$$\mathbf{x}_i^n = \mathbf{r}_k^n + \mathbf{R}_k^n \mathbf{x}_{i,k}^0 \tag{5.23}$$

$\mathbf{x}_{i,k}^0$ in this relation denotes the reference vertex coordinates expressed in the local coordinates of frame k. If vertices lie in regions influenced by multiple interface frames, the pose can be interpolated using the structural shape functions or modal interpolation employed for the coupling. From the updated vertex coordinates $\mathbf{X}^n$, collocation points, vortex filament directions and panel normals are recomputed.

In addition to the instantaneous geometry, the UVLM boundary condition requires the kinematic velocity of the moving surface. For a collocation point $\mathbf{x}_j^n$ associated with interface frame k, the rigid-body contribution to the surface velocity is written as Equation 5.24.

$$\mathbf{V}_{\text{motion},j}^n = \dot{\mathbf{r}}_k^n + \boldsymbol{\omega}_k^n \times \left(\mathbf{x}_j^n - \mathbf{r}_k^n\right) \tag{5.24}$$

Here, $\dot{\mathbf{r}}_k^n$ and $\boldsymbol{\omega}_k^n$ are obtained directly from the Simscape Transform Sensor outputs at the corresponding interface frame.

The additional elastic contributions are included when the local deformation field is provided via modal coordinates and their time derivatives. The velocity $\mathbf{V}_{\text{motion},j}^n$ enters the flow-tangency condition Equation 5.17 and therefore must be updated consistently with the online mesh. After the lattice update, the instantaneous trailing-edge geometry is identified from the updated lattice vertices and used to shed a new wake row. Because wake elements are shed from the instantaneous trailing edge, large wingtip rotations naturally produce a prescribed but non-planar wake even when wake convection is assumed to follow $\mathbf{V}_c \approx \mathbf{V}_\infty$.

5.4.2. Coupling Variables and Scheme

The coupled aeroelastic framework is implemented as a partitioned time-marching scheme in which the structural dynamics are integrated in Simscape and the unsteady aerodynamics are evaluated in PROTEUS. The coupling is realised at a finite set of aerodynamic interface frames that are defined in Section 5.2 and are consistently described in both models. At each coupling time step, Simscape provides the instantaneous pose and kinematics of these interface frames, which are used to update the aerodynamic lattice and to evaluate the motion boundary condition in the UVLM solver. PROTEUS returns aerodynamic loads which are condensed into equivalent interface wrenches and applied back to the structural model, closing the aeroelastic feedback loop.

If $k = 1, \ldots, N_f$ index the coupling interface frames. The structural state passed from Simscape to PROTEUS at time step t_n is collected as Equation 5.25.

$$\mathcal{S}^n := \left\{ \mathbf{r}_k^n,\ \mathbf{R}_k^n,\ \dot{\mathbf{r}}_k^n,\ \boldsymbol{\omega}_k^n \right\}_{k=1}^{N_f} \tag{5.25}$$

In this relation, $\mathbf{r}_k^n \in \mathbb{R}^3$ is the translation of frame k, $\mathbf{R}_k^n \in SO(3)$ is its rotation matrix, and $\dot{\mathbf{r}}_k^n$ and $\boldsymbol{\omega}_k^n$ are the translational and angular velocity of that frame. In the Simscape implementation, these quantities are obtained directly using the Transform Sensor blocks, that measure the relative transform between each interface frame and a common reference frame. In this case, this frame is the world frame that is used to define $\mathbf{V}_\infty$ and the aerodynamic axes.

The separation into pose $\{\mathbf{r}_k^n, \mathbf{R}_k^n\}$ and rates $\{\dot{\mathbf{r}}_k^n, \boldsymbol{\omega}_k^n\}$ reflects the two roles these variables play in the UVLM evaluation: the pose defines the instantaneous aerodynamic geometry through the online mesh update, while the rates define the motion-induced boundary condition through $\mathbf{V}_{\text{motion}}$ in Equation 5.17. In the present implementation, Simscape provides $\mathbf{r}_k^n$, $\mathbf{R}_k^n$, $\dot{\mathbf{r}}_k^n$ and $\boldsymbol{\omega}_k^n$ directly at each coupling interface frame via Transform Sensor outputs. These quantities are passed to PROTEUS to update the aerodynamic lattice and to evaluate the motion-induced boundary condition.

The aerodynamic solver returns equivalent wrenches at the same interface frames as Equation 5.26.

$$\mathcal{A}^n := \left\{ \mathbf{F}_k^n,\ \mathbf{M}_k^n \right\}_{k=1}^{N_f} \tag{5.26}$$

$\mathbf{F}_k^n$ and $\mathbf{M}_k^n$ denote the resultant force and moment about the origin of interface frame k, obtained by integrating panel-wise loads and condensing them to the coupling frames using Equation 5.22. To apply these loads consistently in Simscape, the wrenches are expressed in the same frame convention as used by the External Force and Torque-actuation blocks. When PROTEUS outputs loads in the local interface frame, they are rotated to the world frame as Equation 5.27.

$$\mathbf{F}_{k,\mathrm{w}}^n = \mathbf{R}_k^n \mathbf{F}_k^n, \qquad \mathbf{M}_{k,\mathrm{w}}^n = \mathbf{R}_k^n \mathbf{M}_k^n \tag{5.27}$$

This results in the fact that the applied force direction and the moment arm remain consistent under large FFWT rotations. This is critical for the hinge-dominated motion, where the interface orientation discontinuity must be represented as a true follower-load effect rather than as a small perturbation.

The Simscape structural model is integrated in time using the Simulink solver, while the aerodynamic model is evaluated at discrete coupling instants. In the baseline implementation used in this thesis the coupling is weak, meaning a single pass per time step. At each time step $t_n \rightarrow t_{n+1}$, the coupled update proceeds as:

1. The structural solver advances the structural state and provides the current translations and rotations of the aerodynamic interface frames.
2. The aerodynamic surface lattice is updated accordingly, where panel corner points, control points, and panel normals are recomputed from the current geometry.
3. The aerodynamic influence coefficients are assembled for the updated lattice and the flow-tangency condition is enforced using the motion-induced velocities, in order to solve for the bound circulation.
4. A new wake row is shed from the trailing edge geometry at that time step, and the wake geometry is advanced in time according to the prescribed convection model.

5. Aerodynamic loads are evaluated on the instantaneous lattice and transformed into equivalent interface wrenches, which are in turn mapped to structural nodal loads after which the structural solver proceeds.

To avoid an algebraic loop between structural outputs and aerodynamic loads within the same major solver iteration, the aerodynamic loads are applied using a sample-and-hold convention over the interval $[t_n, t_{n+1})$. This means loads computed at t_n are held constant until the next coupling instant. In practice, this is implemented through the Simulink calling sequence of the aerodynamic S-function, and may include an explicit unit delay if required by the solver settings.

5.5. Validation and Verification

The purpose of this chapter is to establish the credibility of the coupled Simscape-UVLM aeroelastic framework developed to investigate non-linear dynamics of the FFWT, through verification and validation. Here, verification refers to checking whether the implemented coupled model behaves consistently with its underlying mathematical and numerical formulation, whereas validation refers to comparison against trusted reference models or tools. These steps are essential for the model to be used for the final study cases.

5.5.1. Structural Model

To accurately represent the structural dynamics while keeping the ROM suitable for repeated time-domain simulations, the reduction in Simcenter is not assessed on eigenfrequency accuracy alone. Since the structural bandwidth directly influences the numerical stiffness of the system, and thereby the solver step size required for stable and accurate integration, the ROM is also evaluated based on the highest flexible frequencies retained within the modal range that is relevant to the problem at hand. An optimal number of retained modes in the ROM, n_m is then selected through a modal convergence study.

Modal Convergence Analysis

For each exported ROM, a free-free modal analysis (SOL103) is used as a baseline verification step. In the free-free configuration, the first six modes correspond to rigid-body motion and appear as near-zero frequencies; these modes are excluded from the comparison. The remaining elastic modes are compared between ROM and full-order model (FOM). In frequency regions with high modal density, mode ordering by frequency alone can lead to mode swapping; therefore, modes are paired using a Modal Assurance Criterion (MAC) computed on the common physical DOFs prior to computing frequency errors.

Two frequency scales are distinguished, the first of which is the frequency range in which the coupled aeroelastic response is expected to contain dominant dynamics. Secondly, the frequency range over which the structural model is validated to ensure that all globally relevant modes are represented.

For the present wing-only free-free configuration, modes become closely spaced beyond approximately 130 Hz. This behaviour is commonly associated with higher-order and more localized deformation patterns, which are less likely to dominate the global aeroelastic response of the structural model. In addition, Hoyng [18] defines practical dynamic relevance in the demonstrator context in terms of low-order bending dynamics and in his Ground Vibration Testing (GVT) procedure, the measurement bandwidth is selected based on the impact-hammer power spectral density; because the wing has relatively low eigenfrequencies, the measurement bandwidth is set to 204.80 Hz. With 10 seconds acquisition time and 2048 lines, this yields 0.10 Hz resolution.

Based on these considerations, a validation bandwidth of $f_{\text{val}} = 130$ Hz is adopted for eigenfrequency validation of the ROM. This choice captures the set of globally relevant elastic modes prior to the onset of strong modal crowding, while avoiding unnecessary inclusion of dense high-frequency content that primarily increases numerical stiffness in time integration. To mitigate modal truncation effects in forced-response simulations, a higher Craig-Bampton cut-off is used for basis enrichment, selected as $f_{\max} \approx 200$ Hz, which is consistent with the order of magnitude of Hoyng's GVT measurement bandwidth.

To quantify whether the ROM satisfies this condition, a series of ROMs is generated while varying the number of retained fixed-interface modes n_m. The boundary set (ASET) is kept constant at $n_b = 66$ DOFs, and this sequence provides a coarse, monotonic enrichment of the reduction space that is sufficient to identify the region of diminishing returns [84]. The relative eigenfrequency error for the first elastic modes, whose eigenfrequencies are below $f_{\text{val}} = 130$ Hz, is computed using the relation in Equation 5.28.

$$\epsilon_{f,k} = \frac{f_k^{\text{ROM}} - f_k^{\text{FOM}}}{f_k^{\text{FOM}}} \times 100 \qquad k = 1, \ldots, N_e \tag{5.28}$$

The reference solution is obtained from a SOL103 modal analysis of the full FE model under the same boundary conditions applied to the reduced eigenvalue problem in Equation 5.29.

$$\left(\mathbf{K}_{\text{CB}} - \omega^2 \mathbf{M}_{\text{CB}}\right) \boldsymbol{\varphi} = \mathbf{0} \tag{5.29}$$

The results are summarised in Table 5.7 using $\max_k |\epsilon_{f,k}|$ and the mean absolute error $\overline{|\epsilon_f|} = \frac{1}{N_e} \sum_{k=1}^{N_e} |\epsilon_{f,k}|$.

Table 5.7: Modal convergence of the free-free Craig-Bampton ROM.

n_m	**ROM size**	$\max \lvert\epsilon_f\rvert$ **[%]**	$\overline{\lvert\epsilon_f\rvert}$ [%]
20	86	6.32	1.30
40	106	6.32	1.30
60	126	7.34	1.79
80	146	4.05	1.02

The convergence study shows that the eigenfrequency agreement remains unchanged between $n_m = 20$ and $n_m = 40$, indicating that the dominant global dynamics within the validation bandwidth f_{val} are already captured at this level of reduction. Increasing the retained mode count to $n_m = 60$ does not improve the results. Instead, both the maximum and mean absolute frequency errors deteriorate slightly, which can be explained by non-monotonic behaviour in a densely populated spectral region. Modal pairing and the conditioning of the reduced operators become more sensitive to small changes in the basis here. A clear improvement is obtained at $n_m = 80$, where the mean absolute error reduces to 1.02%, giving the best overall agreement with the FOM among the tested reductions, which is why it is selected for the final ROM.

Matrix Property Validation

After determining the amount of retained modes n_m, the fundamental mathematical properties of the exported Craig-Bampton matrices are verified to ensure compliance with physical conservation laws and numerical stability requirements. An analysis is performed regarding the symmetry of the mass and stiffness matrices $\mathbf{M}_{\text{CB}}$ and $\mathbf{K}_{\text{CB}}$ to ensure compliance with Maxwell-Betti reciprocity and energy conservation principles [85]. In practice, this means that matrix symmetry enforces reciprocal work between coordinates, which is required for a conservative elastic system and for obtaining real-valued natural frequencies and mode shapes.

Also, positive definiteness of the mass matrix is checked to ensure that kinetic energy is larger than zero for any non-zero velocity. This confirms that every generalized coordinate is associated with a physically meaningful inertia, avoids zero or negative effective masses, and prevents singular or nearly singular mass matrices that could cause integration failure or ill-conditioned linear solves [66] .

The dimensional consistency is also verified to ensure the correct ASET definition is enforced, confirming the compatibility with the Simscape multibody coupling framework. Since eleven physical nodes are exported with six DOFs each in the Craig-Bampton ROM, it retains $n_b = 66$ physical boundary degrees of freedom (ASET). Additionally, as determined in Subsection 5.1.2, $n_m = 80$ fixed-interface normal modes are retained, resulting in a total of $n_{\text{CB}} = 146$ generalized coordinates in the reduced model, consistent with the ROM size reported in Table 5.7. The mass and stiffness matrices therefore are expected to possess dimensions $\mathbf{M}_{\text{CB}}, \mathbf{K}_{\text{CB}} \in \mathbb{R}^{146\times146}$, where the upper-left 66×66 block represents boundary-to-boundary coupling, the lower-right 80×80 block governs modal dynamics, and off-diagonal blocks capture the interaction between physical interface and elastic deformations. Verification that both matrices maintain these dimensions and appropriate block structure confirms correct superelement assembly and provides assurance that the ROM will interface properly with the Simscape multibody environment.

In addition to definiteness and symmetry, the numerical conditioning of the mass matrix is evaluated through its condition number, defined as Equation 5.30.

$$\kappa(\mathbf{M}_{\text{CB}}) = \frac{\lambda_{\max}(\mathbf{M}_{\text{CB}})}{\lambda_{\min}(\mathbf{M}_{\text{CB}})} \tag{5.30}$$

This number measures the spread of eigenvalues and therefore the sensitivity of linear solves involving $\mathbf{M}_{\mathrm{CB}}$ to floating-point perturbations. A large $\kappa(\mathbf{M}_{\mathrm{CB}})$ is common in Craig-Bampton coordinates, because the reduced coordinate vector mixes physical boundary translations and rotations, which are in disparate units and characteristic magnitudes, with modal coordinates. This leads to a wide range of effective inertial scales. The verification results are summarised in Table 5.8.

Table 5.8: Matrix property validation results.

Property	Criterion	Result
Symmetry ($\mathbf{M}_{\mathrm{CB}}$)	$\lvert\mathbf{M}-\mathbf{M}^T\rvert_F/\lvert\mathbf{M}\rvert_F < 10^{-12}$	7.119×10^{-17}
Symmetry ($\mathbf{K}_{\mathrm{CB}}$)	$\lvert\mathbf{K}-\mathbf{K}^T\rvert_F/\lvert\mathbf{K}\rvert_F < 10^{-12}$	6.457×10^{-13}
Positive definiteness ($\mathbf{M}_{\mathrm{CB}}$)	$\lambda_{\min}(\mathbf{M}_{\mathrm{CB}}) > 0$	1.213×10^{-2}
Positive definiteness ($\mathbf{K}_{\mathrm{CB}}$)	$\lambda_{\min}(\mathbf{K}_{\mathrm{CB}}) > 0$	-1.345×10^{-13}
Dimensional consistency	Size = 146×146	146×146
Condition number	$\kappa(\mathbf{M}_{\mathrm{CB}})$	8.409×10^{4}

Looking at the results of these checks, the symmetry checks confirm that both the reduced mass and stiffness matrices satisfy reciprocity to within numerical precision, with relative Frobenius-norm symmetry errors of 7.119×10^{-17} for $\mathbf{M}_{\mathrm{CB}}$ and 6.457×10^{-13} for $\mathbf{K}_{\mathrm{CB}}$. The minimum eigenvalue of $\mathbf{M}_{\mathrm{CB}}$ is 1.213×10^{-2}, confirming that the reduced mass matrix is positive definite and that the kinetic energy remains strictly positive for any non-zero generalized velocity state. Dimensional consistency is also verified, as both matrices retain the expected size of 146×146.

For the stiffness matrix, the smallest computed eigenvalue is slightly negative, but its relative magnitude is on the order of machine precision. This indicates that the observed negative eigenvalues are numerical artefacts associated with near-zero modes rather than evidence of a physically indefinite stiffness matrix. In practice, $\mathbf{K}_{\mathrm{CB}}$ is therefore interpreted as positive semidefinite, which is consistent with the presence of rigid-body or weakly constrained modes in the reduced structural representation. The reduced mass matrix has a condition number of $\kappa(\mathbf{M}_{\mathrm{CB}}) = 8.409\times10^{4}$, which is sufficiently moderate for stable numerical integration. Overall, these results confirm that the Craig–Bampton reduction preserves the expected matrix structure and yields a numerically consistent basis for the coupled aeroelastic simulations.

Eigenfrequency Comparison

Following the modal convergence study, the selected Craig-Bampton ROM with $n_m = 80$ retained fixed-interface modes is assessed against the FOM by comparing the free-free eigenfrequencies obtained from SOL103. In the free-free configuration, the first six modes correspond to rigid-body motion and appear as near-zero frequencies, so these are excluded from the comparison. The remaining elastic modes are evaluated up to the validation bandwidth $f_{\mathrm{val}} = 130$ Hz, defined on the FOM spectrum. Consequently, all elastic modes with $f_k^{\mathrm{FOM}} \le f_{\mathrm{val}}$ are included, which for the present model corresponds to modes 7-25. The remaining elastic modes are compared in the frequency range relevant to global aeroelastic response, and the relative eigenfrequency error is computed as Equation 5.31.

$$\epsilon_f = \frac{f_{\mathrm{ROM}} - f_{\mathrm{FOM}}}{f_{\mathrm{FOM}}} \times 100\% \tag{5.31}$$

Table 5.9 summarises the eigenfrequency agreement within this bandwidth and overall, the ROM reproduces the low-order elastic eigenfrequencies with good accuracy, seeing the majority of modes exhibit sub-percent deviation, while the maximum deviation within the validation set is 4.05%. The largest errors occur in the lower portion of the elastic spectrum, whereas the densely clustered region near 125–130 Hz is generally captured to within approximately 0.1–0.7%. The maximum error is associated with the third elastic mode, mode 9 in the free-free ordering, which is torsion-dominated. Torsional eigenfrequencies are particularly sensitive to how sectional kinematics are represented in the ROM interface definition and in the present reduction the spanwise interface is represented by a limited set of discrete mesh points, which can bias torsional stiffness and inertia representation relative to bending-dominated modes.

Despite this sensitivity, the observed deviations remain within an acceptable tolerance for the intended time-domain coupled aeroelastic simulations at this stage. The $n_m = 80$ ROM therefore provides an appropriate balance between modal fidelity within f_{val} and computational efficiency for integration in Simscape Multibody.

Table 5.9: Free-free eigenfrequency comparison between the FOM and the ROM.

FOM mode	f_{FOM} [Hz]	f_{ROM} [Hz]	ϵ_f [%]
7	29.2308	29.3987	0.57
8	61.3204	63.3705	3.34
9	76.8134	79.9263	4.05
11	110.1236	110.1323	0.01
12	111.9456	113.9001	1.75
13	119.4949	118.2406	1.05
14	123.7463	119.9675	3.05
15	125.3235	123.7539	1.25
16	125.7314	125.3308	0.32
17	126.8129	125.7542	0.83
18	127.2740	127.1463	0.10
19	127.5420	127.5443	0.00
20	127.9867	128.0266	0.03
22	128.7490	128.7606	0.01
23	129.2334	129.2297	0.00
25	129.3557	129.3743	0.01

Static Response Validation

To verify that the imported flexible wing in Simscape reproduces the same low-frequency (quasi-static) behaviour as the FOM, a linear static analysis (SOL101) is performed on the wing-only FOM with the root clamped. Six independent load cases are applied at the hinge load application node. The load cases consist of three forces (F_x, F_y, F_z) and three moments (M_x, M_y, M_z) defined in the global frame. The resulting hinge translations and small rotations were extracted from the OP2 and are summarised in Table 5.10.

Table 5.10: SOL101 static results FOM.

Subcase	Applied load	u_x [mm]	u_y [mm]	u_z [mm]	θ_x [rad]	θ_y [rad]	θ_z [rad]
1	$F_x = +10$ N	0.9526	-0.04245	-0.006820	-6.4103e-05	1.3207e-05	-0.0009299
2	$F_y = +10$ N	-0.04245	0.008583	-0.0002418	-1.6640e-05	-8.0469e-06	8.6135e-05
3	$F_z = -10$ N	0.006819	0.0002417	-34.4219	-0.04435	0.007815	-3.2322e-05
4	$M_x = 5000$ Nmm	-0.0321	-0.008320	22.1728	0.03938	-0.0008806	0.0001108
5	$M_y = 5000$ Nmm	0.006604	-0.004023	-3.907	-0.0008806	0.05133	-5.2538e-05
6	$M_z = 5000$ Nmm	-0.4650	0.04307	0.01616	0.0001108	-5.2538e-05	0.001152

Taken together, the six load cases define the static hinge flexibility and the associated cross-coupling between translation and rotation. These SOL101 responses serve as the reference against which the Simscape steady-state responses are compared under identical loading and measurement conventions. To create the same boundary conditions in Simscape as the SOL101 responses of the FOM, the root is connected to the World frame via a Weld Joint block, essentially clamping the root. The same force is applied at the hinge frame, of which the results are shown in Table 5.11.

Table 5.11: Simscape incremental static results.

Case	Applied load	u_x [mm]	u_y [mm]	u_z [mm]	θ_x [rad]	θ_y [rad]	θ_z [rad]
1	$F_x = +10$ N	0.9678	-0.04234	-0.07684	-5.3199e-05	0.001255	-0.0009310
2	$F_y = +10$ N	-0.04154	0.02033	-0.5249	-0.0009563	9.7826e-06	8.3417e-05
3	$F_z = -10$ N	-0.07372	0.001640	-34.3204	-0.04423	0.007794	2.2444e-05
4	$M_x = 5000$ Nmm	-0.02844	-0.7115	22.0918	0.03929	-0.0008719	6.7669e-05
5	$M_y = 5000$ Nmm	0.5122	0.001992	-3.9114	-0.0008679	0.05132	-4.6839e-05
6	$M_z = 5000$ Nmm	-0.4655	0.04163	0.01593	0.0001105	-5.03707e-05	0.001151

The six unit-load cases define the hinge compliance matrix and its cross-coupling between translation and rotation. However, for validation of quasi-static stiffness, each unit-load case has a dominant response component. This means that for the F_z response, u_z is the primary response characteristic, as well as θ_y for M_y etc. These are the main drivers for coupled aeroelastic loads because they represent the deformation components that carry the majority of strain energy for the applied generalized load.

Cross-axis terms are retained in Table 5.10 and Table 5.11 for traceability and sign-convention checking, but the numerical validation metric is based on the primary channel error in each subcase, reported in Table 5.12.

Table 5.12: Primary static response errors at the hinge between Simscape and SOL101.

Case	Applied load	Primary response	Δ	Primary % error
1	$F_x = +10$ N	u_x [mm]	+0.01523	+1.599%
2	$F_y = +10$ N	u_y [mm]	+0.01174	+136.826%
3	$F_z = -10$ N	u_z [mm]	+0.1016	−0.295%
4	$M_x = 5000$ Nmm	θ_x [rad]	-8.58×10^{-5}	−0.218%
5	$M_y = 5000$ Nmm	θ_y [rad]	-1.92×10^{-6}	−0.00374%
6	$M_z = 5000$ Nmm	θ_z [rad]	-5.51×10^{-8}	−0.00479%

From this table, it is apparent that there is a very close agreement in static response between the FOM SOL101 and the simulated response of the ROM in Simscape. The F_y unit-load case exhibits a large relative error in the primary response u_y ($\approx$ 137%), but this does not indicate a breakdown of the imported flexible-body stiffness, but results from the fact that the spanwise compliance of the clamped wing is extremely small in the FE reference: $u_y/F_y = 8.58 \times 10^{-4}$ mm/N. Consequently, the SOL101 response under 10 N is only $O(10^{-2})$ mm, which is comparable to numerical noise floors and to small residual kinematic contributions in the multibody representation like frame definitions and solver tolerances. In this ill-conditioned regime, percentage errors become a poor fidelity indicator because very small absolute differences translate into large relative deviations. The F_y case is therefore treated as a secondary check primarily for sign conventions and absence of gross cross-axis coupling, rather than as a calibration target based on relative error in a near-zero response component.

Dynamic Impulse Response

The SOL101 comparison in the previous subsection verifies that the imported flexible body reproduces the static hinge compliance of the full-order model, but it does not yet confirm whether the transient structural dynamics are also captured correctly. Therefore, a dynamic impulse-response validation is introduced as the next step. In Simcenter, the wing-only finite-element model is analysed with the root clamped, consistent with the SOL101 setup, and a short pulse load is applied at the hinge load application node. In Simscape, the same pulse input is applied at the hinge frame.

Although the static validation considered six independent load cases, the dynamic validation is restricted to the two most informative excitation-response pairs. This selection follows directly from the static results, in which each load case is characterized by one dominant hinge response channel. For the structure, the vertical force case F_z is bending-dominated and is therefore validated through the hinge displacement response $u_z(t)$, while the pitching moment case M_y is torsion-dominated and is validated through the hinge rotation response $\theta_y(t)$. Together, these two subcases capture the principal transient behaviour of the structure in bending and torsion, while secondary cross-coupled channels are retained only as supporting checks for sign conventions and qualitative consistency. The result of the first subcase, analysing the bending-dominated response by applying a half-sine pulse of magnitude $F_z = -10N$, is depicted in Figure 5.4.

The plot shows that the Simscape model captures the bending-dominated impulse response in $u_z(t)$ at the hinge location with good overall agreement relative to the Simcenter solution. The first minimum, subsequent maxima, and the general oscillations are reproduced well, indicating that the dominant natural frequency and phase evolution are retained after importing the flexible body into Simscape. A small discrepancy in peak amplitude can be observed, which becomes visible over the later cycles which could be attributed to a slight mismatch in the effective damping. However, since no substantial phase drift is observed and the response envelope remains close to the reference solution, the correlation is considered sufficiently good for validation of the primary structural dynamics in this bending-driven load case. Figure 5.5 shows the torsion-dominated impulse response in $\theta_y(t)$ after applying a half-sine pulse of magnitude $M_y = 5000Nmm$.

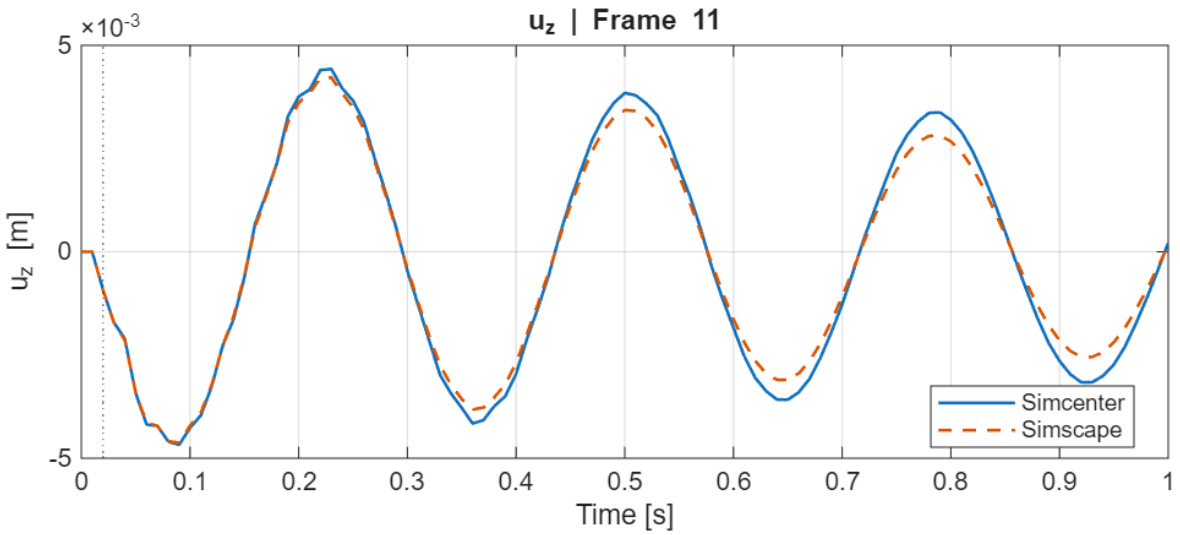

Figure 5.4: Bending-dominated impulse response in $u_z(t)$.

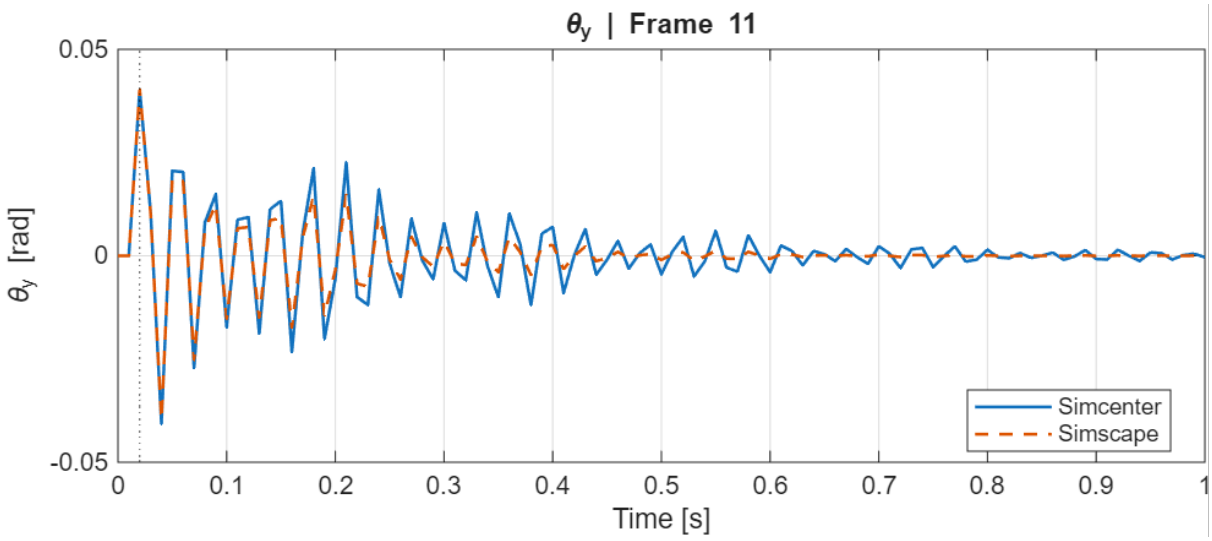

Figure 5.5: Torsion-dominated impulse response in $\theta_y(t)$.

This plot also shows that the torsion-dominated response is reproduced well by the Simscape model relative to the Simcenter benchmark. The agreement in oscillation frequency and phase is strong throughout the response, demonstrating that the relevant rotational structural dynamics are retained in the imported flexible-body representation. Some small differences in peak magnitude can be observed during the initial transient, but these remain limited and do not lead to any meaningful phase divergence. As both the short-term response and the decay trend are matched closely, the correlation is considered satisfactory for validation of the primary torsional dynamics of the structural model.

To support this visual assessment with a concise quantitative comparison, Table 5.13 summarizes the dominant frequencies, correlation coefficients, and the Normalized Root-Mean-Square Error (NRMSE) values for the primary impulse-response pairs considered in the dynamic validation. In this table, f_{SC} and f_{SS} denote the dominant frequencies extracted from the Simcenter and Simscape responses, respectively. corr is the correlation coefficient, quantifying the degree of linear similarity between two response signals, with values close to one, indicating strong agreement in waveform shape and phase.

Table 5.13: Summary of dynamic impulse-response validation results.

Subcase	Response	f_{SC} [Hz]	f_{SS} [Hz]	corr [-]	NRMSE [-]
F_z pulse	u_z	3.03	3.03	0.998	0.065
M_y pulse	θ_y	32.42	32.72	0.966	0.072

The table forms the final confirmation that the two dominant excitation–response channels are reproduced with high accuracy in terms of frequency content, correlation and overall waveform error. Secondary cross-coupled responses are also inspected and found to remain qualitatively consistent, but are not used as primary validation metrics since the quantitative comparison is based on the dominant channel identified for each load case. This validation step therefore complements the preceding static comparison: the SOL101 exercise shows that the imported model reproduces the correct hinge stiffness characteristics, while the impulse-response validation confirms that it also captures the relevant transient bending- and torsion-dominated structural dynamics needed for the subsequent aeroelastic simulations.

