Fluid-structure interaction simulations of the inflated shape of ram-air parachutes

Nicholas Anthony Thomas Fogell

Department of Aeronautics
Imperial College London

Submitted 17th January 2014
Revised 15th November 2014

A thesis submitted for the degree of
Doctor of Philosophy

Per Ardua ad Astra

Declaration

I certify that all material in this thesis which is not my own work has been properly acknowledged.

N. A. T. Fogell

17$^{\text{th}}$ January 2014

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or distribution, researchers must make clear to others the licence terms of this work.

Abstract

Historically, the development of ram-air parachutes, or parafoils, has been mostly empirical. The use of computational methods has the potential to provide a greater level of insight into the underpinning mechanisms of parachute flight than that offered by experiments alone. To this end, a loosely coupled fluid-structure interaction approach has been developed, and is utilised to predict the equilibrium shape of a parafoil during steady glide, thus permitting an analysis of the associated flow-field. In order to reduce computational cost, a novel method for modelling a single cell of a ram-air canopy is derived, which utilises periodic boundary conditions to produce a response representative of a full canopy. The fabric membrane of the structure is modelled using the non-linear Finite Element Method, and the fluid dynamics are characterized by a finite-volume discretization of the steady Reynolds-Averaged Navier-Stokes equations with a k-ε turbulence model.

Two-dimensional flow simulations reveal a strong influence of the ram-air inlet on the lift and drag performance, greater than that caused by the blunt shape of the leading edge cut alone. The results of the three-dimensional FSI analysis show key differences between the inflated parachute structure and the original cut pattern design. These include rounding of the trailing edge and an increase in the parafoil cross-sectional thickness and the size of the inlet. The results highlight a number of instances of interaction between fluid and structure, such as high structural stresses near the inlet, flow separation from the lower leading edge, and a large degree of separation towards the rear of the upper surface. The importance of spanwise forces on the shape of the canopy has been identified, in addition to the significance of internal pressure forces on lift and drag. Finally, the separated nature of the flow means that flow-field results show dependence on the chosen turbulence model, outlining the requirement for a comprehensive experimental dataset to enable the necessary validation of the fluid dynamics simulation approach.

Acknowledgements

First of all I would like to thank Dr Dan Pope of Dstl for arranging this PhD in the first place and the Ministry of Defence for providing the funding, as without them I wouldn't have had the opportunity to pursue such an interesting, challenging and worthwhile project. The continued support of both parties also made a very big difference. I'm also thankful for the support and advice that my supervisors Prof. Lorenzo Iannucci, Prof. Spencer Sherwin and Dr Colin Cotter have provided through what has been a difficult and complicated task. I'm particularly grateful to the help offered by Dr. Rafael Palicios-Nieto in the latter half of my PhD, during which he was very generous with his time despite having no official link to the project.

I was extremely fortunate to have the chance of regularly discussing my work with members of the DGA in Toulouse throughout my PhD and it was a real pleasure. The encouragement and warm hospitality of Yves Coquet, Eric Poquillon and Basile Perrin was always a highlight and I'm very grateful for their help. I'd also like to thank John Cleminson of Dstl who was often present at these meetings and his good humour was always welcome.

A number of people in the Aeronautics department have made my time as a PhD student a lot easier and a lot more fun. Thanks to my office-mates in the Acex building Sara, Eric and Umit who welcomed me to Imperial. My time in a very cold Mechanical Engineering office was made much more pleasant by Jeremy, as well as our esteemed neighbour Professor Lord Robert Winston for the use of the kitchen in his office. Thanks also to Adam and the rest of the "fish-bowl" office. Cheers also to Mark, Dave, Andy, Rushen, Julien, Olly, Henrik, Abeed and a number of other people who I shared a much needed cup of tea or pint of beer with at the end of a long day. The generous sharing of resources and expertise by Alister and Stefano was also greatly appreciated. The support of my friends outside of Imperial also played a big part in keeping motivated and dealing with the challenges of PhD life.

Finally, I'd like to thank my family for their constant support and understanding, and most importantly my girlfriend Lucy for her constant encouragement, unending support and her belief in me, without which this process would have been infinitely more difficult.

'We choose to go to the moon. We choose to go to the moon in this decade and do the other things, not because they are easy, but because they are hard . . . '

John F. Kennedy (September 1962)

Contents

1 Introduction **30**
1.1 Motivation 30
1.2 Research objectives 33
1.3 Thesis outline 33
1.4 Related publications 35
1.5 Thesis contributions 35

2 State of the art in the design and development of ram-air parachutes **36**
2.1 Ram-air parachute design and physics of operation 36
2.2 Experimental studies 42
2.3 Computational studies 46
2.3.1 Simulation of fabric materials 50
2.4 Summary 56

3 Single-cell ram-air parachute model and FSI simulation framework **58**
3.1 Single cell parafoil model 58
3.2 FSI simulation framework 60
3.2.1 Fluid-structure coupling approaches 60
3.2.2 Structural modelling 61
3.2.3 Fluid modelling 65
3.2.4 Loose coupling methodology 74
3.2.5 Assessment of convergence for FSI iterations 78
3.3 Additional modelling issues 79
3.3.1 Treatment of out-of-plane deformation at the trailing edge . . . 79
3.3.2 Treatment of buckling behaviour at the trailing edge 81
3.3.3 Effect of fabric wrinkling on overall deformation 82
3.3.4 Application of Taubin smoothing to deformed structural geometry 84

3.3.5 Effect of pressure sampling resolution on calculation of aerodynamic loading . . . 85

4 Parachute geometry and properties 88
4.1 Parachute geometry . . . 88
4.2 Mesh resolution . . . 90
4.2.1 Structural domain . . . 90
4.2.2 Fluid domain . . . 91
4.3 Flow parameters . . . 95
4.4 Material properties of parachute fabric . . . 96

5 Results 101
5.1 Two-dimensional simulations of flow around ram-air parachute . . . 101
5.1.1 Flow field results . . . 103
5.1.2 Lift and drag comparisons . . . 115
5.1.3 Comparison to other two-dimensional parafoil flow studies . . . 120
5.1.4 Summary and discussion . . . 121
5.2 Initial 3D FSI simulation of infinite span ram-air parachute . . . 124
5.2.1 Prediction of inflated shape . . . 124
5.2.2 Flow field results . . . 129
5.2.3 Summary and discussion . . . 137
5.3 Improved 3D FSI simulation of infinite span ram-air parachute . . . 139
5.3.1 Prediction of inflated shape . . . 140
5.3.2 Flow field results . . . 146
5.3.3 Lift and drag results . . . 153
5.3.4 Summary and discussion . . . 157
5.4 Use of Standard k-ε model for simulation of 3D flow field around inflated parafoil . . . 160
5.4.1 Flow field results . . . 161
5.4.2 Lift and drag coefficient comparison . . . 166
5.4.3 Summary and discussion . . . 168

6 Conclusions 171
6.1 Summary . . . 171
6.2 Primary contributions . . . 178
6.3 Recommendations for future work . . . 179

References **182**

Appendices

A Parafoil cross-section co-ordinates **194**

B Examples of scripts and macros employed in loose coupling framework **198**

B.1 Workflow example for fluid-structure iteration 198

B.2 Scripts for post-processing structural geometry 200

B.3 Post-processing macro for extracting pressure results from Star-CCM+ 203

C Header files for structural simulation **209**

C.1 Example LS-DYNA input files 209

C.1.1 Example LS-DYNA main header file 209

C.1.2 Example LS-DYNA structural header file 210

C.1.3 Example LS-DYNA loading file 212

C.1.4 Example LS-DYNA loading segment set file 212

List of Figures

2.1 Diagram showing key components in the design of a ram-air parachute. 37
2.2 Aerodynamic forces acting on a parafoil canopy during steady glide. . 39
2.3 Front view of forces on a parafoil during steady flight. 41
2.4 Experimental wrinkling pattern in Kapton membrane under-in-plane shear. 51
2.5 Principal status of a membrane - taut, wrinkled and slack. 53
2.6 Principal stress directions and magnitudes for Kapton membrane under in-plane shear. 54
2.7 Principal stresses membrane under in-plane shear with and without membrane wrinkling model. 55

3.1 View of (a) un-deformed half-cell parafoil geometry, (b) half-cell with periodic symmetry applied, and (c) single-cell of BC2 case reflected about xy plane to show unit-cell. 59
3.2 Boundary conditions for structural models in 3D FSI cases. 64
3.3 Schematic showing a 2D side view of the 3D fluid domain and applied boundary conditions. 73
3.4 Diagram outlining transfer of surface pressure data between fluid and structural simulations. 75
3.5 Description of FSI coupling algorithm. 77
3.6 Out of plane displacement seen at the trailing edge when planar constraint is not applied. 80
3.7 Buckling and folding seen at (a) the trailing edge, and the improvement gained by activating the Compressive Stress Elimination (CSE) wrinkling model seen from (b) below and (c) above. 82

3.8 Comparison of results for structural simulation test case. Front (left) and side (right) views of the leading edge deformation: (a) standard membrane elements applied to ribs; (b) CSE wrinkling option applied to ribs. 83
3.9 View of leading edge area of parafoil, showing (a) small mesh distortions on upper surface, and (b) the upper surface area after Taubin smoothing filter is applied. 85
3.10 Comparison of chordwise pressure loadings for the upper surface of the structural model, for two different sampling resolutions, at three spanwise planes. 87

4.1 Cross-section of Ware & Hassell parafoil, and the adjusted cross-section used throughout this study. 89
4.2 Un-deformed half-cell parafoil geometry with main dimensions. 90
4.3 Mesh resolution convergence study for parafoil model. 91
4.4 Side view of fluid domain around an undeformed parafoil showing (a) overview of fluid mesh and (b) schematic of regions where high resolution is prescribed. 92
4.5 Magnified side views of the mesh resolution in the areas surrounding the parafoil. 93
4.6 Figures showing the construction of the boundary layer mesh at (a) the upper leading edge, (b) lower leading edge, and (c) the trailing edge of the parafoil. 94
4.7 Mesh resolution convergence study for (a) lift and (b) drag coefficients of initial undeformed 3D parafoil model. 95
4.8 (a) close up of lower surface fabric under tensile load and (b) view of fabric sample within testing apparatus post-failure. 99
4.9 Stress-strain results from tensile testing of upper surface fabric, in warp and weft directions. 99

5.1 Two-dimensional chordwise cross-sections showing geometry of three different leading edge configurations. 102
5.2 Streamlines for the flow around rounded leading edge model. 103
5.3 Normalized velocity contours for flow around rounded leading edge model. 104
5.4 C_p contours for flow around rounded leading edge model. 105

5.5 Vorticity contours for flow around rounded leading edge model. 105
5.6 Streamlines for flow around parafoil cross-section with closed inlet. . . 106
5.7 Normalized velocity contours for flow around parafoil cross-section with closed inlet. 107
5.8 C_p contours for flow around parafoil cross-section with closed inlet. . . 107
5.9 Vorticity contours for flow around parafoil cross-section with closed inlet.107
5.10 Streamlines for flow around parafoil cross-section with open inlet. . . . 110
5.11 Normalized velocity contours for flow around parafoil cross-section with open inlet. 112
5.12 C_p contours for flow around parafoil cross-section with open inlet. . . 112
5.13 Vorticity contours for flow around parafoil cross-section with open inlet. 112
5.14 Comparison of C_p recorded at the external parafoil surfaces, for three leading edge configurations. 113
5.15 Comparison of wall shear stress recorded at the upper and lower external parafoil surfaces, for three leading edge configurations. 114
5.16 Comparison of aerodynamic force coefficient values for three leading edge configurations. 115
5.17 Lift coefficient values, for three leading edge configurations, separated in terms of structural components. 116
5.18 Comparison of pressure and friction drag contributions, for three leading edge configurations. 118
5.19 Drag coefficient values, for three leading edge configurations, separated in terms of structural components. 119
5.20 Vector plot of pressure coefficient distribution on open ram-air parafoil section. 120
5.21 Pseudo-residuals showing difference between the predicted structural geometry with each FSI iteration. 125
5.22 Resultant 3D geometry of inflated ram-air parafoil cell with infinite canopy boundary conditions. 126
5.23 Chordwise cross-section of the parafoil in three cases: unstressed shape at the rib (cut pattern), inflated shape at the rib, and inflated shape at the centre-line. 127
5.24 3D streamlines for flow around inflated shape of ram-air parafoil cell with infinite canopy boundary conditions. 130

5.25 Chordwise C_p profiles for external surfaces of parafoil at different spanwise positions. 131
5.26 Chordwise cross-section contour plots of flow at 50% span. 132
5.27 Flow contours over the upper surface of canopy. Flow from left to right. 132
5.28 Chordwise cross-sections at different spanwise positions showing constrained streamlines. 133
5.29 Flow contours over lower surface of canopy. Flow from left to right. . . 136
5.30 Convergence of resultant displacement values at several nodal points on the finite element model. 140
5.31 Resultant 3D geometry of inflated ram-air parafoil cell with improved boundary conditions. 141
5.32 Chordwise cross-section of 3D inflated shape taken at various points across the span. 144
5.33 Front isometric view of inflated geometry with contours of maximal principal strain. 145
5.34 3D streamlines for flow around inflated shape of ram-air parafoil with updated boundary conditions. 147
5.35 Chordwise cross-sections at different spanwise positions showing constrained streamlines. 148
5.36 Chordwise cross-sections at different spanwise positions showing contours of C_p. 150
5.37 Flow over upper surface of canopy. Flow from left to right. 151
5.38 Chordwise cross-sections at different spanwise positions showing contours of normalized velocity. 152
5.39 Flow over lower surface of canopy. Flow from left to right. 153
5.40 Lift coefficient values, for both internal and external surfaces, for inflated parafoil cell, separated by structural component. 154
5.41 Drag coefficient values, for both internal and external surfaces, for inflated parafoil cell, separated by structural component. 155
5.42 3D streamlines for flow around inflated shape of ram-air parafoil cell, simulated using S.k-ε turbulence model. 162
5.43 Flow over upper surface of canopy predicted using S.k-ε model. 163
5.44 Chordwise cross-sections at different points across the span, showing constrained streamlines with contours of normalized velocity, predicted using R.k-ε and S.k-ε turbulence models. 164

5.45 Chordwise cross-sections at different points across the span, showing C_p predicted using R.k-ε and S.k-ε turbulence models. 165
5.46 Flow over lower surface of canopy predicted using S.k-ε model. 165
5.47 Comparison of predicted aerodynamic force coefficient values to to experimental measurements for two different RANS turbulence models. . 167
5.48 Comparison of drag coefficient values, by structural component, for inflated parafoil shape simulated using R.k-ε and S.k-ε models. 167
5.49 Comparison of drag coefficient values, by component, for inflated parafoil shape using R.k-ε and S.k-ε models. 168

List of Tables

2.1 Summary of possible criteria for assessing the wrinkle state of a membrane under loading. 54

4.1 Flow parameters applied in all fluid dynamics simulations. 96
4.2 Approximated material properties utilised in periodic symmetry, loaded rib parafoil model, BC1. 96
4.3 Test parameters for fabric material tensile strength tests. 97
4.4 Results for Young's Moduli of nylon fabric in two principle directions, calculated from tensile test results. 98
4.5 Material properties utilised in half-symmetry, unloaded rib parafoil model, BC2. 100

5.1 Lift coefficient for upper surface, for open and closed parafoil model. . 117
5.2 Comparison of drag coefficient for open and closed parafoil model. . . 119
5.3 Comparison of simulation lift and drag results for open-inlet parafoil section to those of other published 2D flow studies. 121
5.4 Comparison of structural dimensions of the simulated ram-air parafoil to wind-tunnel test results. 128
5.5 Comparison of predicted structural dimensions of inflated ram-air parafoil to experimental values. 142
5.6 Drag coefficient results for inflated canopy: pressure and shear components separated by component. 156
5.7 Comparison of simulation results to wind tunnel results, Re=2.1million. 156

A.1 Original Ware & Hassell parafoil cross-section coordinates. 195
A.2 Adjusted parafoil cross-section coordinates used throughout this thesis. 196

A.3 Rounded leading edge geometry cross-section coordinates used within 2D study, Section 5.1. 197

Nomenclature

Symbols

A	cross-sectional area
A_{ref}	reference area
c	chord length
C_{D}	drag coefficient
C_{L}	lift coefficient
C_{p}	pressure coefficient
D	aerodynamic drag
E	Young's modulus
e	engineering strain
F	force
f^B	body forces
f^S	traction force
G	shear modulus
g	acceleration due to gravity

k	turbulent kinetic energy
ΔL	change in length
ℓ	turbulent length scale
L	aerodynamic lift
L/D	lift to drag ratio
L_0	gauge length
L_c	characteristic length
m_p	mass of payload
n	normal unit vector
p	pressure
q_∞	freestream dynamic pressure
R	resultant upwards force acting on canopy
Re	Reynolds number
S_{min}, S_{max}	Principal stresses of a membrane
T	suspension line tension
t	time
u	displacement
V	velocity
v_∞	freestream velocity
V_h	horizontal component of velocity
V_v	vertical component of velocity

v_x, v_y, v_z Cartesian velocity components

v_{turb} turbulent velocity scale

x, y, z Cartesian coordinates

y^+ non-dimensionalised wall distance

Greek Symbols

α angle of attack, degrees

β under-relaxation coefficient

δ_{ij} Kronecker delta

η engineering stress

ε turbulent energy dissipation

$\varepsilon_{min}, \varepsilon_{max}$ Principal strains of a membrane

Γ fluid-structure interface

γ canopy glide slope, degrees

λ_{TS} Taubin smoothing coefficient 2

μ dynamic viscosity

μ_t turbulent eddy viscosity

μ_{TS} Taubin smoothing coefficient 1

ν kinematic viscosity

ν_t turbulent kinematic viscosity

Ω_F fluid domain

Ω_S structural domain

Φ structural geometry as set of nodal coordinates

ρ density

σ Poisson's ratio

θ_{cut} angle of parafoil inlet to horizontal

τ Cauchy stress tensor

Subscripts and superscripts

$(\bar{\bullet})$ mean component of value

$(\bullet')$ fluctuating component of value

$(\bullet)_F$ property of fluid domain

$(\bullet)_S$ property of structural domain

$(\bullet)_{initial}$ initial value

$(\bullet)_{old}$ value from previous simulation

$(\bullet)_{new}$ value from latest simulation

$(\bullet)^*$ under-relaxed value

Abbreviations

CFD Computational Fluid Dynamics

CPU Central Processing Unit

CSD Computational Structural Dynamics

CSE Compressive Stress Elimination

DNS Direct Numerical Simulation

FEM Finite Element Method

FSI Fluid-Structure Interaction

FVM Finite Volume Method

GPS Global Positioning System

GNC Guidance, Navigation and Control

IED Improvised Explosive Device

LDV Laser Doppler Velocimetry

LES Large Eddy Simulation

PIV Particle Image Velocimetry

RANS Reynolds-Averaged Navier-Stokes

R.k-ε Realizable k-epsilon turbulence model

S.k-ε Standard k-epsilon turbulence model

Chapter 1

Introduction

1.1 Motivation

The use of parachutes for the delivery of cargo or stores has been commonplace since World War II, and such operations became known as airdrops [1]. Airdrop missions provided a new method to quickly and directly supply troops with the equipment and supplies that they required. Originally, both personnel and cargo were dropped beneath round ballistic parachutes. These canopies were of simple construction, making them reliable and cheap to produce [2]. The downside of these parachutes however was their susceptibility to the wind, which meant they could easily be blown off course, resulting in large errors between the intended and actual points of impact [3]. The consequences of an airdropped payload landing off course can be serious, and may result in loss of the cargo through damage or capture by opposing forces. Therefore tactics and procedures were changed over time, such as reducing the altitude of drops and calculating the point of release from the aircraft based on recorded or predicted wind data.

The invention of the ram-air parachute canopy, or parafoil, by Jalbert in 1964, permitted the development of high-performance, steerable, gliding parachutes, which offered the potential to solve a number of the issues associated with traditional airdrops [4]. The key capabilities of direction-control and gliding flight possessed by the ram-air parachute made it ideal for use in precision guided airdrop systems, and a number of these have been developed over the years. A typical system consists of a guidance

unit, positioned between the payload and the ram-air canopy, which is used to autonomously guide the system to its pre-determined target. An onboard GPS receiver provides the guidance, navigation and control (GNC) software within the unit with a reading of the current position, and steering is actuated through servo motors attached to the steering lines as appropriate. The combination of onboard GNC, a steerable canopy, and the capability for a gliding canopy to fly against the wind, means that such systems offer the capability to accurately air-drop cargo regardless of the weather conditions, at significant lateral and vertical standoff distances [5].