The verification results presented in this subsection show that the exported Craig–Bampton ROM is numerically well-posed and consistent with the full-order finite-element model in the frequency range relevant to the coupled simulations. The selected reduction ($n_b = 66, n_m = 80$) yields good eigenfrequency agreement within the verification bandwidth $f_{val} = 130$ Hz and preserves the expected symmetry, positive definiteness, and dimensional structure of the reduced mass and stiffness operators. When implemented in Simscape Multibody via the ROFS block, the ROM reproduces the dominant quasi-static hinge compliance channels and, through the impulse-response validation, also captures the principal transient bending and torsional dynamics with good fidelity. Together, these results indicate that the reduced structural model is sufficiently accurate for the subsequent aeroelastic analyses. With the ROM verified, the next section details the Simscape Multibody model assembly, including interface frame definitions, the revolute hinge implementation, and the load and kinematics interfaces used for aeroelastic coupling.

5.5.2. UVLM Solver

After validating the structural component of the Simscape assembly, the coupling with the UVLM solver excluding the force feedback on the structure is separately compared against a well-known, validated VLM solver. XFLR5 [86] is an analysis tool for airfoils, wings and planes and includes wing design and analysis capabilities based on the Lifting Line Theory [11]. Comparing aerodynamic coefficients computed for a rigid configuration at a range of angles of attack α for a prescribed free-stream velocity of $V_\infty = 50 m/s$ allows for tuning of aerodynamic discretization settings such as wake length and panel density. This can be repeated until they are in satisfactory agreement with the reference steady XFLR5 results for the locked-wing configuration. The main purpose of this validation is to check interface consistency and the steady rigid aerodynamic response.

Steady-state Lift Coefficient

The first validation quantity is the steady lift coefficient defined in Equation 5.32.

$$C_L = \frac{L}{q\, S_{\text{ref}}}, \qquad q = \tfrac{1}{2}\rho V_\infty^2 \tag{5.32}$$

To avoid ambiguity in force components, the lift is evaluated in wind axes as presented in Equation 5.33, where F_x and F_z are the total aerodynamic forces in the global Simscape frame produced by the UVLM solver, assuming no spanwise forces.

$$L = -F_x \sin\alpha + F_z \cos\alpha \tag{5.33}$$

PROTEUS provides the UVLM dynamics in discrete-time state-space form as shown in Equation 5.34, where x contains the wake circulations states and u denotes the aerodynamic boundary-condition input vector. In the rigid steady case, this input is fully determines by the free-stream velocity and angle of attack. Also, there is no structural motion and a steady input, so $\dot{\alpha} = 0$.

$$x_{k+1} = Ax_k + Bu_k, \qquad y_k = Cx_k + Du_k \tag{5.34}$$

The equilibrium is then obtained algebraically as presented in Equation 5.35.

$$x_\infty = (I - A)^{-1} Bu \tag{5.35}$$

In the current implementation using PROTEUS' UVLM solver, the steady UVLM problem is assembled in a comparable form in algebraic form as Equation 5.36.

$$K_8\, x_\infty + K_9\, \alpha_{\text{panel}} = 0 \tag{5.36}$$

Where x_∞ denotes the steady circulation state, K_8 is the steady aerodynamic influence operator obtained after assembling the flow-tangency equations on the bound lattice and incorporating the semi-infinite wake through the steady Kutta condition, and K_9 maps the prescribed aerodynamic input α_{panel} into the same boundary-condition residual. The quantity α_{panel} is the effective panel incidence obtained from the local panel normals and the prescribed free-stream direction. Solving this system yields the steady circulation state that satisfies the no-penetration condition for the chosen lattice and wake model as presented in Equation 5.37.

$$x_\infty = -K_8^{-1} \left(K_9\, \alpha_{\text{panel}}\right) \tag{5.37}$$

The UVLM post-processing script is set up such that it provides sectional resultant forces F_{sec}, yielding one resultant per spanwise section. Since these resultants already include the local section width via PROTEUS' internal section-width bookkeeping, the total forces are obtained by summation over sections. Also, no additional multiplication by a spanwise spacing Δy is applied during the processing, as this would double-count the spanwise integration. For the lift coefficient, S_{ref} corresponds to the lifting planform that is used in the aerodynamic model, which is half-wing in symmetry with the appropriate semi-span. This reference-area definition is kept consistent with the XFLR5 configuration.

The wake is then discretized with stream-wise spacing Δx_w and a finite number of wake rows N_w, corresponding to a wake length L_w downstream. The key non-dimensional wake step is defined as Equation 5.38 and the non-dimensional wake length is L_w/c. For steady validation, the wake is not time-marched, but the wake discretization still affects the steady operator through the wake geometry and Kutta enforcement. Therefore, a convergence study is performed by increasing L_w/c (e.g. 10c, 20c, 60c, 120c) until the resulting steady C_L changes negligibly. In addition to this, sensitivity to the wake step $\Delta x_w/c$ is checked by reducing Δx_w while keeping L_w/c fixed and increasing N_w accordingly.

$$U\Delta t/c \;=\; \frac{\Delta x_w}{c} \tag{5.38}$$

In addition to the wake discretization, the panel discretization influences the accuracy of induced velocities and the enforcement of the boundary condition. The aerodynamic lattice is split into spanwise and chordwise panels (N_s, N_c) on the main and tip surfaces. A grid refinement study is carried out by increasing (N_s, N_c) and verifying that C_L converges. In the present configuration, spanwise refinement was prioritized to better resolve the tip effects, and chordwise refinement is used to confirm adequate resolution of the Kutta condition and bound circulation distribution.

During initial validation, numerical sensitivity to wake placement and discretization is traced to a regularization strategy in PROTEUS' internal Aerodynamic Influence Coefficient (AIC) kernel. The original implementation applied a hard cut-off using a fixed core radius $\epsilon = 25mm$, setting induced velocities to zero whenever the distance to a filament was smaller than ϵ. For the present geometry and panel sizes, ϵ was comparable to characteristic panel dimensions, which materially altered the operator and introduced discontinuities with respect to wake location.

To restore a physically consistent and numerically robust desingularization, the hard cutoff was replaced by a smooth vortex-core factor based on the perpendicular distance $D_\perp$ to the filament centreline as presented in Equation 5.39. ϵ is in this case scaled to panel dimensions, which preserves conditioning close to near-singular interactions while maintaining continuity and reducing sensitivity to wake placement.

$$\text{core}(D_\perp) = \frac{D_\perp^2}{D_\perp^2 + \epsilon^2} \tag{5.39}$$

The form as described in Equation 5.37 clarifies why parameters like wake length, wake spacing, panel density and vortex-core regularization affect the extracted steady C_L, since these modify the assembled aerodynamic operator K_8.

A persistent $\sim 10\%$ lift-slope discrepancy relative to the XFLR5 results is found and traced to an aero-structural coordinate mismatch. The ROM includes a physical fixture region, leading to a first structural interface frame at a non-zero spanwise coordinate. A half-wing UVLM model is mirrored about $y = 0$ using symmetry, leading to an artificial centre gap of width $2y_0$ between the mirrored halves, creating additional root-vortex effects and reducing the lift slope. Intentionally introducing the same gap in XFLR5, closes this gap in lift coefficient.

To correct this in the Simscape-UVLM coupling, while keeping the structural world frame unchanged, the aerodynamic lattice is expressed in a local aerodynamic spanwise coordinate $y' = y - y_0$. This ensures that the first lifting station satisfies $y' = 0$ and the symmetry plane coincides with the true aerodynamic root. The resulting forces and moments are then mapped to the original structural frames, so only the aerodynamic geometry is based on y'. The reference area S_{ref} is also updated in the lift coefficient check, in order to represent the true lifting planform, based on the shifted semi-span.

After enforcing the described steady-state algebraic extraction, ensuring the correct load aggregation, converging the wake length and panelling, replacing the PROTEUS' internal AIC hard cut-off by a smooth core model and removing the symmetry-plane gap via the aerodynamic shift in y-direction, the Simscape-UVLM coupling agrees with XFLR5 VLM lift-curve results to within approximately 1% for the rigid benchmark case, as confirmed by the $C_L - \alpha$ plot in Figure 5.6.

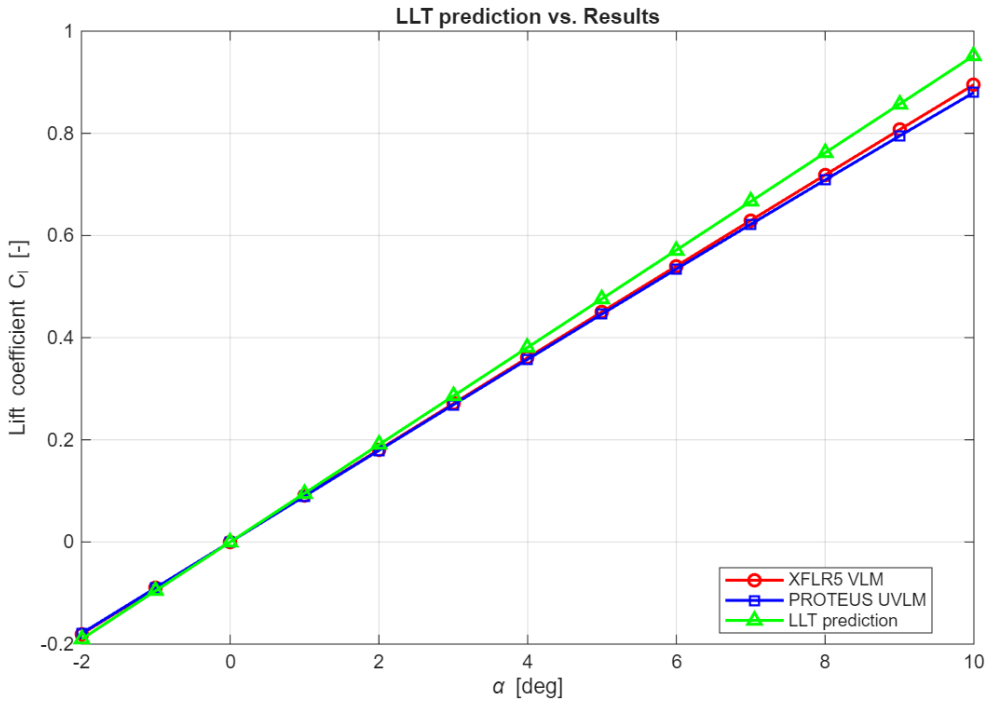

Figure 5.6: C_L-α plot, comparing XFLR5 VLM versus PROTEUS UVLM versus Lifting-Line Theory.

In addition to the XFLR5 VLM reference, a LLT estimate is included in the $C_L - \alpha$ comparison as a lower-order analytical benchmark for the finite-wing lift-curve slope. For a finite wing, LLT predicts a lift-curve slope of the form in Equation 5.40, where a_0 is the two-dimensional airfoil lift-curve slope, AR is the Aspect Ratio and e is the span-efficiency factor.

$$a = \frac{a_0}{1 + \dfrac{a_0}{\pi e AR}}, \qquad C_L = a(\alpha - \alpha_{L=0}) \tag{5.40}$$

Whereas XFLR5 VLM and the present Simscape-UVLM coupling are both lifting-surface methods, LLT models the wing through a spanwise circulation distribution only. It is therefore not used here as the primary validation target, but rather as a physics-based consistency check on the linearity and slope of the lift curve. Agreement between LLT, XFLR5 VLM and PROTEUS UVLM indicates that the definition of angle of attack, lift extraction, symmetry treatment and reference area are implemented consistently. Small differences are expected because LLT does not resolve chordwise loading and approximates tip effects more coarsely than the lifting-surface methods.

Induced Drag Coefficient

The second validation quantity is the induced-drag coefficient. A consistency check is again carried out against the same rigid-wing Benchmark in XFLR5. Since the UVLM formulation is inviscid, the drag predicted by the aerodynamic model should correspond to induced drag only. In contrast to the consistency check on the lift coefficient, using the total force vector for checking small drag values is inaccurate. Therefore, the induced drag is evaluated from the spanwise circulation distribution as an additional validation step, where the starting point is the converged steady bound-circulation distribution obtained from the UVLM solution.

From this distribution a lift coefficient based on the circulation, denoted here as $C_{L,\Gamma}$, is first evaluated and compared with the lift coefficient obtained from the total aerodynamic forces in wind axes. This provides an additional internal consistency check that the recovered circulation field and the force post-processing align. To estimate then the induced drag coefficient, the half-wing circulation distribution is post-processed using a Fourier representation based on the lifting-line theory. The circulation distribution $\Gamma(y)$ is interpreted as the spanwise loading of the finite wing, from which the induced drag coefficient C_{D_i} and span efficiency factor e are obtained. In classical lifting-line form, the induced drag can be written as Equation 5.41.

$$C_{D_i} = \frac{C_L^2}{\pi AR e} \tag{5.41}$$

In the post-processing routine used here, e is obtained from the Fourier representation of the circulation distribution. Particular care is found to be required in the reference-area definition when processing a symmetric half-wing model: The UVLM solution itself is obtained on a half-wing using symmetry, but the induced-drag relation is interpreted at full-wing level. A comparison is made at a reference case of $\alpha = 4$ deg, leading to the results in Equation 5.42.

$$C_{L,\Gamma} \approx 0.3607, \qquad C_{L,\text{lift}} \approx 0.3598 \tag{5.42}$$

The resulting corresponding span-efficiency factor is $e = 0.899$. These results confirm that the lift coefficient obtained from the aerodynamic forces and the lift coefficient obtained from the circulation are correct, and the induced drag coefficient can now be evaluated safely. In Table 5.14, the induced drag coefficient predicted by the present PROTEUS UVLM implementation and the XFLR5 VLM reference over the investigated angle-of-attack range are compared. The same rigid-wing configuration and free-stream velocity $V_\infty = 50$ m/s are used as in the lift validation.

The comparison against the XFLR5 VLM induced-drag coefficient shows satisfactory agreement, where the relative error is a systematic $4 - 5\%$ over the investigated range of α, which can be partly be explained by the fact $C_{D_i} \propto C_L^2$. This amplifies small differences in lift coefficient found before, which in turn propagate as a larger systematic bias in the induced drag comparison. The remaining bias can be associated with the reconstruction of e using lifting-line based Fourier reconstruction.

Table 5.14: Comparison of induced-drag coefficients from XFLR5 VLM and PROTEUS UVLM values.

α [deg]	$C_{D_i,\text{XFLR5}}$ [-]	$C_{D_i,\text{PROTEUS}}$ [-]	Error [%]
-2	0.00079	0.000831	5.19
-1	0.00020	0.000208	4.00
0	0.000000	0.000000	—
1	0.00020	0.000208	4.00
2	0.00079	0.000831	5.19
3	0.00178	0.001867	4.89
4	0.00316	0.003316	4.94
5	0.00492	0.005172	5.12
6	0.00707	0.007434	5.15
7	0.00959	0.010094	5.26
8	0.01248	0.013150	5.37
9	0.01573	0.016592	5.48
10	0.01932	0.020416	5.67

In addition to the comparison against XFLR5, several direct interface-consistency checks regarding the interface are performed for the rigid symmetric benchmark. At $\alpha = 0$, the UVLM solver returns $C_L = 0$, confirming consistency of the incidence definition and force transformation. Reversing the sign of α also flips the sign in lift, and no significant spanwise force or moment components are observed for the locked wing configuration. Together with the agreement between $C_{L,\Gamma}$ and $C_{L,lift}$, these additional verification checks inside the aerodynamic validation provide confidence that the aerodynamic inputs and outputs are transferred consistently in the created Simscape-UVLM interface.

5.5.3. Coupled system

The structural and aerodynamic subsystems are now validated separately, so the next step is to verify the consistency of the coupled model. A direct comparison of the complete Simscape model against a static Simcenter solution is not straightforward, because the Simscape model contains a flexible model represented by a ROM, a rigid outboard wingtip, a revolute hinge that can be either in 'locked' or 'released' state and aerodynamic loads applied at interface frames.

For the first stage, a simulation of the coupled system is executed with the feedback of forces on the structural frames disabled, thereby essentially analysing a rigid, undeformed wing. For this purpose, an AoA of $\alpha = 1$ deg is used, combined with a freestream velocity of $V = 50 m/s$. The choice for these parameters is deliberate: the SOL101 solution in Simcenter/Nastran is a linear solution, so to avoid non-linear effects it is important to keep excitation small for the comparison. The purpose of this exercise is not to reproduce dynamic behaviour, but to verify that the coupled framework transfers a representative aerodynamic load field into the structure in a consistent manner. The forces and moments on the twelve interface frames that are a result of this simulation are combined in the load vector in Equation 5.43.

$$\mathbf{F}_{\text{frozen}} \in \mathbb{R}^{72} \tag{5.43}$$

For the verification, these loads are applied statically in both environments. In Simscape, the UVLM solver is temporarily disabled, and the load vector is instead fed back to the interface frames. In Simcenter/Nastran, a SOL101 load case is created, applying the forces and moments on the same nodes used in the ROM export.

Table 5.15 summarizes the vertical displacement comparison. The interface points are ordered from root-to-tip. The root frame is included for completeness, but this is where the weld joint in Simscape is connected, and is a constraint point in the SOL101 solution, so there is no deflection or rotation present.

Table 5.15: Comparison of vertical displacement SOL101 versus Simscape framework ($V_\infty = 50m/s$, $\alpha = 1°$).

Frame	SOL101 z [mm]	Simscape z [mm]	|Error| [mm]	Error [%]
1	0.000	0.000	0.000	–
2	0.855	0.863	0.008	0.898
3	5.129	5.254	0.125	2.438
4	12.078	12.094	0.016	0.132
5	20.547	20.507	0.040	0.195
6	30.249	30.230	0.019	0.062
7	41.940	41.952	0.013	0.030
8	53.610	53.695	0.085	0.159
9	65.786	65.982	0.196	0.297
10	78.056	78.422	0.367	0.470
11	93.667	94.028	0.361	0.385
12	151.248	148.833	2.415	1.597

The conclusion of this comparison is that overall, the agreement between the Simscape and SOL101 static deflections is good. The mean absolute error is $0.304mm$ and the root-mean-square error is $0.717mm$, while the mean percentage error remains below 1%. The largest absolute deviation occurs at the wingtip, where the Simscape model underpredicts the SOL101 displacement by $2.415mm$, corresponding to a relative error of 1.597%. The maximum relative error of 2.438% occurs at the third interface frame, but the associated absolute deviation there is only $0.125mm$. These results indicate that the Simscape structural implementation reproduces the static bending deformation from the SOL101 reference model with high fidelity.

SOL144

The static deflection of the main wing and wingtip combination is now verified, so the next step in the validation ladder is assessing the coupled static aeroelastic response of the Simscape UVLM-ROM framework against a Simcenter SOL144 reference solution. In this subsonic case, the aerodynamic loading is based on DLM, providing a linearized aeroelastic reference formulation. As discussed in Section 2.3, DLM has well-known limitations regarding non-linear effects, which is the reason the benchmark condition in this case is intentionally restricted to a free-stream velocity of $V_\infty = 20m/s$ and an angle of attack of $\alpha = 1$ deg. Additionally, this case is simulated with the hinge locked in both systems, ensuring that the comparison remains within a low-speed, low-incidence regime and avoiding non-linear effects like large tip rotations.

The SOL144 comparison serves to validate mainly whether the aerodynamic loads are transferred consistently to the structural model, whether the resulting deformations are fed back correctly into the UVLM solver and whether this coupled interaction yields the same equilibrium as the reference SOL144 solution. A simulation of sufficient time to dampen out the oscillations is run in Simscape, leading to a steady state solution regarding the deflections of the interface frames. The results of the comparison are presented in Table 5.16.

Table 5.16: Comparison of Simcenter SOL144 reference solution versus Simscape framework ($V_\infty = 20m/s$, $\alpha = 1°$).

Frame	SOL144 z [mm]	Simscape z [mm]	Error [mm]	Error [%]
1	0.000	0.000	0.000	—
2	0.121	0.130	0.009	7.20
3	0.777	0.813	0.036	4.61
4	1.843	1.886	0.043	2.36
5	3.143	3.207	0.064	2.02
6	4.638	4.730	0.092	1.99
7	6.439	6.560	0.121	1.87
8	8.243	8.397	0.155	1.87
9	10.131	10.327	0.196	1.93
10	12.046	12.288	0.242	2.01
11	14.450	14.730	0.280	1.94
12	23.441	22.722	-0.719	-3.07

Overall, the agreement is very good across the span, excluding the clamped root frame, which is identically zero in both models, the majority of local discrepancies remain close to 2%. Only the second inboard frame shows a somewhat larger deviation of 7.2%, while the difference at the wingtip is limited to −3.1%. This change in sign regarding deviation can be explained by the rigid tip assumption, where the flexible tip in Simcenter bends upwards due to the force distribution and the tip in Simscape remains straight. Otherwise, the interface frames reproduce the SOL144 deformation shape very closely, indicating that the Simscape implementation captures not only the overall magnitude of the response, but also the spanwise distribution of static aeroelastic deflection.

Trend Comparison

As a final validation method a trend reproduction is added, focused on reproducing the locked, free and threshold-released FFWT responses reported by Córcoles [17]. This provides the most direct assessment of whether the present non-linear Simscape–UVLM framework can reproduce established FFWT behaviour before it is used for new design studies. The framework is compared against the dynamic gust-response results for his 'wing B', as depicted in Figure 5.7. This setup is taken as the closest experimental reference configuration for the current thesis.

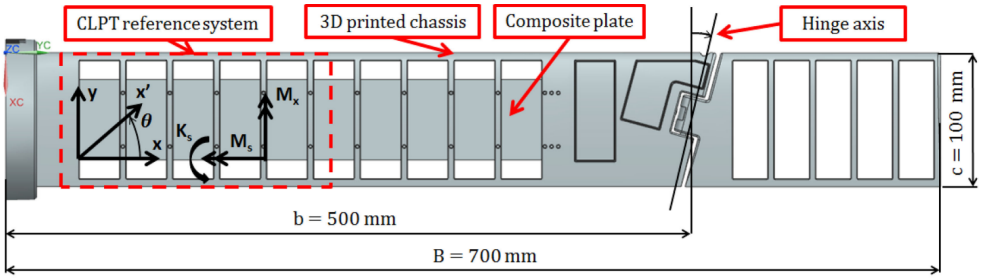

Figure 5.7: CAD representation of Córcoles wind tunnel model without skin with an overview of its dimensions, retrieved from [52].

As becomes apparent from the dimensions annotated in this figure, the model used by Córcoles is smaller in size. A direct one-to-one match in absolute response can therefore not be expected. For this reason, the comparison is interpreted primarily at the level of response trends, namely the ordering of the hinge cases, the sensitivity to release timing, the relative peak WRBM reduction, and the persistence of post-gust oscillations.

To keep the comparison focused and consistent, the final validation is restricted to a single representative gust case at $\alpha = 5°$, gust amplitude $\delta_g = 2.5°$, and gust frequency $f_g = 5$ Hz. This case is selected as an intermediate dynamic condition within the experimental campaign of Córcoles and is sufficiently fast to expose the sensitivity to hinge-release timing, while remaining less extreme than the 8 Hz case. In addition, the $\alpha = 5°$ condition is preferred over $\alpha = 0°$, because it combines the transient gust response with a non-zero static aerodynamic load and is therefore more representative of the operating conditions of interest.

Córcoles defines a couple of release cases, where the hinge-release timing is defined directly from the measured load response during the 1-cosine gust tests. In addition to the locked- and free-hinge reference conditions, four dynamic release cases are considered:

1. Pre-released, in which the hinge is released with the gust-generator trigger before the gust reaches the wing. In the current setup, the hinge is configured to release at gust start.
2. 0% release, defined as the instant at which the gust reaches the wing and the load starts to increase.
3. 50% release, defined as the instant at which the load reaches 50% of its peak value.
4. 100% release, defined as release at the peak load.

These configurations are all retained for the validation because they capture the baseline locked response, the continuously released response, and the delayed-release conditions required to assess the sensitivity of FFWT performance to release timing. Figure 5.8 visualises the definition of the release timings with respect tot the WRBM response. A force ramp-up of one, and a damping period of four seconds is used to make sure a clean baseline is reached before the gust start, hence the five second start time of the plot.

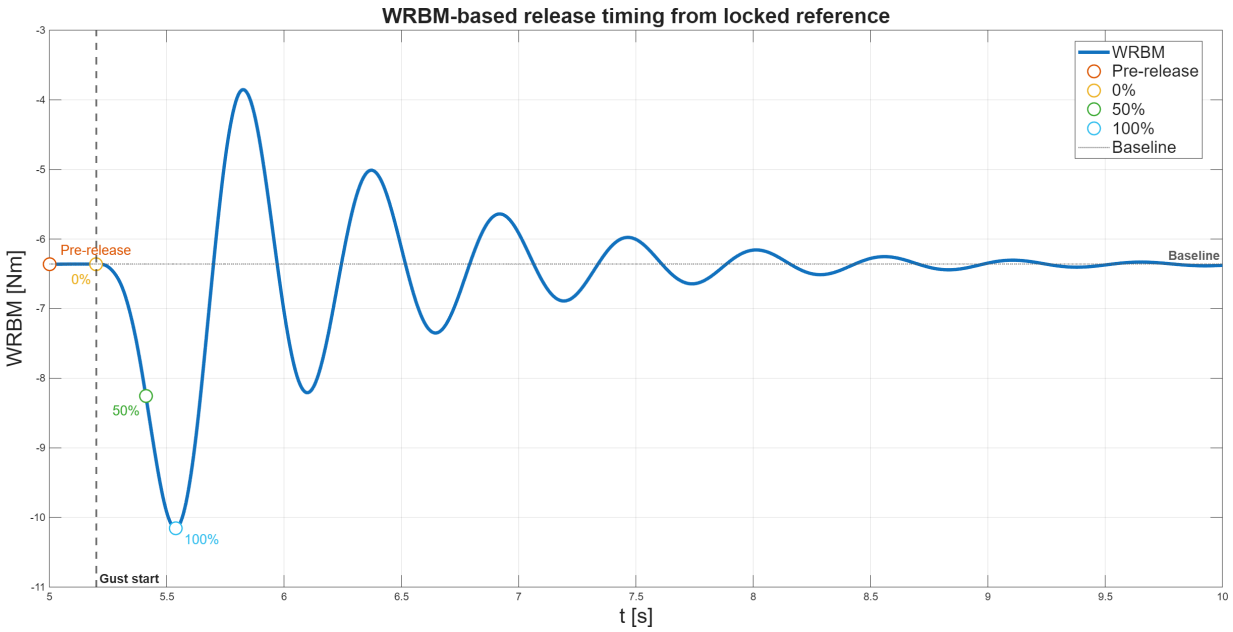

Figure 5.8: Visualisation release timing on the basis of WRBM in Simscape-UVLM framework, as proposed by Córcoles [17].

Followed by Table 5.17 to quantify the exact timing.

Table 5.17: Release timing on the basis of WRBM, as proposed by Córcoles [17].

Release timing	Time [s]
Pre-released	5.000
0%	5.200
50%	5.411
100%	5.538

The comparison is performed on the basis of two WRBM-based quantities. First, the peak WRBM response of each hinge configuration is normalised with respect to the locked-hinge reference case. In this representation, negative values indicate a reduction in peak WRBM relative to the locked configuration, whereas positive values indicate a load increase. Second, the corresponding WRBM time histories are compared to assess whether the release cases follow a physically consistent transient evolution before, during and after the gust encounter. This combination allows the validation to focus on the relative response trends, rather than on absolute load magnitudes, which are not expected to match because of the geometric and structural differences between the present model and the experimental reference.

The main focus of this validation is therefore to observe whether the Simscape-UVLM framework captures:

- the relative ordering of the locked, free, pre-released and threshold-released cases;
- the sensitivity of peak WRBM alleviation to hinge-release timing;
- the deterioration of the alleviation performance when release is delayed toward the peak-load instant;
- the post-gust oscillatory behaviour visible in the WRBM response.

The first comparison is made using the normalised peak WRBM values shown in Figure 5.9. The figure combines the experimental Wing B results of Córcoles with the corresponding recreated Simscape–UVLM results for the same 5 Hz gust encounter. This comparison provides a direct view of how the relative load-alleviation performance changes across the locked, free and prescribed-release hinge configurations.

In the experimental results, the pre-released case gives the strongest reduction, followed by the early-release cases, whereas the 50% and 100% release cases approach the locked-hinge value. The recreated Simscape–UVLM results reproduce this overall sensitivity to release timing.

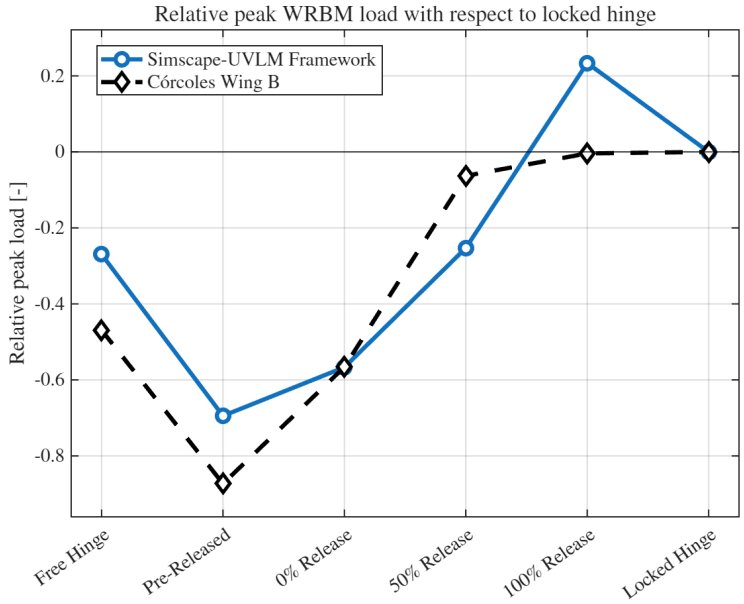

Figure 5.9: Dynamic trend validation against the 5 Hz gust-release experiment of Córcoles for Wing B [17].

The pre-released case again gives the strongest reduction, the 0% release case remains clearly beneficial, and the 100% release case becomes detrimental, producing a peak load larger than the locked-hinge reference.

There are also clear quantitative differences. The free-hinge case in the recreated simulations gives a smaller peak-load reduction than the corresponding experimental result, while the 50% release case remains more beneficial in the Simscape–UVLM framework than it does in the experimental data. This means that the numerical model captures the correct direction of the timing effect, but does not reproduce the full rate at which the benefit collapses between the early-release and near-peak-release cases. The 100% release case, however, does reproduce the important qualitative behaviour that release at or near the locked-reference peak load can remove the benefit of the FFWT and may even increase the peak WRBM. To interpret these differences, the full WRBM time histories of the Simscape–UVLM simulations are presented in Figure 5.10.

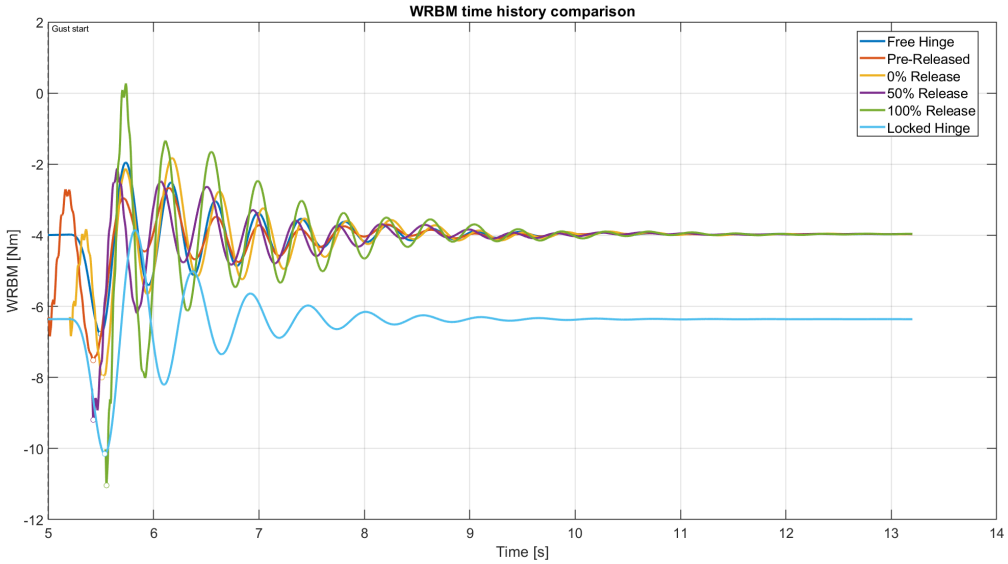

Figure 5.10: WRBM time histories for $5Hz$ gust encounter, obtained in Simscape-UVLM framework.

The main observation from these plots is that the pre-released and 0% release cases do not follow the same initial peak direction as the locked and delayed-release cases. Instead of first developing the large downward WRBM peak observed for the locked, 50%, 100% and free-hinge responses, the early-release cases initially produce an upward WRBM peak. This shows that early release changes the phase of the aeroelastic response, rather than only reducing the magnitude of the locked-hinge peak.

This phase shift explains the observed alleviation mechanism. When the hinge is released before or at gust arrival, the wingtip is already able to move while the gust load is still developing. The resulting motion modifies the instantaneous aerodynamic loading and thereby reduce the downward WRBM excursion, which in the framework used equals upwards bending. The delayed-release cases remain closer to the locked response during the initial gust encounter, so their first dominant response is still governed by the locked-hinge peak before additional post-release oscillations are introduced.

The comparison with Córcoles should therefore be interpreted in terms of effective release phase rather than release label alone. In the present simulations, the release instant is prescribed numerically, whereas the experiment includes the finite response time of the physical release mechanism. As a result, nominally identical release cases may correspond to slightly different phase relations between wingtip motion and gust development. Consequently, the minimum WRBM peak is used as the primary peak-load metric in the following analyses in the results section. This choice avoids interpreting the phase-shifted upward response of the early-release cases as the critical load and instead compares the dominant downward WRBM excursion consistently across release timings.

Overall, the trend validation is considered successful at the intended qualitative level. Although the framework does not reproduce all relative peak-load values exactly, particularly for the free-hinge and 50% release cases, it captures the governing behaviour described by Córcoles [17], where early release provides WRBM alleviation, delayed release reduces this benefit, and release near the peak-load instant can become detrimental. The following simulation stages therefore extend this comparison by generalising the timing study through relative-frequency scaling, before isolating the effects of hinge stiffness, damping and release laws on the basis of the hinge moment.

5.5.4. Concluding Remarks

Taking the various verification and validation results presented in this section into account, it is established that the coupled Simscape-UVLM framework is sufficiently credible for the non-linear FFWT studies carried out in the remainder of this thesis. At structural level, the Craig–Bampton reduction preserves the expected mathematical properties of the reduced operators, reproduces the relevant eigenfrequencies within the selected validation bandwidth, and captures both the quasi-static hinge compliance and the dominant transient bending- and torsion-dominated responses with good fidelity after import into Simscape. This provides confidence that the reduced flexible-body representation is numerically consistent and dynamically representative of the underlying finite-element model.

At aerodynamic level, the rigid-wing UVLM implementation shows good agreement with XFLR5 for the steady lift and induced-drag characteristics after convergence of the wake and panel discretisation and correction of the symmetry-plane offset. These results confirm that the aerodynamic states, reference quantities, force extraction and coordinate transformations are implemented consistently, and that the locked-wing aerodynamic response is captured to a level appropriate for subsequent coupled simulations.

At coupled-system level, the frozen-load comparison demonstrates that the aerodynamic loads are transferred consistently into the structural model, while the SOL144 comparison shows that the coupled framework reproduces the static aeroelastic equilibrium with good agreement in the low-speed, low-incidence regime. Finally, the trend comparison against the dynamic gust-response experiments of Córcoles shows that the framework captures the governing FFWT physics at the level relevant to this thesis: the sensitivity of load alleviation to hinge-release timing, the ordering of the principal release cases, and the transition from beneficial early release to detrimental late release. Although discrepancies remain in some individual cases, these are shown to be physically interpretable in terms of phase sensitivity, actuator delay and unavoidable geometric mismatch between the numerical and experimental configurations.

Overall, the framework is therefore regarded as verified and validated to the level required for the objectives of this thesis. While it is not intended as a point-by-point reproduction of the experimental demonstrator, it has been shown to reproduce the dominant structural, aerodynamic and coupled aeroelastic behaviours that govern the comparative analysis of FFWT concepts. The model can therefore be used with sufficient confidence for the design-space exploration and parametric FFWT studies presented in the following chapters.

5.6. Simulation Campaign

This section presents the simulation campaign used to generate new results with the validated non-linear Simscape-UVLM framework. The objective is to investigate the FFWT response under systematically varied release definitions, hinge properties and release-law formulations.

The campaign is organised into three sequential stages, moving from release-timing interpretation, to hinge-property sensitivity, to hinge-moment-based release design:

1. **Stage I: release-timing extension based on Córcoles**, extending the experimentally observed release-timing trends through systematic variation of WRBM-based release timing, relative gust-to-structural frequency scaling and reduced frequency.
2. **Stage II: hinge-parameter sensitivity analysis**, quantifying the influence of hinge stiffness and hinge damping on load alleviation, wingtip motion and post-gust dynamic behaviour under a common reference excitation.
3. **Stage III: threshold-based release-law analysis**, formulating and assessing hinge-moment-based release laws and associated trigger logic for the selected hinge configurations.

Before defining the individual simulation cases, a common set of performance metrics is introduced to ensure consistent comparison across the three stages.

5.6.1. Performance Metrics

To ensure consistency across the simulation campaign, each simulation is evaluated using a common set of structural, kinematic, and aerodynamic response metrics. The primary structural metric is the WRBM, since this is the principal load quantity used to assess gust-load alleviation. In addition, hinge-angle response, hinge moment, and the associated kinematic change in local wingtip angle of attack are monitored to characterise release timing, aerodynamic unloading, wingtip motion, and post-gust dynamic behaviour.

The primary peak-load metric used in this thesis is the critical WRBM peak in the gust-load direction. In the adopted sign convention, the locked gust response produces a negative WRBM peak. Therefore, the critical WRBM peak for each simulation is defined as the minimum WRBM value reached after gust onset, described in Equation 5.44.

$$M_{\min} = \min_{t_g \le t \le t_{\text{end}}} \text{WRBM}(t) \tag{5.44}$$

The corresponding peak WRBM reduction with respect to the locked-hinge baseline is then defined as Equation 5.45.

$$\Delta_{\text{WRBM}} = \left(1 - \frac{|M_{\min,\text{case}}|}{|M_{\min,\text{locked}}|}\right) 100\% \tag{5.45}$$

Positive values of Δ_{WRBM} indicate that the release case reduces the critical downward WRBM peak relative to the locked baseline. Negative values indicate that the release case reaches a larger downward WRBM peak than the locked configuration. Opposite-sign rebound is not included in this peak-load metric, because it does not represent the maximum bending response in the gust-load direction.

The secondary response metric is the peak hinge angle, defined as Equation 5.46.

$$\theta_{\max} = \max_{t_{\theta,0} \le t \le t_{\text{end}}} |\theta(t)| \tag{5.46}$$

Here, $t_{\theta,0}$ denotes the start of the relevant hinge-motion interval. For locked-free release cases, $t_{\theta,0} = t_{\text{rel}}$. For free-hinge or pre-released cases, $t_{\theta,0} = t_g$ is used so that the reported peak hinge angle corresponds to the gust-driven response. This quantity measures the wingtip-motion demand associated with a given release strategy and is used together with peak WRBM reduction to assess the load-alleviation trade-off. Residual post-release behaviour is assessed qualitatively from the WRBM and hinge-angle time histories.

As a third metric, the kinematic local angle-of-attack relief is used as an auxiliary metric to interpret the local unloading mechanism at the folding wingtip. For small fold angles, the mechanism is commonly interpreted using

the first-order relation described in Section 2.2 and repeated in Equation 5.47.

$$\Delta\alpha_{\text{tip,geom}} \approx -\theta \sin\Lambda \tag{5.47}$$

This relation is useful for explaining the mechanism, but it is not used as the quantitative metric because the simulated fold angles are not necessarily small.

The reported local angle-of-attack relief is therefore obtained from the finite-angle tip-frame geometry. The chordwise and normal velocity components in the folded tip frame are defined as Equation 5.48 and Equation 5.49.

$$u_T = V_\infty \cos\alpha \cos\Lambda \tag{5.48}$$

$$w_T = V_\infty \left(\sin\alpha \cos\theta - \cos\alpha \sin\theta \sin\Lambda\right) \tag{5.49}$$

The corresponding local tip angle of attack is defined as Equation 5.50.

$$\alpha_{\text{tip}}(t) = \text{atan2}(w_T, u_T) \tag{5.50}$$

The spanwise velocity component is treated as sideslip and is therefore not included in the chord-normal incidence angle.

The local angle-of-attack relief is evaluated relative to the locked local tip angle in Equation 5.51.

$$\Delta\alpha_{\text{tip,rel}}(t) = \alpha_{\text{tip,locked}} - \alpha_{\text{tip}}(t) \tag{5.51}$$

With this convention, positive values correspond to nose-down local angle-of-attack relief and therefore to local unloading of the folding wingtip.

Since only unloading generated before the locked-reference WRBM peak can reduce the critical peak load, the useful local relief is reported as Equation 5.52

$$\Delta\alpha_{\text{tip}} = \max_{t_g \le t \le t_{M,\text{locked}}} \Delta\alpha_{\text{tip,rel}}(t) \tag{5.52}$$

The time denotes the time at which the locked-reference WRBM reaches its critical gust-load-direction peak, $t_{M,\text{locked}}$, is defined Equation 5.53.

$$t_{M,\text{locked}} = \arg\min_{t_g \le t \le t_{\text{end}}} \text{WRBM}_{\text{locked}}(t) \tag{5.53}$$

The additional apparent-angle contribution associated with hinge-rate motion is included directly in the coupled UVLM simulation through the moving-surface boundary condition, but is not included in this scalar post-processing metric. The full rotation-matrix derivation is provided in Appendix B. The following additional quantities are extracted where relevant:

- Peak hinge moment and hinge-moment time history.
- If applicable, release time of the hinge.
- Settling time of the hinge motion and, where relevant, of the structural response.

These metrics are selected to reflect the central objectives of the thesis. The critical WRBM peak and its relative reduction provide a direct measure of whether the FFWT release reduces the maximum bending response in the gust-load direction. The peak hinge angle quantifies the wingtip-motion demand associated with this load alleviation. The local angle-of-attack-relief metric indicates whether the folding motion generates aerodynamic unloading at a beneficial instant. The hinge moment and release-timing metrics are included because the later campaign stages explicitly formulate release logic in terms of timing and threshold crossing. Residual post-release behaviour is assessed qualitatively from the WRBM and hinge-angle time histories.

5.6.2. Stage I: Release-timing Extension

The first stage extends the release-timing trends established in the trend comparison in Table 5.5.3. Since the baseline WRBM-based timing hierarchy has already been assessed there, Stage I investigates whether the same early- versus late-release behaviour can be interpreted more generally using frequency-based scaling parameters. The purpose of Stage I is therefore twofold.

Stage I-A: Relative Frequency

First, it examines whether the effectiveness of a given release timing can be linked with the ratio between the gust excitation frequency and the relevant structural natural frequency in Stage I-A, by scaling the gust excitation by structural frequency. The relative-frequency parameter for this stage is defined as Equation 5.54.

$$\hat{f}_g = \frac{f_g}{f_s} \tag{5.54}$$

Here, f_g is the gust excitation frequency and f_s is the structural reference frequency. This parameter is introduced to reinterpret these choices in a more general form, by characterising the excitation through its ratio to the relevant structural frequency. In the experimental campaign of Córcoles [17], the selected gust frequencies for wing B were f_g = 1.5 Hz, 3.0 Hz, 5.0 Hz and 8.0 Hz, corresponding respectively to a flapping-related case, a bending-related case, and two fixed-frequency gust cases. For wing B, the relevant modal frequencies were approximately f_{flap} = 1.33 Hz and f_{bend} = 3.03 Hz. The corresponding relative-frequency values for these cases are presented in Equation 5.55.

$$\hat{f}_{g,\text{flap}} = \frac{1.5}{1.33} \approx 1.13, \qquad \hat{f}_{g,\text{bend}} = \frac{3.0}{3.03} \approx 0.99 \tag{5.55}$$

Taking then the first bending frequency as the structural reference for the higher-frequency cases provides the relative-frequency values in Equation 5.56.

$$\hat{f}_{g,5\text{Hz}} = \frac{5.0}{3.03} \approx 1.65, \qquad \hat{f}_{g,8\text{Hz}} = \frac{8.0}{3.03} \approx 2.64 \tag{5.56}$$

For Stage I-A, f_s is defined on the basis of the modal frequencies identified in the GVT performed by Hoyng [18] for the FFWT demonstrator, since the present reduced-order Simscape model is derived from the full-order model developed for that demonstrator. The GVT results distinguish between the locked- and free-hinge configurations. For the locked hinge, Hoyng identified the first bending mode at f_s = 1.7 Hz, followed by higher modes at 8.3 Hz, 19.5 Hz, 20.3 Hz and 24.9 Hz. For the free-hinge configuration, the first mode is the flapping mode at 0.9 Hz, followed by a first bending-flapping mode at 2.3 Hz and higher modes at 15.1 Hz, 17.4 Hz and 24.4 Hz. This leads to the scaled gust frequencies for Stage I-A presented in Table 5.18. A gust frequency close to the dominant structural mode is expected to produce a stronger dynamic response and a higher sensitivity to release phase, whereas a gust frequency far from the structural frequencies is expected to behave more like an externally imposed transient load.

Table 5.18: Stage I-A: Relative-frequency-matched gust frequencies using Hoyng's GVT frequencies [18].

Case	Reference case	Target ratio $\hat{f}_g$	Present reference f_s [Hz]	Scaled gust frequency f_g [Hz]
I-A1	Córcoles flapping case	1.13	0.9	1.02
I-A2	Córcoles bending case	0.99	1.7	1.68
I-A3	Córcoles 5 Hz case	1.65	1.7	2.81
I-A4	Córcoles 8 Hz case	2.64	1.7	4.49

Each gust-frequency case is evaluated for the same release definitions: free hinge, pre-release, 0%, 50%, and 100% release. For each case, one locked-hinge baseline simulation is first performed to define the WRBM-based release times. The five release cases, namely free hinge, pre-release, 0%, 50%, and 100% release, are then evaluated using those case-specific trigger times. This gives a total of $N_{\text{sim,I-A}} = 4(1 + 5) = 24$ simulations.

Stage I-B: Reduced Frequency

The second part of Stage I reinterprets the Stage I-A gust cases in terms of aerodynamic reduced frequency. Whereas Stage I-A examines the relation between gust excitation and structural response time through $\hat{f}_g$, Stage I-B considers the aerodynamic time scale of the gust through the reduced frequency as defined in Equation 5.57.

$$k_g = \frac{\omega_g c}{2V_\infty} = \frac{\pi f_g c}{V_\infty} \tag{5.57}$$

Here, $\omega_g = 2\pi f_g$ is the gust angular frequency, c is the reference chord, and V_∞ is the freestream velocity. Low values of k_g correspond to a more quasi-steady aerodynamic forcing, whereas higher values indicate a more rapidly varying gust input and therefore a stronger role of unsteady aerodynamic effects.

The four Stage I-A gust cases are reused for this interpretation, since each already corresponds to a known gust frequency and therefore to a corresponding reduced frequency. To extend the reduced-frequency range toward a more quasi-steady regime, one additional low-frequency anchor case is added at $f_g = 0.68$ Hz. This gives the Stage I-B set listed in Table 5.19. The low-k_g anchor is used to assess whether delayed release remains more effective when the gust-induced loading develops more slowly relative to the aerodynamic convective time scale.

Table 5.19: Stage I-B reduced-frequency cases used for the aerodynamic time-scale interpretation. The Stage I-A cases are reused and one additional low-k_g anchor is added.

Case	f_g [Hz]	$\hat{f}_g$ [-]	k_g [-]
I-B0	0.68	0.40	0.047
I-A1	1.02	1.13	0.070
I-A2	1.68	0.99	0.116
I-A3	2.81	1.65	0.194
I-A4	4.49	2.64	0.310

The Stage I-B analysis therefore does not repeat a fully independent reduced-frequency matrix. Instead, it uses the Stage I-A simulations for the cross-case reduced-frequency interpretation and adds the low-frequency anchor case to improve coverage of the low-k_g regime. The additional anchor is evaluated using the same release definitions as Stage I-A: locked, free, pre-release, 0%, 50%, and 100% release. This adds six simulations to the Stage I-A set.

Combining Stage I-A with the additional low-k_g anchor case gives the total number of Stage I simulation runs in Equation 5.58.

$$N_{\text{sim,I}} = N_{\text{sim,I-A}} + N_{\text{sim,I-B0}} = 24 + 6 = 30 \tag{5.58}$$

The outcome of Stage I is a frequency-scaled extension of the trend comparison with the results presented by Córcoles. Stage I-A assesses whether the release-timing hierarchy is better interpreted through structural relative frequency, while Stage I-B assesses whether the same trends are better organised by aerodynamic reduced frequency. The outcome of this stage is used to define the reference excitation and release condition for Stage II. Once the influence of release timing and frequency scaling has been established, the next step is to hold these aspects fixed and isolate the effect of the hinge properties. Stage II therefore shifts the focus from when the wingtip is released to how the released wingtip responds, by varying hinge stiffness and damping under a common reference gust condition.

5.6.3. Stage II: Hinge-parameter Sensitivity Analysis

After establishing the governing release-timing trends in Stage I, the second stage is focused on the influence of the principal hinge parameters under a common reference excitation. It serves as the hinge-design sensitivity study of the campaign, mapping how hinge stiffness and hinge damping affect load alleviation, wingtip motion, and post-gust dynamic behaviour when the release definition is fixed.

The restriction to just the hinge stiffness and damping is deliberate, since the goal is to identify the dominant trends associated with the hinge properties that directly govern the onset, amplitude and decay of wingtip motion. By limiting the number of simultaneously varying parameters, the resulting trends remain interpretable and can subsequently be used to guide the release-strategy investigations intended in Stage III. The objectives of Stage II are defined below.

1. Quantify the influence of the principal hinge parameters on peak WRBM alleviation.
2. Determine how these parameters affect hinge-motion amplitude and post-gust oscillatory behaviour.
3. Assess whether favourable load-alleviation trends are accompanied by acceptable dynamic behaviour.
4. Identify parameter regions that are unfavourable because they lead to delayed motion, excessive oscillation, or incipient instability.

To isolate the first-order influence of hinge stiffness and damping, Stage II is performed at the reference condition selected from Stage I. The selected case is the 0% release case at $f_g = 2.81$ Hz, corresponding to $\hat{f}_g = 1.65$ and $k_g = 0.194$. This case is chosen because it provides clear peak-load alleviation while remaining sufficiently dynamic for hinge-property effects to be visible. The fixed simulation parameters for Stage II are summarised in Table 5.20.

Table 5.20: Stage II reference condition used for the hinge-parameter sensitivity study.

Parameter	Value
Angle of attack, α	5°
Freestream velocity, V_∞	10 m/s
Gust amplitude, δ_g	2.5°
Gust frequency, f_g	2.81 Hz
Relative frequency, $\hat{f}_g$	1.65
Reduced frequency, k_g	0.194
Release condition	0% release
Release time, t_{rel}	5.2 s
Flare angle, Λ	15°
Reference chord, c	0.220 m

The hinge-property sweep is defined using the effective wingtip inertia about the hinge axis. From the Simscape rigid-body properties of the wingtip, the effective inertia about the flared hinge axis is estimated as defined in Equation 5.59.

$$I_{\text{hinge}} = \mathbf{e}_h^{\mathsf{T}} \mathbf{I} \mathbf{e}_h, \qquad \mathbf{e}_h = \begin{bmatrix} \cos\Lambda \\ \sin\Lambda \\ 0 \end{bmatrix} \tag{5.59}$$

This results in $I_{\text{hinge}} \approx 0.0232\ \text{kg m}^2$. The finite hinge cases are then defined by four logarithmically spaced stiffness levels, as shown in Equation 5.60.

$$K_\theta \in \{0.3,\ 1,\ 3,\ 10\}\ \text{N m/rad} \tag{5.60}$$

This is combined with the four damping-ratio levels described in Equation 5.61.

$$\zeta_h \in \{0,\ 0.02,\ 0.10,\ 0.20\} \tag{5.61}$$

For each stiffness level, the corresponding damping coefficient is computed from Equation 5.62.

$$C_\theta = 2\zeta_h \sqrt{K_\theta I_{\text{hinge}}} \tag{5.62}$$

This definition gives the damping level a consistent physical interpretation across stiffness values. The stiffness values are selected on a logarithmic scale to resolve the transition from free-like to locked-like behaviour, while the damping-ratio sweep is used to assess how post-release oscillations are affected by increasing hinge damping.

The Stage II simulation matrix consists of sixteen finite hinge-property cases. The fully locked and fully free hinge responses are retained as bounding references, giving $N_{\text{sim,II}} = 16 + 2 = 18$ simulations. The objective of this stage is not to identify a unique optimum, but to map how the response moves between the free-hinge and locked-hinge limits as stiffness and damping are varied. The resulting response map is used to select representative hinge-property candidates for the Stage III release-strategy study.

5.6.4. Stage III: Threshold-Based Release Strategy

Stage III builds on the hinge-property trends identified in Stage II by introducing a release strategy on the basis of the hinge moment. Whereas Stage II prescribed the release instant at gust arrival and varied the post-release hinge properties, Stage III keeps the selected hinge-property candidates fixed and varies the release threshold. The purpose is to assess whether the favourable hinge-property regions identified in Stage II remain effective when the release instant is determined from a physically motivated hinge-moment criterion.

The operating condition, gust input, and FFWT geometry are kept identical to the Stage II reference condition. The gust frequency is therefore fixed at $f_g = 2.81$ Hz, corresponding to $\hat{f}_g = 1.65$ and $k_g = 0.194$. This keeps the Stage III release-strategy study directly comparable to the Stage II hinge-property study.

Rather than repeating the full hinge-property matrix, three representative hinge-property candidates are carried forward from Stage II. These are selected to represent the main response classes observed in the stiffness–damping study:

1. an *aggressive* candidate, corresponding to the most compliant retained hinge setting and strongest load-alleviation response;
2. a *balanced* candidate, corresponding to a less compliant but still load-alleviating hinge setting;
3. a *conservative* candidate, corresponding to a more restrained hinge setting with reduced hinge-angle demand.

The selected candidates are summarised in Table 6.12.

The release criterion is defined using the third component of the hinge-moment vector, $M_{h,3}$, since this is the component used by the release logic in the aeroelastic model. Because the non-zero angle of attack produces a steady pre-gust hinge load, the trigger variable is defined as the gust-induced hinge-moment increment in Equation 5.63.

$$\Delta M_{h,3}(t) = M_{h,3}(t) - M_{h,3,0} \tag{5.63}$$

where $M_{h,3,0}$ is the pre-gust baseline value of the locked-reference response. The reference hinge-moment increment is then defined as Equation 5.64.

$$\Delta M^{\text{locked}}_{h,3,\text{pk}} = \max_{t \geq t_g} \left| \Delta M^{\text{locked}}_{h,3}(t) \right| \tag{5.64}$$

The release threshold is prescribed as a fraction of this locked-reference peak increment in Equation 5.65.

$$M_{h,\text{thr}} = \eta \, \Delta M^{\text{locked}}_{h,3,\text{pk}} \tag{5.65}$$

In this equation, η is the threshold fraction. The hinge is initially locked and is released once the criterion in Equation 5.66 is met.

$$\left| \Delta M_{h,3}(t) \right| \geq M_{h,\text{thr}}, \qquad t \geq t_g \tag{5.66}$$

After release, the hinge follows the finite stiffness and damping properties of the selected Stage II candidate.

Six threshold fractions are considered, listed in Equation 5.67.

$$\eta \in \{0,\ 0.05,\ 0.10,\ 0.15,\ 0.20,\ 0.25\} \tag{5.67}$$

The $\eta = 0$ case corresponds to release at gust arrival and is retained as the candidate-specific early-release reference. Increasing η delays the release and is used to assess how sensitive each hinge-property candidate is to later triggering within the gust-load build-up.

The Stage III threshold matrix therefore consists of three hinge-property candidates and six threshold fractions, giving $N_{\text{Stage III,thr}} = 3 \times 6 = 18$ finite-hinge release-threshold simulations. The locked response is retained as the common peak-load reference for all candidates. A single fully free reference is not used as a primary Stage III reference, because the aggressive, balanced, and conservative candidates each have different post-release stiffness and damping values. Instead, the $\eta = 0$ cases provide the appropriate candidate-specific early-release references.

In addition to the threshold cases, one prescribed-release anchor is included for each candidate at $t_{\text{rel}} = 5.30$ s. This case is not interpreted as an additional hinge-moment threshold level. It is included because the non-zero thresholds are crossed within a narrow time interval in the locked-reference hinge-moment response. The $t_{\text{rel}} = 5.30$ s case therefore provides an intermediate timing reference between gust-onset (0%) release and the clustered non-zero threshold releases. Including these three prescribed-release anchors gives $N_{\text{Stage III}} = 18 + 3 + 1 = 22$ Stage III simulations, including the locked reference.

The response quantities introduced in the previous stages are retained for the Stage III assessment. The main metrics are:

- the release time associated with each threshold fraction;
- the timing relation between release, gust arrival, and the locked WRBM peak;
- the peak WRBM reduction relative to the locked reference;
- the peak hinge angle after release;
- the post-gust WRBM RMS as a measure of residual dynamic response.

The objective of Stage III is not to repeat the hinge-property optimisation of Stage II, but to test the robustness of the selected hinge-property candidates under hinge-moment-based release timing. The resulting comparison identifies whether the aggressive, balanced, and conservative hinge settings remain favourable when release is governed by the local hinge-moment response rather than by a prescribed gust-arrival release.

6

Results

This chapter presents the aeroelastic response of the numerical FFWT model for the simulation campaign defined in Chapter 5. The results are organised according to the stages of the campaign. First, Stage I examines the sensitivity of the release performance to gust excitation frequency and release timing. Stage II then studies the effect of post-release hinge stiffness and damping. Finally, Stage III evaluates hinge-moment-based release thresholds for selected hinge-property candidates.

6.1. Stage I: Release-Timing Sensitivity

Stage I starts by evaluating the effect of gust frequency and release timing on the gust-load alleviation performance of the FFWT. The response is assessed using two primary quantities: the peak WRBM reduction, which measures the load-alleviation benefit, and the peak hinge angle, which measures the wingtip-motion demand.

6.1.1. Stage I-A: Relative Frequency

Stage I-A investigates whether the release-timing behaviour observed in the validation comparison is preserved when the gust excitation is scaled using the relative-frequency parameter $\hat{f}_g = f_g/f_s$. Four gust frequencies are considered, corresponding to the scaled Córcoles reference cases introduced in Table 5.18.

The trend validation is extended by moving from direct trend reproduction at the original experimental gust frequencies to a modal-scaling interpretation. Instead of asking whether the present model reproduces the absolute experimental response, the analysis examines whether the same release-timing mechanisms persist when the gust excitation is matched by relative frequency. A fixed set of simulation parameters is used to ensure a clean comparison. The gust start time is set to $t_g = 5.2$ s, allowing the aerodynamic forces to settle before the gust is applied. All simulations are run up to $t_{\text{end}} = 13.2$ s, to allow the post-gust oscillations to decay and provide a clear overview of the gust event.

The WRBM reduction values reported below follow the peak-load metric defined in Equation 5.45, so the minimum WRBM value after gust onset is compared against the corresponding locked-hinge value. Positive values therefore indicate peak WRBM alleviation in the gust-load direction, while negative values indicate that the release case reaches a larger downward WRBM peak than the locked baseline. The peak hinge angle is used as a complementary measure of the wingtip-motion demand associated with each release strategy.

The locked-hinge baseline of each gust-frequency case is used to define the prescribed release instants. The resulting release timings are summarised in Table 6.1. As the gust frequency increases, the locked WRBM peak occurs progressively closer to gust onset, reducing the available time for released wingtip motion to influence the critical load build-up.

Table 6.1: Summary of locked-baseline WRBM trigger extraction for all Stage I-A cases.

Case	f_g [Hz]	$\hat{f}_g$ [-]	$t_{\text{rel,pre}}$ [s]	$t_{\text{rel,0\%}}$ [s]	$t_{\text{rel,50\%}}$ [s]	$t_{\text{rel,100\%}}$ [s]
I-A1	1.02	1.13	5.000	5.200	5.731	6.239
I-A2	1.68	0.99	5.000	5.200	5.579	5.823
I-A3	2.81	1.65	5.000	5.200	5.494	5.661
I-A4	4.49	2.64	5.000	5.200	5.426	5.557

I-A1: Gust Frequency $f_g = 1.02$ Hz

The first Stage I-A case corresponds to the relative-frequency-scaled flapping case. The gust frequency is $f_g = 1.02$ Hz, obtained by applying the Córcoles flapping-case relative-frequency ratio to the free-hinge flapping frequency of the present numerical model. The corresponding locked-baseline release instants are listed in Table 6.1.

The WRBM and hinge-angle time histories are shown in Figure 6.1. Because this is the lowest-frequency Stage I-A case, it provides the widest release window and is used to examine how much delayed release can remain effective when the locked load build-up is relatively slow.

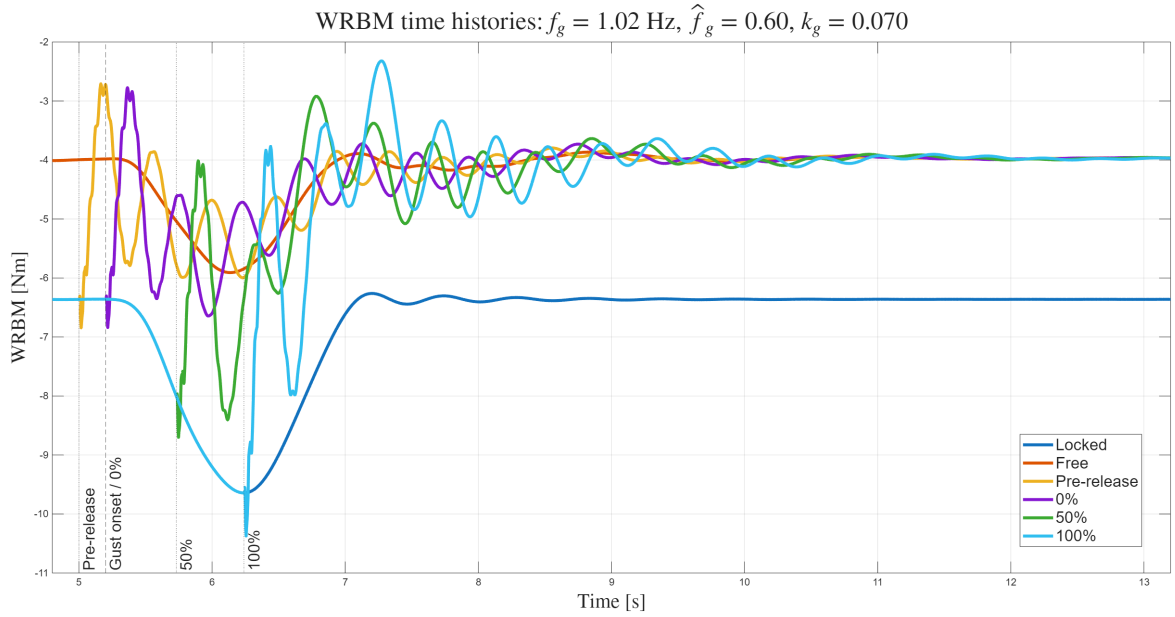

a) WRBM response.

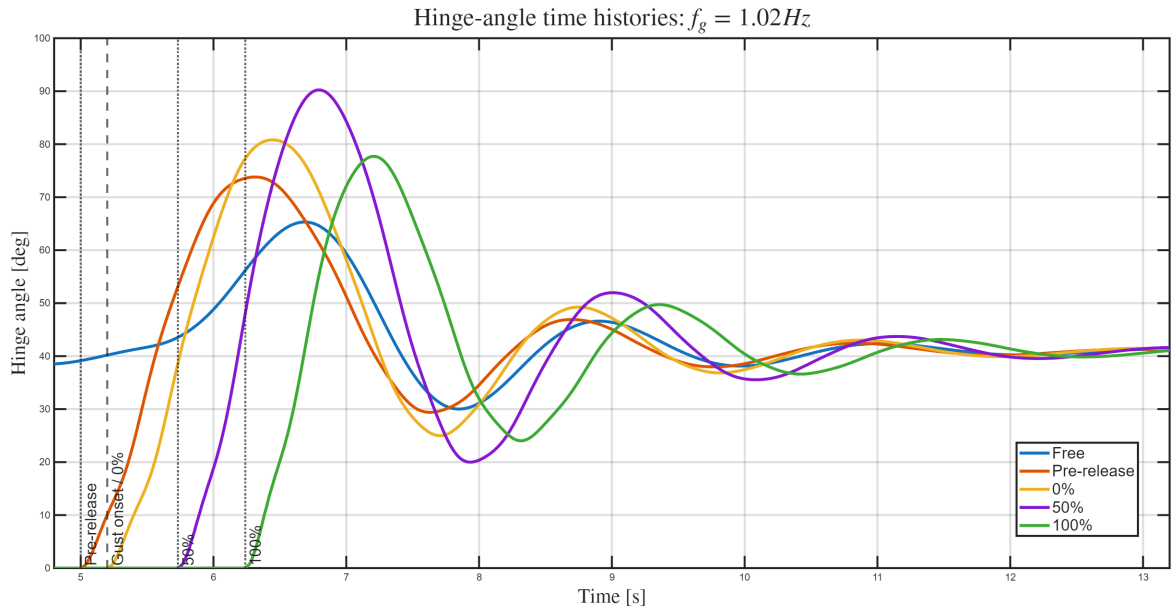

b) Hinge-angle response.