Clearly, autonomous precision guided airdrop systems afford a number of operational advantages, and have been successfully fielded for a number of different applications in the past, both military and humanitarian [6, 7]. More recently, the experience of the conflicts in Iraq and Afghanistan have shown an increased prevalence of attacks on road convoys, particularly in terms of Improvised Explosive Device (IED) use. This has provoked a reassessment of resupply methods, and a renewed interest in the use of precision aerial delivery techniques to provide logistical support, particularly to posts that may be spread over a large geographical area, thereby reducing the reliance on overground lines of supply [8].

The renewed interest in the deployment of precision aerial delivery systems has prompted a large number of experimental trials, and these have shown that whilst the vulnerability of a ram-air based system to adverse weather conditions is reduced, the overall levels of landing accuracy are not as high as would be desired. Moreover, the accuracy levels do not match those of a skilled human parachutist using a similar canopy [9]. Thus, research is required to understand the reason behind the apparent shortfall in accuracy, and to make improvements accordingly.

Historically, parachute design and development has been largely an empirical process, driven by exhaustive programmes of experimental testing [10, 11, 12, 13]. It is acknowledged in the literature that it is unlikely that there will be any giant leaps forward made in parachute design whilst the dominance of testing persists, and it is also an increasingly expensive and time-consuming effort [14, 15, 16, 17]. For these reasons, computational modelling and simulation methods are now sought for use in parafoil design [18].

In addition to a reduction in cost, the use of numerical methods for parafoil development provides a means to move away from “cut-and-try” methods or semi-empirical

relations, and replace them with physics-based formulations [18]. The movement towards a greater use of simulation tools results in an enlargement of the design-space by allowing a greater number of potential configurations to be tested and investigated computationally, in a manner that would be difficult to achieve through conventional experiments. Additionally, computational analysis has the potential to provide a higher range and fidelity of measurements, enabling observations to be made that would otherwise be difficult or impossible to record through experiment alone [19, 20]. Examples include precise fluid velocity and pressure field measurements, or structural stress fields [21].

Due to the high flexibility of the fabric canopy structure, a parafoil in flight exhibits a strong fluid-structure coupling. The shape of the parachute canopy in-flight is highly dependent on the flow field that surrounds it, and conversely the flow behaviour is strongly influenced by the shape of the canopy [22]. Consequently, if numerical methods are to be used to analyse this system, they must be capable of accurately modelling the behaviour of the fluid, the structure, and the interaction between the two [23].

A major goal for the use of modelling and simulation would be to use it to simulate the flow field around an inflated ram-air parachute, during steady glide, thereby providing a prediction of the aerodynamic performance of the canopy. From this, insight can be gained in the influence of the shape of the parafoil over the aerodynamic performance, and therefore a greater understanding can be reached of how the design of the canopy can be optimised to improve the canopy's flight characteristics. Ultimately, the aim would be to reach a point where the technique can be directly employed in the design stage of the parafoil development cycle.

To achieve this goal, the first requirement is for an accurate prediction of the ram-air canopy inflated shape to be acquired, as without this, an assessment of the fluid dynamics around the canopy cannot be attained, and the flight performance cannot be assessed. This challenge can be met by developing and using an appropriate Fluid-Structure Interaction (FSI) simulation technique, such that the aforementioned strong coupling between the fluid and structural forces is accurately captured, enabling a prediction of the equilibrium state of the parachute structure during gliding flight to be calculated. This prediction of the inflated shape can then be used to perform an analysis of the flow around the parafoil. The solution to this challenge forms the main

driver behind the research effort reported within this thesis.

It is to this end that this work aims to investigate, develop and apply an FSI simulation methodology to the simulation of the inflated shape of a representative ram-air parachute canopy during steady flight. This effort will also demonstrate the value of applying these techniques to the study of a parafoil canopy, and advance the state of the art with regard to the use of computational methods in ram-air parachute research and development.

1.2 Research objectives

The key research objectives of this work are as follows:

1. **Development of loosely coupled fluid-structure interaction simulation methodology**
   To develop a methodology by which fluid-structure interaction simulations can be performed in a loosely coupled fashion and applied to the prediction of the inflated shape of a ram-air parachute.
2. **Predict inflated shape of a parafoil and its associated flow field**
   Utilise the aforementioned simulation procedures to predict the inflated shape of a parafoil during steady glide, and simulate the associated flow field around the inflated canopy.
3. **Use simulation results to improve understanding of parafoil performance**
   Analyse simulation results to further the understanding of the relevant factors behind the aerodynamic performance of parafoils, and the principle areas of mutual interaction between the structure and the flow.

1.3 Thesis outline

The chapters of this thesis are broken down in the following manner. Chapter 2 provides a summary of relevant background relating to the design, construction and underlying physics behind ram-air parachutes. Following this, the remainder of the

chapter consists of a survey of the literature, outlining the challenges associated with the study and improvement of parafoil flight performance, and identifying the need for FSI simulations as an underpinning requirement to the successful analysis of parafoil aerodynamics.

Chapter 3 begins by outlining the novel approach used in this work to represent the behaviour of a ram-air canopy. It models a single half-cell of the parafoil structure, with the appropriate boundary conditions such that it behaves as though it was part of a chain of cells, forming a parafoil canopy of infinite span. Beyond that, the key principles behind coupled fluid-structure simulation techniques are described, followed by an explanation of the specific modelling techniques which make up the loosely coupled FSI simulation framework used in this work. This chapter includes an explanation of the two specific boundary condition cases investigated in this thesis, and the necessary adaptations made to the models which ensure an accurate representation of the canopy.

Chapter 4 provides the parafoil geometry adopted for use in the simulations within this work, and provides further details relating to the parafoil simulation such as mesh size and resolution, material properties and flow parameters.

Chapter 5 includes the simulation results, which are separated into four main sections. The first is a 2D fluid dynamics investigation into the influence of a rounded, closed or open leading edge on the flow around a parafoil cross-section, to provide insight into the likely reasons behind the reduction in aerodynamic performance seen in parafoils when compared to an equivalent aerofoil. The results also serve to act as a comparison against which the three-dimensional results reported subsequently can be compared. This is performed using two-dimensional computational fluid dynamics simulations on an undeformed rigid geometry.

The second and third sections of chapter 5 cover the results of FSI simulations using two different approaches to the parafoil boundary conditions, with details of results for both the structural and fluid domains. The fourth and final section within this chapter makes use of the geometry predicted using the loosely coupled FSI methodology, and compares the results of the predicted flow field when predicted using two different turbulence models, namely the Standard and Realizable k-ε models.

Finally, chapter 6 provides a summary of the key conclusions to the work, a summary of the key contributions of this thesis, and suggestions for future work.

1.4 Related publications

The work in this thesis has resulted in the following conference publication and presentation, for which the author was awarded the prize for best student paper.

> Fogell, N., Sherwin, S. J., Cotter, C. J., Iannucci, L., Palacios, R. & Pope, D. J. Fluid-Structure Interaction Simulation of the Inflated Shape of Ram-Air Parachutes. In *22nd AIAA Aerodynamic Decelerator Systems Conference*, Daytona Beach, FL., USA, 25–28 March. 2013.

1.5 Thesis contributions

The primary contributions of this thesis are as follows:

- Presentation of a novel approach for performing the aforementioned FSI analysis in a computationally efficient manner, namely by modelling a single-cell of a parafoil such that it behaves as though it was part of a full canopy.
- Development of a viable methodology for performing loosely coupled FSI simulations of ram-air parachutes.
- Demonstration that the presence of an open inlet has a greater impact on parafoil aerodynamics than that of the blunt nose arising from the associated leading edge cut.
- Successful utilisation of FSI techniques to predict the inflated shape of a ram-air parachute and the surrounding flow field.
- Demonstration of the importance of a spanwise force in determining the shape of the canopy.
- Improvement of derived boundary conditions for a single-cell model by including representation of unloaded ribs and spanwise loading.
- A computational investigation into the main sources of fluid-structure interaction and their influence on the resulting canopy shape or flow field.
- Clear demonstration of the potential benefits of including FSI simulation techniques in parafoil design.

Chapter 2

State of the art in the design and development of ram-air parachutes

In this chapter, a summary of the design and construction of ram-air parachute canopies, as well as the basic physics behind their operation, is provided. In addition, the relevant terminology is introduced. More detailed descriptions of ram-air parafoils and parachutes in general, including more detailed discussions of properties such as flight performance and stability, is provided in the review papers by Lingard [24, 25], Knacke [26] and Cockrell [27]. Once the key terms and concepts have been introduced a review of the literature is presented, outlining noteworthy experimental and computational studies relating to the research and development of ram-air parachutes.

2.1 Ram-air parachute design and physics of operation

The design of the ram-air parachute in 1964 by Jalbert [4], proved to be a significant improvement over the traditional round parachute canopy. The ram-air canopy, or parafoil, is a type of parachute which inflates upon opening to effectively become a low-aspect ratio wing. The parafoil is constructed solely of fabric, with no rigid components, and so is packed, folded and deployed exactly as in the case of a round

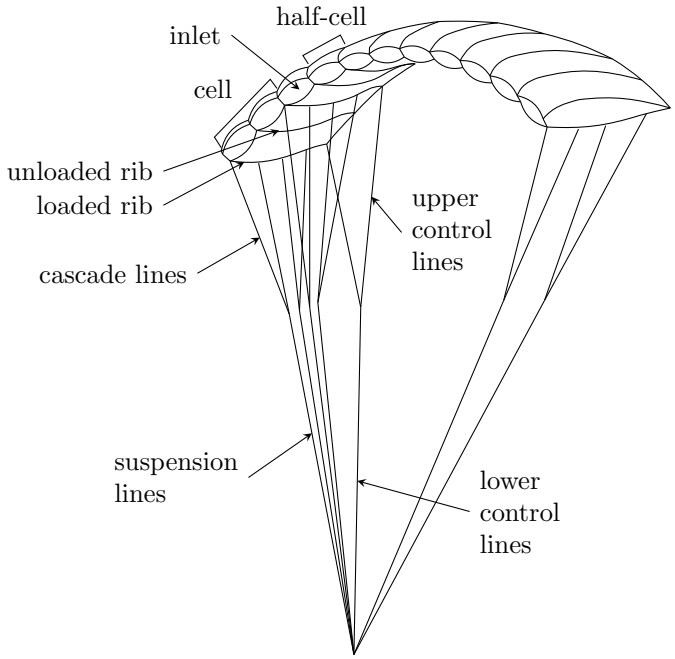


Figure 2.1: Diagram showing key components in the design of a ram-air parachute. Adapted from Lingard [25].

canopy. The design of a typical parafoil is shown in Figure 2.1. The canopy structure consists of two surfaces, one upper and one lower, held together by aerofoil shaped ribs. The ribs are incorporated at set intervals along the span, and are sewn along the full length of the chord, dividing the canopy into sections, known as cells. The entirety of the trailing edge is sewn closed, but the leading edge is open. The opening at the front acts as an inlet, which allows air to enter the cells of the canopy upon deployment, inflating it through a build up of ram-air pressure [27]. To maximise the inflation pressure and hence, form a strong, stable wing shape, the fabric selected for the upper surface is often treated to make it impermeable to air [28]. This also acts to prevent any adverse effects on the aerodynamics that interference between air passing through the upper surface and the the external flow could cause. Typically the material selected for the canopy fabric is nylon, constructed using a ripstop weave, although other materials such as polyester, cotton, and polypropylene have been used [29, 30].

Ram-air canopies typically have a rectangular planform, beneath which the payload is suspended by a series of cords, called suspension lines. The lines are usually attached to the lower canopy surface at every other rib in order to reduce their number, thereby

reducing the level of drag they generate. Suspension lines are most commonly made of nylon, but more recently new man-made fibres such as Spectra® have been used [31, 30]. The lines are attached at multiple, equally spaced points along the length of a rib in order to maintain the shape of the lower surface and of the parafoil as a whole. Ribs that have suspension lines attached are called loaded ribs, and the ribs that do not are called unloaded ribs. The sections between alternating loaded and unloaded ribs are known as half-cells, and two half-cells between a pair of loaded ribs form a cell. Commonly, ribs will have several holes cut into them to allow the ram-air pressure to equilibrate between cells, aiding inflation and preventing collapse of the canopy during flight. This helps to avoid malfunctions such as "end-cell tuck" [26].

The suspension lines, also known as rigging lines, are arranged in a cascaded fashion, shown in Figure 2.1, in an additional effort to reduce the aerodynamic drag force created by the lines, known as line drag [32]. The lines are typically 0.6 – 1.0 spans long, and in modern canopies are all of equal length. This results in the canopies flying with arc-anhedral, which improves stability and manoeuvrability in flight [24]. In addition to the suspension lines, an separate set of lines is attached to the trailing edge, one set on each spanwise half of the canopy. These are the steering lines, used to control the direction and speed of the canopy. Pulling down on the steering lines deforms the trailing edge. Pulling the lines on only one side of the canopy causes a "skid-turn" towards that side, as the deformation of the trailing edge causes that half of the canopy to slow down. Pulling on both sets of control lines at once slows the whole canopy, which is often used to perform a "flare" manoeuvre upon landing. This reduces both forward and vertical velocity, thereby resulting in a soft landing.

Popular cross-sections chosen for parafoils include the Clark-Y aerofoil cross-section, and the more modern NASA LS1-0417 [25]. The angle and position of the inlet has seen a large degree of variation in different parafoil designs, as it is seen to have a significant influence over inflation, as well as lift and drag properties. The incorporation of an aerofoil style cross-section within a ram-air canopy is a key feature that enables the parachute to glide, and thus it behaves very differently to a traditional round, ballistic parachute. The forces acting on a ram-air canopy during steady glide are shown in Figure 2.2. Similarly to a ballistic parachute, the weight of the payload, $m_{\mathrm{p}}g$, where m_{p} is the total mass of the parafoil system including payload, is balanced by an equal and opposite resultant force, R. Unlike a ballistic parachute however, the resultant force in a ram-air canopy consists of components of both the

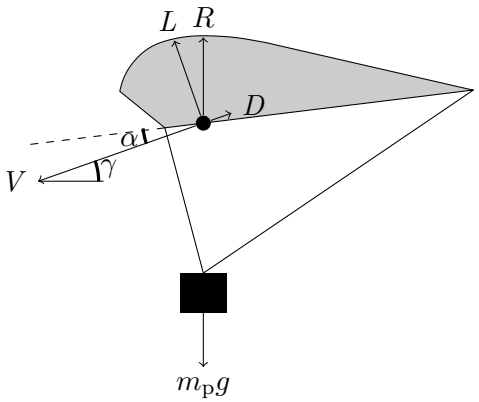


Figure 2.2: Aerodynamic forces acting on a parafoil canopy during steady glide. Forces on lines and aerodynamic forces on payload not shown.

lift force L and drag force D, illustrated in Figure 2.3. The resultant force is defined in Equation (2.1) [26].

$$R = \sqrt{L^2 + D^2} \tag{2.1}$$

A parafoil canopy flies at a particular angle of attack, α, to the system trajectory, according to how the rigging lines are arranged relative to the centre of gravity of the system. The line of the trajectory is denoted by the arrow representing velocity, V in Figure 2.2. The angle to the horizontal that the trajectory of the parachute system takes, is known as the glide slope, denoted by γ. The glide slope is related to the ratio of horizontal to vertical velocities (V_h and V_v respectively), which Equation (2.2) shows to be equal to the ratio of lift to drag [27]. This ratio is an important performance measure for parafoils, as it determines the horizontal range of the canopy, and directly relates to the ability of a canopy to travel into the wind [25]. The lift to drag ratio is given by,

$$\frac{V_\mathrm{h}}{V_\mathrm{v}} = \frac{L}{D} = \frac{C_\mathrm{L}}{C_\mathrm{D}} = \frac{1}{\tan\gamma} \tag{2.2}$$

where the lift and drag coefficients, C_L and C_D are calculated in the same manner as for aerofoils with a reference area A_ref at a freestream dynamic pressure q_∞,

$$C_\mathrm{L} = \frac{L}{q_\infty A_\mathrm{ref}} \qquad\qquad C_\mathrm{D} = \frac{D}{q_\infty A_\mathrm{ref}} \tag{2.3}$$

The ratio L/D is also referred to as the glide ratio, and typical values range from 2:1 to 5:1 for ram-air canopies. The glide ratio is seen to increase with aspect ratio, due to a reduction in induced drag, in the same manner as for rigid wings [33]. However, in the case of a ram-air parachute, there are practical limitations to increasing the aspect-ratio, as higher aspect-ratio canopies are harder to deploy. A higher aspect-ratio also requires more suspension lines, increasing line drag, thus negating the improvement in L/D [24]. This trade-off in design means an optimal value for the aspect ratio is around 3 [25].

The relation of the forces on the canopy shown in Figure 2.2 to the whole canopy-payload system can be seen in Figure 2.3. The force due to the mass of the payload, $m_{\mathrm{p}}g$, is transmitted from the payload to the canopy through the tension, T, in the suspension lines. The lift force generated by the canopy as a whole is a sum of the lift produced by the individual cells of the canopy, the vectors of which are labelled in the figure. As Figure 2.3 clearly demonstrates, the component of the lift produced at each cell in the direction opposing the payload weight varies across the span, due to the arc of the canopy. The arc means that the spanwise component of the lift increases towards the tips of the canopy, which for a parafoil with anhedral rigging may contribute to keeping the canopy open. Another force acting to keep the cells of the parafoil open is the spanwise force acting on the external rib of the cell at each tip of the canopy. This force is generated by the pressure difference between the stagnation pressure inside the cell, and the dynamic pressure outside of the cell [24].

The generation of lift by a parafoil, due to its aerofoil-shaped cross-section, is a major improvement over traditional round parachutes, which only generate a force that opposes the direction of motion. Coupled with the ability to steer and to perform a soft landing, this makes ram-air parachute canopies extremely attractive for use in airdrop operations, as a replacement for the traditionally used round parachute [6]. The incorporation of ram-air parachutes into airdrop operations has led to the development of precision aerial delivery systems, which are autonomous and guided. The use of a ram-air canopy enables the unit to be steered, under control of a GPS guided system, to guide it to its pre-determined destination point. The level of control offered by such a system reduces the effect of the wind on airdrop accuracy, which traditional ballistic canopies are extremely susceptible to [25]. Removing this restriction also has further benefits, as “low and slow” airdrop operations, where items are dropped at low altitudes and low velocity to improve accuracy, are no longer required. The

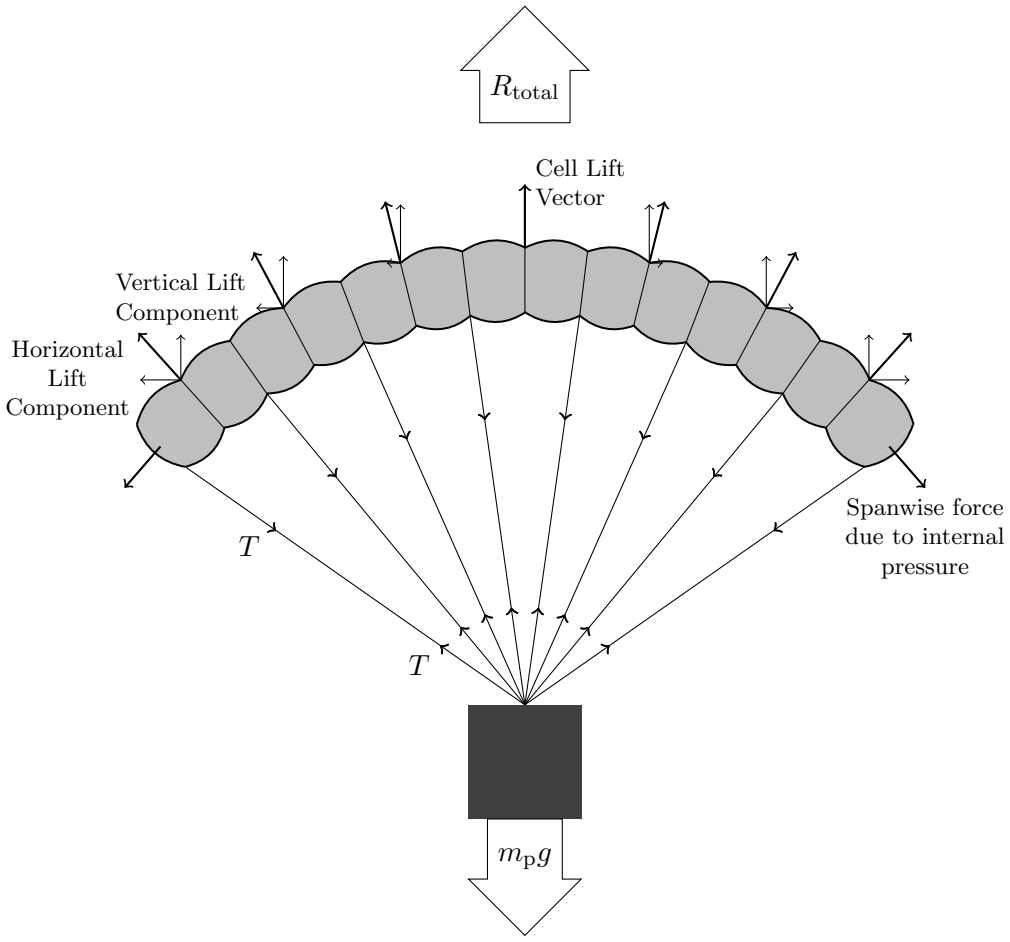


Figure 2.3: Front view of forces on a parafoil during steady flight. Adapted from Sobieski [34]

gliding ability of the canopy means that payloads can be dropped from greater vertical and horizontal distances, reducing the risk to the aircraft and making the airdrop operation less detectable to opposition forces [8]. Therefore, precision aerial delivery systems offer a number of significant advantages, which have motivated a number of experimental and computational efforts, including this one, to understand and improve the performance of ram-air parachutes, with a view to bettering the accuracy and reliability of automated precision airdrop operations.