Figure 6.1: Stage I-A1 response for $f_g = 1.02$ Hz, $\hat{f}_g = 1.13$ and $k_g = 0.070$.

The corresponding peak-load and peak-motion metrics extracted from these time histories are summarised in Table 6.2.

Table 6.2: Stage I-A response summary for $f_g = 1.02$ Hz.

Case	$M_{\min}$ [Nm]	$\lvert M_{\min}\rvert$ [Nm]	Δ_{WRBM} [%]	$\theta_{\max}$ [°]
Locked	−9.644	9.644	0.0	0.0
Free	−5.908	5.908	38.7	65.31
Pre-release	−5.997	5.997	37.8	73.81
0% release	−6.841	6.841	29.1	80.82
50% release	−8.702	8.702	9.8	90.25
100% release	−10.372	10.372	−7.6	77.70

The locked configuration reaches a downward peak WRBM of −9.644 Nm at $t = 6.239$ s. The free and pre-release cases provide the largest alleviation, reducing the critical peak by 38.7% and 37.8%, respectively. Release at gust onset remains effective, although the reduction decreases to 29.1%. Delaying the release to the 50% WRBM threshold leaves only a modest reduction of 9.8%, while the 100% release becomes detrimental and increases the downward peak by 7.6% relative to the locked baseline.

The hinge-angle response in Figure 6.1b shows that the magnitude of the wingtip motion alone does not determine the alleviation performance. The 50% release produces the largest peak hinge angle, $\theta_{\max} = 90.25°$, but only achieves a 9.8% WRBM reduction because the motion starts late in the locked load build-up. Similarly, the 100% release still produces a large hinge motion, but this motion occurs after the locked critical peak has already formed. The effectiveness of the release is therefore governed primarily by phase relative to the WRBM build-up, rather than by peak hinge deflection alone.

This lowest-frequency case supports the release-window interpretation of Stage I-A. The relatively slow WRBM build-up provides enough time for the wingtip motion to develop after gust onset, which explains why the pre-release, 0% and even 50% release cases remain beneficial. However, the case also shows that a wider release window does not make all delayed releases effective. Release at the locked peak-load instant is still too late: by the time the wingtip starts moving, the critical downward WRBM peak has already developed. The subsequent hinge motion therefore acts primarily as a dynamic perturbation, increasing both the peak load and the post-gust oscillation level. This reproduces the timing mechanism identified in the validation comparison: early release provides load alleviation, delayed release progressively loses effectiveness, and release near the locked peak can become detrimental.

I-A2: Gust Frequency $f_g = 1.68$ Hz

The second Stage I-A case corresponds to the relative-frequency-scaled bending case. The gust frequency is $f_g = 1.68$ Hz, obtained by applying the Córcoles bending-case relative-frequency ratio to the locked-hinge first bending frequency of the present numerical model. The corresponding locked-baseline release instants are listed in Table 6.1.

The WRBM and hinge-angle time histories are shown in Figure 6.2. As in the previous case, the WRBM response is used to assess the peak-load alleviation, while the hinge-angle response indicates whether the wingtip motion develops early enough to influence the critical load build-up. The main peak-response quantities extracted from these time histories are collected in Table 6.3.

Table 6.3: Stage I-A2 response summary for $f_g = 1.68$ Hz.

Case	$M_{\min}$ [Nm]	$\lvert M_{\min}\rvert$ [Nm]	Δ_{WRBM} [%]	$\theta_{\max}$ [°]
Locked	−9.983	9.983	0.0	0.0
Free	−6.047	6.047	39.4	62.64
Pre-release	−6.989	6.989	30.0	80.24
0% release	−7.226	7.226	27.6	85.13
50% release	−8.972	8.972	10.1	82.35
100% release	−10.753	10.753	−7.7	70.41

The locked configuration reaches a downward peak WRBM of −9.983 Nm at $t = 5.823$ s. The free-hinge case gives the largest reduction, lowering the critical peak by 39.4%. The pre-release and 0% release cases remain effective, with reductions of 30.0% and 27.6%, respectively. The 50% release case provides only limited alleviation, reducing the peak by 10.1%, while the 100% release is detrimental and increases the downward peak by 7.7% relative to the locked baseline.

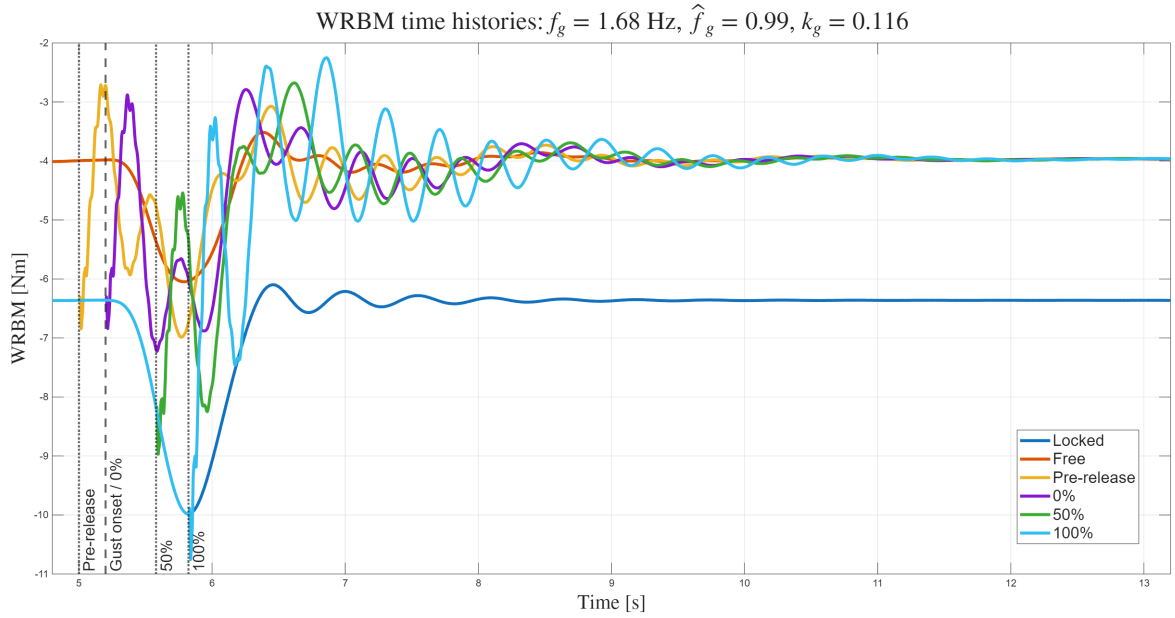

a) WRBM response.

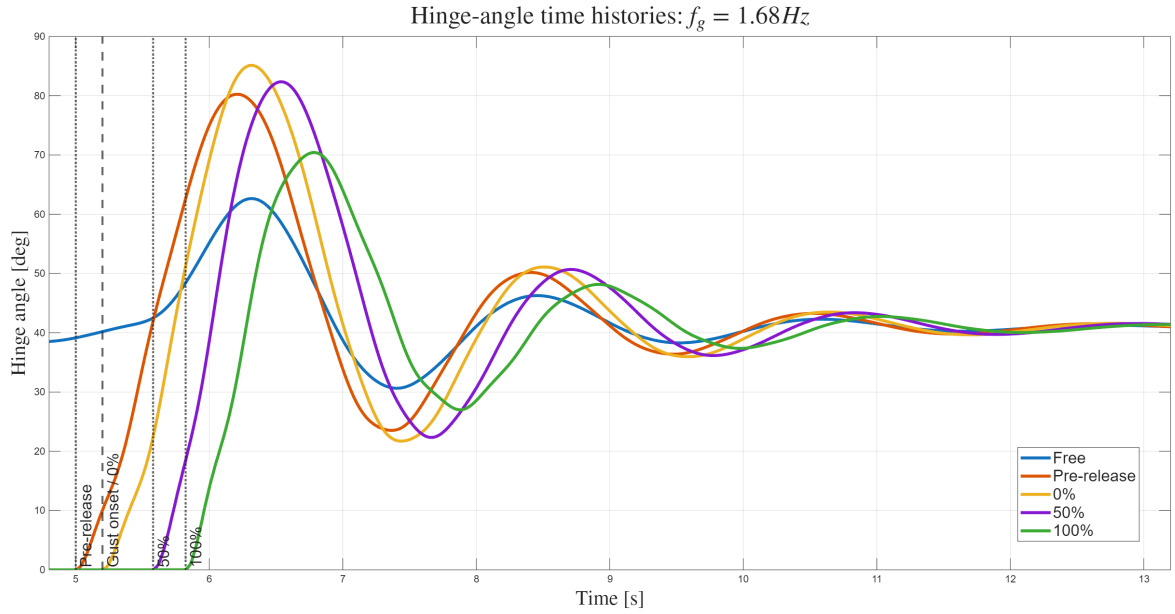

b) Hinge-angle response.

Figure 6.2: Stage I-A2 response for $f_g = 1.68$ Hz, $\hat{f}_g = 0.99$ and $k_g = 0.116$.

The hinge-angle response shows again that the peak hinge angle is not a sufficient indicator of alleviation performance. The 0% and 50% release cases reach comparable peak hinge angles, with $\theta_{\max} = 85.13°$ and $82.35°$, respectively, but their WRBM reductions differ substantially. The 0% release still allows the wingtip motion to develop during the WRBM build-up, whereas the 50% release occurs later in the locked response and therefore has less time to unload the outboard wing before the critical peak is reached. The 100% release also produces a large hinge motion, but this motion occurs after the locked peak-load instant and therefore contributes mainly to the post-release transient.

Since this case lies close to the locked first-bending frequency scaling, with $\hat{f}_g \approx 0.99$, it represents the Stage I-A case in which the gust excitation is most directly aligned with the locked first-bending reference. The response therefore provides a useful check on whether the release-window mechanism remains valid near the structural frequency scale. The results confirm this: early release remains beneficial, but the benefit decreases rapidly as the release is delayed. Compared with the lower-frequency case, the pre-release case no longer matches the free-hinge response as closely, which indicates increased sensitivity to the precise release phase. The 50% threshold release still provides peak-load alleviation, but with only a small margin, while release at the locked peak-load instant is again too late and produces both a larger critical WRBM peak and a stronger post-gust response.

I-A3: Gust Frequency $f_g = 2.81$ Hz

The third case corresponds to the scaled version of the Córcoles 5 Hz fixed-frequency case. Here, the locked-hinge first bending frequency is used as the structural reference frequency, resulting in $f_g = 2.81$ Hz. The corresponding locked-baseline release instants are listed in Table 6.1.

The WRBM and hinge-angle time histories are shown in Figure 6.3. Compared with the two lower-frequency cases, the locked WRBM peak occurs sooner after gust onset, reducing the time available for post-release wingtip motion to influence the critical load build-up.

Table 6.4: Stage I-A3 response summary for $f_g = 2.81$ Hz.

Case	$M_{\min}$ [Nm]	$\lvert M_{\min}\rvert$ [Nm]	Δ_{WRBM} [%]	$\theta_{\max}$ [°]
Locked	−10.581	10.581	0.0	0.0
Free	−6.624	6.624	37.4	60.89
Pre-release	−6.643	6.643	37.2	78.74
0% release	−8.405	8.405	20.6	78.95
50% release	−9.392	9.392	11.2	70.19
100% release	−11.449	11.449	−8.2	62.87

As summarised in Table 6.4, the locked configuration reaches a downward peak WRBM of −10.581 Nm at $t = 5.661$ s. The free-hinge and pre-release cases provide the largest peak-load alleviation, reducing the downward peak to −6.624 Nm and −6.643 Nm, respectively. This corresponds to reductions of 37.4% and 37.2%. The 0% release case remains beneficial, but the reduction decreases to 20.6%, showing that release at gust onset is already less effective for this faster gust excitation. The 50% release case gives only limited alleviation, with a reduction of 11.2%, while the 100% release case is detrimental and increases the downward peak by 8.2% relative to the locked baseline.

The hinge-angle response again shows that the timing of the motion is more important than the peak hinge angle alone. The pre-release and 0% release cases reach nearly identical peak hinge angles, with $\theta_{\max} = 78.74°$ and $78.95°$, respectively, but their WRBM reductions differ substantially. In the pre-release case, the wingtip motion is already available when the gust load starts to build up, whereas in the 0% release case the hinge motion must still develop during the shorter load-build-up interval. The 50% and 100% release cases occur even later in the locked response, so their hinge motion increasingly takes place after the critical WRBM excursion has already formed.

At $\hat{f}_g = 1.65$, the gust excitation is faster than the locked first-bending reference frequency. The available release window is therefore shorter than in the lower-frequency cases, and the response becomes more sensitive to release phase. The results support the same release-window mechanism: pre-release and free-hinge motion provide the strongest peak-load alleviation, release at gust onset remains beneficial but loses effectiveness, and delayed threshold releases provide only limited relief. The detrimental behaviour of the 100% release case confirms that releasing at the locked peak-load instant is too late to alleviate the critical transient; instead, it introduces additional hinge and structural motion after the load has already developed.

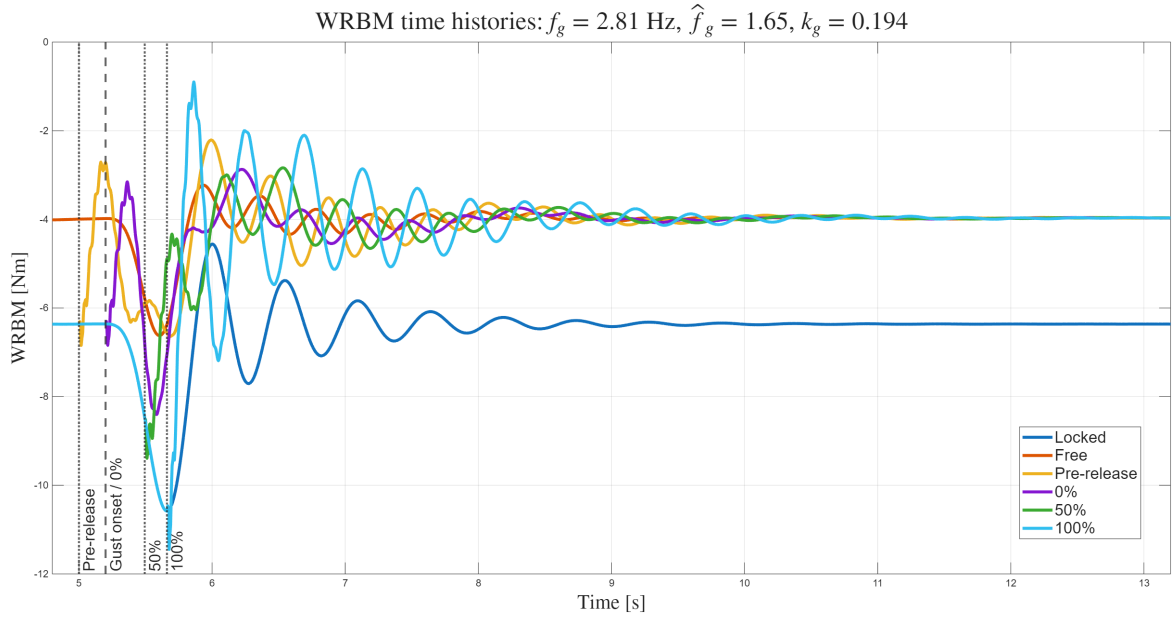

a) WRBM response.

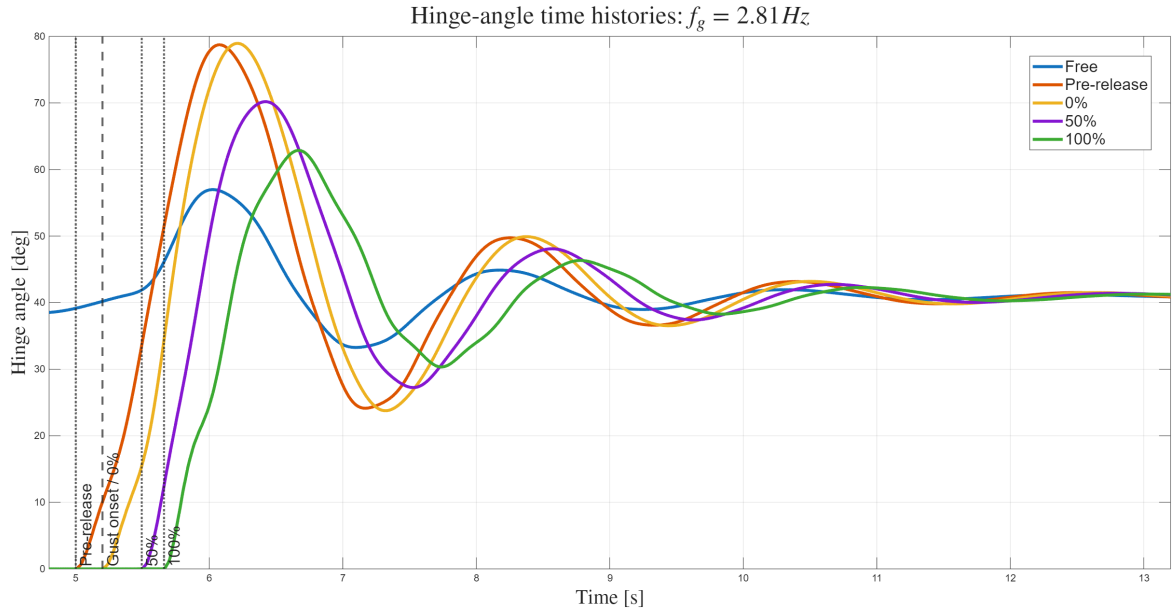

b) Hinge-angle response.

Figure 6.3: Stage I-A3 response for $f_g = 2.81$ Hz, $\hat{f}_g = 1.65$ and $k_g = 0.194$.

I-A4: Gust Frequency $f_g = 4.49$ Hz

The fourth and final case of Stage I-A corresponds to the scaled version of the Córcoles 8 Hz fixed-frequency case. With the locked-hinge first bending mode as structural reference, this gives $f_g = 4.49$ Hz. This case has the largest relative-frequency value in Stage I-A and is used to assess whether the response becomes less governed by the structural frequency scale and more by the externally imposed gust transient. The corresponding locked-baseline release instants are listed in Table 6.1.

The WRBM and hinge-angle time histories are shown in Figure 6.4. Compared with the lower-frequency cases, the locked WRBM peak occurs shortly after gust onset, so the release cases provide a direct test of whether the wingtip can move quickly enough to influence the critical load build-up.

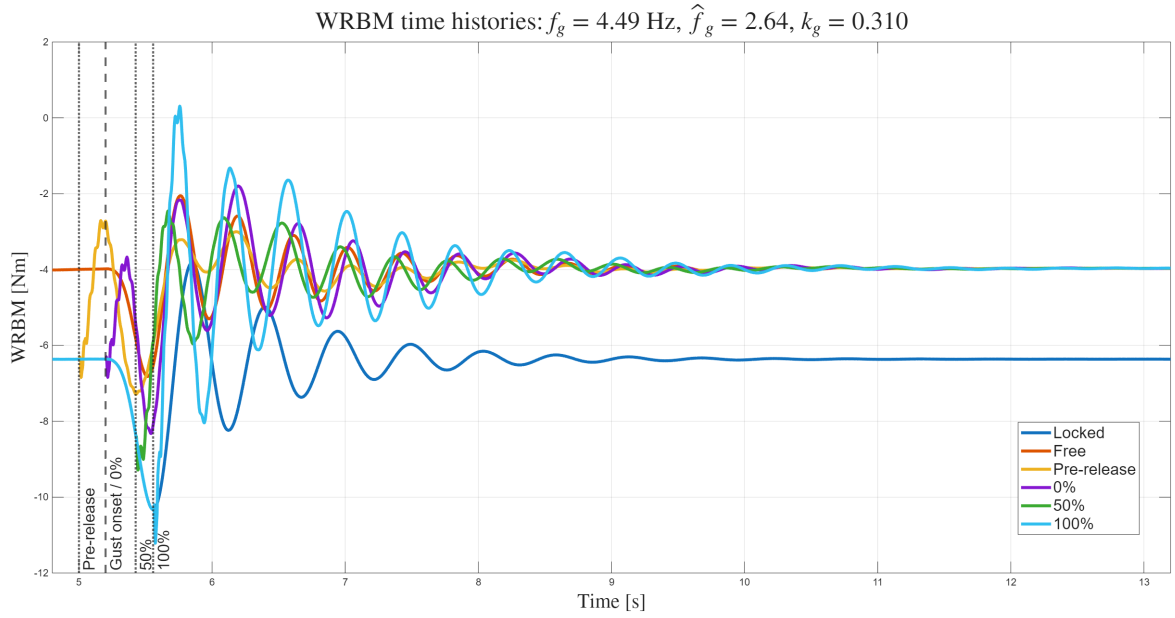

a) WRBM response.

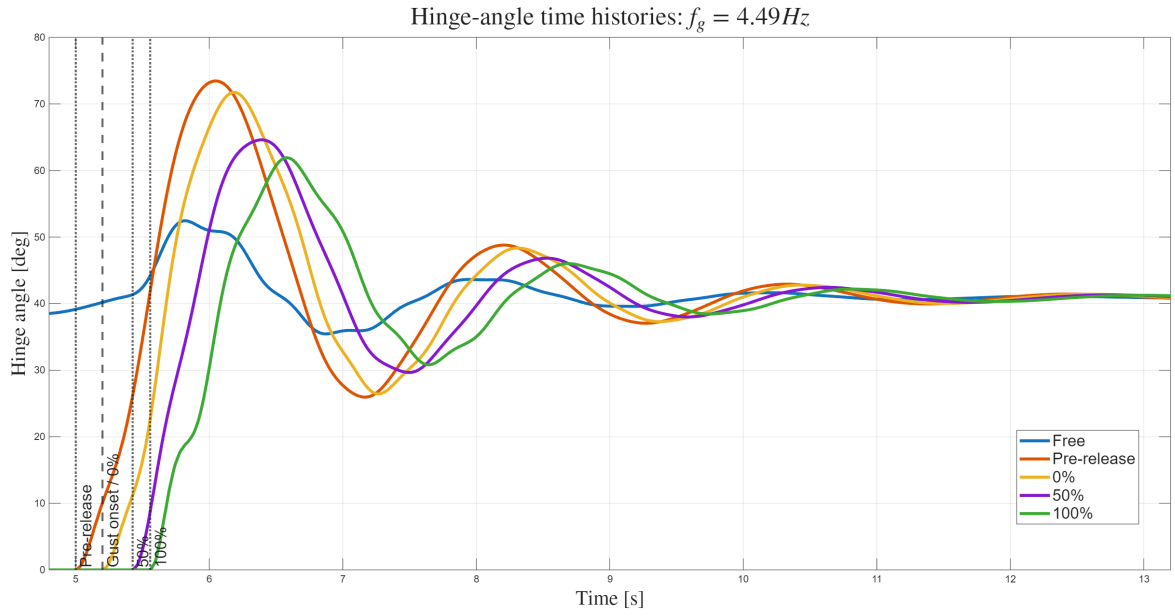

b) Hinge-angle response.

Figure 6.4: Stage I-A4 response for $f_g = 4.49$ Hz, $\hat{f}_g = 2.64$ and $k_g = 0.310$.

The critical WRBM peaks, corresponding load reductions and peak hinge angles extracted from these time histories are summarised in Table 6.5.

Table 6.5: Stage I-A4 response summary for f_g = 4.49 Hz.

Case	$M_{\min}$ [Nm]	$\lvert M_{\min}\rvert$ [Nm]	Δ_{WRBM} [%]	$\theta_{\max}$ [°]
Locked	−10.327	10.327	0.0	0.0
Free	−6.819	6.819	34.0	60.89
Pre-release	−7.288	7.288	29.4	73.45
0% release	−8.326	8.326	19.4	71.74
50% release	−9.289	9.289	10.0	64.61
100% release	−11.215	11.215	−8.6	61.92

For this case, the locked configuration reaches a downward peak WRBM of −10.327 Nm at t = 5.557 s. The free-hinge case gives the largest peak-load reduction, lowering the critical peak by 34.0%. The pre-release case gives the strongest design-relevant released response, with a reduction of 29.4%. Release at gust onset remains beneficial, but the reduction decreases to 19.4%, indicating that the hinge motion has less time to develop before the locked peak is reached. The 50% release case gives only limited alleviation, with a reduction of 10.0%, while the 100% release case is detrimental and increases the downward peak by 8.6% relative to the locked baseline.

The hinge-angle response confirms that the reduction in performance is governed mainly by release phase rather than by the magnitude of the wingtip motion alone. The pre-release and 0% release cases reach similar peak hinge angles, with $\theta_{\max}$ = 73.45° and 71.74°, respectively, but the pre-release case gives a substantially larger WRBM reduction. This indicates that, for this faster gust excitation, even release at gust onset leaves insufficient time for the hinge motion to develop fully before the critical load is formed. The 50% and 100% release cases start even later in the locked response, so their hinge motion occurs increasingly after the relevant load-build-up phase.

Consistent with the release-window interpretation, the $\hat{f}_g$ = 2.64 case is the least forgiving Stage I-A case. The useful timing window is shortest, so the wingtip must already be free, or be released very early, to provide substantial unloading during the critical WRBM build-up. Release at gust onset remains beneficial, but is less effective than in the lower-frequency cases. The 50% release still gives limited peak-load relief, while the 100% release is again too late and mainly introduces an additional post-release transient. This confirms that delayed release becomes increasingly unattractive at high relative frequency, primarily because it excites residual hinge–wing dynamics after the critical load has already developed.

Cross-Case Comparison

The cross-case comparison shows how the release-window duration changes across the Stage I-A gust-frequency cases. The release timings show that increasing the gust frequency reduces the time interval between gust onset and the WRBM-threshold crossings. This is quantified in Table 6.6, where the time from gust onset to the locked WRBM peak decreases from 1.039 s in I-A1 to 0.357 s in I-A4. The corresponding interval between the 50% threshold release and the locked WRBM peak decreases from 0.508 s to 0.131 s. Higher gust-frequency cases therefore provide a narrower window in which the wingtip can rotate early enough to unload the outboard wing before the critical WRBM peak is formed.

Since the I-A1 case uses the free-hinge flapping frequency as structural reference while I-A2–I-A4 use the locked first-bending frequency, the ordering in f_g is not identical to the ordering in $\hat{f}_g$. The release-window trend should therefore be interpreted primarily as a consequence of the faster gust excitation and the associated earlier locked WRBM peak. Consequently, the comparison between the 0%, 50% and 100% release cases is not only a threshold comparison, but also a phase-sensitivity comparison relative to the locked WRBM response.

Table 6.6: Release-window duration relative to the locked WRBM peak for the Stage I-A gust-frequency cases.

Case	f_g [Hz]	$\hat{f}_g$ [-]	$t_{\mathrm{peak,locked}} - t_{\mathrm{rel},0\%}$ [s]	$t_{\mathrm{peak,locked}} - t_{\mathrm{rel},50\%}$ [s]
I-A1	1.02	1.13	1.039	0.508
I-A2	1.68	0.99	0.623	0.244
I-A3	2.81	1.65	0.461	0.167
I-A4	4.49	2.64	0.357	0.131

Although the free-hinge case is included in the time-history comparisons as a reference for the maximum passive unloading response, it is omitted from the cross-case summary plots. Stage I-A is concerned with practically relevant locked-to-released strategies relative to the locked baseline.

Continuous free-hinge operation would not represent the intended in-flight configuration, since the wingtip would not remain locked for span efficiency outside the gust encounter. The free case is therefore retained as a physical reference, but not treated as a release strategy in the trend comparison. The resulting cross-case peak WRBM reductions are shown in Figure 6.5.

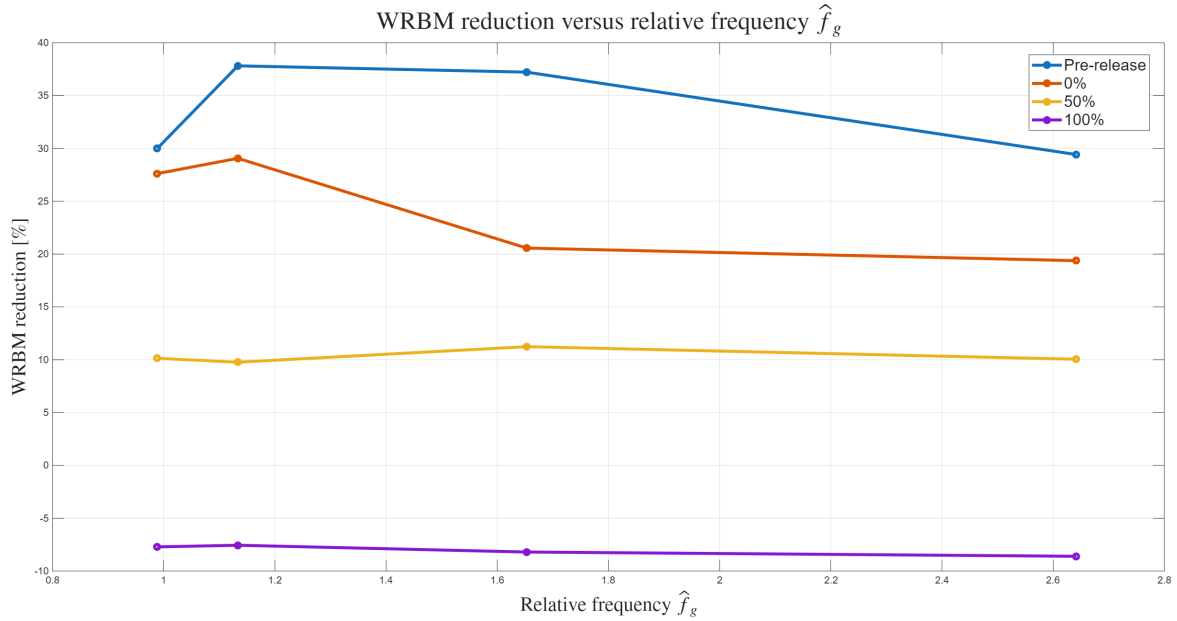

Figure 6.5: Peak WRBM reduction in the gust-load direction as a function of relative frequency for the Stage I-A release cases.

The figure shows that the release strategies separate into three groups. The pre-release and 0% release cases provide the largest design-relevant peak-load reductions, although the 0% release loses effectiveness as the gust excitation becomes faster. The 50% release case remains consistently beneficial, but only weakly, with peak-load reductions close to 10% for all four cases. The 100% release case is consistently detrimental, increasing the critical downward WRBM peak by approximately 8% relative to the locked baseline. This confirms that release at the locked peak-load instant occurs too late to contribute to the initial unloading mechanism. The corresponding peak hinge-angle trend is shown in Figure 6.6.

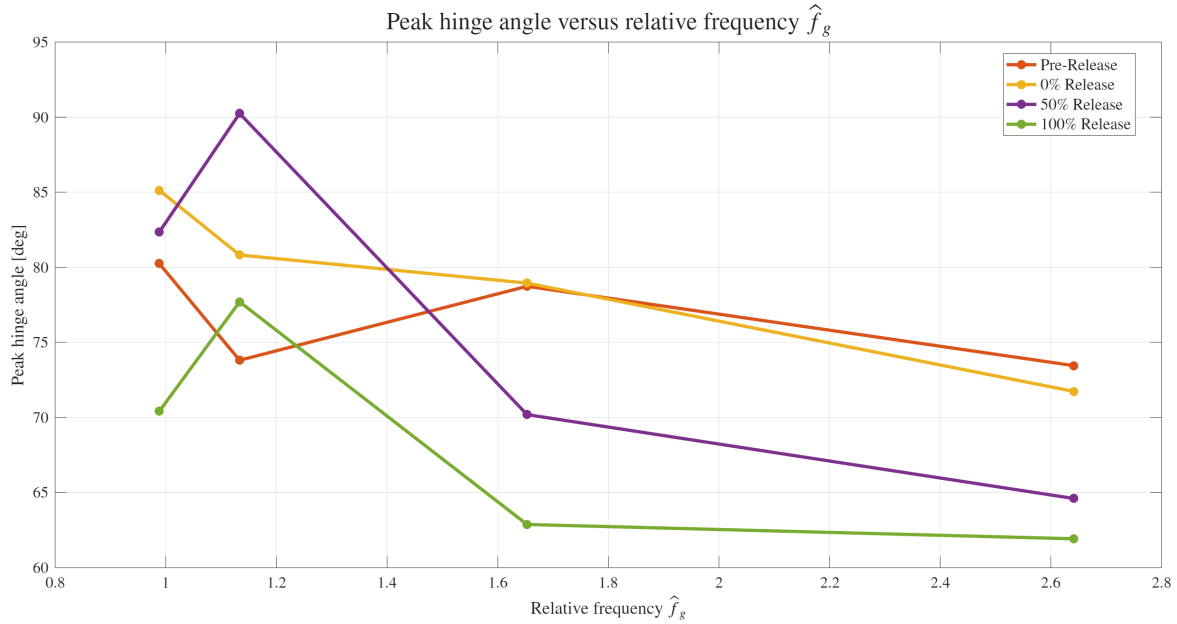

Figure 6.6: Peak hinge angle as a function of relative frequency for the Stage I-A release cases.