However, due to the proprietary nature of the majority of parafoil development, a large proportion of previous research has not been published. Despite this, a number of notable studies on the behaviour of ram-air parachutes can be found in the literature. It

is evident from these studies that the vast majority of parachute research and development has been undertaken through the use of experimental testing, a view supported by the recent review paper by Potvin et al. [14]. Numerous shortcomings of experimental testing have been revealed or identified through previous test programmes, and this has encouraged the application of computational simulation techniques to the study of ram-air parachutes. In the next sections, several notable examples of both experimental and computational studies are discussed, and the challenges related to both types of approach are examined.

2.2 Experimental studies

The use of aircraft drop-tests for testing and certification of ram-air parachute canopies has a long history, and is continued today, for the obvious reason that the performance of such a safety-critical piece of equipment must be tested to achieve a very high level of confidence in the design. However, in terms of research and development, studying the design factors behind the performance of parafoils through drop tests is not easy. Taking experimental measurements such as recording the shape of the canopy or observing the flow around it in flight is either impossible or extremely difficult, and the need to deploy the parachute from a folded configuration means that possibilities for applying instrumentation to the canopy fabric itself are limited [35].

More recent tests have made use of small-scale powered models, as a substitute to traditional style drop-tests. The use of a motorised payload provides more flexibility in the test, providing greater flight duration and repeatability, whilst also removing the need for an aircraft from which to perform the drop, thereby drastically reducing cost and logistical burden. This technique has been utilised to test novel control methods such as canopy spoilers [36] and variable glide slope mechanisms [37, 38], in the hope of improving the accuracy of automated guided systems. Whilst more adaptable than full-scale airdrops from an aircraft, the results attained through this type of experiment are still limited to overall flight characteristics, such as lift to drag ratios and turn rates. Similarly, tests in which a reduced-scale canopy is towed by a moving vehicle have also been trialled [39]. Examination of the changes in structure or behaviour of the flow during flight is still extremely challenging in either of these settings.

As an alternative to drop-testing, wind-tunnel testing provides a more controlled environment within which to observe a parafoil in flight, enabling more detailed measurements to be taken. The large scale of typical parachutes means that studies must often incorporate reduced-scale versions of the canopy, although a number of full-scale wind-tunnel studies have been undertaken [21]. An important example of this is the work of Ware & Hassell [33], who utilised the 30ft $\times$ 60ft wind tunnel located at NASA Langley to study a range of flight characteristics, including the effect of aspect-ratio and angle of attack on the performance of a ram-air canopy. The design tested by Ware & Hassell is taken as the parafoil design for all of the simulations described in the following chapters.

Further full-scale tests are described by Nicolaides [40], Geiger & Wailes [41], and Sanger & Ware [42], all of which provide measurements of the integrated forces about the canopy, and their dependence on angle of attack. The results all show a high level of drag and low level of lift, relative to the values generally observed for an aerofoil, resulting in a range of lift to drag ratios of 2–6. The results show a large influence of the aspect ratio on the lift and drag coefficients, in the same manner as for a wing.

It is clear from these reports, however, that the use of such large canopies in the tests creates a number of problems. The rigging is made difficult by the long length of the lines, and the angle of attack is hard to control and measure accurately as a consequence. The large number of lines required for a large canopy also contribute a significant amount of drag, but they may be too long to be included in the flow stream, yet must still be considered in the final results. Most importantly however, the full scale tests have not shown the potential for more detailed measurements of the parafoil in flight, such as flow visualisation or measurement of the inflated shape. The size of the canopy seemingly makes this too difficult [43].

A more modern example of a full-scale test is described by Askins, Zell and Ross [44], who have developed a method for installing a chordwise strip of pressure tappings within the fabric of a canopy, such that the pressure at the external parafoil surface can be recorded. Due to the lightweight and flexible nature of the fabric construction of a parachute, this approach is limited to one chordwise strip of 64 tappings. Any more than this, and the weight of the instrumentation will be great enough to affect the flight of the canopy. This highlights the difficulties involved in making comprehensive measurements on the lightweight structure of a parachute canopy even within the more manageable setting of the wind-tunnel.

To address the problems presented by large-scale wind-tunnel tests, scaled models of the parafoils can be used instead for a more manageable experiment. However, scaled wind-tunnel studies still raise a number of associated difficulties. One of the unique challenges presented by parachutes is their inflatable, flexible nature. This poses several obstacles to acquiring a detailed assessment of their behaviour in flight. As already mentioned, the flexibility of the fabric makes attaching and supporting instrumentation difficult without influencing the shape of the structure. The flexibility also presents a problem when trying to accurately record the angle of attack for example, which could be influenced by oscillations of the canopy or changes in the geometry over a range of angle of attack. The biggest problems however, relate to the accurate measurement of the inflated geometry, to allow the influence of the aerodynamics on the structure to be recorded and analysed; and the successful visualisation of the flow around the inflated geometry, to enable the relationship between the inflated structure and the associated flow to be understood.

A solution to the problem of accurately measuring the inflated shape of the parafoil is to make use of stereo-correlation imaging, or photogrammetry. This technique uses multiple cameras to image the structure, and a set of known calibration points are marked on the canopy to be tracked in the camera footage. The images can then be processed to provide a three-dimensional reproduction of the canopy shape. Matos et al. [45] successfully applied this approach to the wind-tunnel testing of a model parafoil, and were able to accurately reproduce the three-dimensional shape of the inflated canopy. A similar approach using stereo-correlated laser-scanning is used by Lee & Li [46]. The measured 3D geometries allow quantitative comparisons of the difference between the inflated and uninflated shapes, particularly in terms of the cross-section, which is shown very obviously to expand between the ribs as the ram-air pressure causes spanwise scalloping of the parafoil surfaces. Knowledge of this difference provides a feedback to the design of the uninflated shape, such that the aerodynamics of the inflated structure can be more carefully controlled.

Whilst solutions for the problem of measuring the inflated shape of the canopy have been found, there is a lack of examples available in the literature that provide information on the flow field around the ram-air parafoil. A rare insight is provided in a report by Speelman et al. [47], who performed a wind-tunnel study of reduced-scale parafoils, for Reynolds numbers of around 1.5×10^5–3×10^5. The smaller size of the canopy allowed wool tufts to be attached at various points on the canopy upper

surface, that were photographed from above, providing information on the variation of the flow in the chordwise and spanwise directions. Observation of the wool tufts revealed the presence of upper surface separation even at angles of attack as low as 2°, for a parafoil geometry very similar to that used by Ware & Hassell [33]. Interestingly, the separated region was shown to be present at the centre-line at first, moving forward from the trailing edge, and spanwise away from the centre, as the angle of attack increased to a maximum of 10°. The results demonstrate a tendency for the flow to separate at a mid-chord position, even at relatively low angles of attack. Unfortunately, the wool tufts were not used by Speelman et al. to investigate the nature of the flow around the lower surface.

Aside from the difficulties associated with the nature of parafoil flight behaviour, the use of scaled flexible models in the wind-tunnel also results in an added complication, due to the need to consider the effect of aeroelastic scaling. Whilst Reynolds numbers may be matched in the wind tunnel to ensure the similitude of the fluid dynamic behaviour, the dynamic pressure will not be equal to the full-scale case, and so the response of the structure to this loading may also not match the full-scale behaviour. Therefore the stiffness of the structure may need to be adjusted to account for this [48, 21].

A less complex alternative to using small-scale model canopies in the wind tunnel is to use purpose built rigid models that reflect the key features of the parafoil design being studied, an approach adopted by Furfaro, Hailey and Behr [49]. Their study used models constructed of Styrofoam and balsa wood to compare the effect of the presence of the leading edge cut, the influence of an open leading edge over a closed one, and the effect of spanwise scalloping of the parafoil. Whilst not incorporated into this particular experiment, the use of a rigid model provides more opportunities to make observations on the flow, through the use of pressure tappings in the surfaces of the model for example. A rigid parafoil structure is also a more reliable base around which to perform flow visualisation, via techniques such as oil-flow visualisation, Particle Image Velocimetry (PIV) or Laser Doppler Velocimetry (LDV). The rigid structure ensures that the measurements of the flow field are not influenced or affected by any oscillation or translations in the fabric of the canopy. However, recreating a realistic 3D geometry for an inflated parafoil, instead of a simplified 2D prism geometry is a challenge, and so flow visualisation around a canopy in steady flight is difficult. Consequently, to the author's knowledge, there are no examples

of experimental visualisation of the flow around a ram-air parachute available in the literature. This places a clear restriction on the level of sophistication to which the design of parafoil canopies can reach, as the relationship between the inflated design and its corresponding flow field cannot be clearly understood.

2.3 Computational studies

Acknowledgement of the aforementioned challenges and constraints involved in the experimental study of full-scale parafoil behaviour during flight makes the application of computational methods to the problem an attractive and promising alternative. The use of Computational Fluid Dynamics (CFD) and Computational Structural Dynamics (CSD) techniques offer the opportunity to study and observe the flow field and associated structural stress and deformation fields as a whole.

The flow around two-dimensional parafoil cross-sections has been examined using CFD simulations by Mittal et al. [50], Balaji et al. [51], and Mohammadi & Johari [52]. In these studies the Reynolds-Averaged Navier-Stokes (RANS) equations were solved for the flow around rigid aerofoil and parafoil cross-sections. Mittal et al. and Balaji et al. modelled the flow as unsteady, and Mohammadi & Johari modelled it as steady. These papers provide valuable insight into the difference between the flow behaviour seen around a parafoil cross-section with a leading edge cut, and their baseline aerofoil equivalent. The analyses showed a tendency for the sharp leading edges, present on the upper and lower surfaces of the canopy as a consequence of the ram-air inlet, to cause the flow to separate. As all three studies used different parafoil cross-section geometries, including those based on Clark-Y and LS(1)-0417 profiles, the magnitude and position of the separation was seen to vary between different designs. The size and location of the leading edge cut was also seen to be highly influential on the resulting flow field. In all of these cases, the parafoil cross-sections produced less lift, typically $\approx$ 15%–20%, and created around double the amount of drag than their counterpart aerofoils with closed, smooth leading edges. This is attributed to the blunt shape of the leading edge which results from the cut in the leading edge made to create the ram-air inlet [52, 51].

These 2D CFD studies provide the first examples of the likely flow behaviour around a parafoil in the literature, improving the understanding behind the aerodynamic

performance of parafoils, and demonstrating the value of CFD studies. They also strongly suggest that the effect of the ram-air inlet cannot be ignored in terms of the aerodynamic performance of parafoils, and their improved design could offer benefits in terms of increased lift to drag ratios. However, two-dimensional studies are limited in terms of predicting the performance of the parafoil overall. The reason for this is that if the aerodynamics of a canopy are to be properly assessed, the three-dimensional nature of the structure needs to be taken into account. This is because the inflation of the parachute cells, which provide the rigidity required for the canopy to keep its shape, causes the cells to billow outwards. Consequently, a large proportion of the canopy will have a cross-section that deviates significantly from that of the cut pattern, as seen by the front view of a parafoil in Figure 2.3. Therefore, for a reliable prediction of aerodynamic performance to be made a three-dimensional CFD simulation of the flow around the canopy is required, and the shape of the parafoil used in this analysis must be accurate.

Due to the strong fluid-structure interaction inherent in the flight of a flexible parafoil canopy, the inflated shape is highly dependent on the balance between the aerodynamic forces on the structure, and the corresponding behaviour of the flow around the structure that results. This means that the inflated shape is not known a priori. Due to the previously stated necessity for obtaining an accurate shape for use in the 3D CFD analysis this must either be measured experimentally, using the techniques described above, or calculated numerically using FSI simulation techniques. A shortcut to this is to take the approach used by Eslambolchi & Johari [53], who approximated a 3D geometry of an inflated ram-air canopy by scaling dimensions from photographs of the parachute in flight. The 3D geometry was used for a CFD study, using the steady RANS description of the fluid. The results revealed a three-dimensional flow field, as expected, with pressure and velocity fields varying across the scalloped upper surfaces. Flow separation from the lower leading edge, previously observed in the 2D CFD simulations, is also observed, but is shown to be influenced by the curvature of the lower surface in the 3D case.

The results of Eslambolchi & Johari clearly show the need for the flow around parafoils to be considered in three dimensions, and provides a valuable insight into the flow field around an inflated 3D parafoil geometry which had not previously been achieved either experimentally or computationally. Significantly, this both demonstrates the utility of applying CFD techniques to the study of parafoils and highlights the importance

of acquiring and using an accurate canopy geometry. The process of approximating a geometry from scaling photographs without the use of photogrammetry is difficult and time-consuming, and for a true predictive capability to be achieved, that is useful for design studies and performance assessments, confidence in the parafoil shape being studied needs to be high. In light of the difficulties of acquiring the shape experimentally, this reiterates the value in developing methods by which the inflated geometry can be predicted directly through the use of FSI simulation techniques.

One method of analysing the deformation inherent in the inflated canopy, and thereby predicting the inflated shape, is to use a strongly coupled, partitioned approach to model the equilibrium between the structural and fluid dynamics. This method is employed by Kalro & Tezduyar [20, 54], where an approximation of a ram-air parachute is simulated using bespoke, Finite Element Method (FEM) based CSD and CFD solvers. In this case, the canopy has a closed leading edge, with a prescribed internal pressure, and therefore any flow features associated with the open leading edge will be missed, along with any associated effect these may have on the deformation of the parachute. The work is extended to include prescribed shape changes at the trailing edge by Tezduyar et al. [55], but the representation of the remainder of the parafoil geometry was unchanged.

A similar approach is adopted by Ibos et al. [56], who combine the CFD code N3S with the CSD code SAMCEF [57]. The coupling software is designed to allow transient analysis of the parafoil canopy in flight, by solving the unsteady RANS equations for the flow, and iterating at each time-step to find the equilibrium at the fluid-structure interface, in a strongly coupled fashion. This methodology is used by Ibos et al. [56] to simulate the equilibrium inflated shape of a ram-air canopy. Unlike the work of Kalro & Tezduyar discussed above, the analysis of the ram-air canopy by Ibos et al. includes the presence of the ram-air inlet, but a slight simplification of the trailing edge of the canopy is required due to the required resolution. The reported lift and drag values produced by the simulation for the inflated canopy compare well to experiment, but no indication as to the accuracy of the predicted structure is provided. Furthermore, there is also no additional insight into the flow field around the canopy as the fluid dynamic results are not analysed in detail.

Given the disadvantages of the complexity involved in developing a strongly coupled framework such as the ones described herein [58], it is not obvious from the results

of these studies whether this approach is advantageous. In particular the potential for transient, time-dependent analyses permitted by the use of a strong coupling, is not used in any of the studies, making the reasons for selecting a strongly coupled approach unclear.

A more recent example of an FSI analysis of a ram-air canopy is presented by Oye [59], which makes use of a loose coupling between a RANS based fluid solver and an FEM based structural solver. The inflated shape of a full canopy is predicted, for which the aerodynamic performance is assessed through CFD. The lift values predicted by the model for the inflated parafoil compare well to the experiments, but the drag force is over predicted. As in previous studies, detailed results on the nature of the flow field or the behaviour of the structure are lacking. Another similarity between this study and previous ones is the need to approximate the canopy in some way to enable a successful completion of the simulations. In this case, extra chordwise stiffening elements are required in the ribs in the structural model of the canopy, to provide extra chordwise support. Overall, the outcomes of this study suggest that for the prediction of the inflated shape of a parafoil during steady glide, a loosely coupled approach is likely to be equally as successful as a strongly coupled approach, but may be more straightforward in terms of implementation, and less intensive in terms of computational resources.

A more simplified approach can be found in the work of Chatzikonstantinou & Puskas [60], Chatzikonstantinou [61], and more recently Altmann [62], who utilise a vortex panel code coupled to a finite element solver. This provides an approximation to the inflated shape, as well as predictions of the aerodynamic characteristics. This approach is limited by the vortex panel code, which is applied to a canopy with a closed leading edge. Therefore the effect of separation at the leading edge is ignored and the internal pressure must be prescribed. Despite this, the importance of considering the inflated shape is made clear by the results, which show a significant difference between the cross-section at the mid-cell point and at the ribs. Other observations, such as the unwanted deformation seen at the trailing edge also point to the effect that these differences will have on the fluid flow.

2.3.1 Simulation of fabric materials

An aspect of structural modelling that is particularly important to consider when simulating parachutes is the replication of the behaviour of fabric materials. An accurate representation of the structural properties of fabric is a challenge, principally due to their inherent lack of flexural stiffness [63]. A lack of flexural stiffness means that fabric membrane structures can only react to loads through in-plane resistance [64]. As a result, such structures naturally find the optimal shape for the given constraints. For inflatable or pneumatic envelope type structures such as a parachute this optimal shape typically requires a degree of curvature in order to resolve the distributed perpendicular pressure loading into in-plane tensile stresses [64]. Such a structure will often undergo large displacements and rotations in the process of reaching the final configuration [65]. To recreate this behaviour computationally, the numerical model for the fabric must be capable of coping with the resulting geometrical non-linearities, and the fabric's inability to support compressive stresses [63]. Compressive stresses in fabrics result in wrinkling, and so a strategy to successfully consider the effects of wrinkling is required. Additionally, more sophisticated problems may require a more detailed fabric constitutive model to take into account the many effects related to the directional nature of a woven fabric, including the interaction between the woven filaments and the evolution of the fabric anisotropy with time [65, 66].

Whilst anisotropic material behaviour and the intricacies related to the interaction of fabric filaments under loading are both beyond the scope of this current work, the inherent propensity for wrinkling in loaded fabric structures must be considered. As already stated, wrinkles such as those observed in the membrane material under in-plane shear shown in Figure 2.4 occur due to the extremely low or non-existent bending stiffness of a fabric. This means that a fabric membrane will readily buckle under even a very small amount of compression. This buckling behaviour manifests itself as a local instability in the form of small out-of-plane displacements occurring in a perpendicular direction to that of the applied compressive stress. The presence of these displacements, commonly known as wrinkles, serve to remove any compressive stresses that may appear in the structure [67]. As such, the presence of wrinkles in membrane structures will have a definite effect on the resultant deformation of the structure as a whole, and therefore their influence must be incorporated into any simulation of an inflatable membrane structure [67].

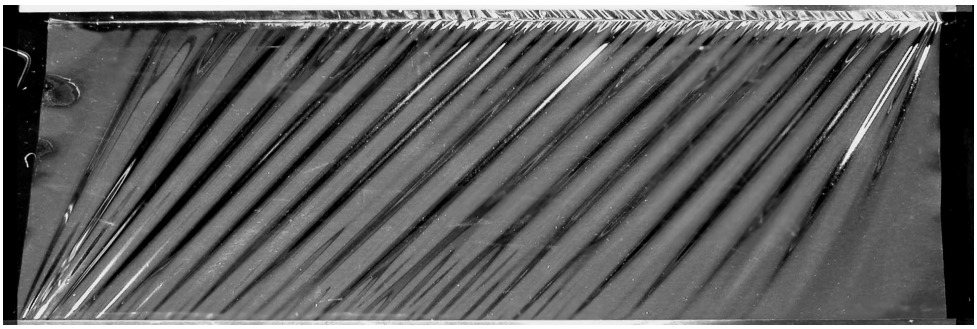


Figure 2.4: Experimental result showing wrinkle pattern in a 0.0025mm thick Kapton membrane under a mid-plane shear displacement of 3mm [68].

The average size of structural wrinkles is a function of the bending stiffness, and so typically the size of the wrinkles are small relative to the scale of the structure [69]. Therefore, if the formation of individual wrinkles is to be replicated accurately in a simulation either an extensively refined mesh of lower order elements or a reference mesh of higher order elements is required. A number of examples where a refined mesh of elements is used to recreate the wrinkle patterns in membrane structures can be found in the literature. Tessler, Sleight and Wang [70], Flores and Onate [71], and Wong and Pellegrino [72] all focus on the replication of the wrinkle pattern in a square or rectangular membrane film under in-plane shear and comparison against experimental results such as those in Figure 2.4 [68]. The studies show that shell elements must be used to provide the correct coupling between the membrane and bending deformations if the correct wrinkling patterns are to be predicted, despite the bending stiffness being very low. The results also demonstrate the need for a careful choice of element formulation if convergence issues and accuracy problems are to be avoided, and an element size that is smaller than the wavelength of a wrinkle [72].

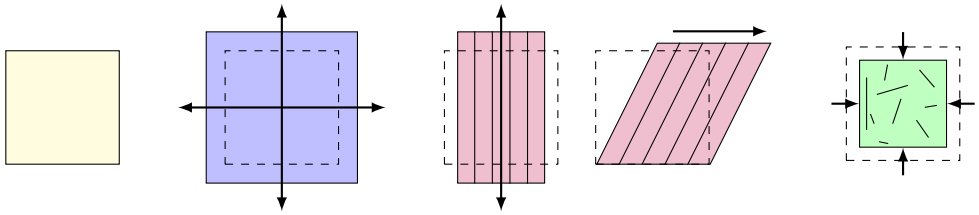
Tessler et al. and Wong & Pellegrino found that the introduction of imperfections into the model was a further requirement of an accurate prediction of the amplitude and frequency of fabric wrinkles. The applied imperfections are in the form of small displacements in a pattern that was either pseudo-random or based on the superposition of the modal eigenvectors of the membrane. Imperfections are required to initiate the formation of wrinkling and ensure the necessary coupling between membrane and bending deformations [70]. This is similar to the conditions found in a real membrane where localised buckling occurs in a random fashion, determined by the distribution of initial imperfections in the material [69].

The results from the studies highlighted above compare well to experimental results, but they highlight the difficulties associated with this approach, such as the need for very high mesh resolutions, the need for careful selection of element formulations and the requirement for carefully tuned imperfections to be incorporated into the mesh. The high level of detail relating to the amplitude and position of individual wrinkles provided by this method is often not required for many engineering applications. Therefore the computational expense and the intricacies of the approach are not justified, and alternative, less challenging approaches are required [70].

The alternative to simulation of individual wrinkles in the fabric is to "enrich" the elements used in the structural simulation through the use of a wrinkling model. The concept here is not to predict the occurrence of the wrinkles themselves, but to allow the average stress and displacement field to be accurately described without the need for the wrinkles to be resolved [67]. The advantage to this is that element size is no longer limited by the size of the wrinkles, thereby greatly reducing the computational cost without any loss in accuracy. This approach is no less precise than the method described previously, as the global stress field is correctly described even if no information about the actual wrinkles is provided. It is also noted in the literature that any attempt to directly model the size and position of wrinkles would always be based on arbitrary positions due to the lack of data on the imperfections that cause them, and hence no predictive capability is deemed to be lost [69].

The element enrichment can be applied in two main ways, both of which aim to adapt the solution procedure for membrane elements such that fictitious compressive stresses that would be present in a standard membrane element are eliminated, thereby replicating the effect wrinkling has on the behaviour of the membrane structure.

The first possibility is to manipulate the formulation of the membrane element itself, by superimposing a term on the displacement field related to the formation of wrinkles, creating an effective deformation gradient that describes the shortening of the average plane of the membrane in the presence of compressive stresses [69]. This approach is adopted by Kang & Im [73], Lu et al. [63] and Raible et al. [74], and the results compare well to a range of different experimental results, proving the effectiveness of the element enrichment technique. Whilst successful, the main drawback to the deformation gradient based wrinkling model is that it is inherently complicated to implement.



(a) undeformed membrane. (b) taut membrane. (c) wrinkled membrane. (d) slack membrane.

Figure 2.5: Principal status of a membrane - taut, wrinkled and slack. Adapted from Raible et al. and Jarasjarungkiat at al. [74, 67].

A second method for "enriching" elements is to modify the stress-strain relationship of the material, by altering the constitutive law as to not permit compressive stresses. This has the advantage of being applicable to any element type, thereby avoiding the complications for the previous method and making implementation more straightforward, including into commercial Finite Element codes as a "user-material" [64]. It also has the benefit of being easily adaptable to include orthotropic material behaviour [67].

The wrinkling algorithm based on modification of the constitutive law consists of two key steps that occur during each time-step. Firstly, the membrane must be assessed to identify what state it is in; whether regions of the structure should be considered to be taut, wrinkled or slack. Examples of these states are provided in Figure 2.5. Three different sets of criteria can be used to make this assessment, based upon a consideration of the principal stresses, principal strains or a combination of the two. It has been found that the most effective method is to use a combination of the principal stresses and strains proposed by Roddeman [75], known as the mixed criterion. The criteria for all three cases are given in Table 2.1.

Once the assessment stage is complete, any areas identified as being wrinkled or slack are corrected in the second stage by applying the constitutive relation to remove compressive stresses. In an isotropic membrane under uniaxial stress the wrinkles will form in a direction perpendicular to that of the first principal stress [67]. Any compressive stresses present in that direction are therefore considered fictitious and can be removed. An example of the direction and magnitude of the principal stresses in a wrinkled membrane is shown in Figure 2.6, where the result of a numerical simulation by Wong & Pellegrino, that recreates the test in Figure 2.4, is presented.

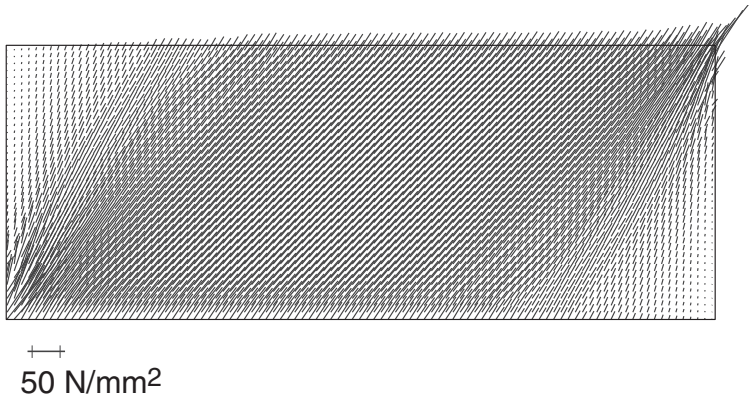


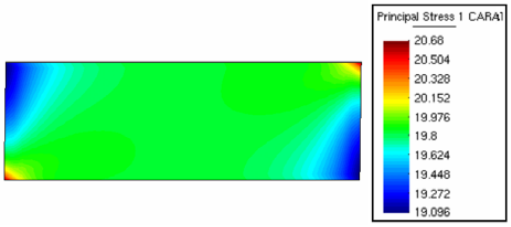
Figure 2.6: Principal stress directions and magnitudes in a thin-shell simulation of a 0.0025mm thick Kapton membrane under a mid-plane shear displacement of 3mm [68].

The removal of compressive stresses is based upon the analytical Tension Field theory, which treats the fabric as a *true membrane* of zero bending stiffness. This is in contrast to the small but finite bending strength required to simulate wrinkles using shells as discussed previously [64].

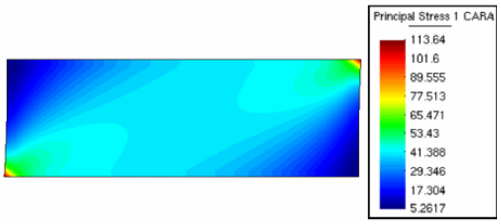
Several examples of this type of wrinkling model can be found in the literature, such as the work of Rossi et al. [69, 76], Jarasjarungkiat, Wüchner and Bletzinger (Figure 2.7) [67], and Liu, Jenkins and Schur [64]. The studies show the successful application of the material modification wrinkling model to a number of different examples, including a rectangular membrane under in-plane shear. The results show excellent correlation to the equivalent shell model with considerably lower mesh density, and a clear improvement on the overly stiff behaviour displayed by membrane

Table 2.1: Summary of possible criteria for assessing the wrinkle state of a membrane under loading [67, 74].

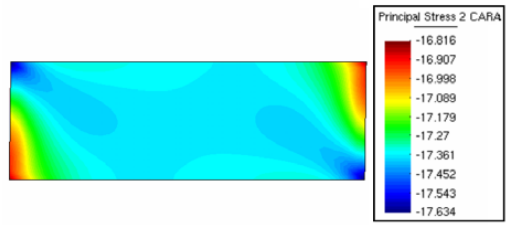
Principal stress criterion	Principal strain criterion	Mixed criterion	Membrane state
$S_{min} > 0$	$\varepsilon_{min} > 0$	$S_{min} > 0$	Taut
$S_{min} \leq 0$ and $S_{max} > 0$	$\varepsilon_{min} \leq 0$ and $\varepsilon_{max} > 0$	$S_{min} \leq 0$ and $\varepsilon_{max} > 0$	Wrinkled
$S_{min} \leq 0$ and $S_{max} \leq 0$	$\varepsilon_{min} \leq 0$ and $\varepsilon_{max} \leq 0$	$\varepsilon_{max} \leq 0$	Slack



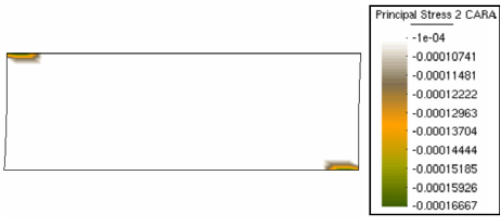
(a) Maximal principal stress without wrinkling model.



(b) Maximal principal stress with wrinkling model.



(c) Minimal principal stress without wrinkling model.



(d) Minimal principal stress with wrinkling model.

Figure 2.7: Principal stresses in a membrane element simulation of a 0.0025mm thick Kapton membrane under a mid-plane shear displacement of 3mm [67].

elements without a wrinkling model. This can be seen clearly in the comparison of Figures 2.7(a) and (b) where the central band of stress observed in Figure 2.6 is more closely replicated by the membrane elements with the wrinkling model applied in Figure 2.7(b). Compressive stresses are shown to be efficiently removed from the material which ensures the correct global stress and deformation fields are generated within the structure [67]. The successful removal of compressive stresses by such a wrinkling model is demonstrated in a comparison of the results of Jarasjarungkiat et al. in Figures 2.7(c) and (d), where the compressive stresses present in the standard membrane model of a membrane under in-plane shear in (c) are non-existent when the wrinkling model is applied in (d).

The material modification approach is believed to be the basis of the wrinkling algorithm incorporated within the fabric constitutive model in the structural dynamics solver LS-DYNA [77]. As discussed in Sections 3.3.2 and 3.3.3 the wrinkling algorithm within LS-DYNA is employed for all structural simulations within this current work. The effectiveness of this method is examined by Pruett, Accorsi and Charles, who carried out three experiments on structures constructed of parachute fabrics, to act as validation studies [65]. The tests consisted of an inflated cylinder under

torsion, an inflated airbag, and an inverted round canopy filled with water. A digital photogrammetry system was used to measure the position and displacement of the fabric structures, which therefore enabled the calculation of in-plane strains [78]. Quantitative comparisons of the results with simulations performed using LS-DYNA showed good agreement, and the wrinkling model was shown to accurately replicate the changes in the displacement and stress fields caused by wrinkling of the material. This suggests that the fabric constitutive model within LS-DYNA is suitable for capturing the effects of wrinkling likely to occur in parachute structures.

2.4 Summary

The above review of the challenges surrounding the analysis of ram-air parachute behaviour and performance during flight has established a number of major advantages in the application of computational methods to the problem. In particular, the use of Fluid-Structure Interaction simulation techniques offer the potential for the design and analysis of ram-air canopies to have significantly less reliance on experimental testing. However, despite the individual efforts found in the literature, a comprehensive application of FSI simulation techniques to ram-air parachutes has not been successfully realised.

Whilst there are examples of the prediction of an inflated shape of a parafoil canopy resulting from the use of both loosely and strongly coupled FSI simulations, none of them provide any detailed results regarding their predictions for the behaviour of the fluid or the structure. This makes it difficult to assess the merits of the particular study, and to judge the suitability of the approach. More importantly, it also means that even though the methodology may have successfully predicted the inflated shape, the step of using this shape to provide further insight into the behaviour of ram-air canopies has not been made. In particular, the predicted shape can be used to further the understanding of the aerodynamic performance of the canopy in terms of the key flow features, as well as the mechanisms behind the interaction between the fluid flow and the canopy structure.

Therefore, this study aims to progress beyond the state of the art by developing a loosely coupled framework for the prediction of the inflated equilibrium shape of a

parachute during steady glide. This shape will then be used for analysis of the fluid flow around the structure, to provide a prediction of the flow behaviour around a typical ram-air parachute, and enable an analysis of the key flow features responsible for the aerodynamic performance of a typical canopy. The FSI simulations will also offer insight into the interactions between the flow field and the shape of the canopy, and how this leads to the steady equilibrium shape. To achieve this goal, a novel approach to the simulation of the parafoil will be adopted, using a single-cell of a ram-air canopy in periodic symmetry, to represent the complete canopy as a chain of cells. Full details of the loosely coupled FSI framework, as well as the single-cell boundary conditions, are provided in the following chapter.

Chapter 3

Single-cell ram-air parachute model and FSI simulation framework

This chapter introduces the approach used to represent the behaviour of the ram-air parachute in this study. A single half-cell of a parafoil is modelled such that it behaves as though it was part of a full canopy of infinite span, formed of a chain of half-cells. Following this the underlying physics behind the simulation techniques used for the fluid and structural simulations are outlined, and the coupling methodology is defined. In addition a description of the boundary conditions for simulation of both the fluid and structural domains is provided, followed by a discussion of the necessary modelling approaches and adaptations that are required to ensure an accurate representation of the canopy.

3.1 Single cell parafoil model

Due to the high computational cost involved in performing large CFD simulations, such as that of a full parachute canopy, only a single half-cell of the parafoil is modelled. In order to make this approach relevant to to a full canopy, the structural and fluid boundary conditions have been set-up to represent periodic symmetry, thereby simulating a cell that behaves as one repeating part of a flat (zero-degrees anhedral)

canopy of infinite span. The undeformed single-cell, a representation of the full symmetry case and a view of the periodic unit-cell of the BC2 case is seen in Figure 3.1. This is a departure from the design of many modern canopies, such as those tested by Ware & Hassell [33] upon which the design chosen for this thesis was based, which are all rigged to have arc-anhedral. However, it still provides a good approximation to the flight behaviour of a generic ram-air canopy, and is close to the behaviour observed for the centre cells of an anhedral rigged parafoil.

The fluid boundary conditions are set up with the side boundaries of the fluid domain in the same plane as the ribs, to ensure periodic boundary conditions are enforced for both the fluid and structural domain. It is important to note that a consequence of the infinite span boundary conditions is that any edge (wing tip) effects seen on a full parachute canopy are neglected, and therefore the effects of induced drag seen in real-life parachutes is ignored. No suspension lines are included in the model; representative constraints that enforce a constant angle of attack are used instead which are discussed in Section 3.2.2.

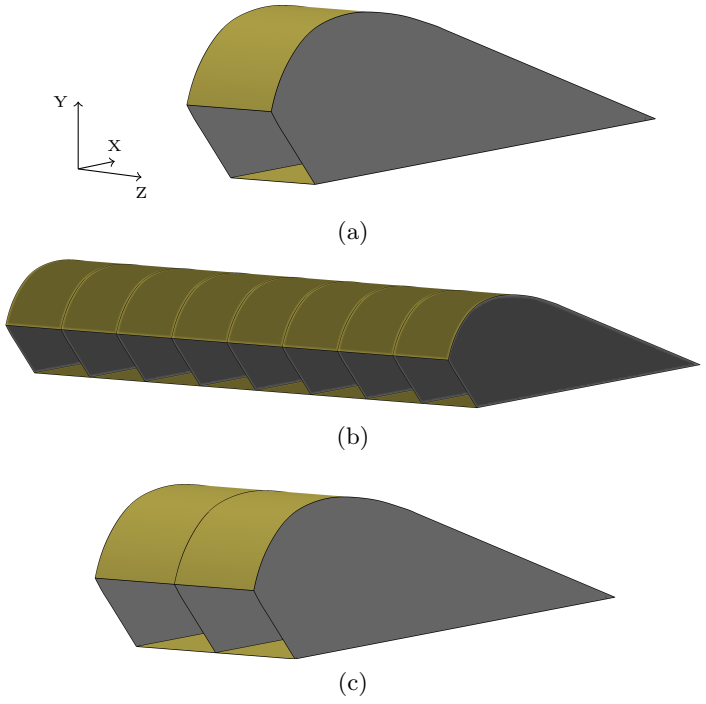


Figure 3.1: View of (a) un-deformed half-cell parafoil geometry, (b) half-cell with periodic symmetry applied, illustrating the infinite span of full-symmetry canopy, and (c) single-cell of BC2 case reflected about xy plane to show unit-cell consisting of two loaded ribs and one central unloaded rib.

3.2 FSI simulation framework

Due to the aforementioned inter-dependence between aerodynamic and structural forces acting on a typical parachute canopy, simulating the inflated shape of the single-cell model above must involve a coupled fluid-structure modelling approach. In the remainder of this section, the selected level of coupling is described, and full details of the FSI simulation framework used during this study are provided.

3.2.1 Fluid-structure coupling approaches

As already stated, complex multi-physics problems such as the behaviour of a lightweight structure during flight requires the consideration of two separate domains, and the interaction between them at their interface. In this case, the two domains are the structural domain, Ω_S, and the fluid domain, Ω_F, both of which are governed by a separate, independent system of equations, which interact at the fluid-structure interface, Γ. In this case, the interface is at the upper and lower surfaces of the parachute fabric.

In order to combine the two systems, they are coupled, such that the outputs of each model become the inputs for their respective counterparts. This coupling process therefore relies on the exchange of information between the two domains. Depending on the problem, this exchange will happen in various different ways, requiring different levels of complexity [79, 80].

The most complex method of coupling is the monolithic coupling approach, where the two (or more) systems are combined into one large system of equations, which is solved simultaneously. This is known as a strongly coupled approach, with solutions for both domains solved at each time-step. This makes the monolithic approach very useful for problems where the interaction between the fluid and structural forces is very strong [81], and where transient effects are important. With this style of coupling, there is less need for transfer of information between the two systems. However, monolithic approaches are very computationally intensive, due to the number of equations which have to be solved at the same time. They are also difficult from an implementation perspective because the solvers for each domain must be written with such a coupling method in mind from the beginning [79].

Alternatively, two pre-existing solvers can be used, using a partitioned approach. In this case, the individual domains are treated as subsystems, which are solved separately at each time-step. The equations of each subsystem are formulated such that the solution of each field is separate, but coupled through shared terms. The values of the shared terms must be exchanged between the two systems to achieve the coupling. For a strong coupling, a small time step must be used, or a number of iterations between the two systems are required at each time step, to ensure convergence of the FSI solution [82]. If a strong coupling is achieved, the outcome is the same as it would be from a monolithic solver, with the advantage of solving each system by a method best suited to the physics of each one [81, 83].

This study is focused on predicting the equilibrium flying shape of a parafoil, and therefore transient, time-dependent effects on the structure can be neglected. This means that a strong coupling is not required, and both systems can be coupled whilst maintaining separate time-steps. This is known as a loosely coupled approach, and is the method chosen for this study. A loosely coupled, partitioned approach is adopted, using two separate solvers for the fluid and structure domains, coupled at the fluid structure interface [84].

With the general outline of a loosely coupled FSI coupling method presented, the details of the loosely coupled framework developed within this study are now described, starting with the composite fluid and structural solvers. The specific methodology by which the loose coupling is achieved is then discussed in further detail.