The figure shows that hinge-motion magnitude alone does not determine the load-alleviation performance, since several delayed-release cases generate large hinge-angle excursions while providing only limited WRBM reduction. This is because the wingtip motion occurs too late relative to the locked WRBM build-up. The cross-case comparison therefore shows that release effectiveness is governed primarily by phase, while the peak hinge angle represents the motion demand required by each strategy. The corresponding peak-load reductions are summarised in Table 6.7, which compares the prescribed release strategies across the four relative-frequency-scaled gust cases.

Table 6.7: Cross-case comparison of peak WRBM reduction in the gust-load direction for the Stage I-A release strategies.

Case	$\hat{f}_g$ [-]	Pre-release [%]	0% release [%]	50% release [%]	100% release [%]
I-A1	1.13	37.8	29.1	9.8	-7.6
I-A2	0.99	30.0	27.6	10.1	-7.7
I-A3	1.65	37.2	20.6	11.2	-8.2
I-A4	2.64	29.4	19.4	10.0	-8.6

Stage I-A conclusion

Across Stage I-A, the release-timing mechanism observed in the Córcoles trend comparison is preserved under relative-frequency scaling. Pre-release gives the strongest peak-load reduction in most cases, with Δ_{WRBM} between approximately 29% and 38%. Release at gust onset remains beneficial for all cases, but its effectiveness decreases as the available time between gust onset and the locked WRBM peak becomes shorter. The 50% release gives only modest relief of approximately 10%, while the 100% release consistently increases the critical downward WRBM peak by approximately 8%.

The results therefore support a release-window interpretation. The FFWT is effective only when the wingtip motion develops before the critical WRBM peak is formed. Large hinge-angle excursions alone do not guarantee load alleviation, because delayed release can generate substantial motion after the useful unloading window has closed. The relative-frequency study therefore indicates that release logic should be biased toward early, phase-aware release rather than peak-load-triggered release.

6.1.2. Stage I-B: Reduced Frequency

Stage I-B reinterprets the Stage I-A release-timing results in terms of the gust reduced frequency, $k_g = \frac{\pi f_g c_{\mathrm{ref}}}{V_\infty}$. Whereas Stage I-A focused on the relation between release timing and the structural response time, Stage I-B is meant to examine whether the aerodynamic time scale of the gust helps explain the local unloading mechanism at the folding wingtip. The four Stage I-A gust cases are therefore reused, but the emphasis is shifted from WRBM and hinge-angle time histories to the angle-of-attack relief of the tip-frame, generated by the folding motion.

To extend the reduced-frequency range, an additional low-frequency case with $f_g = 0.68$ Hz is included. This case acts as a low-k_g anchor and is used to assess whether the release mechanism approaches a more quasi-steady behaviour when the gust excitation is slow relative to the convective time scale. The full set of Stage I-B cases is listed in Table 6.8, combining the four Stage I-A gust cases with the additional low-frequency anchor used for the reduced-frequency interpretation.

Table 6.8: Stage I-B gust-frequency cases used for the reduced-frequency interpretation.

Case	f_g [Hz]	$\hat{f}_g$ [-]	k_g [-]
I-B0	0.68	0.40	0.047
I-A1	1.02	1.13	0.070
I-A2	1.68	0.99	0.116
I-A3	2.81	1.65	0.194
I-A4	4.49	2.64	0.310

Low Reduced-Frequency Anchor: $f_g = 0.68$ Hz

The additional $f_g = 0.68$ Hz case extends the Stage I-A dataset toward the low-k_g range. With $k_g = 0.047$, this case lies close to the quasi-steady boundary and is therefore used to assess how the release mechanism behaves when the gust input develops slowly relative to the convective time scale. The WRBM and hinge-angle time histories are presented in Figure 6.7. The locked WRBM develops over an even longer time interval than in the Stage I-A cases, so the prescribed threshold releases occur with a wider physical time margin before the locked peak is reached.

The WRBM response shows that the delayed release cases remain more effective than would be expected for the higher-k_g cases. The corresponding peak-response metrics are summarised in Table 6.9.

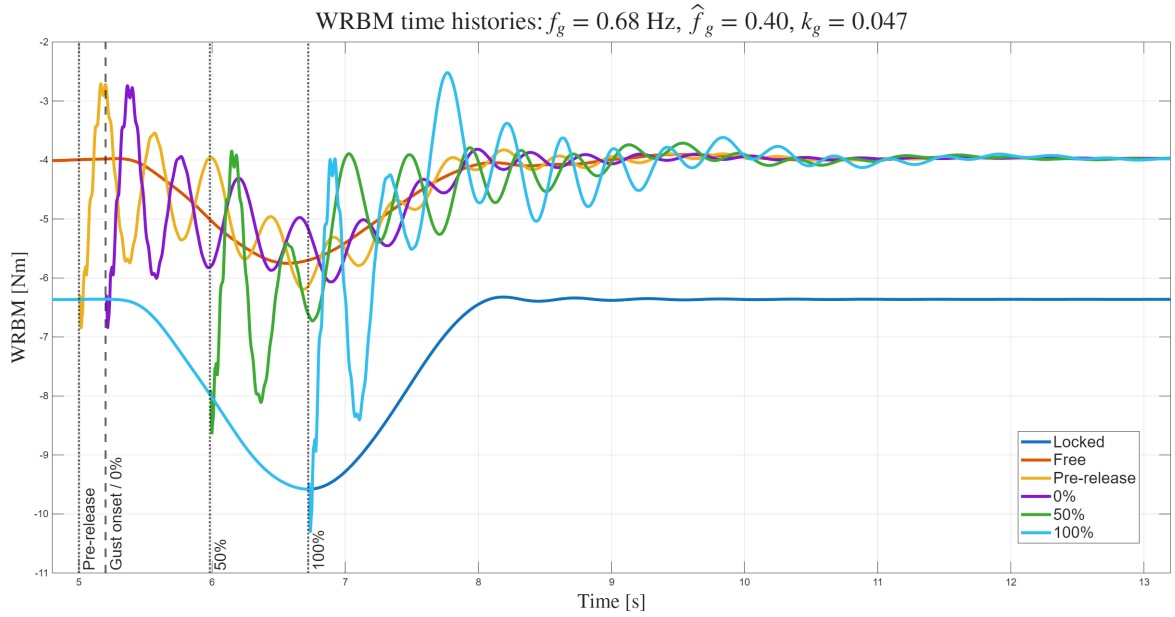

a) WRBM response.

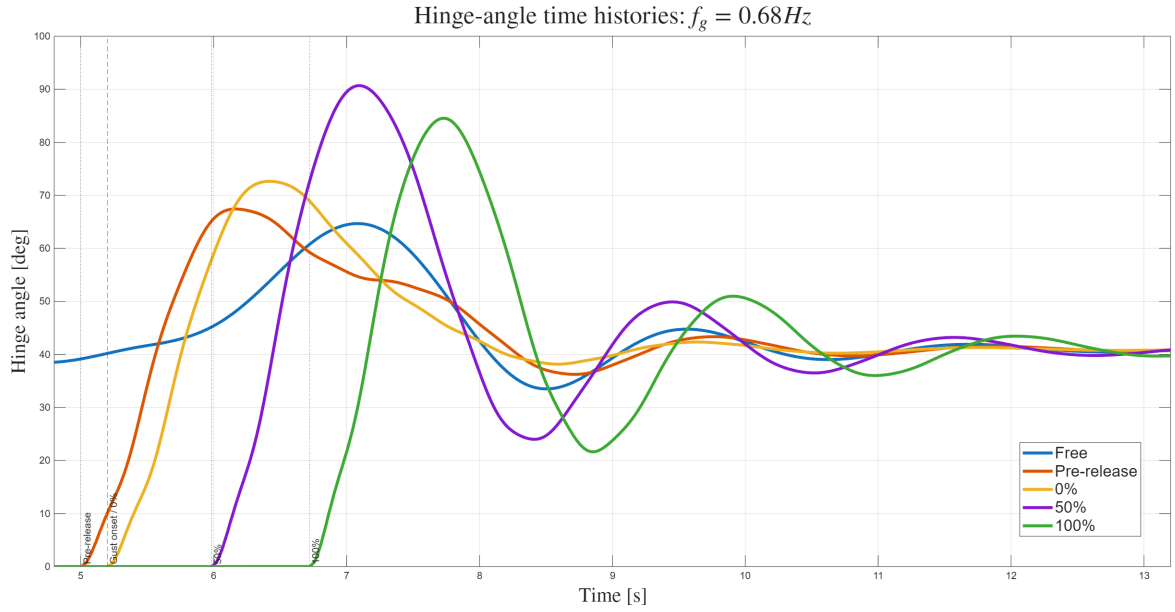

b) Hinge-angle response.

Figure 6.7: Stage I-B low reduced-frequency anchor response for $f_g = 0.68$ Hz, $\hat{f}_g = 0.40$ and $k_g = 0.047$: (a) WRBM response and (b) hinge-angle response.

Table 6.9: Stage I-B low reduced-frequency anchor response summary for $f_g = 0.68$ Hz.

Case	$M_{\min}$ [Nm]	Δ_{WRBM} [%]	$\theta_{\max}$ [°]	$\Delta\alpha_{\mathrm{tip}}$ [°]
Locked	−9.579	0.00	0.36	0.00
Free	−5.751	39.96	64.68	15.91
Pre-release	−6.190	35.38	67.42	17.18
0% release	−6.841	28.58	72.64	18.06
50% release	−8.627	9.94	90.69	18.06
100% release	−10.298	−7.51	84.54	0.00

This follows directly from the slower load build-up: even when the release is triggered at a WRBM threshold, the wingtip still has time to rotate and generate local unloading before the critical load is fully developed. The hinge-angle response confirms this interpretation, since the delayed cases produce substantial wingtip motion before the response has fully decayed into the post-gust oscillatory regime.

Local Angle-of-Attack Relief

The local unloading mechanism is assessed using the finite-angle tip-frame angle-of-attack relief metric defined in Subsection 5.6.1. This metric reports positive values when the folding motion produces nose-down local incidence relief, and only the relief generated before the locked-reference WRBM peak is treated as useful for peak-load alleviation. The selected time histories in Figure 6.8 show that the reduced-frequency effect is expressed mainly through the timing of the local unloading response.

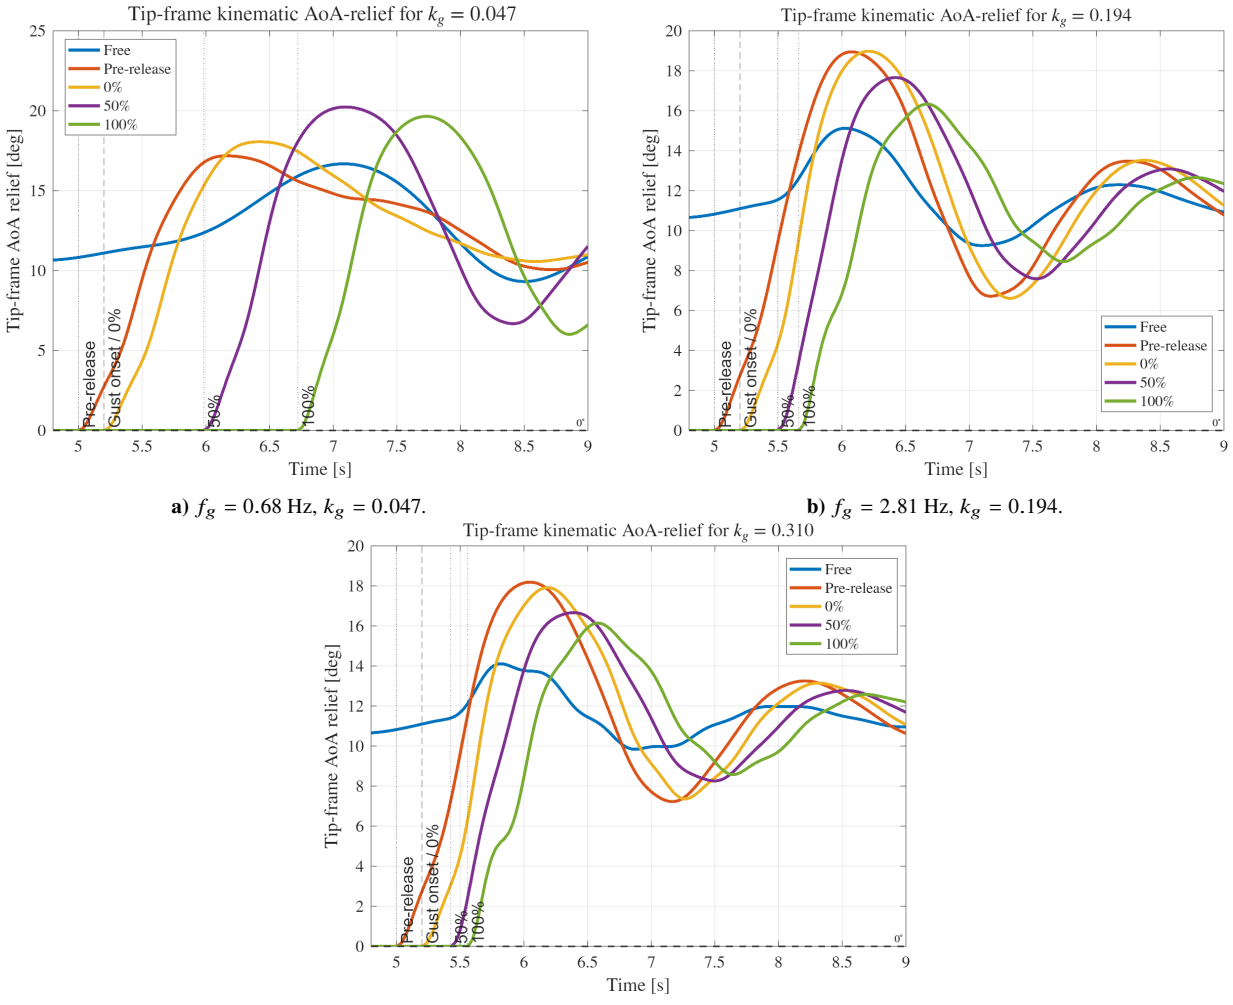

a) $f_g = 0.68$ Hz, $k_g = 0.047$.

b) $f_g = 2.81$ Hz, $k_g = 0.194$.

c) $f_g = 4.49$ Hz, $k_g = 0.310$.

Figure 6.8: Selected tip-frame angle-of-attack relief time histories for Stage I-B.

At low reduced frequency, the AoA-relief curves develop over a longer time interval, so even delayed release cases can still generate positive relief before the locked WRBM peak is reached. At intermediate and high reduced frequency, the relief curves become more compressed in time, and delayed cases increasingly generate their largest relief after the useful pre-peak window has already closed.

The comparison also shows that the largest local AoA relief does not necessarily correspond to the largest WRBM reduction. Some delayed cases generate substantial local AoA relief in the full time history, but only the relief generated before the locked WRBM peak contributes directly to peak-load alleviation. This supports the Stage I-A interpretation: the relevant mechanism is not only how much local unloading the folding motion can generate, but whether that unloading is available early enough to reduce the critical root load.

Reduced-Frequency Trend Comparison

The WRBM peak-reduction trends have already been discussed in Stage I-A. The reduced-frequency comparison is therefore used here to examine the local unloading mechanism through the useful tip-frame angle-of-attack relief defined in Subsection 5.6.1. The resulting trend is shown in Figure 6.9.

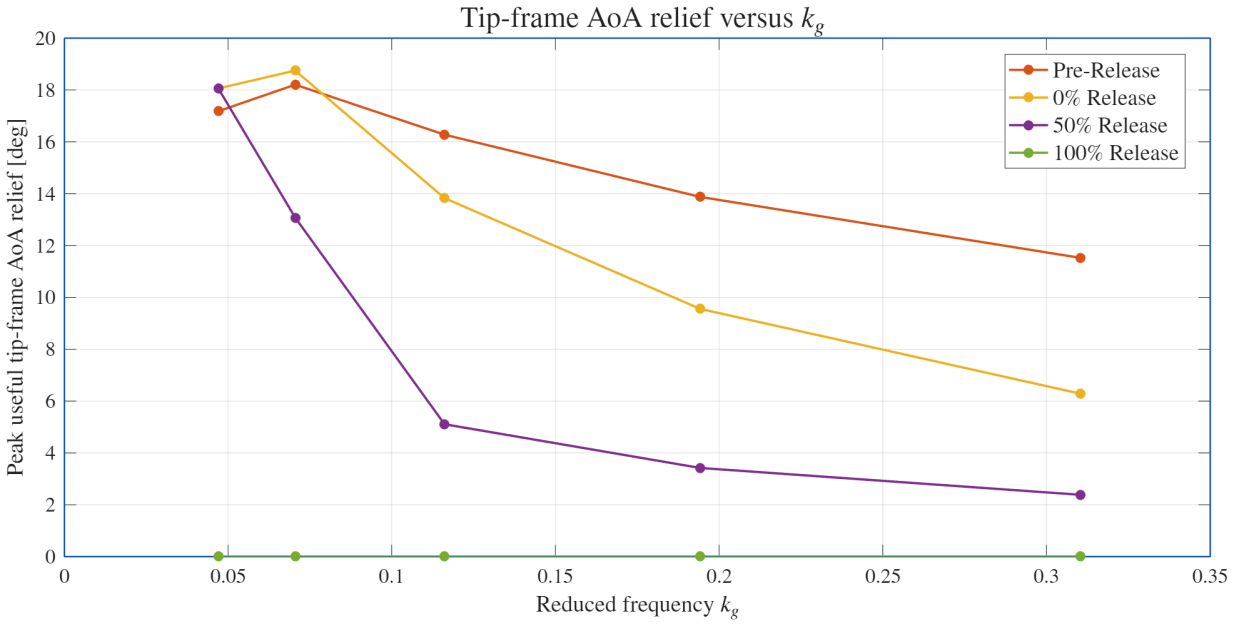

Figure 6.9: Useful tip-frame angle-of-attack relief as a function of reduced frequency for the Stage I-B release cases. The reported value corresponds to the maximum positive relief generated before the locked-reference WRBM peak.

The plot shows that the useful local AoA relief is governed by both reduced frequency and release timing. The pre-release and 0% release cases produce the largest useful relief, while the 50% release loses useful relief rapidly as k_g increases. The 100% release case remains close to zero useful relief across the full range, because release at the locked-reference WRBM peak leaves no time for unloading to develop before that peak forms.

This explains why reduced frequency complements, rather than replaces, the relative-frequency interpretation of Stage I-A. The reduced frequency characterises the aerodynamic time scale of the gust, while the release-window analysis determines whether the folding motion is phased correctly with the structural load build-up. The FFWT mechanism is therefore governed by both the amount of local unloading generated by the tip and the time at which this unloading becomes available.

Stage I-B Conclusion

Stage I-B shows that the reduced-frequency interpretation provides additional insight into the local unloading mechanism of the FFWT. The added $f_g = 0.68$ Hz case extends the dataset toward the low-k_g range, where the gust-induced loading develops more gradually. In this case, delayed releases can remain partly effective because the wingtip still has sufficient time to generate useful local AoA relief before the critical WRBM peak is formed. At higher reduced frequency, the available unloading window becomes shorter, so delayed release produces less useful pre-peak relief and contributes less effectively to peak-load alleviation.

The Stage I results are therefore used to define the reference condition for the hinge-parameter study in Stage II. The selected case is the 0% release at $f_g = 2.81$ Hz, corresponding to $\hat{f}_g = 1.65$ and $k_g = 0.194$. This case represents a practical triggered-release condition rather than an idealised free-hinge or pre-released configuration, while still providing clear peak WRBM alleviation. It also lies outside the near-quasi-steady low-k_g range, so the response remains sufficiently dynamic for assessing the influence of post-release hinge stiffness and damping.

6.2. Stage II: Hinge-Parameter Study

Stage II investigates how the post-release hinge stiffness and hinge damping affect the FFWT response. The gust input and release timing are fixed using the reference condition selected from Stage I, so that the influence of the hinge properties can be isolated from the release-timing effects studied previously.

The selected reference condition is the 0% release case at $f_g = 2.81$ Hz, corresponding to $\hat{f}_g = 1.65$ and $k_g = 0.194$, placing it outside the near-quasi-steady low-k_g range considered in Stage I-B. The response therefore remains sufficiently dynamic for the effects of hinge stiffness and damping to be meaningful. Stage II consequently keeps the gust input, release instant, and geometric parameters fixed, and varies only the post-release hinge properties. The fixed simulation parameters used throughout Stage II are listed in Table 6.10.

Table 6.10: Fixed simulation parameters for Stage II.

Parameter	Symbol	Value
Angle of attack	α	5°
Freestream velocity	V_∞	10 m/s
Gust amplitude	δ_g	2.5°
Gust frequency	f_g	2.81 Hz
Relative frequency	$\hat{f}_g$	1.65
Reduced frequency	k_g	0.194
Release definition	–	0% release
Release time	t_{rel}	5.2 s
Flare angle	Λ	15°
Reference chord	c	0.220 m
Hinge-axis inertia	I_{hinge}	0.0232 kg m^2

The finite hinge-property matrix used in Stage II follows the simulation campaign defined in Subsection 5.6.3. It consists of four hinge-stiffness levels, $K_\theta \in \{0.3,\ 1,\ 3,\ 10\}$ N m/rad, combined with four damping-ratio levels, $\zeta_h \in \{0,\ 0.02,\ 0.10,\ 0.20\}$. The fully locked and fully free hinge responses are retained as bounding references.

The objective of Stage II is not to identify a single optimum hinge setting, but to map the transition between free-like and locked-like behaviour as stiffness and damping are varied. The results are presented through representative time histories, stiffness–damping heatmaps, and a final trade-off interpretation used to select the Stage III candidate cases.

Representative Time-History Response

The full set of simulations was evaluated using the same metrics. The present section therefore focuses on representative time histories and summary maps, since these are sufficient to identify the dominant stiffness–damping trends. The selected figures are chosen to isolate the two main effects: the transition from free-like to locked-like behaviour with increasing stiffness, and the influence of damping at fixed stiffness.

Figure 6.10 shows the stiffness sweep at fixed damping ratio $\zeta_h = 0.10$. The WRBM response shows that the lowest stiffness case remains closest to the free-hinge response and gives the strongest reduction of the critical downward WRBM peak. As K_θ is increased, the response moves progressively toward the locked-hinge limit. This transition is even clearer in the hinge-angle response: low stiffness allows large wingtip rotation, whereas high stiffness strongly restricts the post-release motion.

The stiffness sweep therefore shows the primary Stage II trade-off. A compliant hinge is beneficial for peak-load alleviation because it allows the wingtip to rotate and generate local unloading during the gust encounter. However, the same compliance also produces large hinge-angle excursions. A stiff hinge avoids large tip motion, but it also removes the mechanism responsible for the load reduction.

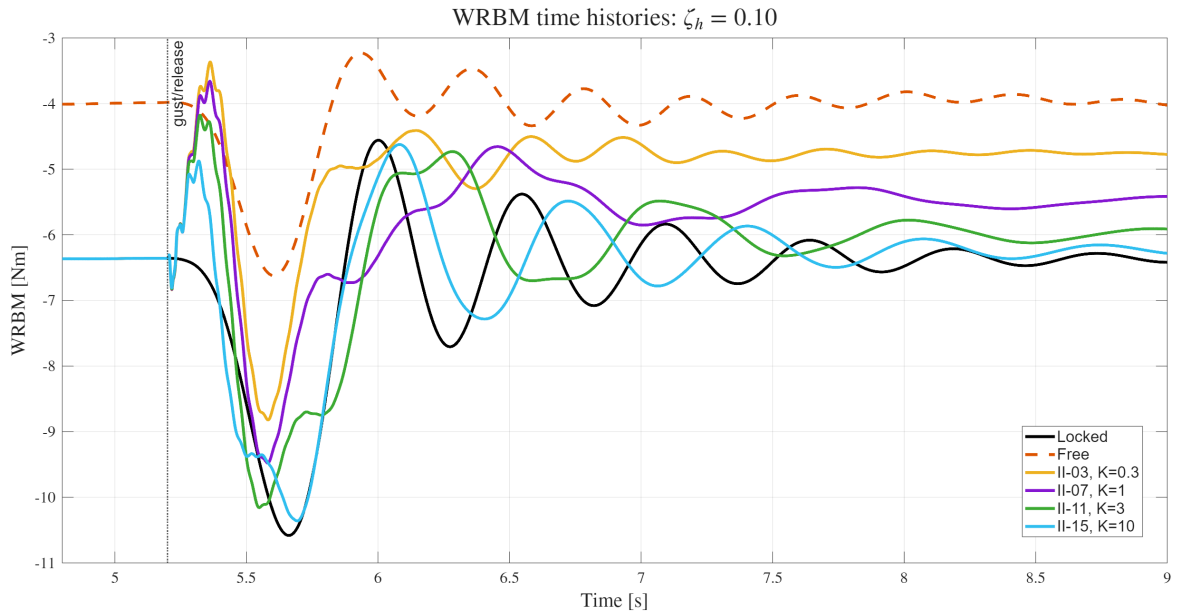

a) WRBM response.

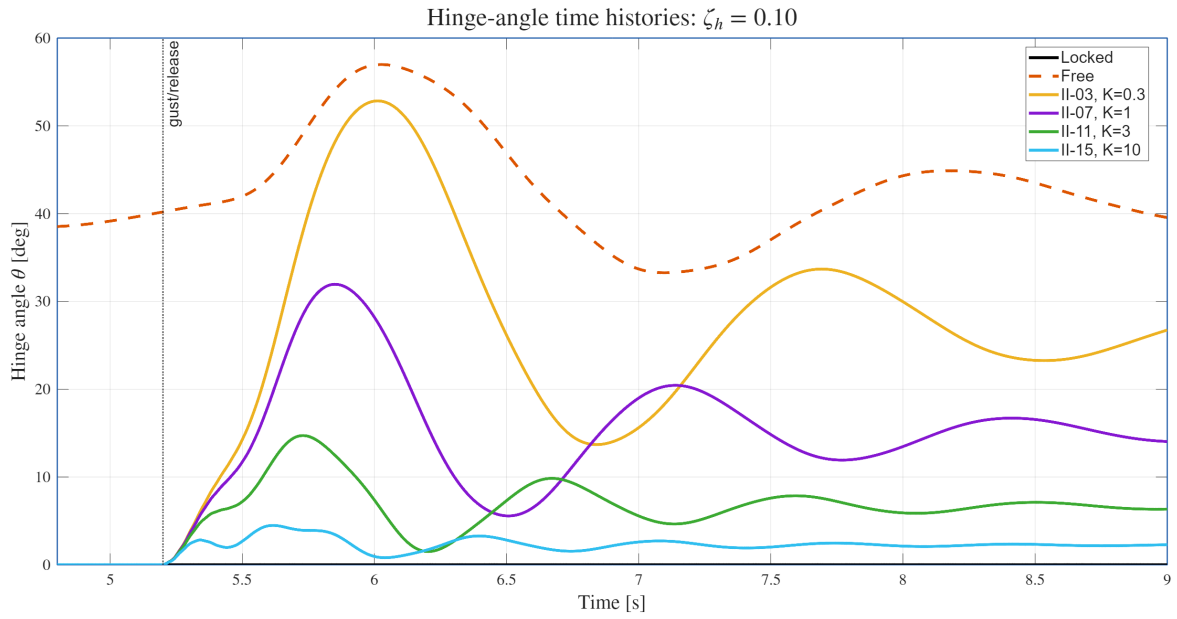

b) Hinge-angle response.

Figure 6.10: Representative Stage II time histories showing the influence of hinge stiffness at fixed damping ratio $\zeta_h = 0.10$: (a) WRBM response and (b) hinge-angle response. The locked and free responses are included as bounding references.

Figure 6.11 shows the damping sweep at fixed hinge stiffness $K_\theta = 1$ N m/rad. The first WRBM peak is only weakly affected by damping, while the later oscillatory response and hinge-angle decay are more visibly modified. Increasing ζ_h reduces the persistence of the hinge motion, but it also slightly limits the folding response. Damping should therefore not be interpreted only as a stabilising addition: it improves the post-release dynamic quality, but can also reduce part of the motion that produces local aerodynamic unloading.

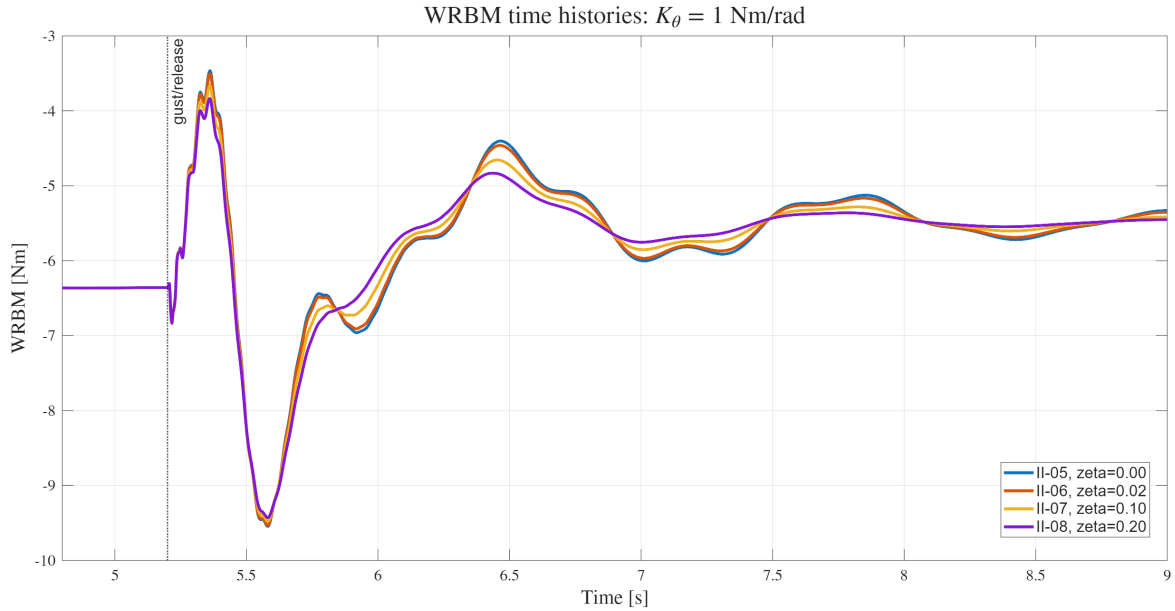

a) WRBM response.

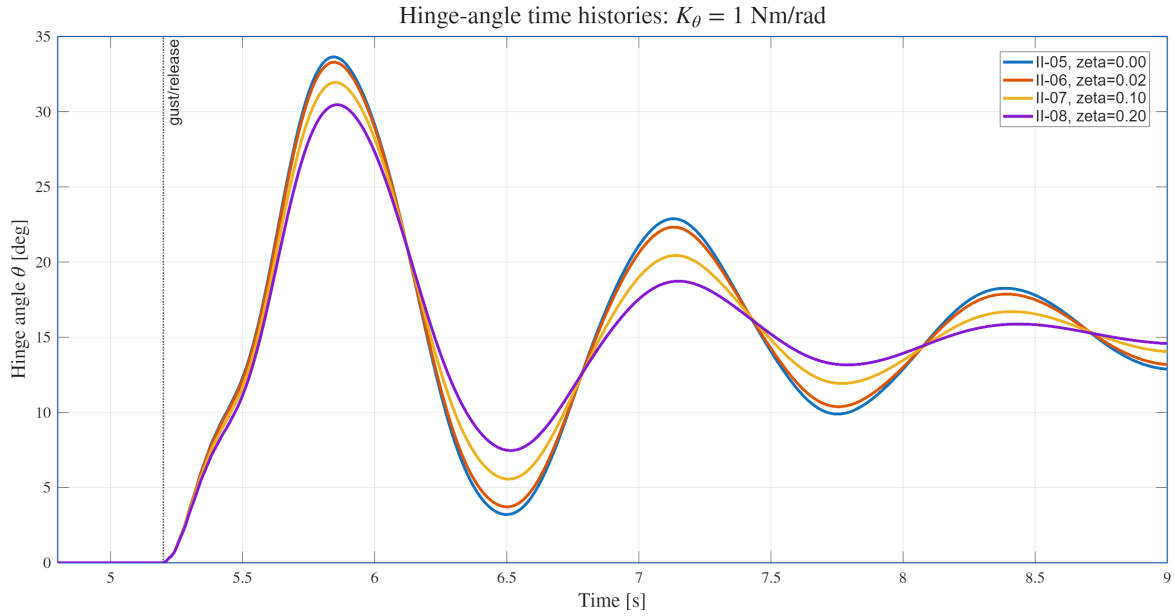

b) Hinge-angle response.

Figure 6.11: Representative Stage II time histories showing the influence of hinge damping at fixed hinge stiffness $K_\theta = 1$ N m/rad: (a) WRBM response and (b) hinge-angle response.

The damping sweep therefore confirms that damping mainly modifies the post-release dynamic response rather than the first WRBM peak. Increasing ζ_h reduces the persistence of the hinge motion, but it also slightly limits the folding motion that contributes to local aerodynamic unloading. The damping level must therefore be interpreted together with the peak-load reduction, peak hinge angle, and qualitative post-release time-history behaviour.

Hinge-property Maps

The full Stage II matrix is summarised using two primary metrics: peak WRBM reduction and peak hinge angle. The WRBM reduction quantifies the load-alleviation benefit relative to the locked reference, while the peak hinge angle quantifies the wingtip-motion demand required to obtain this benefit.

The peak-load performance of the full hinge-property matrix is evaluated in Figure 6.12 using the WRBM reduction metric introduced previously. Positive values indicate that the finite hinge case reduces the critical WRBM excursion relative to the locked reference, while negative values indicate a more severe response than the locked case.

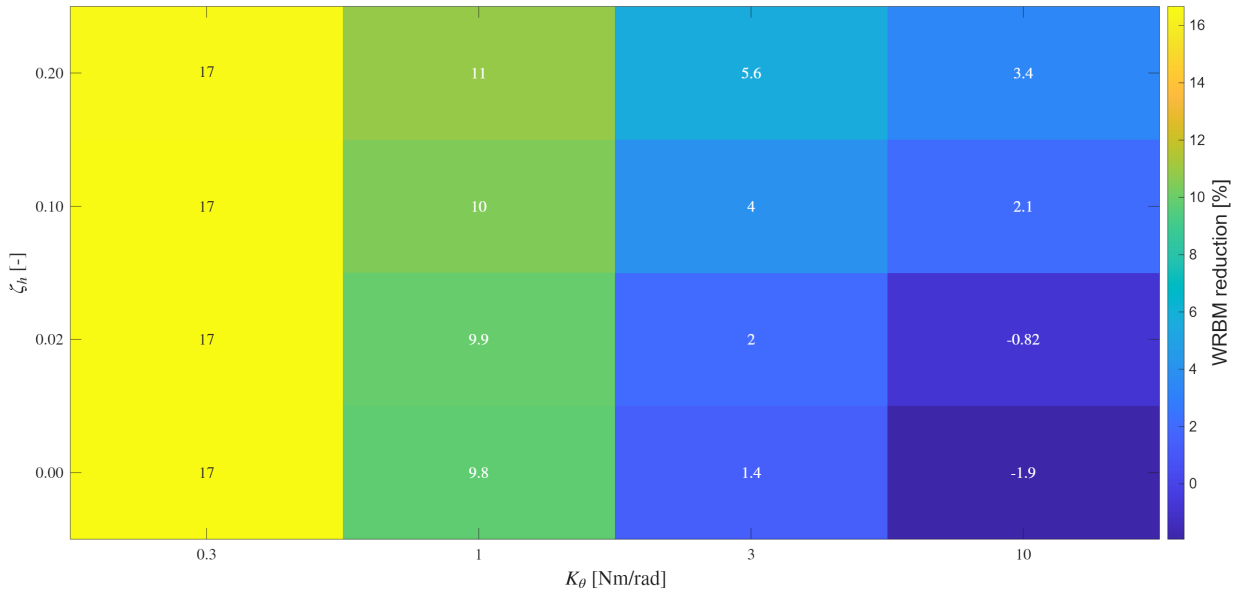

Figure 6.12: Stage II peak WRBM reduction over the hinge stiffness–damping plane. Positive values indicate alleviation relative to the locked reference.

The heatmap shows a clear stiffness dependence in the peak-load metric. The lowest stiffness level, $K_\theta = 0.3\,\mathrm{N\,m/rad}$, gives the largest WRBM reduction, approximately 17%, and is almost insensitive to damping. At $K_\theta = 1\,\mathrm{N\,m/rad}$ the reduction decreases to approximately 10–11%. For $K_\theta = 3\,\mathrm{N\,m/rad}$ the peak-load benefit becomes small, while the stiffest cases at $K_\theta = 10\,\mathrm{N\,m/rad}$ approach the locked-hinge limit and can even become slightly detrimental at low damping. To interpret this trend further, the peak hinge-angle map is presented in Figure 6.13.

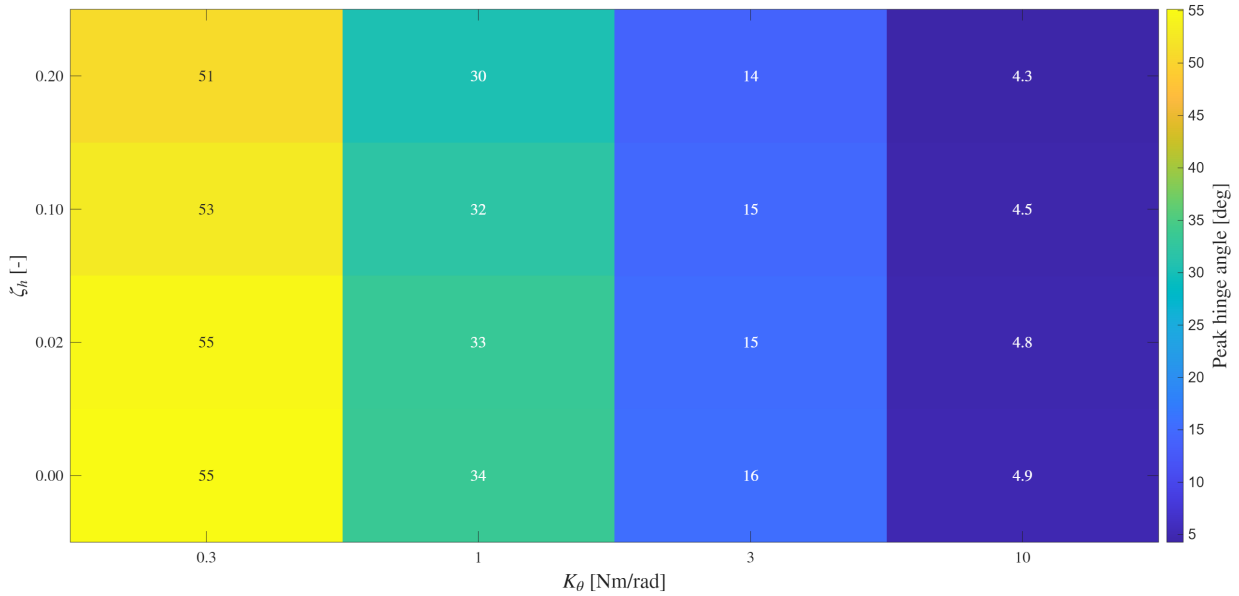

Figure 6.13: Stage II peak hinge angle over the hinge stiffness-damping plane.

Consistent with the Stage I interpretation, the dynamic-response maps show why the WRBM reduction map cannot be used on its own. The peak hinge-angle map decreases strongly with stiffness. The most compliant cases reach hinge angles of roughly 51–55°, while the $K_\theta = 10\,\mathrm{N\,m/rad}$ cases remain below approximately 5°. Damping reduces the peak hinge angle only moderately compared with the effect of stiffness. Therefore, stiffness is the dominant parameter controlling the amplitude of the wingtip motion.

Overall, the heatmaps show that increasing damping generally improves the qualitative post-release response, but does not change the dominant stiffness trend. The best peak-load performance is obtained at low stiffness, whereas the smallest hinge motion is obtained at high stiffness. The intermediate stiffness range does not automatically provide the best compromise, because $K_\theta = 3\,\mathrm{N\,m/rad}$ produces relatively low load alleviation while still allowing a non-negligible hinge response.

Stage II Conclusion

The Stage II results define a trade-off rather than a single optimum. To make this trade-off explicit, Table 6.11 lists one representative case for each stiffness level. The highest damping ratio, $\zeta_h = 0.20$, is used for this comparison because it gives the dynamically cleanest response within each stiffness column and therefore provides a consistent basis for comparing the effect of stiffness.

Table 6.11: Representative Stage II hinge-property trade-off cases at $\zeta_h = 0.20$.

Case	K_θ [Nm/rad]	ζ_h [-]	Δ_{WRBM} [%]	$\theta_{\max}$ [°]
II-04	0.3	0.20	17.0	51
II-08	1.0	0.20	11.0	30
II-12	3.0	0.20	5.6	14
II-16	10.0	0.20	3.4	4.3

The representative cases show that hinge stiffness is the dominant parameter controlling the compromise between peak-load alleviation and hinge-motion demand. The most compliant hinge, $K_\theta = 0.3\,\mathrm{N\,m/rad}$, gives the strongest WRBM reduction, but requires a large peak hinge angle of approximately 51°. Increasing the stiffness to $K_\theta = 1\,\mathrm{N\,m/rad}$ reduces the WRBM reduction to approximately 11%, while the peak hinge angle remains of the same order. At $K_\theta = 3\,\mathrm{N\,m/rad}$, the hinge-angle demand is reduced considerably, but the WRBM reduction drops to 5.6%. The stiffest case, $K_\theta = 10\,\mathrm{N\,m/rad}$, behaves increasingly like the locked-hinge reference: the hinge motion is small, but the load-alleviation mechanism is largely suppressed.

Hinge damping improves the post-release dynamic response, but cannot recover the load-alleviation benefit once the hinge becomes too stiff. A useful hinge-property selection must therefore balance WRBM reduction against hinge-angle magnitude, while retaining acceptable qualitative post-release behaviour in the time histories. Based on this trade-off, three representative hinge-property candidates are selected for the Stage III release-threshold study, as listed in Table 6.12.

Table 6.12: Stage II hinge-property candidates selected for the Stage III release-threshold study.

Stage III candidate	Stage II case	K_θ [Nm/rad]	K_{free} [Nm/deg]	C_{free} [Nms/deg]
Aggressive	II-04	0.3	0.005236	0.000582
Balanced	II-08	1.0	0.017453	0.001063
Conservative	II-12	3.0	0.052360	0.001842

The aggressive setting retains the strongest load-alleviation behaviour, the balanced setting provides a less compliant but still effective response, and the conservative setting reduces the hinge-angle demand more strongly. The $K_\theta = 10\,\mathrm{N\,m/rad}$ case is not retained because it behaves too close to the locked-hinge reference and provides limited load alleviation. These three candidates are carried forward into Stage III to assess whether the identified hinge-property behaviour remains robust when the release threshold is varied around the fixed $f_g = 2.81$ Hz gust condition.

6.3. Stage III: Release Strategy

Stage III investigates the sensitivity of the selected FFWT hinge-property candidates to the release strategy. In Stage II, the release instant was fixed at the 0% condition, corresponding to gust arrival at $t = 5.2$ s, while the post-release hinge properties were varied. Stage III reverses this emphasis: the hinge properties are kept fixed, while the release threshold is varied.

The gust input, freestream condition, geometry and reference frequency parameters remain identical to the Stage II reference condition. The selected gust case is therefore still $f_g = 2.81$ Hz, with $\hat{f}_g = 1.65$ and $k_g = 0.194$. This keeps the release-strategy study directly comparable to the Stage II hinge-property results and ensures that changes in the response can be attributed to the release logic rather than to a change in gust input.

The three hinge-property candidates selected at the end of Stage II are carried forward unchanged. These are the aggressive, balanced and conservative candidates listed in Table 6.12. For each candidate, the release threshold is varied using

$$\eta \in \{0,\ 0.05,\ 0.10,\ 0.15,\ 0.20,\ 0.25\}.$$

The locked response is retained as the common peak-load reference, while the $\eta = 0$ cases provide candidate-specific early-release references. The release thresholds are defined from the incremental locked-reference hinge moment, using the formulation introduced in Subsection 5.6.4. Before comparing the candidate responses, the release-threshold definition is first examined using the locked-reference hinge-moment response. This is shown in Figure 6.14.

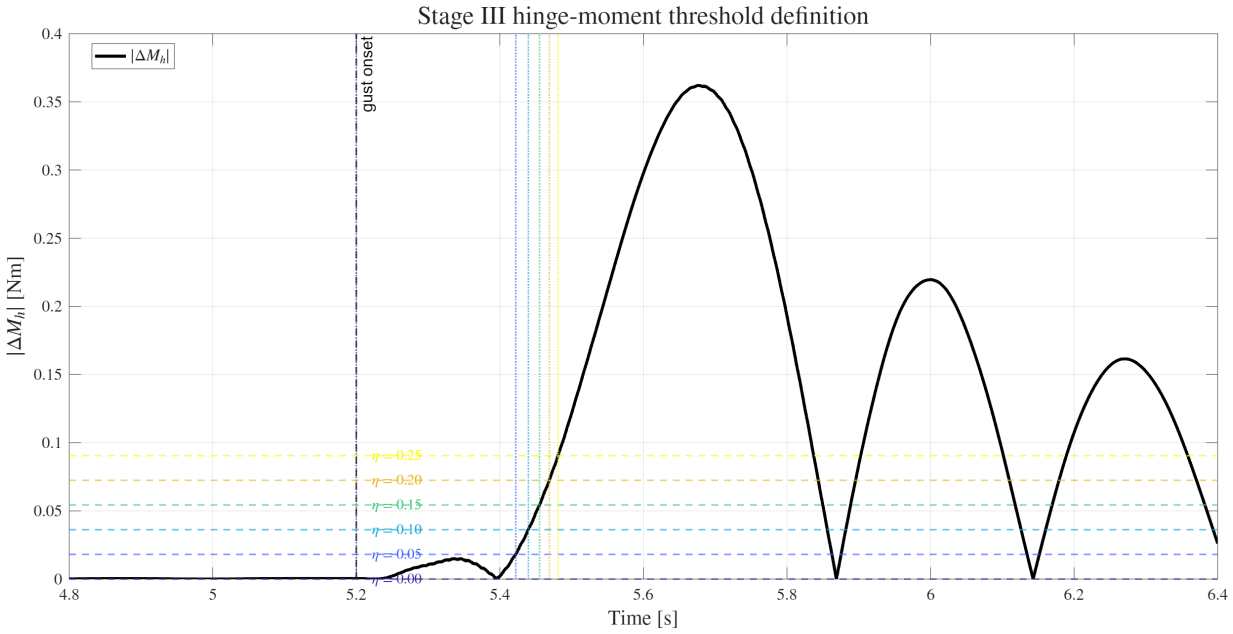

Figure 6.14: Definition of the Stage III release thresholds based on the locked-reference incremental hinge moment.

The release times extracted from this locked-reference hinge moment are summarised in Table 6.13. As expected, increasing the threshold fraction delays release relative to gust onset. The resulting release times therefore provide the phase reference for interpreting the WRBM and hinge-angle responses.

Table 6.13: Stage III release times obtained from the locked-reference $M_{h,3}$ threshold criterion.

η [-]	0	0.05	0.10	0.15	0.20	0.25
$M_{h,\text{thr}}$ [Nm]	0.000	0.018	0.036	0.054	0.072	0.091
t_{rel} [s]	5.200	5.422	5.440	5.455	5.469	5.481
$t_{\text{rel}} - t_g$ [s]	0.000	0.222	0.240	0.255	0.269	0.281

The important observation from Figure 6.14 and Table 6.13 is that the non-zero threshold cases do not produce a broad release-timing sweep. While the $\eta = 0$ case releases at gust onset, the $\eta = 0.05$-0.25 cases release between 5.422 s and 5.481 s, giving a total spread of only 0.059 s. This clustering occurs because the locked-reference $|\Delta M_{h,3}|$ response first shows a small initial excursion, then returns toward a lower value, and subsequently crosses all selected threshold levels on the same steep rising branch.

The Stage III threshold sweep should therefore be interpreted as a local sensitivity study around the hinge-moment-triggered release window, rather than as a broad release-timing study. This distinction is important, because the non-zero threshold cases are all released substantially later than the gust-onset (0%) case, but only slightly separated

from each other. To bridge the larger timing gap between $\eta = 0$ and the first non-zero threshold case, an additional prescribed-release anchor is included at $t_{rel} = 5.30$ s for each hinge-property candidate. This anchor is not treated as an additional hinge-moment threshold level; it is used only to assess whether the response changes primarily between gust-onset release and the clustered threshold-triggered releases.

Time-Histories

Representative time histories are shown for the balanced candidate in Figure 6.15. The aggressive and conservative candidates show the same qualitative threshold-timing behaviour as the balanced case. Their responses are therefore summarised through the threshold-sensitivity metrics rather than repeated as full time-history plots.

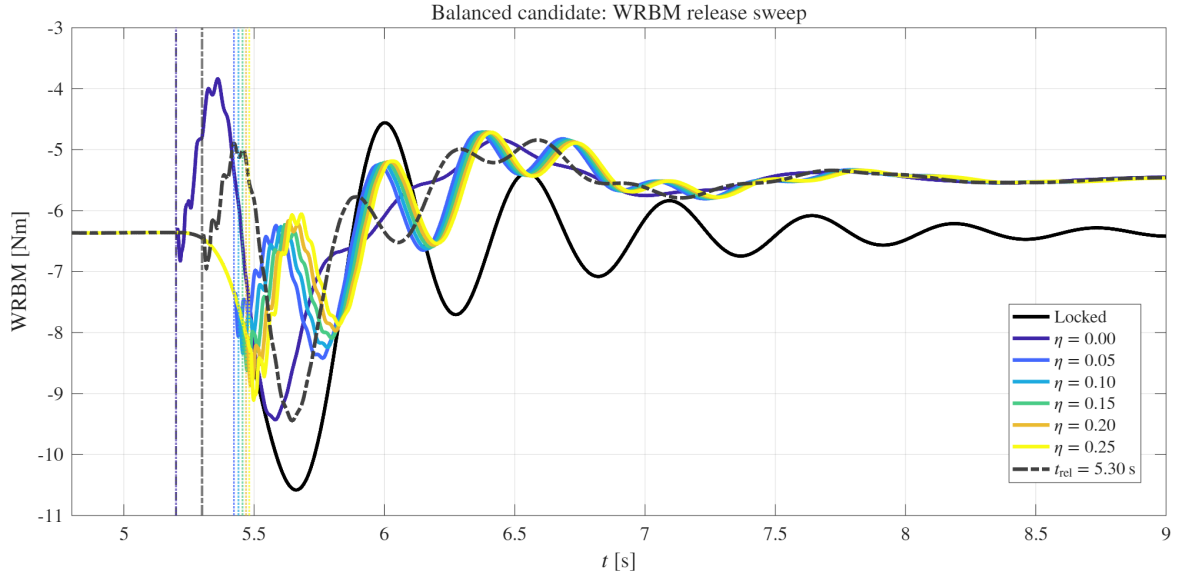

a) WRBM response.

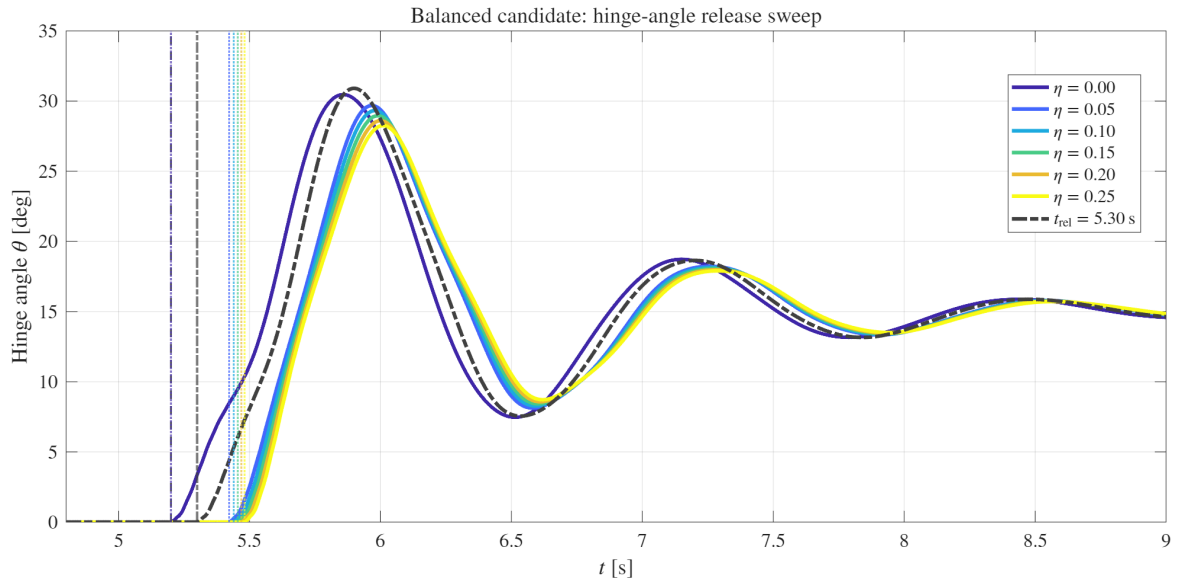

b) Hinge-angle response.

Figure 6.15: Representative Stage III threshold sweep for the balanced hinge-property candidate: (a) WRBM response and (b) hinge-angle response.

The balanced response shows that the threshold-triggered cases do not simply interpolate between gust-onset release and the locked response. The $\eta = 0$ case gives a WRBM reduction of 10.9%, while the additional prescribed anchor at $t_{rel} = 5.30$ s gives a similar reduction of 10.7%. In contrast, the first non-zero threshold cases give a much larger reduction: $\eta = 0.05$ gives 20.4% and $\eta = 0.10$ gives 21.2%. The corresponding hinge-angle peaks remain close to

30°, so the improved WRBM reduction is not caused by a larger hinge-angle excursion. Instead, it is caused by a more favourable release phase relative to the developing WRBM peak.

For thresholds above $\eta = 0.10$, the WRBM reduction decreases progressively, from 18.6% at $\eta = 0.15$ to 13.9% at $\eta = 0.25$. At the same time, the peak hinge angle decreases slightly from 29.0° to 28.3°, while the time histories indicate a more pronounced post-release transient. This indicates that later release reduces the available time for beneficial unloading before the critical WRBM peak, while still exciting residual structural response after release. The balanced candidate therefore has its best peak-load performance for a low but non-zero threshold, rather than for release exactly at gust onset.

The same mechanism is observed for the aggressive candidate, but with larger hinge motion and stronger load alleviation. The aggressive $\eta = 0$ case reduces the WRBM peak by 16.6%, while the $\eta = 0.05$ threshold gives the largest reduction in the full Stage III matrix, equal to 23.8%. Increasing the threshold beyond this point progressively reduces the benefit, with the reduction falling to 13.8% at $\eta = 0.25$. The peak hinge angle remains large for all aggressive cases, decreasing only from approximately 51° for $\eta = 0$ to 45° for $\eta = 0.25$. The aggressive candidate therefore remains the most effective load-alleviation setting, but it does so at the cost of the largest wingtip motion.

The conservative candidate behaves differently. For this case, increasing the threshold improves the WRBM reduction from 5.6% at $\eta = 0$ to 10.9% at $\eta = 0.25$. This trend shows that the more restrained hinge does not benefit from very early release in the same way as the aggressive and balanced candidates. The additional prescribed anchor at $t_{\text{rel}} = 5.30$ s is particularly ineffective for the conservative candidate, giving only 1.5% WRBM reduction. The conservative hinge therefore requires release closer to the main load build-up to generate useful unloading, but even then its maximum load reduction remains below that of the aggressive and balanced candidates. The full set of peak-response metrics for the three hinge-property candidates is summarised in Table 6.14.

Table 6.14: Stage III response summary for the aggressive, balanced, and conservative hinge-property candidates.

Candidate	Release case	t_{rel} [s]	Δ_{WRBM} [%]	$\theta_{\max}$ [°]
Aggressive	$\eta = 0$	5.200	16.6	50.8
Aggressive	$t_{\text{rel}} = 5.30$ s	5.300	18.0	49.6
Aggressive	$\eta = 0.05$	5.422	23.8	47.2
Aggressive	$\eta = 0.10$	5.440	21.1	46.7
Aggressive	$\eta = 0.15$	5.455	18.5	46.3
Aggressive	$\eta = 0.20$	5.469	16.1	45.9
Aggressive	$\eta = 0.25$	5.481	13.8	45.5
Balanced	$\eta = 0$	5.200	10.9	30.5
Balanced	$t_{\text{rel}} = 5.30$ s	5.300	10.7	30.9
Balanced	$\eta = 0.05$	5.422	20.4	29.7
Balanced	$\eta = 0.10$	5.440	21.2	29.3
Balanced	$\eta = 0.15$	5.455	18.6	29.0
Balanced	$\eta = 0.20$	5.469	16.2	28.6
Balanced	$\eta = 0.25$	5.481	13.9	28.3
Conservative	$\eta = 0$	5.200	5.6	14.1
Conservative	$t_{\text{rel}} = 5.30$ s	5.300	1.5	15.0
Conservative	$\eta = 0.05$	5.422	6.1	14.5
Conservative	$\eta = 0.10$	5.440	7.4	14.3
Conservative	$\eta = 0.15$	5.455	8.6	14.0
Conservative	$\eta = 0.20$	5.469	9.8	13.8
Conservative	$\eta = 0.25$	5.481	10.9	13.5

Threshold-Sensitivity Metrics

The full Stage III matrix is summarised in Figure 6.16 using two primary metrics: peak WRBM reduction and peak hinge angle. These metrics show that the threshold effect depends strongly on the selected hinge-property candidate.

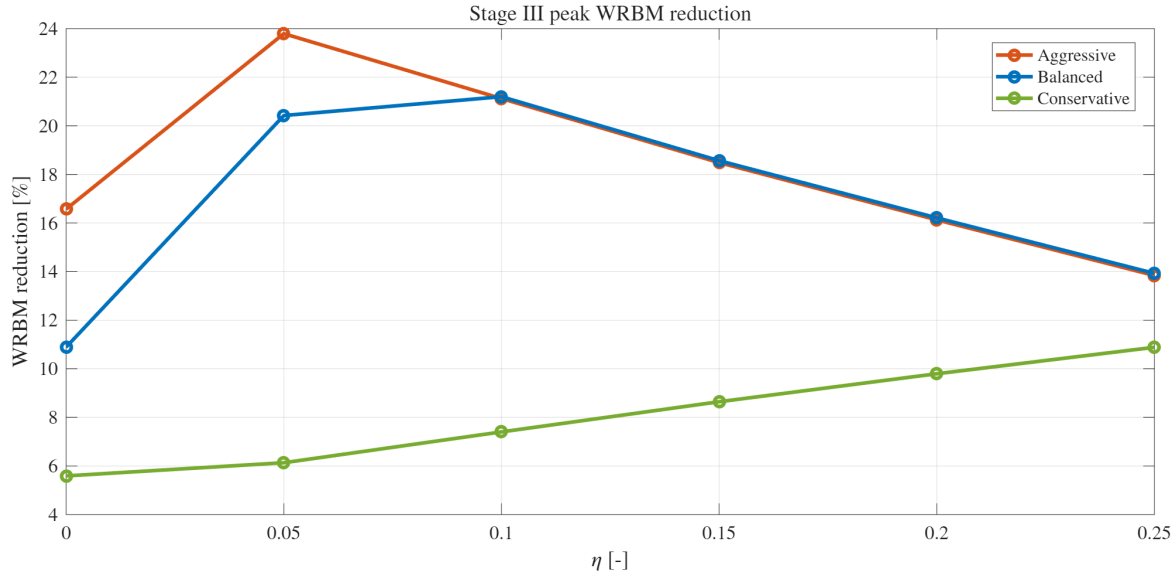

a) Peak WRBM reduction.

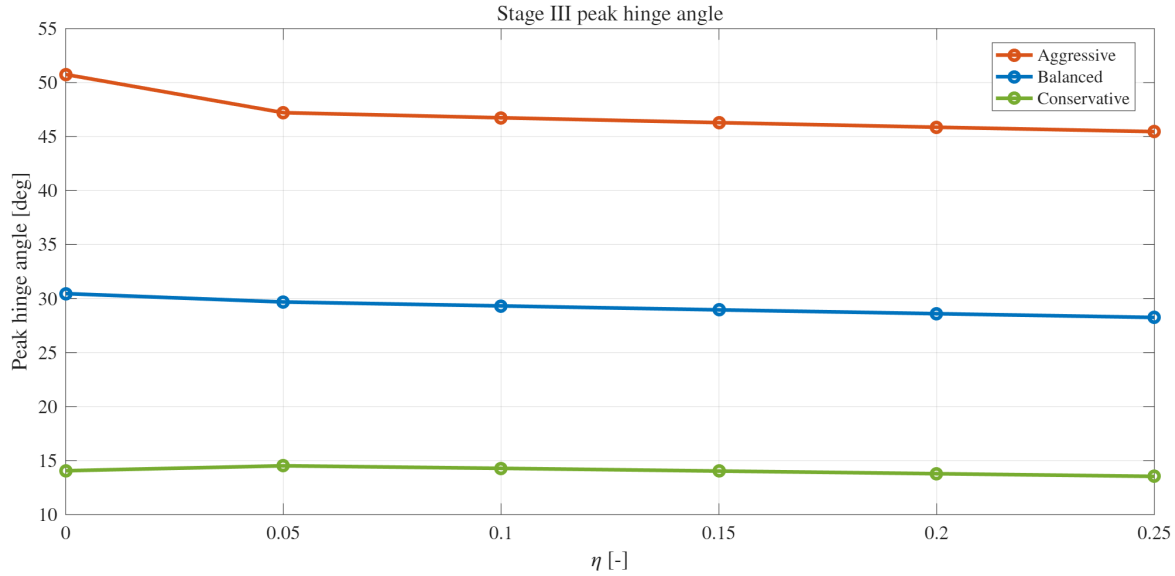

b) Peak hinge angle.

Figure 6.16: Stage III threshold sensitivity of the aggressive, balanced, and conservative hinge-property candidates.

The WRBM-reduction trend shows that the aggressive and balanced candidates have a clear optimum at low non-zero threshold. For the aggressive candidate, the maximum reduction is obtained at $\eta = 0.05$, with $\Delta_{\text{WRBM}} = 23.8\%$. For the balanced candidate, the maximum reduction occurs at $\eta = 0.10$, with $\Delta_{\text{WRBM}} = 21.2\%$. In both cases, increasing the threshold further reduces the peak-load benefit, because the wingtip motion is initiated too late to affect the full WRBM build-up. The conservative candidate follows the opposite trend over the tested range: its WRBM reduction increases with η, reaching 10.9% at $\eta = 0.25$. This confirms that the release phase interacts with the post-release hinge stiffness, and that the same threshold fraction does not have the same effect for all hinge-property candidates.

The peak hinge-angle trend is governed mainly by the hinge stiffness selected in Stage II. The aggressive candidate produces the largest hinge excursions, with peak angles between approximately 45° and 51°. The balanced candidate remains near 28°–31°, while the conservative candidate stays near 14°–15°. Increasing η generally reduces the peak hinge angle for the aggressive and balanced candidates, because the wingtip has less time to rotate before and during the main load transient. For the conservative candidate the variation is smaller, since the larger stiffness already restricts the motion.

To visualise the resulting design compromise more directly, Figure 6.17 plots the peak WRBM reduction against the corresponding peak hinge angle for the Stage III threshold-release cases. This representation collapses the threshold sweep onto the two primary design quantities and makes the trade-off between load alleviation and hinge-motion demand explicit. The aggressive candidate occupies the high-alleviation, high-motion region, whereas the conservative candidate remains in the low-motion but low-alleviation region. The balanced candidate lies between these extremes and retains most of the peak-load reduction of the aggressive setting while requiring substantially smaller hinge excursions.