3.2.2 Structural modelling

Governing equations and solver overview

The equation governing the behaviour of the structural domain Ω_S is the Lagrangian description of conservation of linear momentum,

$$\rho_S \frac{\partial^2 \boldsymbol{u}}{\partial t^2} = \nabla \cdot \boldsymbol{\tau} + \rho_S \boldsymbol{f}^B \tag{3.1}$$

where ρ_S is the material density, $\boldsymbol{u}$ is a vector of structural displacements, t is the time, $\boldsymbol{\tau}$ is the Cauchy stress tensor, and $\boldsymbol{f}^B$ is a vector of body forces [85]. The equation is completed by the use of a constitutive equation for the material. This can

be linear or non-linear. The boundary conditions needed to solve 3.1 are,

$$\boldsymbol{u} = \boldsymbol{u}_S \quad \text{on } S_u \tag{3.2}$$

$$\boldsymbol{\tau} \cdot \boldsymbol{n} = \boldsymbol{f}^S \quad \text{on } S_f \tag{3.3}$$

where S_u and S_f represent the parts of the boundary with prescribed displacements, $\boldsymbol{u}_S$, and tractions, $\boldsymbol{f}^S$, respectively; and $\boldsymbol{n}$ is a unit outward normal vector to the boundary. These are defined more specifically below.

The solution of (3.1) is achieved through discretisation of the domain using the Finite Element Method (FEM). This is accomplished in this project using the commercial code LS-DYNA [86], which has been selected as the structural solver due to the strong background it has with parachute simulations, examples of which can be found in papers by Coquet et al. [58] and Tutt et al. [87]. The parachute fabric is modelled using three and four-noded membrane elements, with one through-thickness integration point, to which Rayleigh damping is applied for stability, based on the findings of Tanner, Cruz and Braun [88]. Additional global hourglass damping is also applied to aid stability. The material of the canopy is treated as linear elastic and isotropic, with full material properties provided in Section 4.4.

Two different time integration approaches have been adopted for the structural FEM studies within this study. Initially, an explicit time-stepping scheme was used, Section 5.2, due to concerns raised in [88] about the use of the fabric material model in the implicit solver within LS-DYNA. The explicit solver uses a central difference scheme. For the results in Section 5.3, the applicability of the fabric constitutive model for the implicit case was confirmed, and the non-linear implicit solver was therefore adopted instead. In the implicit case, the time integration is performed with the Newmark scheme and the non-linearity due to large geometrical displacements is solved with a BFGS quasi-Newton scheme (more details are found in [89]).

Use of an explicit solver for pseudo-static structural simulations such as those performed as part of this analysis, requires care in order to maintain pseudo-static conditions during the simulation. In practice, this means applying loads slowly by steadily increasing the load to the full value over time and ensuring that the level of kinetic energy is kept very low within the structure. The use of the explicit solver brings a penalty in terms of run times, but it is a reliable method that successfully reaches a solution to the non-linear structural problem. Subsequent use of the implicit solver

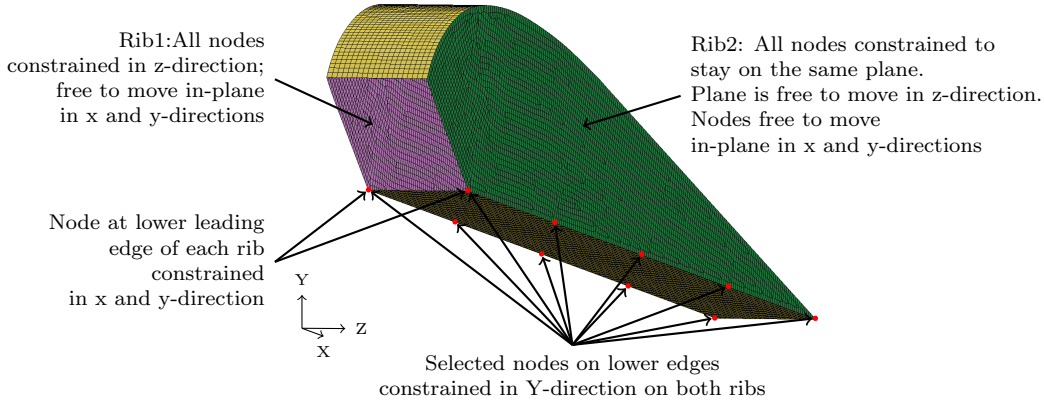
vastly reduced run-times, and therefore greatly reduced the overall run-time for each FSI iteration cycle. The gradual application of loads was also adopted for the implicit structural analyses, to aid the convergence of the initial load-steps.

The capability present in LS-DYNA for modelling fabric materials is of benefit to this study, and has been validated by Pruett, Accorsi and Charles [65]. The parachute fabric is modelled using the fabric material constitutive model (`*MAT_FABRIC`) incorporated within LS-DYNA [77]. This constitutive model includes an option to prevent the generation of any compressive stresses in the material, which prevents the material buckling under compression, and instead the elements reduce in size, offering no resistance to any compression loadings. As discussed in detail in Section 2.3.1 this is a means by which the constitutive model can represent the behaviour of fabric wrinkling, without explicitly modelling the wrinkles. This model is activated in these simulations as a means to represent the effect of wrinkling. The practical benefits to incorporating this model into the structural simulations of parafoil canopies is discussed in detail in Sections 3.3.2 and 3.3.3.

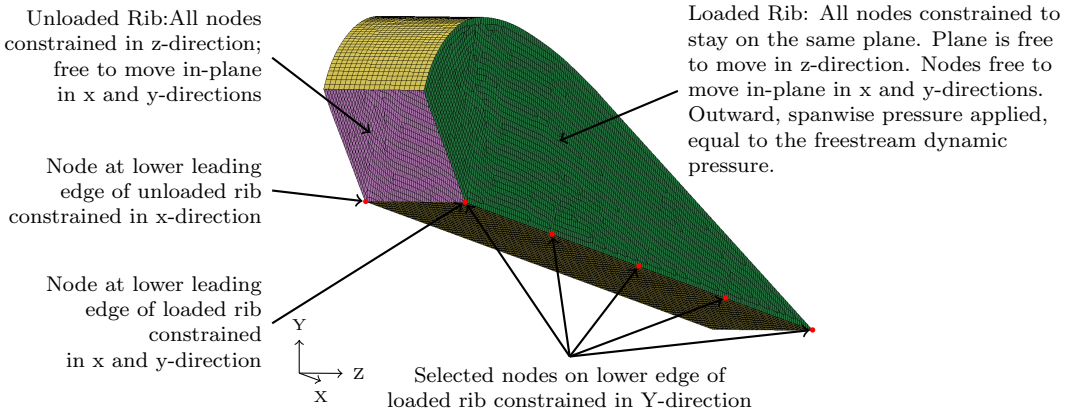
Structural boundary conditions

As already stated, the objective of this FSI study is to simulate a single cell of a parafoil canopy, such that it behaves as though it was part of a larger canopy. To achieve this, suitable boundary conditions have had to be determined. Within this study two general approaches have been followed, the boundary conditions for both are labelled in Figure 3.2.

In the first instance, the parafoil cell is set up in periodic symmetry as though it is a single-cell in an infinitely wide canopy, Figure 3.2(a). Due to the periodicity vertical constraints are applied to both ribs, making both of them loaded ribs. These vertical constraints are at points along the lower edge of both ribs, and they recreate the effect of the suspension lines. The pressure forces seen on the inner surface of the ribs in the fluid dynamics model are not applied to the structural model. This is because it is assumed that an equal and opposite force, originating from the neighbouring cell, would be applied on the other side of the rib. The forces on each side of the rib therefore cancel each other out. No other lateral forces are applied either, as this would not be consistent with the approximation to an infinite span canopy. A full symmetry representation of this case, known as BC1 or the infinite-span case, is seen



(a) Boundary Conditions for infinite span case - BC1



(b) Boundary Conditions for finite span case - BC2

Figure 3.2: Boundary conditions for structural models in 3D FSI cases with FEM mesh shown. z-axis runs spanwise, and x-axis runs chordwise.

in Figure 3.1(b). The results produced from this approach are discussed in Section 5.2. The second approach is an evolution of the first, with some aspects of the finite nature of the true parafoil canopy structure taken into account to produce a more realistic representation. The boundary conditions for this case, known as the finite-span case or BC2, are found in Figure 3.2(b). An outward spanwise force is prescribed in the positive z-direction on the loaded rib of the cell. This replicates the lateral force present on the external rib of the end cell due to the pressure difference between the stagnated air inside the cell and the freestream air outside it [24]. The magnitude is equal to the freestream dynamic pressure, q_∞. The presence of unloaded ribs in the true physical canopy is also incorporated, by removing the vertical constraint from one

of the ribs allowing the whole rib to move freely in the y-direction. To provide extra stability this rib is still constrained in the chordwise direction at the lower leading edge, which is an approximation that would not be the case in real life. The result is a full-cell modelled in half-symmetry which, similarly to the previous conditions, forms the unit cell for a periodic model of a ram-air canopy. A view of the un-deformed unit cell for this case is seen in Figure 3.1(c).

In both cases, the nodes on both ribs are constrained so that they do not move out of plane relative to each other, thereby maintaining a flat surface. This is not strictly the case in real life, as the ribs of a parafoil are not rigid, but it approximates the ideal case where forces on the ribs are equal in each cell. This condition keeps the ribs flat and parallel, which is necessary for periodic symmetry to be achieved and ensures the ribs are always in the same plane as the side boundary walls of the fluid domain. This makes the inclusion of the structural geometry into the fluid simulations consistent and straightforward.

Instead of suspension lines, nodal constraints in the vertical direction are applied at various points along the lower edges of the parafoil ribs, clearly labelled in Figure 3.2. In all cases, each constraint is spread over a group of four adjacent nodes in order to prevent localised mesh distortions. This replicates the design of a real-life ram-air canopy, which would either use reinforcements known as "line attachment tapes" at the suspension line attachment points, or an extra piece of material called a "flare", to distribute the load from the lines into the fabric of the loaded rib [90]. The nodal constraints maintain a constant angle of attack. Rigid-body motion of the structure is prevented using nodal constraints on the lower leading edge.

3.2.3 Fluid modelling

Governing equations and turbulence closure treatment

The behaviour of a fluid under motion is described by the Navier-Stokes equations, given here for an unsteady, incompressible flow of a Newtonian fluid in three dimensions, with body forces such as gravity ignored [91]. They comprise of the continuity equation describing mass conservation, (3.4), and the three momentum equations in (3.5), that describe the momentum conservation in each spatial dimension, x, y, z.

$$\frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} + \frac{\partial v_z}{\partial z} = 0 \tag{3.4}$$

$$\rho_F \left(\frac{\partial v_x}{\partial t} + v_x \frac{\partial v_x}{\partial x} + v_y \frac{\partial v_x}{\partial y} + v_z \frac{\partial v_x}{\partial z} \right) = -\frac{\partial p}{\partial x} + \mu \left(\frac{\partial^2 v_x}{\partial x^2} + \frac{\partial^2 v_x}{\partial y^2} + \frac{\partial^2 v_x}{\partial z^2} \right) \tag{3.5a}$$

$$\rho_F \left(\frac{\partial v_y}{\partial t} + v_x \frac{\partial v_y}{\partial x} + v_y \frac{\partial v_y}{\partial y} + v_z \frac{\partial v_y}{\partial z} \right) = -\frac{\partial p}{\partial y} + \mu \left(\frac{\partial^2 v_y}{\partial x^2} + \frac{\partial^2 v_y}{\partial y^2} + \frac{\partial^2 v_y}{\partial z^2} \right) \tag{3.5b}$$

$$\rho_F \left(\frac{\partial v_z}{\partial t} + v_x \frac{\partial v_z}{\partial x} + v_y \frac{\partial v_z}{\partial y} + v_z \frac{\partial v_z}{\partial z} \right) = -\frac{\partial p}{\partial z} + \mu \left(\frac{\partial^2 v_z}{\partial x^2} + \frac{\partial^2 v_z}{\partial y^2} + \frac{\partial^2 v_z}{\partial z^2} \right) \tag{3.5c}$$

where v_x, v_y, v_z are the three components of velocity, t is time, ρ_F is the fluid density, ν is the kinematic viscosity, related to the dynamic viscosity, μ, through $\nu = \mu/\rho$, and p is the pressure [92].

The appropriate Neumann and Dirichlet boundary conditions are applied as required,

$$v_x = v_\infty,\ v_y = v_z = 0 \quad \text{at an inlet} \tag{3.6}$$

$$\boldsymbol{v} \cdot \boldsymbol{n} = 0,\ \frac{\partial v_n}{\partial n} = 0 \quad \text{at a symmetry boundary} \tag{3.7}$$

$$v_x = v_y = v_z = 0 \quad \text{at a non-slip wall} \tag{3.8}$$

$$p = p_{outlet} \quad \text{at an outlet} \tag{3.9}$$

where $\boldsymbol{n}$ is a unit vector normal to the boundary.

Simulating a fluid flow requires the solution of the governing equations, which for certain simple cases is possible analytically, but for more complex flows must be performed numerically. The nature of the fluid flow is also influenced by the Reynolds number (Re), which is the ratio between the viscous and inertial forces, defined for an object with a characteristic length L_c as $Re = (\rho_F v_\infty L_c)/\mu$ [92]. For low Reynolds number flows the viscous forces dominate the inertial forces, and the flow remains laminar due to the viscous dissipation of naturally occurring disturbances. For higher Reynolds number flows the inertial forces are great enough to amplify any disturbances, and the flow becomes turbulent [93].

For a turbulent flow the fluid motion consists of a range of small disturbances, thereby making the flow inherently unstable and resulting in flow properties that vary in a random and chaotic way. The disturbances form a collection of mutually interacting

and continuously evolving eddies, the scale of which covers several orders of magnitude [94]. The flow regime around ram-air parachutes is turbulent and so the effects of turbulence on the behaviour of the flow must be included in the fluid dynamic simulations.

The only method by which a high-fidelity solution of the Navier-Stokes equations for non-trivial conditions can be achieved is to use Direct Numerical Simulation (DNS) for which, due to the need to resolve the full range of the aforementioned turbulent eddy scales, the mesh resolution requirements are very high [95]. The range of eddy size increases in line with $Re^{3/4}$, hence the required mesh-resolution rises by $Re^{9/4}$, resulting in an increase of CPU time in line with Re^3 [95].

An alternative to this is the more economical approach of the Large Eddy Simulation (LES) method, which instead of aiming to resolve the full range of turbulent eddy scales, resolves only the large to medium scale eddies allowing the use of a coarser grid. The effect of the smallest eddy scale range, which are responsible for the dissipation of turbulent energy through viscous destruction of small-scale motion, is modelled. Whilst lower than for DNS, the computational cost of LES is still high, especially if an accurate representation of the boundary layer is required [95].

The Reynolds-Averaged Navier-Stokes (RANS) method provides a further alternative, where the entire range of turbulent eddy scales is incorporated into a statistical model which represents the effects of turbulence. This method is the least computationally intensive approach, and for this reason has been chosen for this study as an acceptable solution to the trade-off between accuracy and computational cost. This is a particularly pertinent point for this study because of the FSI iterations, which inherently involve multiple runs of a CFD simulation for each FSI result.

The derivation of the RANS equations begins with simplifying the description of the fluctuating turbulent flow properties. This is achieved by describing the flow properties, such as velocity or pressure, in terms of a combination of their mean values and their corresponding statistical fluctuations [95]. For example, the velocity $\boldsymbol{v}(t)$ can be decomposed into a steady mean value $\bar{\boldsymbol{v}}$ and a fluctuating component $\boldsymbol{v}'(t)$ which is superimposed onto it, giving:

$$\boldsymbol{v}(t) = \bar{\boldsymbol{v}} + \boldsymbol{v}'(t) \tag{3.10}$$

Incorporating the mean and fluctuating components from equation (3.10) into the continuity and momentum equations in the Navier-Stokes equations, (3.4) and (3.5), and time-averaging the results, provides the steady Reynolds-Averaged Navier-Stokes equations [94]. This allows the time-averaged properties of the flow to be solved for, i.e. the mean velocities, mean pressures, mean stresses, for a steady representation of the flow. The RANS equations below are described for an incompressible flow, which is acceptable for parafoils as their typical flow velocity is less the 0.3 times the speed of sound in air.

$$\frac{\partial \bar{v}_x}{\partial x} + \frac{\partial \bar{v}_y}{\partial y} + \frac{\partial \bar{v}_z}{\partial z} = 0, \tag{3.11a}$$

$$\begin{aligned}\rho_F \left(\bar{v}_x \frac{\partial \bar{v}_x}{\partial x} + \bar{v}_y \frac{\partial \bar{v}_x}{\partial y} + \bar{v}_z \frac{\partial \bar{v}_x}{\partial z} \right) = &-\frac{\partial \bar{p}}{\partial x} + \mu \left(\frac{\partial^2 \bar{v}_x}{\partial x^2} + \frac{\partial^2 \bar{v}_x}{\partial y^2} + \frac{\partial^2 \bar{v}_x}{\partial z^2} \right) \\ &- \left(\frac{\partial}{\partial x}(\rho_F \overline{v_x'^2}) + \frac{\partial}{\partial y}(\rho_F \overline{v_x' v_y'}) + \frac{\partial}{\partial z}(\rho_F \overline{v_x' v_z'}) \right)\end{aligned} \tag{3.11b}$$

$$\begin{aligned}\rho_F \left(\bar{v}_x \frac{\partial \bar{v}_y}{\partial x} + \bar{v}_y \frac{\partial \bar{v}_y}{\partial y} + \bar{v}_z \frac{\partial \bar{v}_y}{\partial z} \right) = &-\frac{\partial \bar{p}}{\partial x} + \mu \left(\frac{\partial^2 \bar{v}_y}{\partial x^2} + \frac{\partial^2 \bar{v}_y}{\partial y^2} + \frac{\partial^2 \bar{v}_y}{\partial z^2} \right) \\ &- \left(\frac{\partial}{\partial x}(\rho_F \overline{v_x' v_y'}) + \frac{\partial}{\partial y}(\rho_F \overline{v_y'^2}) + \frac{\partial}{\partial z}(\rho_F \overline{v_y' v_z'}) \right)\end{aligned} \tag{3.11c}$$

$$\begin{aligned}\rho_F \left(\bar{v}_x \frac{\partial \bar{v}_z}{\partial x} + \bar{v}_y \frac{\partial \bar{v}_z}{\partial y} + \bar{v}_z \frac{\partial \bar{v}_z}{\partial z} \right) = &-\frac{\partial \bar{p}}{\partial x} + \mu \left(\frac{\partial^2 \bar{v}_z}{\partial x^2} + \frac{\partial^2 \bar{v}_z}{\partial y^2} + \frac{\partial^2 \bar{v}_z}{\partial z^2} \right) \\ &- \left(\frac{\partial}{\partial x}(\rho_F \overline{v_x' v_z'}) + \frac{\partial}{\partial y}(\rho_F \overline{v_y' v_z'}) + \frac{\partial}{\partial z}(\rho_F \overline{v_z'^2}) \right)\end{aligned} \tag{3.11d}$$

The RANS equations above introduce a number of extra terms, related to the velocity fluctuations. These are additional viscous terms known as the Reynolds stresses [96]. They form a tensor called the Reynolds stress tensor, in 3D:

$$-\rho_F \begin{pmatrix} \overline{v_x'^2} & \overline{v_x' v_y'} & \overline{v_x' v_z'} \\ \overline{v_x' v_y'} & \overline{v_y'^2} & \overline{v_y' v_z'} \\ \overline{v_x' v_z'} & \overline{v_y' v_z'} & \overline{v_z'^2} \end{pmatrix} \tag{3.12}$$

thereby producing six new independent variables to solve. In order for the equations to be solved, the unknown Reynolds stresses must be associated with the properties of the mean flow, in order for their values to be determined. This is known as the "closure" problem, and a great number of models relating the properties of the flow fluctuations to the behaviour of the mean flow have been developed. These models belong to several different groups, but the brief discussion here is limited to linear eddy-viscosity models, which is the type used in the simulations described herein. Information on other approaches and their derivation is found in review papers such as those by Leschziner [95] and Leschziner & Drikakis [96], as well as the work of Pope [94].

The starting point for models of these type is the proposal by Boussinesq that the Reynolds stresses could be linked to the mean rates of deformation, in a similar manner to viscosity [94]. This stress-strain relationship for an incompressible fluid, using Einstein notation, is

$$-\rho\overline{v_i v_j} = \mu_t \left(\frac{\partial V_i}{\partial x_j} + \frac{\partial V_j}{\partial x_i} \right) - \frac{2}{3}\rho k \delta_{ij} \tag{3.13}$$

where μ_t is introduced as an unknown eddy viscosity, $\boldsymbol{V}_i$ is the velocity vector, x_i is the spatial coordinate, and δ_{ij} is the Kronecker delta. A further parameter is introduced, k, which is the turbulent energy, defined $k = \frac{1}{2}(\overline{v}_x{}^2 + \overline{v}_y{}^2 + \overline{v}_z{}^2)$. This is present in the second term on the right hand side of (3.13), which is an adjustment included to ensure that normal stresses sum up to $2k$ in instances of zero strain [95].

The relationship proposed by Boussinesq in equation 3.13 expresses the fact that the level of turbulent activity, or mixing, is closely associated with the strain action within the flow. This relation can be used, if μ_t can be determined, to calculate the Reynolds stresses thereby closing the RANS equations. In order to achieve this it is useful to acknowledge the fact that from dimensional analysis, the (kinematic) turbulent viscosity written as $\nu_t = \mu_t/\rho$, depends on a linear combination of a turbulent length scale, ℓ, and a turbulent velocity scale, v_{turb}.

$$\nu_t \propto \ell v_{turb} \tag{3.14}$$

The simplest approach to determine a value for ν_t therefore is to use take advantage of an algebraic model of the form of (3.14). An example of this is the mixing-length

model of Prandtl and von Karman, which related the turbulent velocity scale to the cross-flow velocity gradient, and prescribed a value for the turbulent length scale based on the length-scale of the larger, most energetic eddies in a shear layer, as they are the largest contributors to mixing within the shear layer [96].

The aforementioned algebraic approach where the length scale is prescribed, may be appropriate for the simulation of very simple shear layers if the length scale is chosen carefully. However, in terms of applications to more complex flows this approach is too restrictive. Algebraic models are limited by their neglect of any transport of turbulence, which implies that the state of turbulence is entirely dictated by the local relationship between turbulence-energy production and dissipation. However, this is rarely the case and turbulence cannot generally be described purely by local relationships [96].

More realistic models consider the turbulent transport through the use of one (one-equation models) or two (two-equation models) differential equations to account for the evolution of the turbulence scales used to determine the eddy-viscosity. One-equation models still have a limitation, in that one of the turbulence scales is prescribed and not described by a transport equation. This precludes their use in flows containing separated flow for example [95]. Therefore, a two-equation linear eddy-viscosity model is used for the CFD simulations throughout this work.

The two-equation model selected for this study is the k-ε model, which as stated uses two transport equations to describe the turbulence length and velocity scales independently. The velocity scale selected is the root of the turbulence energy, k, and a surrogate for the length scale is used, the turbulent energy dissipation ε.

The results of the transport equations for k and ε are then utilised in a rephrased equation for the eddy viscosity

$$\nu_t = C_\mu \frac{k^2}{\varepsilon} \tag{3.15}$$

from which the Reynolds stresses can then be calculated, and the RANS equations closed [97].

The majority of fluid dynamic simulations within this work actually make use of a slightly adjusted version of the k-ε model, known as the Realizable k-ε model. This makes an adjustment for the fact that the standard k-ε model does not meet certain mathematical constraints on the normal stresses, and thereby non-physical results can

occur. The model has an adjusted equation for ε, and the constant C_μ present in the standard k-ε model is no longer a constant, but is calculated from the mean flow and turbulent properties. More detail on the formulation of the model is found in [93] and [98].

The selection of the turbulence model was based upon a number of factors. The lack of any comprehensive experimental data for the flow around a parafoil, even for an undeformed cross-section, made making an empirically justified choice extremely difficult. Preliminary simulations of a NACA0012 aerofoil were useful in terms of verifying the flow simulation set-up, and ensuring good mesh convergence, as well as showing good agreement for lift and drag when using the popular turbulence models. The literature suggested that the k-ω SST model and k-ε series of model were more suited to flows involving separation, when compared to models such as the Spallart-Almaras [99].

For complicated flow problems of this kind, the literature favoured use of the k-ω SST model, but once this was applied to the simulation of flow around deformed parafoil geometries, such as the ones seen after multiple FSI iterations, convergence was difficult to achieve. This prompted a move towards the k-ε family of models, from which the Realizable k-ε model (R.k-ε) was chosen. This was based in part on the conclusions of a study of flow over a backward facing step found in Tu [93], where the R.k-ε model was shown to provide the most accurate prediction of the reattachment length for the separation bubble downstream of the step. Similar results were observed for a backward facing step by Collie et al. [100], which also showed the R.k-ε model to provide better predictions and be more robust than the standard model. The Realizable k-ε model has also been used successfully in the simulation of flow around yacht sails, which typically exhibit flow separation [101, 102].

Given the lack of experimental data, the means for selecting the turbulence model in this case provides a justifiable first approximation. In order to provide a measure on how the selection of the turbulence model influences the results, a CFD study of the flow around the predicted structural geometry from Section 5.3, using the standard k-ε turbulence model, has been performed and detailed in Section 5.4. Such an analysis is not possible without an inflated geometry upon which to base it, and so the circular nature of the problem supports the need for a turbulence model that works to be chosen initially, such that the dependency on the choice of model can be investigated.

Wall treatment

When simulating external flows around an object, the interaction of the flow with the solid boundary of the structure is an important component of the flow field. This interaction, of course, results in the development of a boundary layer which for high Re flows is likely to be turbulent. In order for the influence of turbulence to be considered in the boundary-layer, the behaviour of the flow in direct vicinity to the wall must be taken into account. This region, adjacent to the wall, is known as the "viscous sub-layer". Within this layer, the viscous action is much stronger than in the rest of the boundary layer [103]. Correct simulation of the flow in this region requires either an incorporation of wall laws, based on an empirical relation known as the "law of the wall", or an adjustment to the turbulence model to extend its use down to the wall, often described as a "low-Reynolds number" version of the model [93]. In this study a "two-layer" approach is used, which as an alternative to using a modified version of the Realizable (or Standard) k-ε model, couples a specifically calibrated one-equation model near the wall to the k-ε model for the rest of the flow [104, 103]. Regardless of the approach, the objective for all models is to reflect the greater influence of viscosity that causes a higher rate of turbulent decay, and the inviscid wall-blocking, which combine in the viscous sub-layer to reduce the mixing in the direction normal to the wall [96].

Fluid dynamics solver overview

STAR-CCM+ has been selected for the fluid simulations as it is a tried and tested commercial CFD solver. The steady viscous incompressible RANS equations are solved in 3D, using either the Realizable (or Standard in Section 5.4) k-ε turbulence model for closure. The steady state model has been chosen because computationally it is much less intensive, although the steady solution does exclude transient effects. As the flow is shown to separate at the leading edge by Mittal et al. [50] and Balaji et al. [51], as well as in our own simulations, this does remove some detail. However, it is accepted as a first approximation and time-dependent effects offer a possibility for future work. The Finite Volume Method (FVM) discretised equations are solved using the SIMPLE algorithm [105, 106].

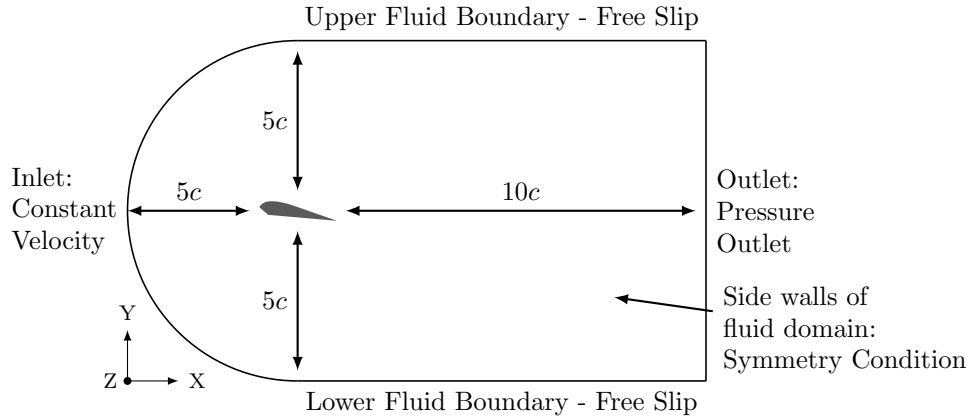


Figure 3.3: Schematic showing a 2D side view of the 3D fluid domain and applied boundary conditions. The depth (z-direction) dimension is determined by the results of the structural solver

Fluid domain boundary conditions

A schematic of the D-shaped fluid domain is given in Figure 3.3 in which the applied boundary conditions are labelled. At the surface of the parafoil, interface Γ, a no-slip boundary condition is applied, and all surfaces are treated as impermeable and rigid. The boundary conditions on the rest of the domain are uniform velocity for the inflow, pressure-outlet at the outflow boundary, free-slip on the upper and lower fluid walls, and a symmetry condition on the two side wall boundaries. As stated above the side walls of the fluid domain extend directly from the ribs, therefore any flow features that would move across the span of an actual finite-span parachute canopy, or any edge effects seen at the cells on the edge of the canopy, are not captured. Such features, however, would not be expected for the infinite-span case that is being simulated.

3.2.4 Loose coupling methodology

To couple the two domains, the following conditions are upheld at the fluid-structure interface, Γ

$$\boldsymbol{\tau}^S \cdot \boldsymbol{n} = \boldsymbol{\tau}^F \cdot \boldsymbol{n} \tag{3.16}$$

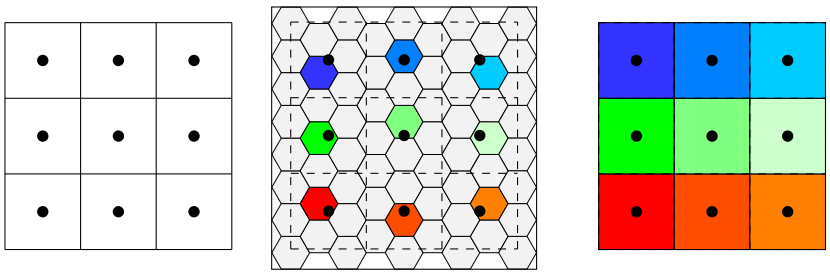
$$\boldsymbol{u}^S = \boldsymbol{u}^F \qquad \text{on } \Gamma \tag{3.17}$$

$$\boldsymbol{v}^S = \boldsymbol{v}^F = 0 \qquad \text{on } \Gamma \tag{3.18}$$

where $\boldsymbol{\tau}$ is the Cauchy stress tensor, $\boldsymbol{n}$ is a unit vector normal to the fluid-structure interface, $\boldsymbol{u}$ is the displacement of the domain, and $\boldsymbol{v}$ is the velocity of the domain [81]. Superscripts S and F denote properties of domains Ω_S and Ω_F respectively.

In order to realize the coupling between the two and to enforce the conditions above at interface Γ within a loosely coupled framework, data must be exchanged between the fluid and structural solvers. This is achieved through the transfer of two key parameters. Firstly, a deformed geometry, Φ, is output from the structural solver, to be read into the fluid solver. Once the flow around this geometry is computed the second parameter, the surface pressures, p_F, on the external parafoil surfaces are output from the fluid solver. This is used as loading conditions for the next structural simulation.

In order to aid convergence an under-relaxation step is performed on the structural geometry during each FSI iteration step, creating a biased combination of the starting and deformed geometries. The under-relaxation step is described by Equation 3.19, where Φ represents the structural geometry as a whole, and β is a constant between 0 and 1 that determines the bias between the contributions of the two geometries. The under-relaxed geometry Φ^*, for $\beta < 1$, provides a geometry which has a lesser degree of deformation than the result exported from the structural solver, Φ_{new}. This step is carried out to avoid large deformations in the structure from producing large changes in the aerodynamic loadings produced by the CFD simulation step, which could slow down the convergence or even cause it to diverge [107, 108]. This process was seen in preliminary studies to be an extremely valuable stage of the iteration cycle, as without it the differences in the structural geometry, Φ, between FSI cycles remained large and in some cases were seen to grow, signifying a divergence in the result. Due to this the under-relaxation step has been used in all of the FSI simulations performed



(a) (b) (c)

Figure 3.4: Diagram outlining transfer of surface pressure data between fluid and structural simulations. In (a) centroids of structural mesh elements are located and (b) used as sampling points to extract surface pressures from the results of the fluid simulation. Pressures are then applied (c) as pressure loadings to the elements for the next structural simulation. The difference in size between structural and fluid domain mesh sizes is exaggerated for clarity.

within this thesis. The value used for the under-relaxation coefficient is $\beta = 0.5$.

$$\Phi^* = \Phi_{old} + \beta(\Phi_{new} - \Phi_{old}) \tag{3.19}$$

The external aerodynamic loading seen on the structure is applied in the structural simulations as a pressure, as illustrated by Figure 3.4. This pressure data is extracted from the results of the fluid solver by sampling at different locations on the upper and lower parafoil surface boundaries, Figure 3.4(b). The locations of the sampling points correspond to the locations of the element centroids of the structural model, Figure 3.4(a), the coordinates of which are extracted automatically from the model by using LS-PrePost [109]. This means that each structural element on the upper and lower surfaces has an individual loading applied, Figure 3.4(c). This permits any span-wise differences in the flow field to be incorporated into the FSI analysis, providing a realistic loading. However, the pressure loading is not distributed between the nodes of the element in a way that ensures the conservation of work done, and consequently the loading is not strictly a consistent loading [110]. This simplified loading method is a satisfactory initial approach that is straightforward to implement and has been judged to provide a good level of accuracy, as demonstrated in Section 3.3.5.

The pressure internal to the parafoil structure is seen to be constant throughout the inside of the parafoil half-cell, and therefore is applied as a uniform loading to all elements on the upper and lower surfaces. A very slight variation is seen near the

leading edges of the upper and lower surfaces, but this variation is neglected for the sake of limiting complexity. As stated already and labelled in Figure 3.2, the pressure on the internal surfaces of the ribs is ignored and the loading associated with this is either disregarded, as in the BC1 loaded rib case, or prescribed, as in the BC2 unloaded rib case.

The data transfer and processing is handled by MATLAB [111], which is also used to automatically perform steps within other software packages, such as LS-PrePost and HyperMesh [112]. The surface mesh for the fluid boundary is automatically combined with the under-relaxed structural geometry using HyperMesh, and the combined geometry is exported in .stl format to be imported into the fluid dynamics software, STAR-CCM+ [113]. The construction of the CFD model and mesh, and the export of pressure data from the completed fluid simulation, is performed in STAR-CCM+ commanded by macros written in the Java programming language [114]. Note that the CFD model is re-meshed at every iteration. The coupling process is described in more detail in Figure 3.5, and specific examples of the code-base and specific implementation details used for the loose coupling are provided in Appendix B.

The first fluid-structure iteration is initiated by performing a structural simulation in which the applied loading is an estimated value for the internal pressure, $p_{initial}$ based on the stagnation pressure of the freestream flow, and the starting geometry is the cut pattern of the parafoil shown in Figure 3.1(a). In this first case the second rib is not allowed to move freely, and is instead gradually moved in the spanwise direction in a prescribed manner to avoid the large deformations which would otherwise be seen. The result of this simulation then provides the starting point of all further structural simulations, which are restarted with updated loading conditions from the fluid simulations. This way simulation times are lowered, whilst the susceptibility of the simulation to instabilities and mesh distortions occurring due to the geometrically non-linear nature of a large deformation problem with a light, flexible material, is reduced.

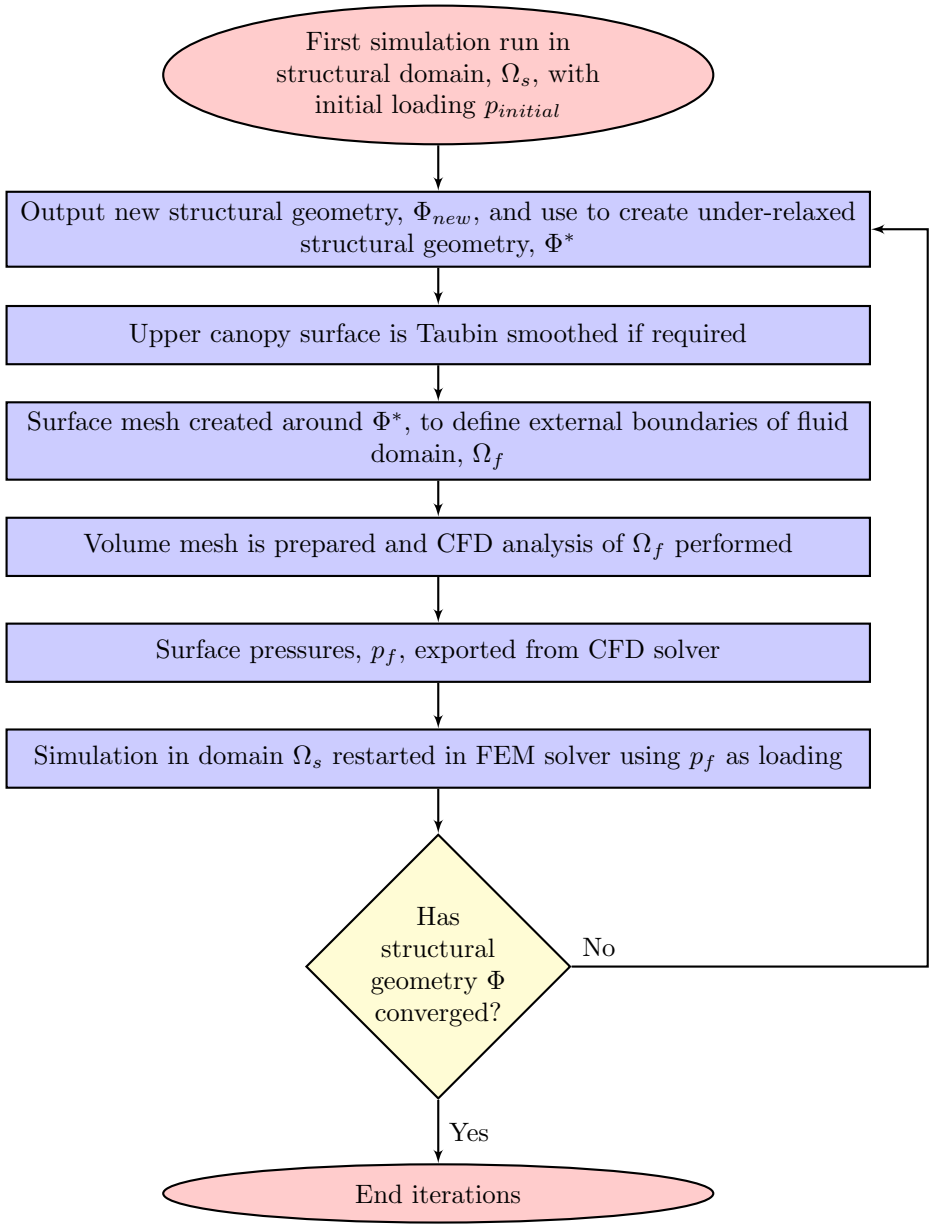


Figure 3.5: Description of FSI coupling algorithm.

3.2.5 Assessment of convergence for fluid-structure interaction iterations

As the loosely coupled FSI approach utilised in this study is iterative, it is important to have a metric by which the proximity of the result to the equilibrium state can be judged. In other words, a method by which the convergence of the iterations can be assessed is required. For all of the FSI simulations in this study the convergence criteria were based on the difference in resultant deformation between the results of the structural simulation step of each cycle. The difference between the position of selected points on the geometry was measured for the predicted geometry at each stage, and compared.

For the loaded ribs case BC1, three points along the centre line of the upper surface were used to judge the change in the structural shape between iterations. One was at the upper leading edge, one was at the thickest point on the parafoil, and one was in the region of the trailing edge. These points were chosen because the upper surface exhibited the largest amounts of deformation, and this was greatest at the centre-line of the geometry. The specific three points were chosen to provide a view of the structure as a whole, whilst reducing the complexity and ensuring that the convergence was judged upon the points where the maximum difference in the nodal co-ordinate was observed. This aided in assessing the susceptibility of the different parts of the structure to oscillations.

For the unloaded rib case BC2, the number of selected points was increased due to the different modes of deformation observed. The additional degrees of freedom created by relaxing the vertical constraint on rib 1, introduced a new parameter to be observed. The spanwise force that is applied in this case also meant that the lateral position of rib 2 was worthy of observation. Therefore, in addition to the three aforementioned points on the upper surface the resultant deformation at the lower leading edge of both ribs was compared, as well as three points along the lower surface for completeness. Again, specific points were chosen so that the convergence was judged against the largest differences in nodal position between iterations. It also provided a clear view of how the balance of forces changed over consecutive iterations. The specific levels of convergence achieved in each case are provided at the beginning of the relevant section in Chapter 5.