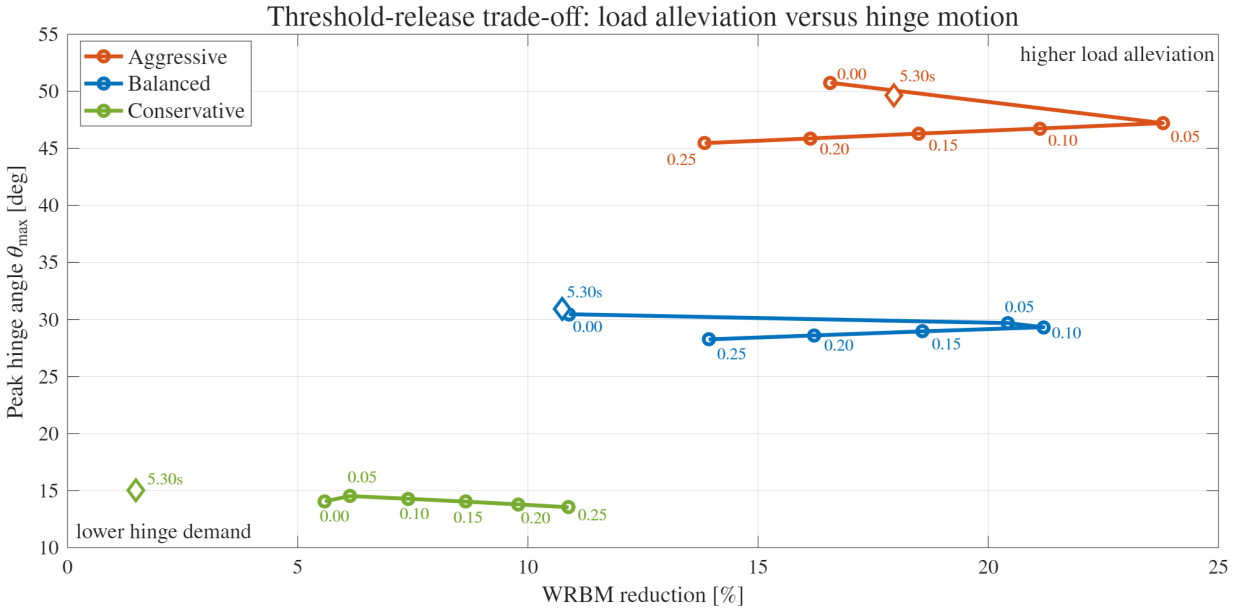

Figure 6.17: Stage III trade-off between peak WRBM reduction and peak hinge angle for the aggressive, balanced and conservative hinge-property candidates. Positive WRBM reduction indicates load alleviation relative to the locked reference, while lower peak hinge angle indicates lower hinge-motion demand.

Together, these trends show that the best release threshold is not obtained by simply maximising the delay or by releasing as early as possible. Instead, the favourable threshold is one that initiates hinge motion during the load build-up phase, while still leaving sufficient time for the folding motion to generate local unloading before the critical WRBM peak. The aggressive candidate provides the highest peak-load reduction, but does so at the cost of large hinge excursions. The balanced candidate therefore represents the more useful compromise, since it retains most of the load-alleviation benefit while requiring substantially smaller hinge motion.

The most relevant Stage III trade-off cases are summarised in Table 6.15. The aggressive $\eta = 0.05$ case gives the strongest peak-load alleviation, but requires a large hinge-angle excursion of approximately 47°. The balanced $\eta = 0.10$ case gives almost the same WRBM reduction, while reducing the peak hinge angle to approximately 29°. The conservative $\eta = 0.25$ case gives the smallest hinge motion, but its WRBM reduction remains only about half of the best aggressive and balanced cases.

Table 6.15: Representative Stage III trade-off cases.

Candidate	Release case	t_{rel} [s]	Δ_{WRBM} [%]	θ_{max} [°]
Aggressive	$\eta = 0.05$	5.422	23.8	47.2
Balanced	$\eta = 0.10$	5.440	21.2	29.3
Conservative	$\eta = 0.25$	5.481	10.9	13.5

From a pure peak-load perspective, the aggressive $\eta = 0.05$ case is the best-performing Stage III configuration. However, the balanced $\eta = 0.10$ case provides a more attractive compromise, because it retains nearly the same peak WRBM reduction while substantially reducing the hinge-angle demand. The conservative candidate is dynamically more restrained in terms of hinge angle, but it does not recover enough peak-load alleviation to be competitive with the aggressive and balanced candidates under the present gust condition.

Stage III Conclusion

Stage III shows that hinge-moment-based release can improve the response relative to simple gust-onset release, but only when the threshold places the release within a favourable load-build-up phase. Because the non-zero threshold cases are clustered in a narrow timing window, the sweep should be interpreted as a local phase-sensitivity study rather than as a broad release-timing study.

The aggressive candidate gives the largest peak WRBM reduction, reaching 23.8% for $\eta = 0.05$, but requires peak hinge angles above 45°. The conservative candidate limits the hinge motion to about 14°, but suppresses too much of the folding mechanism and reaches only 10.9% WRBM reduction in its best case. The balanced candidate with $\eta = 0.10$ gives the most favourable compromise, retaining a high WRBM reduction of 21.2% while limiting the peak hinge angle to approximately 29°.

Overall, the Stage III results show that the release threshold should be tuned using the phase of the hinge-moment response, rather than by assuming that the earliest possible release or the largest delay is optimal.

7

Conclusions

This thesis investigated the non-linear aeroelastic response of a flared folding wingtip subjected to discrete vertical gust excitation. The central objective was to develop and apply a time-domain Simscape–UVLM framework capable of modelling a flexible main wing, a rigid flared folding wingtip, and a configurable hinge-release mechanism. The framework was then used to assess whether suitable combinations of hinge properties and release logic can reduce the critical wing-root bending moment while limiting hinge-motion demand and avoiding clearly adverse post-release behaviour.

The resulting framework couples main-wing flexibility, rigid wingtip motion, flared hinge kinematics, unsteady aerodynamic loading and gust excitation in the time domain. This makes it possible to study not only the basic FFWT unloading mechanism, but also the timing sensitivity, hinge-motion demand and residual oscillatory response that determine whether the concept is useful as a practical gust-load alleviation device.

Research Questions

Before answering the research questions formulated at the start of this thesis, they are repeated here for the reader's convenience. The main research question was defined as follows:

> Which combinations of FFWT design parameters and hinge-release settings reduce wing-root bending moment relative to a locked-hinge baseline, while keeping post-gust motion and aeroelastic behaviour within acceptable limits?

This main question was supported by the following sub-research questions:

1. How accurately can the framework reproduce existing FFWT reference cases in terms of wing-root bending moment, hinge-angle dynamics and aeroelastic modes?
2. How can hinge-moment-based release strategies be formulated and implemented within the framework?
3. How do key FFWT design parameters affect dynamic behaviour, load alleviation performance and post-gust tip motion for standard vertical discrete gusts?

The first sub-research question addresses whether the developed framework provides a sufficiently credible basis for the subsequent release and hinge-parameter studies. The numerical model was verified through structural, aerodynamic and coupled aeroelastic checks, and was then validated at trend level against existing FFWT gust-release experiments. Although the present model was not intended to reproduce the experimental setup point-by-point, it captured the governing release-timing hierarchy observed in previous work: early release provides load alleviation, delayed release reduces the benefit, and release near the locked peak-load instant can become detrimental. The framework can therefore be considered suitable for comparative non-linear FFWT release studies, provided that its results are interpreted as trend-level predictions within the selected modelling assumptions.

Building on this framework, the second sub-research question concerns the formulation and implementation of hinge-moment-based release logic. This was achieved by extracting an incremental hinge-moment response from the locked-reference simulation and defining release thresholds as fractions of the locked-reference peak hinge moment. These thresholds were then used to trigger the transition from locked to released hinge behaviour in the time-domain simulation. The Stage III results show that this formulation is suitable for comparing threshold-based release strategies and can improve the response relative to simple gust-onset release, but only when the

threshold causes the wingtip to start moving during the favourable load-build-up phase. The threshold value is therefore not sufficient on its own, but its effect must be interpreted through the release phase relative to the WRBM build-up.

The third sub-research question then determines which design parameters and release settings are favourable once such release logic is available. The Stage I frequency-scaled release-trend assessment showed that release effectiveness is governed primarily by whether the folding motion develops before the critical WRBM peak is formed. Pre-release produced peak WRBM reductions of approximately 29 – 38%, while gust-onset release remained beneficial but became less effective as the gust excitation became faster. Release at the 50% WRBM threshold provided only modest alleviation of approximately 10%, and release at the 100% threshold was consistently too late, increasing the critical WRBM peak by approximately 8% relative to the locked reference. The reduced-frequency interpretation confirmed that local angle-of-attack relief only contributes to load alleviation if it is generated early enough relative to the structural load build-up.

The Stage II hinge-parameter study showed that hinge stiffness is the dominant design parameter controlling the trade-off between WRBM reduction and hinge-angle demand. Low stiffness increases load alleviation but also increases peak hinge angle, while high stiffness limits wingtip motion but suppresses the unloading mechanism. Hinge damping improves the qualitative post-release response and hinge-angle decay, but cannot recover the load-alleviation benefit once the hinge is too stiff. The Stage III threshold study further showed that release timing and hinge stiffness cannot be selected independently: the same hinge-moment threshold can lead to different performance depending on the post-release hinge properties.

Together, these answers lead to the answer to the main research question. Within the scope of the present model and gust cases, the FFWT reduces WRBM most effectively when a sufficiently compliant post-release hinge is combined with release logic that triggers during the gust-load build-up, before the locked-reference WRBM peak is formed. The aggressive low-stiffness hinge setting gives the largest peak-load reduction when released at a favourable phase, reaching 23.8%, but requires large hinge-angle excursions of approximately 47°. The conservative setting limits hinge motion to approximately 14°, but suppresses too much of the folding mechanism to provide comparable load alleviation, with a maximum WRBM reduction of 10.9%. The balanced hinge setting with a low non-zero hinge-moment threshold provides the most favourable compromise, reaching a WRBM reduction of 21.2% while limiting the peak hinge angle to approximately 29°.

The results therefore support the potential of the FFWT as a gust-load alleviation device, provided that it is treated as a dynamically phased aeroelastic system rather than as a simple load-threshold device. Releasing the hinge too late can reduce or even reverse the benefit, while excessive hinge stiffness prevents the wingtip from generating sufficient local unloading. These conclusions should therefore be interpreted within the selected geometry, discrete gust cases, hinge-property range and modelling assumptions used in the present framework.

8

Recommendations

Based on the conclusions of this thesis in Chapter 7, future work should focus on improving the physical realism of the model, extending the release logic, and testing the identified design trends under a broader set of operating conditions.

Model Fidelity

The first recommendation is to improve the aerodynamic model to also extend beyond the present attached-flow UVLM assumptions. Large wingtip rotations may introduce local separation or other viscous effects that are not captured by UVLM. Future work should therefore consider viscous corrections, higher-fidelity aerodynamic modelling, or a free-wake formulation to assess whether the prescribed-wake and attached-flow assumptions remain valid during rapid folding motion.

Finally, the structural model should be extended to include wingtip flexibility. The present model treats the main wing as flexible and the folding wingtip as rigid. Including wingtip flexibility would provide a more complete representation of the structural load path and could influence the predicted local angle-of-attack relief, hinge loads and coupling with the main-wing modes.

Hinge Mechanism and Design Constraints

The hinge model should be extended from idealised stiffness and damping to a more realistic mechanism representation. Future models should include effects such as friction, backlash, mechanical stops, release latency, lock-disengagement dynamics and re-locking. These effects are important because they influence both the actual release instant and the post-release hinge-angle response. Mechanical stops are especially relevant for compliant hinge settings, where the present simulations predict large peak hinge angles.

Future optimisation studies should also treat hinge motion as an explicit design constraint. The present results show that the largest WRBM reductions can require substantial wingtip motion. Therefore, peak WRBM reduction should not be maximised alone; peak hinge angle, hinge rate, hinge moment and residual post-release oscillations should be included as constraints or penalty terms.

Release Logic and Control

Future release logic should be designed around the phase of the gust-induced load build-up rather than around peak-load detection alone. The present results show that release at or near the locked-reference peak-load instant is consistently too late. A useful release law must therefore trigger early enough for the wingtip motion to generate local unloading before the critical WRBM peak forms.

A promising next step is to develop hinge-moment-based release logic that is aware of the phase. Instead of using only a fixed hinge-moment threshold, the release criterion could include hinge-moment rate, an estimated WRBM response, hinge-state information or gust-preview information. The release strategy should also be extended to include re-locking, since practical operation requires controlling the full release–motion–recovery sequence without introducing secondary load peaks or excessive hinge moments.

Once the passive release framework has been extended, active or preview-based control should be investigated. Preview information from pressure probes, accelerometers or lidar-based gust sensing could trigger release earlier and more robustly than a purely local threshold.

Active hinge torque or control surfaces could then be used to shape the wingtip motion after release, reducing peak loads while limiting hinge-angle demand and suppressing post-release oscillations.

Expanded Simulation Campaign

The simulation campaign should be extended to a broader set of gust and operating conditions. Future studies should include variations in gust amplitude, gust gradient length, freestream velocity and baseline angle of attack. This would show whether the identified balanced hinge-property candidate and low-threshold release strategy remain favourable outside the tested condition.

Continuous turbulence and manoeuvre loading should also be included. Discrete vertical gusts are useful for controlled release-timing analysis, but an operational FFWT would encounter stochastic atmospheric excitation and combined flight-load cases. These simulations would allow the fatigue-relevant response, repeated release events, manoeuvre loads and asymmetric behaviour to be assessed.

Finally, full-aircraft integration effects should be evaluated. The present model isolates the wing and folding-tip mechanism, which is appropriate for studying the local aeroelastic response. However, practical implementation requires assessing aircraft-level loads, flight mechanics, control allocation, certification constraints and operational requirements.

Priority for Future Work

The most direct continuation of this thesis is to combine a broader gust library with a finite-latency hinge-release model and explicit hinge-angle, hinge-rate and hinge-moment constraints. This would test whether the balanced hinge-property candidate and low-threshold release strategy remain favourable once practical mechanism limitations are included. After that, preview-based and actively controlled release strategies can be evaluated on the same numerical basis.

Part IV

Appendix

A

Python ROM Extraction and Verification

This appendix lists the Python script used to extract the reduced Craig–Bampton mass and stiffness matrices from the Nastran `.op4` superelement export. The script reads the `MXX` and `KXX` matrices, converts them to dense NumPy arrays, verifies their symmetry and dimensions, checks the definiteness of the reduced mass matrix, checks the near-positive-semidefinite character of the stiffness matrix, computes the natural frequencies from the generalised eigenvalue problem, and finally writes the matrices to a MATLAB-compatible `.mat` file for use in the Simscape workflow.

```
#!/usr/bin/env python3
"""
Extract and verify Craig-Bampton ROM matrices from a Nastran OP4 file.

The script performs the following steps:
1. Reads the reduced stiffness and mass matrices, KXX and MXX.
2. Converts sparse or matrix-like objects to dense NumPy arrays.
3. Checks symmetry, dimensions, mass positive definiteness and stiffness
   near-positive-semidefiniteness.
4. Solves the generalised eigenvalue problem

       K phi = omega^2 M phi

   to obtain the natural frequencies of the reduced model.
5. Exports the verified matrices to a MATLAB-compatible .mat file.

The script is intended as a traceable preprocessing step between the
Simcenter/Nastran superelement export and the Simscape implementation.
"""

from pathlib import Path
import sys

import numpy as np
from pyNastran.op4.op4 import OP4
from scipy import linalg, sparse
from scipy.io import savemat


def to_dense(matrix):
    """Convert a sparse or array-like matrix to a dense NumPy array."""
    if sparse.issparse(matrix):
        return matrix.toarray()
    return np.asarray(matrix)


def print_header(title, width=72):
    """Print a formatted section header."""
    print("\n" + "=" * width)
    print(title)
    print("=" * width)


def status(flag):
    """Return a compact pass/fail string."""
    return "PASS" if flag else "FAIL"


def check_symmetry(M, K, tol=1e-11):
```

```
    """
    Check relative Frobenius-norm symmetry errors of M and K.

    Returns
    -------
    passed : bool
        True if both matrices satisfy the prescribed tolerance.
    err_M : float
        Relative symmetry error of the mass matrix.
    err_K : float
        Relative symmetry error of the stiffness matrix.
    """
    err_M = np.linalg.norm(M - M.T, ord="fro") / np.linalg.norm(M, ord="fro")
    err_K = np.linalg.norm(K - K.T, ord="fro") / np.linalg.norm(K, ord="fro")

    pass_M = err_M < tol
    pass_K = err_K < tol

    print_header("SYMMETRY CHECK")
    print(f"Mass matrix M:      {err_M:.3e}  {status(pass_M)}")
    print(f"Stiffness matrix K: {err_K:.3e}  {status(pass_K)}")
    print(f"Tolerance:          {tol:.1e}")

    return pass_M and pass_K, err_M, err_K


def check_dimensions(M, K, expected_size):
    """
    Check whether M and K are square and have the expected ROM size.
    """
    same_shape = M.shape == K.shape
    square = M.shape[0] == M.shape[1] and K.shape[0] == K.shape[1]
    expected = M.shape == (expected_size, expected_size)

    passed = same_shape and square and expected

    print_header("DIMENSION CHECK")
    print(f"Expected size: {expected_size} x {expected_size}")
    print(f"M shape:       {M.shape[0]} x {M.shape[1]}")
    print(f"K shape:       {K.shape[0]} x {K.shape[1]}")
    print(f"Result:        {status(passed)}")

    return passed


def check_matrix_spectrum(A, name, tol=1e-10, require_positive_definite=True):
    """
    Inspect the eigenvalue spectrum of a symmetric matrix.

    For the mass matrix, strict positive definiteness is required.
    For the stiffness matrix, near-positive-semidefiniteness is accepted,
    because free-free or weakly constrained structural models can contain
    near-zero rigid-body modes and small negative numerical artefacts.
    """
    eigvals = np.linalg.eigvalsh(A)

    lam_min = float(np.min(eigvals))
    lam_max = float(np.max(eigvals))

    if require_positive_definite:
        passed = lam_min > tol
        criterion = f"lambda_min > {tol:.1e}"
    else:
        passed = lam_min > -tol
        criterion = f"lambda_min > -{tol:.1e}"

    cond_num = lam_max / lam_min if lam_min > tol else np.inf

    print_header(f"SPECTRUM CHECK: {name}")
    print(f"Minimum eigenvalue: {lam_min:.6e}")
    print(f"Maximum eigenvalue: {lam_max:.6e}")
    print(f"Condition number:   {cond_num:.6e}")
    print(f"Criterion:          {criterion}")
    print(f"Result:             {status(passed)}")

    return passed, lam_min, lam_max, cond_num


def compute_natural_frequencies(M, K, n_modes=20, eig_tol=1e-10):
    """
    Compute natural frequencies from the generalised eigenvalue problem.

    The problem solved is

```

```
        K phi = lambda M phi,

    where lambda = omega^2. Near-zero and negative eigenvalues below
    eig_tol are removed before converting to frequencies.
    """
    M = np.asarray(M, dtype=np.float64)
    K = np.asarray(K, dtype=np.float64)

    print_header("MODAL ANALYSIS")
    print("Solving K phi = omega^2 M phi ...")

    eigvals, _ = linalg.eigh(K, M)
    eigvals = np.sort(eigvals)

    n_total = len(eigvals)
    n_filtered = int(np.sum(eigvals <= eig_tol))

    eigvals = eigvals[eigvals > eig_tol]
    eigvals = eigvals[:min(n_modes, len(eigvals))]

    omega_rad_s = np.sqrt(eigvals)
    freq_hz = omega_rad_s / (2.0 * np.pi)

    print(f"Total eigenvalues:             {n_total}")
    print(f"Filtered near-zero eigenvalues: {n_filtered}")
    print(f"Frequencies displayed:          {len(freq_hz)}")

    print("\n{:<8s} {:>16s} {:>18s} {:>18s}".format(
        "Mode", "f [Hz]", "omega [rad/s]", "lambda"
    ))
    print("-" * 68)

    for i, (f, omega, lam) in enumerate(zip(freq_hz, omega_rad_s, eigvals), start=1):
        print(f"{i:<8d} {f:>16.6f} {omega:>18.6f} {lam:>18.6e}")

    return freq_hz, omega_rad_s, eigvals


def read_op4_matrices(op4_path):
    """
    Read KXX and MXX from a Nastran OP4 file.

    Parameters
    ----------
    op4_path : str or Path
        Path to the OP4 file.

    Returns
    -------
    M : ndarray
        Reduced mass matrix.
    K : ndarray
        Reduced stiffness matrix.
    """
    op4_path = Path(op4_path)

    if not op4_path.is_file():
        raise FileNotFoundError(f"OP4 file not found: {op4_path}")

    op4 = OP4()
    matrices = op4.read_op4(str(op4_path))

    try:
        K = to_dense(matrices["KXX"].data)
        M = to_dense(matrices["MXX"].data)
    except KeyError as exc:
        available = ", ".join(matrices.keys())
        raise KeyError(
            f"Required matrix {exc} not found. Available matrices: {available}"
        ) from exc

    return M, K


def export_rom_to_mat(
    op4_path,
    mat_path,
    expected_size,
    n_modes_display=20,
    symmetry_tol=1e-11,
    eig_tol=1e-10,
):
    """
    Extract, verify and export Craig-Bampton ROM matrices.
```

```
218
219     The matrices are exported only if the core checks pass:
220     symmetry, dimensional consistency and mass positive definiteness.
221     The stiffness matrix is allowed to be near-positive-semidefinite.
222     """
223     print_header("ROM MATRIX EXTRACTION AND VERIFICATION")
224     print(f"Input OP4 file: {op4_path}")
225     print(f"Output MAT file: {mat_path}")
226
227     try:
228         M, K = read_op4_matrices(op4_path)
229     except Exception as exc:
230         print(f"\nERROR: Failed to read OP4 file: {exc}")
231         return False
232
233     print("\nSuccessfully extracted MXX and KXX.")
234     print(f"M shape: {M.shape}")
235     print(f"K shape: {K.shape}")
236
237     sym_pass, sym_err_M, sym_err_K = check_symmetry(
238         M, K, tol=symmetry_tol
239     )
240
241     dim_pass = check_dimensions(
242         M, K, expected_size=expected_size
243     )
244
245     mass_pass, m_min, m_max, m_cond = check_matrix_spectrum(
246         M,
247         name="mass matrix M",
248         tol=eig_tol,
249         require_positive_definite=True,
250     )
251
252     stiff_pass, k_min, k_max, k_cond = check_matrix_spectrum(
253         K,
254         name="stiffness matrix K",
255         tol=eig_tol,
256         require_positive_definite=False,
257     )
258
259     modal_pass = True
260     try:
261         freq_hz, omega_rad_s, eigvals = compute_natural_frequencies(
262             M,
263             K,
264             n_modes=n_modes_display,
265             eig_tol=eig_tol,
266         )
267     except Exception as exc:
268         print(f"\nWARNING: Modal analysis failed: {exc}")
269         modal_pass = False
270         freq_hz = np.array([])
271         omega_rad_s = np.array([])
272         eigvals = np.array([])
273
274     core_pass = sym_pass and dim_pass and mass_pass
275
276     print_header("VALIDATION SUMMARY")
277     print(f"Symmetry check:              {status(sym_pass)}")
278     print(f"Dimension check:             {status(dim_pass)}")
279     print(f"Mass positive definiteness:  {status(mass_pass)}")
280     print(f"Stiffness semidefiniteness:  {status(stiff_pass)}")
281     print(f"Modal analysis:              {status(modal_pass)}")
282     print(f"Core export checks:          {status(core_pass)}")
283
284     if not core_pass:
285         print("\nCore checks failed. Matrices are not exported.")
286         return False
287
288     validation = {
289         "symmetry_error_M": sym_err_M,
290         "symmetry_error_K": sym_err_K,
291         "lambda_min_M": m_min,
292         "lambda_max_M": m_max,
293         "condition_number_M": m_cond,
294         "lambda_min_K": k_min,
295         "lambda_max_K": k_max,
296         "condition_number_K": k_cond,
297         "frequencies_hz": freq_hz,
298         "frequencies_rad_s": omega_rad_s,
299         "eigenvalues": eigvals,
300     }
301
```

```
    try:
        savemat(
            mat_path,
            {
                "M": M,
                "K": K,
                "validation": validation,
            },
        )
    except Exception as exc:
        print(f"\nERROR: Failed to write MAT file: {exc}")
        return False

    print(f"\nSuccessfully saved ROM matrices to: {mat_path}")
    return True


if __name__ == "__main__":
    success = export_rom_to_mat(
        op4_path="sol103_wing_only_sv2-wing-cmsv2_0.op4",
        mat_path="wing_SECMS.mat",
        expected_size=146,
        n_modes_display=20,
    )

    sys.exit(0 if success else 1)
```

Listing A.1: Python script for ROM matrix extraction and verification.

B

Kinematic Local Angle-of-Attack Derivation

This appendix derives the kinematic local angle-of-attack metric used in this thesis to interpret the local unloading mechanism of the flared folding wingtip. The metric is used as a post-processing quantity. It does not replace the aerodynamic loads computed by the UVLM solver, but provides an interpretable measure of how the folding motion changes the local incidence of the wingtip.

B.1. Coordinate Frames and Sign Convention

A right-handed wing-fixed frame W is used. The x_W-axis is aligned with the undeformed chordwise direction, the y_W-axis points spanwise outboard, and the z_W-axis points upward, normal to the undeformed lifting surface. The folding wingtip is described by a local tip frame T. The hinge line is flared by an angle Λ, and the fold angle is denoted by θ. In the adopted convention, positive θ corresponds to upward folding of the wingtip.

The local velocity components in the folded-tip frame are obtained by transforming the relative velocity from the wing frame to the tip frame. The transformation is treated as a passive coordinate transformation: it expresses a vector given in the wing frame in the folded-tip frame. The wing-to-tip transformation is written as

$$\mathbf{V}_T = \mathbf{R}_{W\to T}\mathbf{V}_W, \qquad \mathbf{R}_{W\to T} = \mathbf{R}_x(-\theta)\mathbf{R}_z(\Lambda). \tag{B.1}$$

The rotation matrices are

$$\mathbf{R}_z(\Lambda) = \begin{bmatrix} \cos\Lambda & -\sin\Lambda & 0 \\ \sin\Lambda & \cos\Lambda & 0 \\ 0 & 0 & 1 \end{bmatrix}, \tag{B.2}$$

and

$$\mathbf{R}_x(-\theta) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\theta & \sin\theta \\ 0 & -\sin\theta & \cos\theta \end{bmatrix}. \tag{B.3}$$

Substitution gives

$$\mathbf{R}_{W\to T} = \begin{bmatrix} \cos\Lambda & -\sin\Lambda & 0 \\ \cos\theta\sin\Lambda & \cos\theta\cos\Lambda & \sin\theta \\ -\sin\theta\sin\Lambda & -\sin\theta\cos\Lambda & \cos\theta \end{bmatrix}. \tag{B.4}$$

This sign convention is chosen such that positive upward folding about a positive flare angle produces a negative local normal velocity component in the folded-tip frame, corresponding to local nose-down incidence relief.

B.2. Velocity Components in the Folded-Tip Frame

The freestream velocity in the wing frame is written as

$$\mathbf{V}_{\infty,W} = V_\infty \begin{bmatrix} \cos\alpha_0 \\ 0 \\ \sin\alpha_0 \end{bmatrix}, \tag{B.5}$$

where α_0 is the reference angle of attack in the wing frame. If the fold-rate contribution is included, the relative velocity seen by a point on the folding wingtip is obtained by subtracting the local body velocity induced by the fold motion. For a point located a distance L from the hinge axis in the local spanwise direction, this can be written in the tip frame as

$$\mathbf{V}_{\mathrm{rel},T} = \mathbf{R}_{W\to T}\mathbf{V}_{\infty,W} - \boldsymbol{\omega}_T \times \mathbf{r}_T, \tag{B.6}$$

with

$$\boldsymbol{\omega}_T = \begin{bmatrix} \dot\theta \\ 0 \\ 0 \end{bmatrix}, \qquad \mathbf{r}_T = \begin{bmatrix} 0 \\ L \\ 0 \end{bmatrix}. \tag{B.7}$$

The fold-rate contribution is therefore

$$\boldsymbol{\omega}_T \times \mathbf{r}_T = \begin{bmatrix} 0 \\ 0 \\ L\dot\theta \end{bmatrix}. \tag{B.8}$$

The local relative velocity in the folded-tip frame is denoted by

$$\mathbf{V}_{\mathrm{rel},T} = \begin{bmatrix} u_T \\ v_T \\ w_T \end{bmatrix}, \tag{B.9}$$

where u_T, v_T, and w_T are the chordwise, spanwise, and normal velocity components in the folded-tip frame. Substitution gives

$$u_T = V_\infty \cos\alpha_0 \cos\Lambda, \tag{B.10}$$

$$v_T = V_\infty \left(\cos\alpha_0 \cos\theta \sin\Lambda + \sin\alpha_0 \sin\theta\right), \tag{B.11}$$

and

$$w_T = V_\infty \left(\sin\alpha_0 \cos\theta - \cos\alpha_0 \sin\theta \sin\Lambda\right) - L\dot\theta. \tag{B.12}$$

The quasi-static contribution responsible for the FFWT incidence-relief mechanism is

$$-V_\infty \cos\alpha_0 \sin\theta \sin\Lambda. \tag{B.13}$$

For positive upward folding and positive flare angle, this term is negative. It therefore reduces the local normal velocity component w_T, corresponding to a local nose-down incidence change.

B.3. Local Flow Angle, Sectional Angle of Attack and Sideslip

Two related local angles can be defined from the tip-frame velocity components. The first is the local flow-elevation angle,

$$\alpha_{\text{flow}} = \operatorname{atan2}\left(w_T, \sqrt{u_T^2 + v_T^2}\right). \tag{B.14}$$

This angle describes the inclination of the local velocity vector relative to the local x_T-y_T plane. It includes the spanwise velocity component v_T in the in-plane velocity magnitude and is therefore useful as a geometric description of the three-dimensional local flow direction.

However, the sectional aerodynamic angle of attack relevant for interpreting local lift and drag is normally defined in the local chord-normal plane. The spanwise velocity component is then interpreted as sideslip rather than as part of the two-dimensional sectional angle of attack. For this reason, the local sectional angle of attack used in this thesis is defined as

$$\alpha_{\text{loc}} = \operatorname{atan2}(w_T, u_T). \tag{B.15}$$

The local sideslip angle can be defined separately as

$$\beta_{\text{loc}} = \operatorname{atan2}(v_T, u_T). \tag{B.16}$$

The kinematic local angle-of-attack relief used in the main text is based on the sectional definition in Equation B.15. It is evaluated relative to the unreleased reference condition as

$$\Delta\alpha_{\text{tip,kin}} = \alpha_{\text{loc}}(\theta, \dot{\theta}) - \alpha_{\text{loc}}(0, 0). \tag{B.17}$$

With this definition, positive upward folding about a positive flare angle produces a negative value of $\Delta\alpha_{\text{tip,kin}}$, corresponding to local incidence relief.

B.4. Quasi-Static Limit

For the quasi-static case with $\alpha_0 = 0$ and $\dot{\theta} = 0$, the tip-frame velocity components reduce to

$$u_T = V_\infty \cos\Lambda, \qquad v_T = V_\infty \cos\theta \sin\Lambda, \qquad w_T = -V_\infty \sin\theta \sin\Lambda. \tag{B.18}$$

The local flow-elevation angle becomes

$$\alpha_{\text{flow}} = \operatorname{atan2}\left(-\sin\theta \sin\Lambda, \sqrt{\cos^2\Lambda + \cos^2\theta \sin^2\Lambda}\right), \tag{B.19}$$

whereas the sectional local angle of attack becomes

$$\alpha_{\text{loc}} = \operatorname{atan2}\left(-\sin\theta \sin\Lambda, \cos\Lambda\right). \tag{B.20}$$

For a fully folded tip, $\theta = 90°$, both definitions reduce to

$$\alpha_{\text{flow}} = \alpha_{\text{loc}} = \operatorname{atan2}\left(-\sin\Lambda, \cos\Lambda\right) = -\Lambda. \tag{B.21}$$

Thus, for a flare angle of $\Lambda = 15°$, the quasi-static local incidence at a fold angle of 90° becomes

$$\alpha_{\text{loc}} = -15°. \tag{B.22}$$

This result provides an important consistency check. The kinematic incidence change remains bounded by the flare-angle scale instead of increasing without bound as the fold angle approaches 90°.