3.3 Additional modelling issues

The development of the loosely-coupled FSI simulation approach in this study required addressing a number of modelling issues related to the specific characteristics of ram-air parachutes, before FSI results could be successfully acquired. In this section some of these issues are described, and the techniques employed to overcome them are explained. It is hoped that their discussion may assist further research efforts and provoke consideration of better solutions should they be available, as well as making clear the justification for the selection of the various techniques employed throughout the FSI modelling process described herein. The examples here are all based on the model geometry described in Chapter 4, but the problems are related to the modelling of all parafoils and are not specific to the particular geometry.

3.3.1 Treatment of out-of-plane deformation at the trailing edge

The application of the periodic symmetry conditions discussed in Section 3.2.2 to the structural simulation affects the deformation seen at the rear portion of the parafoil cell. Without the planar constraints the parafoil structure in the vicinity of the trailing edge tends to bulge out in the spanwise direction, making the trailing edge area noticeably wider than the rest of the structure. This can be seen in Figure 3.6 and it is interesting to note that this mode of deformation appears to be the natural shape for the unconstrained parafoil. This suggests that all rectangular planform canopies are likely to exhibit similar behaviour, with the cell widening towards the trailing edge.

The behaviour observed in Figure 3.6 demonstrates the obvious need for the ribs to be constrained to achieve periodic symmetry (this is also advantageous when applying the spanwise loading to the loaded rib in the BC2 unloaded rib model, Section 5.3). The application of these constraints has a number of knock-on effects.

Firstly, when the planar constraint is applied the fabric of the upper and lower surfaces bulges around the rib, therefore crossing the plane of symmetry. In order to prevent this the rib is extended in the vertical direction in the affected area, thereby creating a “stopper” or end-plate that prevents the material crossing the symmetry boundary. With the inclusion of an appropriate contact scheme into the simulation, the stoppers successfully prevent the symmetry conditions from being broken. The stoppers were

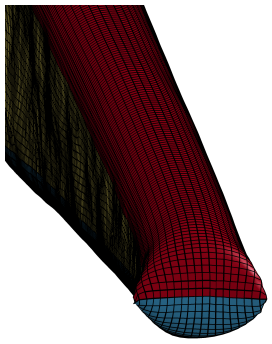


Figure 3.6: Out of plane displacement seen at the trailing edge when planar constraint is not applied.

required for the results of the loaded rib model, Section 5.2, but were not needed for the unloaded rib model BC2, due to the application of an additional spanwise loading force.

The stoppers are constructed out of the same fabric material as the rest of the parafoil, which prevents them having any influence on the shape of the ribs from which they extend. To enable the stoppers and canopy materials to interact a contact algorithm was required. This was chosen to be the `*CONTACT_AUTOMATIC_SINGLE_SURFACE` algorithm in LS-DYNA, which is a penalty method based contact approach. As both of the components are made of a soft fabric material, the `SOFT=2` option in the contact algorithm was activated, which changes the contact logic from node-segment based to being segment-segment based, and alters the penalty stiffness to be proportional to the element masses instead of the element bulk moduli [77]. This change was necessary to ensure reliable contact behaviour.

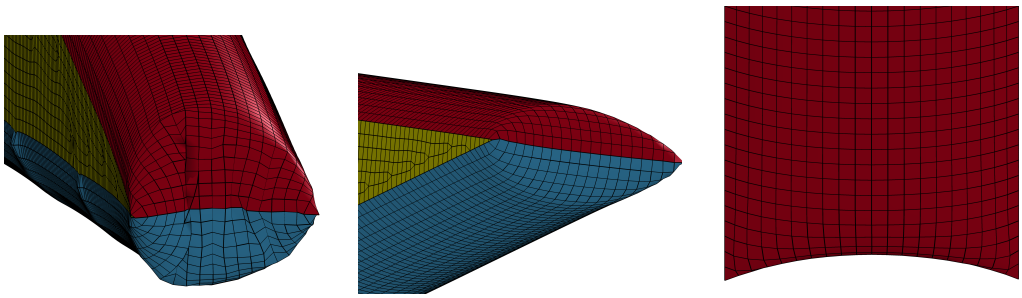
When an element on the upper or lower canopy surface makes contact with the stopper it is in the same plane as the periodic boundary that makes up the side wall of the fluid domain. Therefore when in contact with the stoppers an element no longer has a fluid flow around it. When the fluid geometry is created for the exported structural geometry these elements are incorporated into the side wall fluid boundary, thus slightly reducing the surface area of the canopy exposed to the flow. This transfer of elements from the parafoil to the fluid boundary had a very small effect on the fluid results but was necessary to allow for a consistent fluid mesh to be created at every step. It is also important to note that this is the explanation behind the irregular outlines of the rib observed in the results of Section 5.2.

3.3.2 Treatment of buckling behaviour at the trailing edge

A second consequence of the planar constraint on the ribs was that the trailing edge was compressed more than it would be if the rib was free to move out of plane. this compression caused the elements around the seam of the trailing edge to buckle and fold over themselves, an example of which can be seen in Figure 3.7(a). This poses a problem when exporting the deformed geometry to the CFD solver, as the folded and distorted elements create a very difficult surface for the volume meshing algorithm to deal with.

This issue prevents a fluid volume mesh from being created around the structure; a major obstacle in the process of performing the FSI iterations. To avoid any manual adjustment or correction of the deformed structural mesh, the capabilities of the fabric material model within LS-DYNA were explored. The key benefit of using this model, described above, is the option within the `*MAT_FABRIC` constitutive relations to eliminate compressive stresses. When this option is activated the solver does not permit any compressive stress to be sustained within the material and the element itself reduces in size. This reflects the true behaviour of fabrics, which provide no resistance to compressive forces, producing in-plane deformation [115]. Due to the obvious incompressible nature of a fabric material this in-plane deformation is only possible due to some of the material also moving out of plane, thus producing wrinkles.

The out-of-plane behaviour seen in fabric wrinkling is difficult to simulate directly in finite element simulations without the need for a very high level of element resolution. The compressive stress elimination algorithm (CSE) in LS-DYNA [77] however, provides a means through which the in-plane deformation is facilitated without the need for simulating the wrinkles themselves [65]. The result of activating the CSE option in the elements in the trailing edge region is an obvious reduction in the amount of buckling behaviour, and the reduction in size of the membrane elements under compressive loads produces a smooth trailing edge, shown in Figures 3.7(b) and (c). This solves the problem of creating an acceptable CFD mesh around the deformed shape of the structure, while providing a realistic characterization of the physics of the problem.



(a) (b) (c)

Figure 3.7: Buckling and folding seen at (a) the trailing edge, and the improvement gained by activating the Compressive Stress Elimination (CSE) wrinkling model seen from (b) below and (c) above.

3.3.3 Effect of fabric wrinkling on overall deformation

In addition to the LS-DYNA fabric wrinkling model being applied to the trailing edge, the effect it has on the structural response as a whole was considered. To be consistent and avoid any false boundaries in the simulation of the parafoil fabrics the CSE option was applied to the whole of the upper and lower surfaces in line with the results from the trailing edge. However, the suitability of this to the cell ribs had to be determined.

To test the effect the wrinkling model would have on the ribs' behaviour, a simplified test case was run, using the half-cell model with only internal pressure loadings applied. To further simplify the model a constraint on the amount of span-wise movement was added. A comparison of the structural response of the structure with and without the fabric wrinkling model can be seen in Figure 3.8.

The test results show that the overall shape of the parafoil cell is altered by the inclusion of the wrinkling model in the representation of the ribs. The comparison in Figure 3.8 shows a large enough effect on the results for it to be significant; with Figure 3.8(b) showing a reduction in the length of the nose of the parafoil under wrinkling. This results in a blunter, more rounded leading edge and a steeper angle at the inlet. The results in Figure 3.8(a) show a stiffer response, with the ribs maintaining a shape close to that of the cut pattern, and the upper leading edge moving into a peak shape. The difference in geometry between the two is great enough that it could affect the characteristics of the flow over the upper surface.

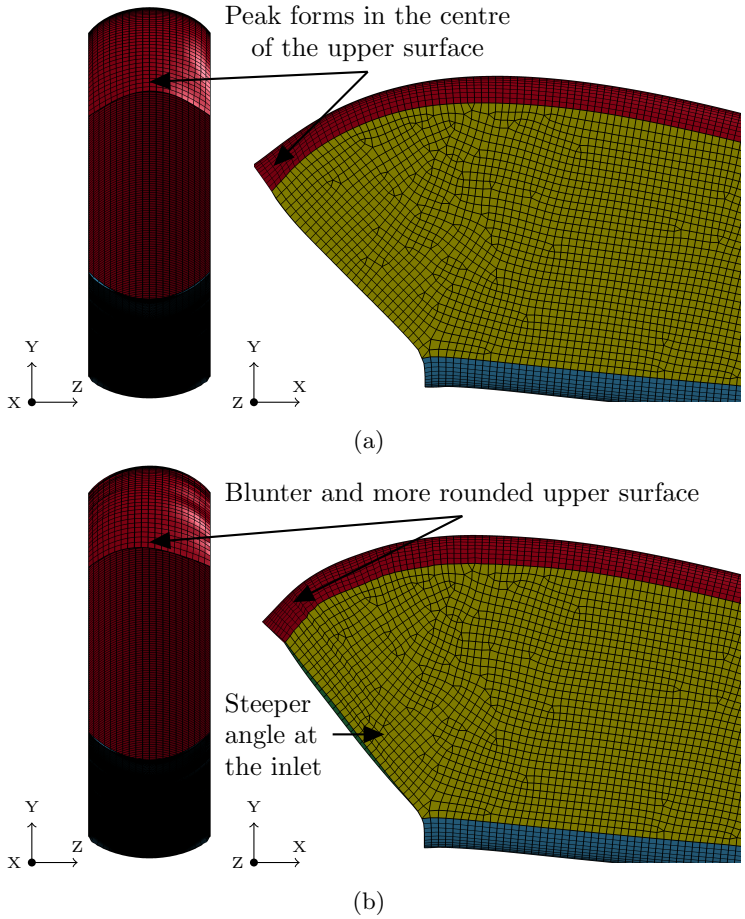


Figure 3.8: Comparison of results for structural simulation test case. Front (left) and side (right) views of the leading edge deformation: (a) standard membrane elements applied to ribs; (b) CSE wrinkling option applied to ribs.

For all of the FSI studies the CSE wrinkling algorithm is utilised for the fabric material in every part of the structure, including the ribs, as the results in Figure 3.8 show the structural response to be overly stiff without it. However, the effect of the wrinkling model may be over-estimated in this case, as the geometry used does not include any stitching or reinforcement that may be included in the design of parafoils to control this type of deformation. Furthermore, the material properties used in the model are based on characterization of fabric from a different model of ram-air parachute, because details on the fabric used in this particular parafoil geometry were not available. This is also likely to have an effect on the predicted degree of deformation.

3.3.4 Application of Taubin smoothing to deformed structural geometry

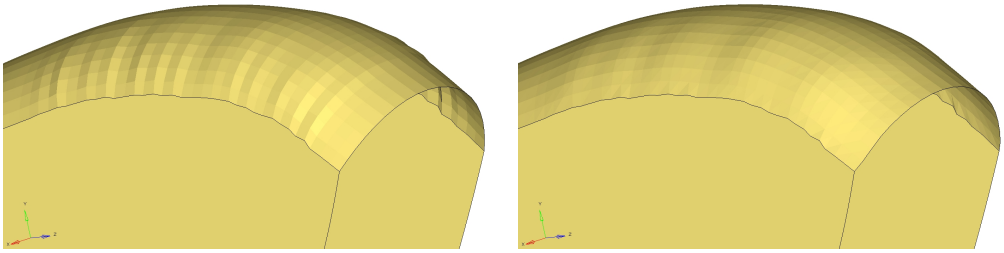
Whilst a controlled loading approach, described in Section 3.2.2, reduces the mesh distortion and element hourglassing in the structural results, there are still some distortions which persist [116]. This is the case on the upper surface towards the leading edge, in the area where the ribs are subject to the biggest amount of planar deformation. The problem arises from the wrinkling algorithm being activated in the ribs, which allows the elements in the ribs to compress as discussed in 3.3.2 above. This compression affects the position of the elements at the seam between the rib and the upper surface, and the change in position of the rib nodes connected to the upper surface creates distortion in some of the upper surface elements.

The impact of this is that the front part of the upper surface is not completely smooth. Although these small individual distortions have no effect in the shape of the parafoil as a whole, they do have a knock-on effect when the geometry is transferred to the fluid solver. The distorted elements create small raised bumps in the surface which act as small obstructions to the flow, which produce noticeable spikes in the surface pressure because the leading edge area where they are located is an area of fast, attached flow. This also means that the coupling of the aerodynamic forces to the structure may magnify this issue even further and may have an adverse effect on the convergence of the FSI iterations.

In order to reduce these small bumps in the parafoil surface and avoid any localised deformations, the shape of the upper surface geometry is smoothed before the fluid simulation step. A Taubin smoothing algorithm is used because it has been shown

to be very effective in removing high-frequency noise from surfaces without changing the underlying geometry [117, 118]. The smoothed surface is then recombined with the rest of the structure and transferred to the fluid solver. The surface pressures for the smoothed surface show reduced fluctuations but the larger scale bulges created by the reaction of the fabric to the structural loads are still represented.

The Taubin smoothing was performed on the upper surface only, using the MeshLab program [119]. The degree of smoothing is determined by two parameters, λ_{TS} and μ_{TS}, as well as by the number of smoothing cycles used. During this study the parameters were chosen as $\lambda_{TS} = 0.5$ and $\mu_{TS} = -0.53$, and the number of smoothing steps was chosen as 10. The difference between the deformed structure and the smoothed structure can be seen in Figure 3.9 below.



(a) without smoothing (b) with smoothing

Figure 3.9: View of leading edge area of parafoil, showing (a) small mesh distortions on upper surface, and (b) the upper surface area after Taubin smoothing filter is applied.

The Taubin smoothing step is only required for the simulations of the loaded ribs case, with BC1 boundary conditions, Section 5.2, due to the fabric material properties used for that case, in particular the value of the Poisson's ratio, $\sigma \neq 0$. Taubin smoothing was not used for the results in Section 5.3. Further information on the material properties used for the two different cases is provided in Section 4.4.

3.3.5 Effect of pressure sampling resolution on calculation of aerodynamic loading

A key aspect of the FSI iteration cycle developed herein, Section 3.2, is the extraction of the aerodynamic pressures on the parafoil surfaces from the fluid simulation and the mapping of these to the structural simulation as loading conditions. As already stated,

the most efficient way to apply the loading in LS-DYNA is to apply a pressure loading which is applied to each element independently. The pressure is then converted within the FEM calculation process into an applied load at each node of the element via the shape functions of the particular element.

The resolution of the mesh of the fluid domain is required to be much finer than the structural mesh of the parafoil, and this therefore presents a question as to whether the pressures extracted from the CFD model and applied to the CSD model accurately capture the variations in pressure seen in the fluid domain. In order to test this a comparison was made between two different sampling resolutions.

Firstly, the surface pressures were sampled from the fluid model at the locations equivalent to the centroid of each structural element. For comparison, the same fluid simulation was sampled at five equally spaced points within each structural element, and averaged to provide a single pressure value at that element. This was performed using the results of a fluid simulation around the second FSI iteration of the BC2 unloaded rib case to ensure that the parafoil geometry included three-dimensional spanwise deformation, and thereby the flow produced pressure gradients in both the chordwise and spanwise directions. This provides a more realistic pressure field upon which to test the required pressure sampling resolution compared to the undeformed geometry.

The calculated pressure loadings were compared for chordwise strips of elements at five different spanwise locations. The results for a sample of three of the strips are found in Figure 3.10. Strips at 25% and 75% of span were also compared, showing even closer matches between the sampling methods. These are not included in the figure for clarity reasons. Pressures have only been plotted over the first half of the chord because beyond this the pressures calculated by the two methods are identical.

The results in Figure 3.10 show that the two pressure sampling resolutions produce very similar loadings. The pressure distributions compare very well, with differences only observed in the first 15% of the chord, beyond which the results of the two sampling methods are identical. The main points of difference between the two occur predominantly at points of local maxima or minima. The largest deviation is seen at the spanwise position of 95% span, in the region near the upper leading edge, and is a localised difference of approximately 4%. Given that this maximum value is a reasonably small error, and that it is at a localised point on one strip out of

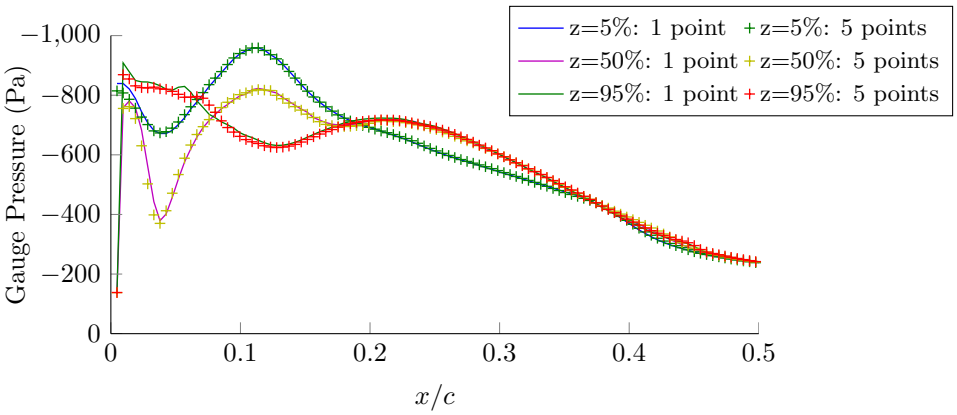


Figure 3.10: Comparison of chordwise pressure loadings for the upper surface of the structural model, for two different sampling resolutions, sampled from fluid simulation results at three spanwise planes. Loadings generated by two different sampling methods: pressure recorded at 1 point per structural element; pressure averaged from 5 sampling points per structural element.

5, the single pressure sampling approach was deemed to be satisfactory for use in all FSI simulations. This was also preferable due to the significant increase in runtime associated with sampling at five points for each element, which was seen as disproportionate for the very slight increase in accuracy seen.

Further details of the specific computational approach used to extract the pressure values in Star-CCM+ are found in Appendix B.

Chapter 4

Parachute geometry and properties

In this section the specific details of the parameters chosen for the CFD and FSI simulations are discussed in terms of the selected parafoil geometry, the resolution of the fluid and structural models, associated material properties, and the properties for the flow around the canopy structure.

4.1 Parachute geometry

The geometry of the ram-air canopy used within all of the studies within this thesis is based upon the cross-section of a parachute canopy used in wind-tunnel tests documented by Ware & Hassell [33]. This design was chosen because it was, to the author's knowledge, the only accurate description available in the literature at the time, and the associated experimental test data provides confidence that the parafoil design is sound both structurally and aerodynamically. The test data also provides an opportunity for basic validation, but it is important to note that an exacting computational replication of the wind-tunnel tests was neither an achievable nor desired objective for this programme of work.

The coordinates for the parafoil cross-section are provided in Appendix A. The section used throughout this study includes a slight modification from the one provided

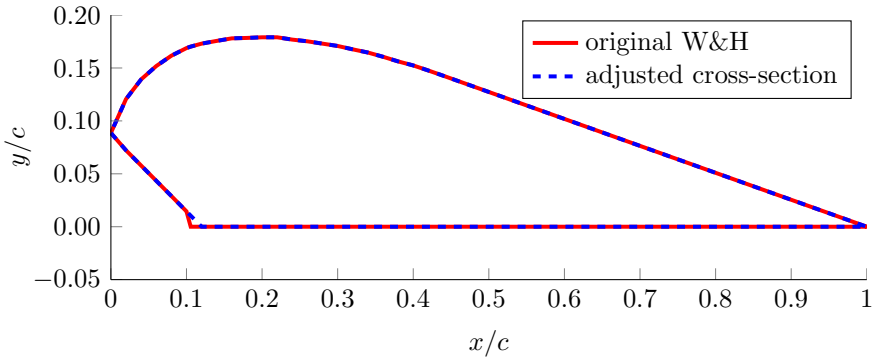


Figure 4.1: Cross-section of Ware & Hassell parafoil, and the adjusted cross-section used throughout this study.

in [33], which is adjusted to remove a kink in the join between the inlet face and the lower leading edge. The original and adjusted parafoil cross-sections are plotted in Figure 4.1.