B.5. Relation to the Small-Angle Interpretation

For small fold angles, the local incidence change induced by a flared hinge is commonly interpreted using the first-order relation

$$\Delta\alpha_{\mathrm{tip,kin}} \approx -\theta \sin\Lambda. \tag{B.23}$$

This expression is useful for explaining the sign and first-order scaling of the FFWT unloading mechanism. It follows directly from the flow-elevation definition when the fold angle is small. In that case,

$$\sin\theta \approx \theta, \tag{B.24}$$

and the normal velocity component scales as

$$w_T \approx -V_\infty \theta \sin\Lambda. \tag{B.25}$$

For the sectional definition used in the reported metric, the corresponding small-fold-angle limit for $\alpha_0 = 0$ is

$$\alpha_{\mathrm{loc}} \approx -\theta \tan\Lambda. \tag{B.26}$$

For the flare angle used in this thesis, $\Lambda = 15°$, the difference between $\sin\Lambda$ and $\tan\Lambda$ is small. The relation $-\theta \sin\Lambda$ is therefore retained in the main text as the conventional first-order interpretation of the FFWT unloading mechanism, while the reported local angle-of-attack relief is computed using the full tip-frame velocity formulation in Equation B.15.

B.6. Comparison with the Original Scalar Relation

The local incidence relief of a flared folding wingtip is often introduced through the scalar relation

$$\Delta\alpha_{\mathrm{scalar}} = -\arctan\left(\tan\theta \sin\Lambda\right). \tag{B.27}$$

For small fold angles, this relation gives the same qualitative behaviour as the tip-frame velocity formulation. However, because it contains $\tan\theta$, it becomes singular as θ approaches 90°. It therefore predicts an artificial growth of the local incidence relief toward −90°. The tip-frame velocity formulation does not show this behaviour. Instead, it remains bounded and gives $\alpha_{\mathrm{loc}} = -\Lambda$ for $\theta = 90°$, $\alpha_0 = 0$, and $\dot{\theta} = 0$.

This difference is illustrated in Table B.1 for $\Lambda = 15°$, $\alpha_0 = 0°$, and $\dot{\theta} = 0$.

Table B.1: Comparison of kinematic local angle-of-attack relief predicted by the original scalar relation, the local flow-elevation definition, and the sectional tip-frame velocity formulation for $\Lambda = 15°$, $\alpha_0 = 0°$, and $\dot{\theta} = 0$. Negative values indicate local nose-down incidence change.

θ [°]	Original scalar relation [°]	Flow-elevation definition [°]	Sectional definition [°]
0	0.00	0.00	0.00
10	−2.61	−2.58	−2.66
20	−5.38	−5.08	−5.24
30	−8.50	−7.44	−7.63
40	−12.25	−9.58	−9.77
50	−17.14	−11.44	−11.60
60	−24.15	−12.95	−13.06
70	−35.42	−14.08	−14.13
80	−55.73	−14.77	−14.78
90	−90.00	−15.00	−15.00

The flow-elevation and sectional tip-frame definitions remain bounded and approach the flare angle in the fully folded quasi-static limit. The scalar expression agrees reasonably well only for small fold angles. At large fold angles, it predicts increasingly large local incidence relief due to the singular behaviour of $\tan\theta$, whereas the rotation-matrix formulation gives a physically interpretable limit set by the flare angle.

The final metric used in this thesis is therefore based on the sectional tip-frame velocity formulation in Equation B.15. This definition is preferred because the local sectional incidence is determined by the chordwise and normal velocity components seen by the folded wingtip section. The spanwise velocity component is retained separately through the local sideslip angle, rather than being included in the sectional angle of attack.

C

Full Simulation Matrix

This appendix summarises the complete simulation matrix used in the staged FFWT gust-release campaign. The matrix is organised according to the three stages introduced in Section 5.6. Stage I varies gust frequency and prescribed release timing, Stage II fixes the gust input and release instant while varying the post-release hinge properties, and Stage III fixes selected hinge-property candidates while varying the hinge-moment-based release threshold.

C.1. Campaign Overview

Table C.1: Overview of the simulation campaign.

Stage	Description	Number of simulations
Stage I-A	Relative-frequency release-timing study	24
Stage I-B	Additional low-reduced-frequency anchor case	6
Stage I total	Release-timing and frequency-scaling study	30
Stage II	Hinge stiffness–damping sensitivity study	18
Stage III	Hinge-moment-based threshold-release study	22
Full campaign	Stages I–III combined	70

C.2. Stage I: Release-Timing and Frequency-Scaling Matrix

Stage I uses prescribed release definitions. For each gust-frequency case, a locked-hinge baseline is first evaluated. The locked WRBM response is then used to define the release instants for the pre-release, gust-onset release, 50% WRBM-trigger release and 100% WRBM-trigger release cases. The free-hinge case is included as a bounding reference and is free from the start of the simulation.

Table C.2: Stage I gust-frequency cases.

Case	f_g [Hz]	$\hat{f}_g$ [-]	k_g [-]
I-B0	0.68	0.40	0.047
I-A1	1.02	1.13	0.070
I-A2	1.68	0.99	0.116
I-A3	2.81	1.65	0.194
I-A4	4.49	2.64	0.310

Table C.3: Stage I release definitions.

Release case	Description	Release definition
Locked	Locked-hinge baseline	No release
Free	Free-hinge reference	Hinge free from start of simulation
Pre-release	Released before gust arrival	$t_{\mathrm{rel}} = 5.000$ s
0% release	Release at gust arrival	$t_{\mathrm{rel}} = t_g = 5.200$ s
50% release	WRBM-triggered release	$t_{\mathrm{rel}} = t(\lvert\mathrm{WRBM}\rvert = 0.5\lvert\mathrm{WRBM}_{\mathrm{locked,pk}}\rvert)$
100% release	Release at locked WRBM peak	$t_{\mathrm{rel}} = t(\lvert\mathrm{WRBM}\rvert = \lvert\mathrm{WRBM}_{\mathrm{locked,pk}}\rvert)$

Table C.4: Stage I-A locked-baseline release times used for the prescribed release cases.

Case	f_g [Hz]	$\hat{f}_g$ [-]	k_g [-]	$t_{\mathrm{rel,pre}}$ [s]	$t_{\mathrm{rel},0\%}$ [s]	$t_{\mathrm{rel},50\%}$ [s]	$t_{\mathrm{rel},100\%}$ [s]
I-A1	1.02	1.13	0.070	5.000	5.200	5.731	6.239
I-A2	1.68	0.99	0.116	5.000	5.200	5.579	5.823
I-A3	2.81	1.65	0.194	5.000	5.200	5.494	5.661
I-A4	4.49	2.64	0.310	5.000	5.200	5.426	5.557

The exact 50% and 100% WRBM-trigger release times for I-B0 are obtained from the locked response of the low-reduced-frequency anchor case in the same way as for the Stage I-A cases. If these values are extracted from the simulation data, they can replace the trigger-definition entries in Table C.5.

C.3. Stage II: Hinge-Property Sensitivity Matrix

Stage II uses the Stage I-A3 gust case as the reference condition. The gust input, freestream condition and release instant are kept fixed, while the post-release hinge stiffness and damping ratio are varied. The fixed condition is $f_g = 2.81$ Hz, $\hat{f}_g = 1.65$, $k_g = 0.194$ and $t_{\mathrm{rel}} = 5.200$ s.

Table C.5: Additional Stage I-B low-reduced-frequency anchor case.

Case	f_g [Hz]	$\hat{f}_g$ [-]	k_g [-]	$t_{\mathrm{rel,pre}}$ [s]	$t_{\mathrm{rel},0\%}$ [s]	50% release	100% release
I-B0	0.68	0.40	0.047	5.000	5.200	WRBM-triggered	Locked-peak-triggered

Table C.6: Stage II finite hinge-property simulation matrix.

Case	f_g [Hz]	$\hat{f}_g$ [-]	k_g [-]	K_θ [Nm/rad]	ζ_h [-]
II-01	2.81	1.65	0.194	0.3	0.00
II-02	2.81	1.65	0.194	0.3	0.02
II-03	2.81	1.65	0.194	0.3	0.10
II-04	2.81	1.65	0.194	0.3	0.20
II-05	2.81	1.65	0.194	1.0	0.00
II-06	2.81	1.65	0.194	1.0	0.02
II-07	2.81	1.65	0.194	1.0	0.10
II-08	2.81	1.65	0.194	1.0	0.20
II-09	2.81	1.65	0.194	3.0	0.00
II-10	2.81	1.65	0.194	3.0	0.02
II-11	2.81	1.65	0.194	3.0	0.10
II-12	2.81	1.65	0.194	3.0	0.20
II-13	2.81	1.65	0.194	10.0	0.00
II-14	2.81	1.65	0.194	10.0	0.02
II-15	2.81	1.65	0.194	10.0	0.10
II-16	2.81	1.65	0.194	10.0	0.20

Table C.7: Stage II bounding reference simulations.

Case	f_g [Hz]	$\hat{f}_g$ [-]	k_g [-]	Hinge condition	Release definition
II-L	2.81	1.65	0.194	Locked	No release
II-F	2.81	1.65	0.194	Free	Free-hinge reference

C.4. Stage III: Hinge-Moment-Based Release Matrix

Stage III keeps the gust input and operating condition identical to Stage II, but carries forward three hinge-property candidates. The release threshold is based on the locked-reference increment of the third hinge-moment component, $M_{h,3}$. The threshold is defined as

$$M_{h,\mathrm{thr}} = \eta \Delta M_{h,3,\mathrm{pk}}^{\mathrm{locked}}, \tag{C.1}$$

and release occurs when

$$|\Delta M_{h,3}(t)| \geq M_{h,\mathrm{thr}}, \qquad t \geq t_g. \tag{C.2}$$

Table C.8: Stage III hinge-property candidates selected from Stage II.

Candidate	Stage II case	K_θ [Nm/rad]	ζ_h [-]	Description
Aggressive	II-04	0.3	0.20	Most compliant retained setting
Balanced	II-08	1.0	0.20	Intermediate load-alleviating setting
Conservative	II-12	3.0	0.20	More restrained hinge setting

In Table C.10, the prescribed-release anchor at $t_{\mathrm{rel}} = 5.300$ s is not treated as an additional hinge-moment threshold. It is included only as an intermediate timing reference between gust-onset release and the clustered non-zero threshold-triggered cases.

Table C.9: Stage III hinge-moment threshold levels and corresponding release times.

η [-]	$M_{h,\mathrm{thr}}$ [Nm]	t_{rel} [s]	$t_{\mathrm{rel}} - t_g$ [s]
0.00	0.000	5.200	0.000
0.05	0.018	5.422	0.222
0.10	0.036	5.440	0.240
0.15	0.054	5.455	0.255
0.20	0.072	5.469	0.269
0.25	0.091	5.481	0.281

Table C.10: Stage III full threshold-release simulation matrix.

Case	Candidate	f_g [Hz]	$\hat{f}_g$ [-]	k_g [-]	K_θ [Nm/rad]	ζ_h [-]	Release definition
III-L	Locked reference	2.81	1.65	0.194	–	–	No release
III-A-00	Aggressive	2.81	1.65	0.194	0.3	0.20	$\eta = 0$
III-A-530	Aggressive	2.81	1.65	0.194	0.3	0.20	$t_{\mathrm{rel}} = 5.300$ s
III-A-005	Aggressive	2.81	1.65	0.194	0.3	0.20	$\eta = 0.05$
III-A-010	Aggressive	2.81	1.65	0.194	0.3	0.20	$\eta = 0.10$
III-A-015	Aggressive	2.81	1.65	0.194	0.3	0.20	$\eta = 0.15$
III-A-020	Aggressive	2.81	1.65	0.194	0.3	0.20	$\eta = 0.20$
III-A-025	Aggressive	2.81	1.65	0.194	0.3	0.20	$\eta = 0.25$
III-B-00	Balanced	2.81	1.65	0.194	1.0	0.20	$\eta = 0$
III-B-530	Balanced	2.81	1.65	0.194	1.0	0.20	$t_{\mathrm{rel}} = 5.300$ s
III-B-005	Balanced	2.81	1.65	0.194	1.0	0.20	$\eta = 0.05$
III-B-010	Balanced	2.81	1.65	0.194	1.0	0.20	$\eta = 0.10$
III-B-015	Balanced	2.81	1.65	0.194	1.0	0.20	$\eta = 0.15$
III-B-020	Balanced	2.81	1.65	0.194	1.0	0.20	$\eta = 0.20$
III-B-025	Balanced	2.81	1.65	0.194	1.0	0.20	$\eta = 0.25$
III-C-00	Conservative	2.81	1.65	0.194	3.0	0.20	$\eta = 0$
III-C-530	Conservative	2.81	1.65	0.194	3.0	0.20	$t_{\mathrm{rel}} = 5.300$ s
III-C-005	Conservative	2.81	1.65	0.194	3.0	0.20	$\eta = 0.05$
III-C-010	Conservative	2.81	1.65	0.194	3.0	0.20	$\eta = 0.10$
III-C-015	Conservative	2.81	1.65	0.194	3.0	0.20	$\eta = 0.15$
III-C-020	Conservative	2.81	1.65	0.194	3.0	0.20	$\eta = 0.20$
III-C-025	Conservative	2.81	1.65	0.194	3.0	0.20	$\eta = 0.25$

References

[1] K. Calvin, D. Dasgupta, G. Krinner, et al. *IPCC, 2023: Climate Change 2023: Synthesis Report*. Tech. rep. July 2023. DOI: 10.59327/IPCC/AR6-9789291691647.001. URL: https://www.ipcc.ch/report/ar6/syr/.

[2] Unfccc. *Adoption of the Paris Agreement*. Tech. rep. 2015.

[3] T. Hague. *Climate Agreement*. Tech. rep. 2019.

[4] P. Netherlands Environmental Assessment Agency. *Climate and Energy Outlook of the Netherlands 2025*. Tech. rep. 2025. URL: www.pbl.nl/kev.

[5] IEA. *Aviation - IEA*. 2025. URL: https://www.iea.org/energy-system/transport/aviation.

[6] D. S. Lee, D. W. Fahey, A. Skowron, et al. "The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018". In: *Atmospheric Environment* 244 (Jan. 2021), p. 117834. ISSN: 1352-2310. DOI: 10.1016/J.ATMOSENV.2020.117834. URL: https://www.sciencedirect.com/science/article/pii/S1352231020305689.

[7] European Commission. *ReFuelEU aviation - Mobility and Transport - European Commission*. URL: https://transport.ec.europa.eu/transport-modes/air/environment/refueleu-aviation_en.

[8] *Notification document European Commission notification Balanced Approach procedure for Schiphol*. Tech. rep. 2023.

[9] Airbus. *First ZEROe engine fuel cell successfully powers on*. URL: https://www.airbus.com/en/newsroom/stories/2024-01-first-zeroe-engine-fuel-cell-successfully-powers-on.

[10] "Challenge the future Flight and Orbital Mechanics Lecture slides". In: ().

[11] J. Anderson. *Fundamentals of Aerodynamics*. 5th ed. McGraw-Hill, Feb. 2011. ISBN: 9780073398105.

[12] *Induced Drag Coefficient | Glenn Research Center | NASA*. URL: https://www1.grc.nasa.gov/beginners-guide-to-aeronautics/induced-drag-coefficient/.

[13] D. Scholz. "Definition and discussion of the intrinsic efficiency of winglets". In: *INCAS Bulletin* 10.1 (2018), pp. 117–134. ISSN: 22474528. DOI: 10.13111/2066-8201.2018.10.1.12.

[14] *Winglets Save Billions of Dollars in Fuel Costs | NASA Spinoff*. URL: https://spinoff.nasa.gov/Spinoff2010/t_5.html.

[15] *Trans-Sonic Truss-Braced Wing May Help Reduce Fuel Consumption - NASA*. URL: https://www.nasa.gov/image-article/trans-sonic-truss-braced-wing-may-help-reduce-fuel-consumption/.

[16] BC Airplanes. *777-9 Airplane Characteristics for Airport Planning*. Tech. rep. Boeing, 2020.

[17] X. C. Córcoles. "An experimental study on flared folding wingtips". PhD thesis. 2022.

[18] J. P. Q. Hoyng. "Design and Manufacturing of a Demonstrator for the Flared Folding Wingtip Concept". PhD thesis. 2025.

[19] A. Ghosh. "Development of an Aeroelastic Model for a Flared Folding Wing Tip An exploration into the multibody framework of PROTEUS". PhD thesis. 2025.

[20] A. R. Collar. "The Expanding Domain of Aeroelasticity". In: *The Aeronautical Journal* 50.428 (Aug. 1946), pp. 613–636. ISSN: 0368-3931. DOI: 10.1017/S0368393100120358. URL: https://www.cambridge.org/core/journals/aeronautical-journal/article/abs/expanding-domain-of-aeroelasticity/489D68FBDE5D9801BB5615F4F405C069.

[21] Z. Wu, Y. Cao, and M. Ismail. "Gust loads on aircraft". In: *The Aeronautical Journal* 123.1266 (Aug. 2019), pp. 1216–1274. ISSN: 0001-9240. DOI: 10.1017/AER.2019.48. URL: https://www.cambridge.org/core/journals/aeronautical-journal/article/gust-loads-on-aircraft/C37F9E9233E81557F8435CAA406DABB6.

[22] Easa. *Easy Access Rules for Large Aeroplanes (CS-25) (Amendment 27)*. Tech. rep. URL: http://eur-lex.europa.eu/.

[23] H. H. Khodaparast, M. I. Friswell, J. Cooper, et al. *The Effects of Structural Nonlinearity on the Dynamic Response to Aeroelastic Gust Models The effects of structural nonlinearity on the dynamic re-sponse to aeroelastic gust models*. Tech. rep. 2016. URL: https://www.researchgate.net/publication/308761782.

[24] A. P. Ricciardi, M. J. Patil, R. A. Canfield, et al. “Evaluation of Quasi-Static Gust Loads Certification Methods for High-Altitude Long-Endurance Aircraft”. In: *https://doi.org/10.2514/1.C031872* 50.2 (Mar. 2013), pp. 457–468. ISSN: 15333868. DOI: 10.2514/1.C031872. URL: /doi/pdf/10.2514/1.C031872?download=true.

[25] R. E. Bartels. “Development, Verification and Use of Gust Modeling in the NASA Computational Fluid Dynamics Code FUN3D”. In: (2012). URL: http://www.sti.nasa.gov.

[26] B. Hevko and Y. Bondar. “Comparison of Two Methods to Calculate External Loads in Continuous Turbulence”. In: *Aviation* 26.3 (Oct. 2022), pp. 160–168. ISSN: 18224180. DOI: 10.3846/aviation.2022.17788.

[27] J. C. Yeager. “Implementation and Testing of Turbulence Models for the F18-HARV Simulation”. In: (1998).

[28] K. Balunov, V. Chedrik, and F. Ishmuratov. “Aeroelastic Tailoring of Composite Wing by using Multidisciplinary Design Environment”. In: (2018).

[29] S. Guo, J. E. De Los Monteros, and Y. Liu. “Gust Alleviation of a Large Aircraft with a Passive Twist Wingtip”. In: *Aerospace 2015, Vol. 2, Pages 135-154* 2.2 (Apr. 2015), pp. 135–154. ISSN: 2226-4310. DOI: 10.3390/AEROSPACE2020135. URL: https://www.mdpi.com/2226-4310/2/2/135/htm%20https://www.mdpi.com/2226-4310/2/2/135.

[30] C. V. Jutte and B. K. Stanford. “Aeroelastic Tailoring of Transport Aircraft Wings: State-of-the-Art and Potential Enabling Technologies”. In: (2014). URL: http://www.sti.nasa.gov.

[31] J. Sodja, N. P. M Werter, and R. De Breuker. “Design of a flying demonstrator wing for manoeuvre load alleviation with cruise shape constraint”. In: (). DOI: 10.2514/6.2018-2153. URL: https://doi.org/10.2514/6.2018-2153.

[32] J. de la Cierva. *Rotor-Equipped Aircraft*. July 1945.

[33] T. W. Finch and D. W. Briggs. “NACA Research Memorandum: Preliminary Results of Stability and Control Investigation of the Bell X-5”. In: *DATE: F EBRUARY* 24 (1956).

[34] Jan Irving and Robert Davies. *Wing tip device*. 2004.

[35] Thomas Wilson, Martin Herring, John Pattinson, et al. *An aircraft wing with a moveable wing tip device for load alleviation*. 2016.

[36] H. Gu, F. Healy, S. Jayatilake, et al. “Flight dynamics of aircraft incorporating the semi-aeroelastic hinge”. In: *Aerospace Science and Technology* 147 (Apr. 2024), p. 109026. ISSN: 1270-9638. DOI: 10.1016/J.AST.2024.109026. URL: https://www.sciencedirect.com/science/article/pii/S1270963824001597.

[37] Arévalo and Félix. *Update on AlbatrossONE Semi Aeroelastic Hinge Small Scale Flying Demonstrator Project*. Tech. rep. 2022. URL: https://www.researchgate.net/publication/361450092.

[38] M. E. Renzelmann and A. D. Marx. *Application Data 63 Continuation-in-part of Ser*. Tech. rep. 828. 1992, p. 479.

[39] T. Wilson, J. Kirk, J. Hobday, et al. *Small scale flying demonstration of semi aeroelastic hinged wing tips*. Tech. rep. 2019. URL: https://www.researchgate.net/publication/333559132.

[40] Airbus. *Airbus launches extra high performance wing demonstrator to fortify decarbonisation ambition*. Tech. rep.

[41] A. Castrichini, V. Hodigere Siddaramaiah, D. E. Calderon, et al. “Preliminary Investigation of Use of Flexible Folding Wing-Tips for Static and Dynamic Loads Alleviation”. In: *Aeronautical Journal* 121 (2017), pp. 73–94. DOI: 10.1017/aer.2016.108. URL: http://www.bristol.ac.uk/red/research-policy/pure/user-guides/brp-terms/.

[42] R. C. M. Cheung, D. Rezgui, J. E. Cooper, et al. *Testing of a Hinged Wing-Tip Device for Gust Loads Alleviation*. Tech. rep. 2018.

[43] R. C. M. Cheung, D. Rezgui, J. E. Cooper, et al. *Testing of Folding Wingtip for Gust Load Alleviation of a Flexible High Aspect Ratio Wing*. Tech. rep. 2020.

[44] E. Albano and W. P. Hodden. “A doublet-lattice method for calculating lift distributions on oscillating surfaces in subsonic flows.” In: *https://doi.org/10.2514/3.5086* 7.2 (May 1969), pp. 279–285. ISSN: 00011452. DOI: 10.2514/3.5086. URL: /doi/pdf/10.2514/3.5086?download=true.

[45] A. Castrichini, V. Hodigere Siddaramaiah, D. E. Calderon, et al. “Nonlinear folding wing tips for gust loads alleviation”. In: *Journal of Aircraft* 53.5 (2016), pp. 1391–1399. ISSN: 15333868. DOI: 10.2514/1.C033474.

[46] A. Castrichini, J. E. Cooper, T. Wilson, et al. “Nonlinear negative stiffness wingtip spring device for gust loads alleviation”. In: *Journal of Aircraft* 54.2 (2017), pp. 627–641. ISSN: 15333868. DOI: 10.2514/1.C033887.

[47] J. Pattinson, T. Wilson, and M. Herring. *High fidelity simulation of the folding wing tip for loads alleviation*. Tech. rep. 2015. URL: https://www.researchgate.net/publication/327156018.

[48] R. M. Ajaj. “Flight dynamics of transport aircraft equipped with flared-hinge folding wingtips”. In: *Journal of Aircraft* 58.1 (2021), pp. 98–110. ISSN: 15333868. DOI: 10.2514/1.C035940.

[49] G. Dussart, S. Yusuf, and M. Lone. "Identification of In-Flight Wingtip Folding Effects on the Roll Characteristics of a Flexible Aircraft". In: *Aerospace 2019, Vol. 6, Page 63* 6.6 (May 2019), p. 63. ISSN: 2226-4310. DOI: 10.3390/AEROSPACE6060063. URL: https://www.mdpi.com/2226-4310/6/6/63/htm%20https://www.mdpi.com/2226-4310/6/6/63.

[50] F. Healy, R. Cheung, T. Neofet, et al. "Folding Wingtips for Improved Roll Performance". In: *AIAA Scitech 2021 Forum*. American Institute of Aeronautics and Astronautics Inc, AIAA, 2021, pp. 1–24. ISBN: 9781624106095. DOI: 10.2514/6.2021-1153.

[51] D. Sanghi, C. E. Cesnik, and C. Riso. "Roll Maneuvers of Very Flexible Aircraft with Flared Folding Wingtips". In: *AIAA SciTech Forum and Exposition, 2023*. American Institute of Aeronautics and Astronautics Inc, AIAA, 2023. ISBN: 9781624106996. DOI: 10.2514/6.2023-0186.

[52] X. C. Córcoles, C. Mertens, A. Sciacchitano, et al. "Effect of Wing Stiffness and Folding Wingtip Release Threshold on Gust Loads". In: *Journal of Aircraft* 60.6 (Nov. 2023), pp. 1917–1936. ISSN: 15333868. DOI: 10.2514/1.C037108.

[53] Airbus. *AlbatrossOne: Revolutionising Aircraft Wing Design*. 2019. URL: https://www.youtube.com/watch?v=Rq2poT9pW3w.

[54] A. Castrichini, T. Wilson, and J. E. Cooper. *On the Dynamic Release of the Semi Aeroelastic Wing-Tip Hinge Device*. Tech. rep. 2018. URL: https://www.researchgate.net/publication/328137819.

[55] D. Balatti, H. H. Khodaparast, M. I. Friswell, et al. "Experimental and numerical investigation of an aircraft wing with hinged wingtip for gust load alleviation". In: *Journal of Fluids and Structures* 119 (May 2023), p. 103892. ISSN: 0889-9746. DOI: 10.1016/J.JFLUIDSTRUCTS.2023.103892. URL: https://www.sciencedirect.com/science/article/pii/S0889974623000609.

[56] H. Gu, F. Healy, L. Constantin, et al. "Aeroelastic Scaling of a High-Aspect-Ratio Wing Incorporating a Semi-Aeroelastic Hinge". In: *AIAA Journal* (Apr. 2024), pp. 1–13. ISSN: 0001-1452. DOI: 10.2514/1.j063646.

[57] M. Narimani, H. Haddadpour, and S. H. Pourtakdoust. "Model predictive control of a flared folding wingtip for gust load alleviation". In: *Aerospace Science and Technology* 159 (Apr. 2025). ISSN: 12709638. DOI: 10.1016/j.ast.2025.109992.

[58] A. Fallier. *Siemens PLM LMS Virtual.Lab brochure*.

[59] MSC Software Corporation. *MSC Nastran Dynamic Analysis User's Guide Dynamic Analysis User's Guide*. Tech. rep. 1992. URL: http://msc-documentation.questionpro.com.

[60] C. Conti, F. Saltari, F. Mastroddi, et al. "Quasi-steady aeroelastic analysis of the semi-aeroelastic hinge including geometric nonlinearities". In: *Journal of Aircraft* 58.5 (2021), pp. 1168–1178. ISSN: 15333868. DOI: 10.2514/1.C036115.

[61] P. Mastracci, F. Saltari, F. Mastroddi, et al. "Unsteady Aeroelastic Analysis of the Semi Aeroelastic Hinge Including Local Geometric Nonlinearities". In: *AIAA Journal* 60.5 (2022), pp. 3147–3165. ISSN: 1533385X. DOI: 10.2514/1.J061108.

[62] J. Mills and R. Ajaj. "Flight dynamics and control using folding wingtips: An experimental study". In: *Aerospace* 4.2 (June 2017). ISSN: 22264310. DOI: 10.3390/aerospace4020019.

[63] R. C. Cheung, C. Wales, D. Rezgui, et al. "Modelling of folding wing-tip devices for gust loads alleviation". In: *AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, 2018*. American Institute of Aeronautics and Astronautics Inc, AIAA, 2018. ISBN: 9781624105326. DOI: 10.2514/6.2018-0462.

[64] *Nastran Superelements and Modules User's Guide*. 2023. URL: https://help-be.hexagonmi.com/bundle/MSC_Nastran_2023.1_Superelements_and_Modules_User_Guide/raw/resource/enus/MSC_Nastran_2023.1_Superelements_and_Modules_User_Guide.pdf.

[65] R. R. Craig and M. C. Bampton. "Coupling of substructures for dynamic analyses." In: *https://doi.org/10.2514/3.4741* 6.7 (May 1968), pp. 1313–1319. ISSN: 00011452. DOI: 10.2514/3.4741. URL: /doi/pdf/10.2514/3.4741?download=true.

[66] R. R. Craig and A. J. Kurdila. "Fundamentals of Structural Dynamics". In: (2006). URL: https://download.e-bookshelf.de/download/0000/5929/18/L-G-0000592918-0002363113.pdf.

[67] P. P. Friedmann, G. A. Lesieutre, and D. Huang Frontmatter. *Structural Dynamics*. 2023. ISBN: 9781108909617. DOI: https://doi.org/10.1017/9781108909617. URL: www.cambridge.org.

[68] P. Benner, S. Gugercin, and K. Willcox. "A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems *". In: 57.4 (2015), pp. 483–531. DOI: 10.1137/130932715. URL: http://www.siam.org/journals/sirev/57-4/93271.html.

[69] *Superelement User's Guide*. URL: https://simcompanion.hexagon.com/customers/s/article/superelement-user-s-guide-doc9185.

[70] *Direct Matrix Input*. URL: https://nexus.hexagon.com/documentationcenter/en-US/bundle/MSC_Nastran_2021/page/Nastran_Combined_Book/linear/matrix/TOC.Direct.Matrix.Input1.xhtml.

[71] *Simcenter 3D Motion Flexible Bodies*. URL: https://www.aerofem.com/assets/files/Siemens-PLM-Simcenter-3D-Motion-Flexible-Bodies_EN.pdf.

[72] "LMS Virtual.Lab Motion". In: (). URL: www.realworldready.com.

[73] S. Doyle. *pyNastran: Nastran BDF/F06/OP2/OP4 File Reader/Writer/Viewer*. 2024.

[74] D. Balbuena. *Modeling flexible bodies using the Reduced Order Flexible Solid block » Guy on Simulink - MATLAB & Simulink*. Mar. 2021. URL: https://blogs.mathworks.com/simulink/2021/03/22/modeling-flexible-bodies-using-the-reduced-order-flexible-solid-block/.

[75] MathWorks. *File Solid - Solid with properties derived from CAD file - MATLAB*. URL: https://nl.mathworks.com/help/sm/ref/filesolid.html.

[76] MathWorks. *Revolute Joint - Joint with one revolute primitive - MATLAB*. URL: https://nl.mathworks.com/help/sm/ref/revolutejoint.html.

[77] *SciPy - savemat() Function*. URL: https://www.tutorialspoint.com/scipy/scipy_savemat_function.htm.

[78] MathWorks. *External Force and Torque - Apply external force and/or torque to connected frame - MATLAB*. URL: https://nl.mathworks.com/help/sm/ref/externalforceandtorque.html.

[79] J. Murua, R. Palacios, J. Michael, et al. "Applications of the Unsteady Vortex-Lattice Method in Aircraft Aeroelasticity and Flight Dynamics". In: *Progress in Aerospace Sciences* (2012).

[80] R. De Breuker, M. M. Abdalla, and Z. Gürdal. *Energy-Based Aeroelastic Optimisation of a Morphing Wing*. Tech. rep. 2007.

[81] N. Werter. "Aeroelastic Modelling and Design of Aeroelastically Tailored and Morphing Wings". In: (2017). DOI: 10.4233/UUID:74925F40-1EFC-469F-88EE-E871C720047E. URL: https://research.tudelft.nl/en/publications/aeroelastic-modelling-and-design-of-aeroelastically-tailored-and-/.

[82] N. P. Werter, R. De Breuker, and M. M. Abdalla. "Continuous-Time State-Space Unsteady Aerodynamic Modeling for Efficient Loads Analysis". In: *AIAA Journal* 56.3 (Oct. 2018), pp. 905–916. ISSN: 00011452. DOI: 10.2514/1.J056068. URL: /doi/pdf/10.2514/1.J056068?download=true.

[83] R. Palacios, J. Murua, and R. Cook. "Structural and Aerodynamic Models in Nonlinear Flight Dynamics of Very Flexible Aircraft". In: *AIAA Journal* 48.11 (Apr. 2010), pp. 2648–2659. ISSN: 00011452. DOI: 10.2514/1.J050513. URL: /doi/pdf/10.2514/1.J050513.

[84] C. Papadimitriou and D. C. Papadioti. "Component mode synthesis techniques for finite element model updating". In: *Computers & Structures* 126.1 (Sept. 2013), pp. 15–28. ISSN: 0045-7949. DOI: 10.1016/J.COMPSTRUC.2012.10.018. URL: https://www.sciencedirect.com/science/article/abs/pii/S0045794912002520.

[85] R. D. Cook, D. S. Malkus, M. E. Plesha, et al. *Concepts and Applications of Finite Element Analysis*. 2002, p. 733. ISBN: 978-0-471-35605-9. URL: https://www.wiley.com/en-us/Concepts+and+Applications+of+Finite+Element+Analysis%2C+4th+Edition-p-9780471356059.

[86] A. Deperrois. *XFLR5 v6 User Guide: Analysis of Foils and Wings Operating at Low Reynolds Numbers*. 2013.