The cross-section shares a similar outline to the Clark-Y aerofoil, particularly in terms of the flat lower surface. However, the cross-section differs from the Clark-Y in a number of ways including the larger maximum thickness of $0.18c$, compared to $0.12c$ of the Clark-Y, and closer inspection confirms that it is a completely independent aerofoil section and is not a scaled version of a Clark-Y. In fact, the parafoil cross-section does not seem to be based at all on an equivalent standard aerofoil shape. The leading edge cut of the parafoil is at an angle of 36° to the freestream.

For the two-dimensional fluid dynamic simulations detailed in Section 5.1 the cross-section was the only part of the design required, but for the results of the three-dimensional FSI simulations detailed in Sections 5.2–5.4, a full 3D structural geometry was required. The spanwise width of the three-dimensional model was estimated from diagrams of the inflated canopy provided in [33], as the uninflated dimension was not explicitly stated. The dimensions for the geometry of the half-cell can be seen in Figure 4.2. Reinforcements and stitching within the parachute have been neglected, owing to a lack of details of the parafoil's design and construction.

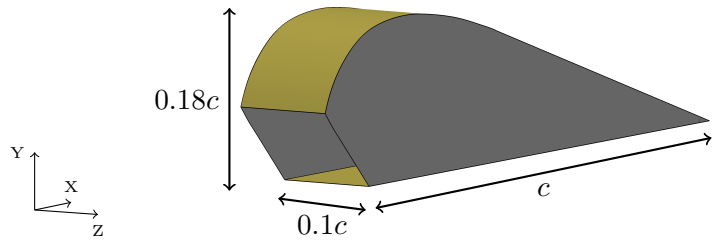


Figure 4.2: Un-deformed half-cell parafoil geometry with main dimensions.

4.2 Mesh resolution

Given the selected parafoil geometry described above, the mesh resolutions for the simulation of the structural and fluid domains were selected accordingly.

4.2.1 Structural domain

The upper and lower surfaces are discretised using four-sided membrane elements, approximately 5mm (0.5% of chord) square. The ribs are meshed identically using a combination of three and four-sided membrane elements, of a similar size to the upper and lower surfaces. The mesh size was chosen very carefully, in order to achieve a good level of accuracy whilst maintaining a manageable time-step. This was extremely important for the structural simulations of the loaded ribs model, BC1, which used an explicit solver. Limiting the smallest element size was necessary due to the direct impact this had on the time-step, which was of the order of 2×10^{-6} seconds, and the effect this subsequently had on the simulation run time. For the FSI simulations of the unloaded ribs case, BC2, an implicit solver was used and therefore the time-step was no longer directly dependent on the mesh resolution. For the implicit solver, a variable time-step was used, automatically adjusted by the solver, which was of the order of $\approx 0.02 - 0.4$ seconds. For consistency, the same structural mesh was used in both instances. Typical run-times for the explicit solver were of the order of 24 hours, whereas the implicit solver took around 20 minutes.

Convergence of the structural mesh resolution was tested by simulating the response of the parafoil cell to inflation via an internal loading. The loading was applied to the internal upper and lower surfaces and was equal to the stagnation pressure (645Pa) observed at a velocity of 33m/s. An external spanwise loading equal to the freestream

dynamic pressure of 645Pa was applied to one of the ribs as in the BC2 case described in Section 3.2.2, Figure 3.2. All other boundary conditions were the same as those defined for the BC2 case. Figure 4.3 compares three key dimensions of the parafoil cell at different mesh densities, and shows minimal dependence of these dimensions on the mesh resolution. The values measured at the chosen mesh density are denoted by the $*$ symbol. The chosen mesh resolution can be seen from a number of viewpoints in Figures 3.2, 3.6, 3.7 and 3.8.

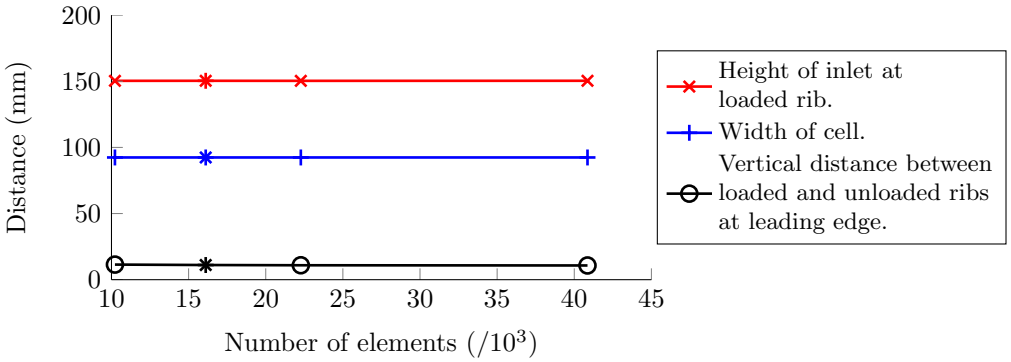


Figure 4.3: Mesh resolution convergence study for structural simulation of parafoil model. Values relating to the chosen mesh resolution are denoted using the $*$ symbol.

4.2.2 Fluid domain

As shown in Figure 3.3, a D-shaped domain of size 16 by 10 chord lengths was used. This is based on the work of Mohammadi and Johari [52], and the domain size was seen to be large enough such that the presence of the boundaries had no effect on the results. For the three-dimensional simulations discussed in Sections 5.2–5.4 the undeformed depth of the mesh is $0.1c$, a value that changes with successive FSI iterations, in line with the value for the width of the cell determined in the structural simulations. Figure 4.4(a) provides a side view of the fluid domain and its discretisation.

An unstructured polygonal mesh is used to discretise the vast majority of the domain, including the internal volume of the canopy, with the only exception being the boundary layer mesh adjacent to the upper and lower surfaces. Full consideration of the flow within the interior region of the parafoil is shown by the results of the 2D study in Section 5.1 to be necessary in order to accurately assess the flow around and downstream of the lower leading edge. Figure 4.4(a) shows the large area of high resolution

used in the volume surrounding the parafoil, which is achieved by prescribing the size of the elements within three specified regions of the domain. The three regions are shown by boxes marked by dashed lines in Figure 4.4(b); one surrounding the parafoil which has the finest resolution, one in the wake region directly behind the parafoil with slightly larger elements, and a larger volume that includes both of the other two regions with a further reduced resolution. The volumes extend through the full depth of the 3D models. The elements within the higher resolution areas can be seen in more detail in the close-up views of the mesh surrounding the parafoil in Figure 4.5.

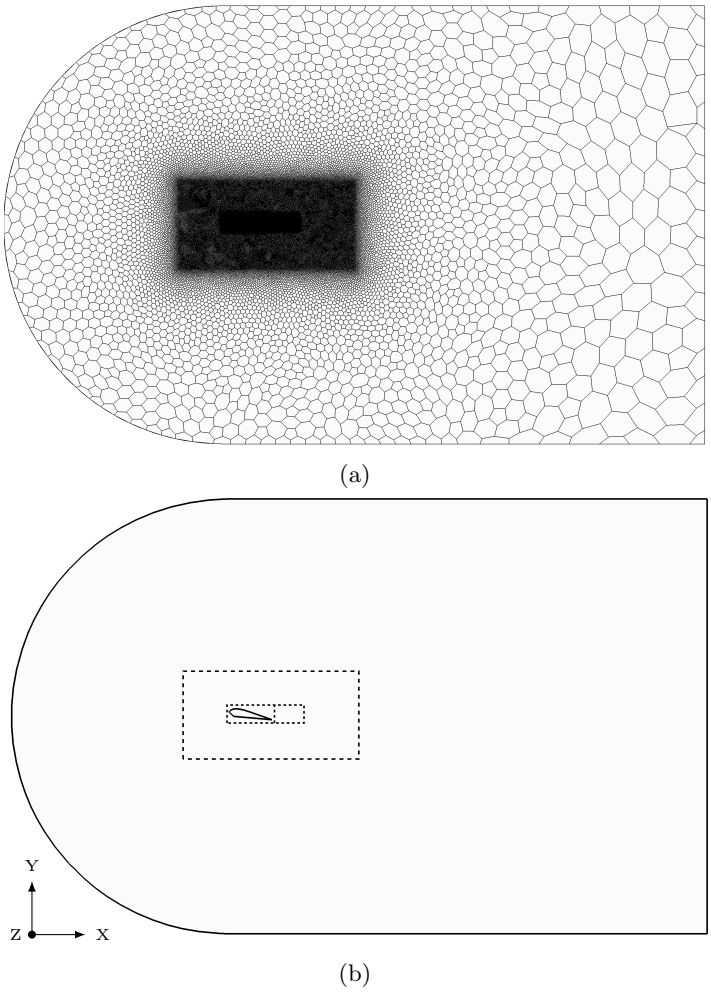
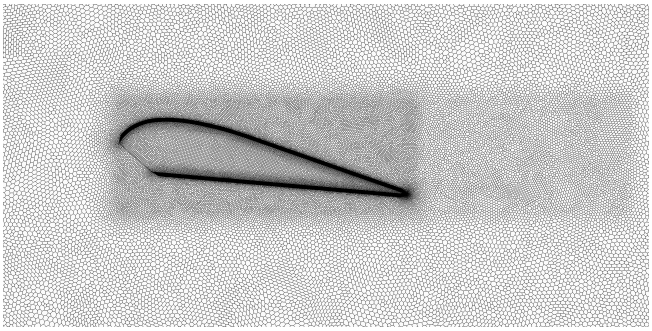


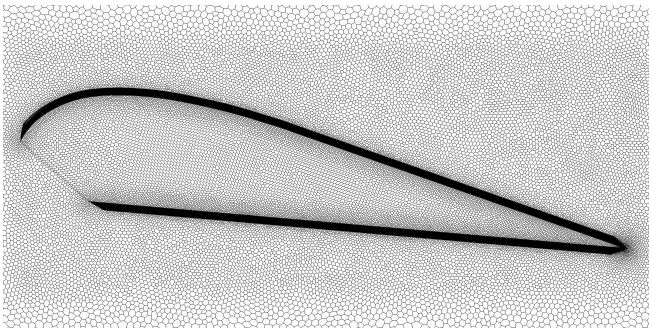
Figure 4.4: Side view of fluid domain around an undeformed parafoil showing (a) overview of fluid mesh and (b) schematic of regions where high resolution is prescribed.

As well as the specified increase in resolution in the proximity of the parafoil structure,

Figures 4.5(a) and (b) also show the thin hexahedral boundary-layer type mesh used to discretise the region in immediate vicinity (1.2% of chord length) to the upper and lower external surfaces of the canopy. This value was reached during preliminary tests, by using the contours of turbulent viscosity to judge the thickness of the boundary layer over the forward part of the upper surface. The first cell adjacent to these surfaces is of sufficient size to place the first node at a non-dimensional wall distance, or y^+, of $y^+ \approx 1$, necessary for the viscous sub-layer to be resolved by the two-layer near-wall turbulence model [93]. Fifty elements are applied in the thickness direction of the boundary layer mesh. Due to the low velocities seen within the parafoil, a boundary layer mesh is not applied to interior surfaces. The surfaces of the parafoil are modelled as solid, impermeable surfaces of zero thickness. The construction of the boundary layer mesh at the upper and lower leading edges, and at the trailing edge is shown in Figures 4.6(a), (b) and (c) respectively.



(a)



(b)

Figure 4.5: Magnified side views of the mesh resolution in the areas surrounding the parafoil.

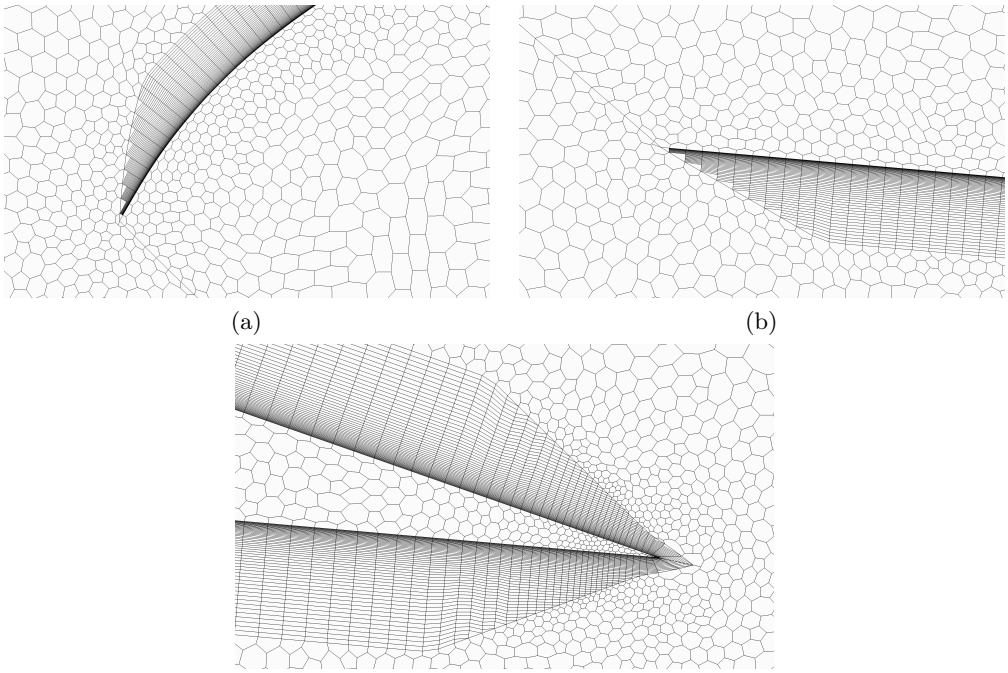


Figure 4.6: Figures showing the construction of the boundary layer mesh at (a) the upper leading edge, (b) lower leading edge, and (c) the trailing edge of the parafoil. A boundary layer mesh is generated on the external surfaces of the parafoil only.

As seen in Figure 4.4(a), to reduce the total element count for the mesh the element sizes in the domain are graded, allowing for larger element sizes in the regions far from the parafoil. The more rounded spanwise shape of the deformed parafoil for the infinite span BC1 case, Section 5.2, resulted in large fluid mesh sizes, with a total element count of $\approx$ 9–10 million. The flatter shape of the deformed parafoil for the BC2 model, Section 5.3, resulted in smaller fluid meshes, with an element count of $\approx$ 6 million. All models were meshed on a PC with 12 Gb of RAM, for which 10 million elements was the extreme upper limit in terms of computational resources, resulting in frequent problems in software stability. Meshes of $\approx$ 6 million elements were more appropriate, whilst still making the most of the available resources. The 3D mesh was shown to be mesh independent, demonstrated by the results for lift and drag coefficient in Figure 4.7 where the results for the undeformed 3D mesh at the chosen resolution are marked by the $*$ symbol.

For all cases convergence of the fluid simulation was judged by ensuring that a sufficient number of iterations was carried out for each flow simulation such that the normalized residuals decreased by at least three orders of magnitude, and the lift and drag forces had stabilised. A typical CFD simulation, run on 16 CPUs, took between 15–25 hours to achieve convergence. A full FSI iteration cycle therefore took between 17–27 hours when using the implicit structural solver, resulting in a total run time of around two weeks for 10 FSI iteration cycles.

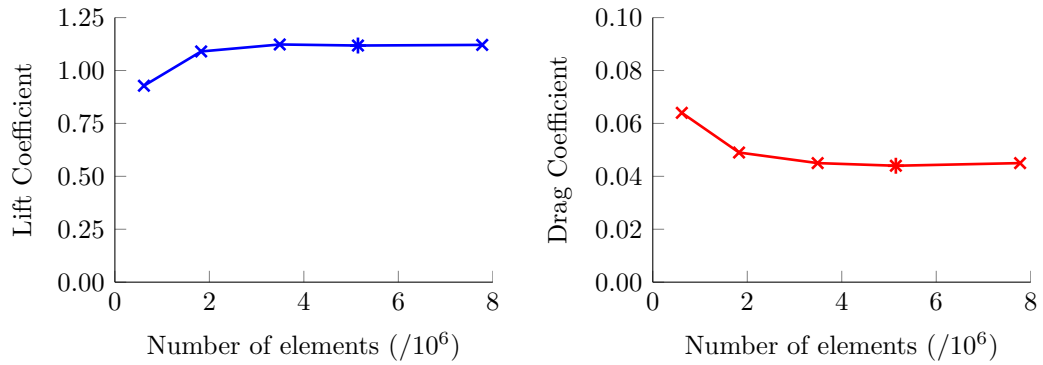


Figure 4.7: Mesh resolution convergence study for (a) lift and (b) drag coefficients of initial undeformed 3D parafoil model. Values relating to the chosen mesh resolution are denoted using the $*$ symbol.

4.3 Flow parameters

In each of the 2D and 3D simulations, the angle of attack is fixed at $\alpha = 5°$ relative to the freestream. Unlike the convention for aerofoils, where the angle of attack is measured against the chord line, the convention for parafoils is for α to be measured against the lower surface. This convention is also adopted here. The material simulated for the fluid in all cases is air, with the properties equal to those of Table 4.1. The same inlet velocity and Reynolds number is applied in every case, based on a characteristic length of a chord length of 1 metre. The Reynolds number has been selected to be representative of flight conditions, and to match the value used by Ware & Hassell to allow for a comparison to be made between the simulation and experimental results.

Table 4.1: Flow parameters applied in all fluid dynamics simulations.

Parameter	Value
Fluid	air
Reynolds number	2.1×10^6
Velocity (m/s)	33
Fluid Density (kg/m^3)	1.18415
Dynamic Viscosity (Ns/m^2)	1.855×10^{-5}
Dynamic Pressure (Pa)	644.8
Angle of attack (degrees)	5

4.4 Material properties of parachute fabric

The material properties used for the first three-dimensional FSI simulations, for which the boundary conditions represent an infinite span canopy, Section 5.2, are provided in Table 4.2. These were based on initial measurements and approximations based on values found in the literature, as material properties were not readily available for a ram-air parachute. For the second FSI case, Section 5.3, where the boundary conditions are adjusted to include the effects of the finite canopy span, the material properties of the fabric were updated.

The updated material properties for the BC2 case, Section 5.3, were based on tensile testing of samples of fabric from an AerAzur [120] tandem personnel ram-air canopy. This particular canopy was constructed from two different types of Nylon material; the lower canopy surfaces and the ribs are made of a ripstop nylon, whilst the upper

Table 4.2: Approximated material properties utilised in periodic symmetry, loaded rib parafoil model, BC1 (Section 5.2).

Parameter	Value
Young's Modulus E_1 (MPa)	100
Young's Modulus E_2 (MPa)	100
Density (kg/m^3)	3100
Poisson's ratio	0.3
Shear Modulus (MPa)	38.5
Thickness (μm)	30.5

canopy surface is constructed of a ripstop nylon fabric which is coated to ensure zero permeability of air through the material.

The details of the tensile tests are given in Table 4.3 below. The two woven fabrics are constructed of nylon yarns running in two principal directions; the warp direction which runs continuously through the entire length of the fabric roll, and the weft or fill yarn direction that runs perpendicularly across the fabric, interlacing with the warp yarn [121]. The yarns are therefore at 90° to each other, illustrated by the rip-stop pattern seen in the close-up view of a fabric sample during testing in Figure 4.8(a). Each of the two fabrics was tested in the weft and warp directions, and each test was repeated five times using samples cut from the same piece of fabric. The sample preparation and test procedure were all in accordance with ISO 13934-1:2013 [122]. Due to their extremely small thickness, the fabric samples were held in pneumatically driven grips, seen with the compressed air-line attached in Figure 4.8(b).

Table 4.3: Test parameters for fabric material tensile strength tests.

Parameter	Value
Extension rate (mm/min)	100
Sample width (mm)	50
Sample gauge length (mm)	200
Number of tests per fabric type	5

All of the tensile tests were recorded using a high-speed camera, to enable photogrammetric measurement of the strains in the samples. A close-up view of a still taken from the high-speed video footage of a test on the lower-surface fabric can be seen in Figure 4.8(a) showing the fiducial markers used for the photogrammetry. The use of high-speed video was motivated by the need to ensure that there was no measurement error introduced as a consequence of the inclusion of the grips into the load path. The results of the photogrammetry confirmed that no error resulted from the use of the grips.

However, comparison of the photogrammetry results with the stress-strain curves generated from the testing machine revealed that the fabric in all of the tests had slipped. The stress-strain curves for the weft and warp directions of the upper surface fabric material are in Figure 4.9, where the values are engineering stress η and engineering

strain e, defined in (4.1),

$$\eta = \frac{F}{A} \qquad\qquad e = \frac{\Delta L}{L_0} \tag{4.1}$$

where F is the uniaxial force applied to the sample, A is the cross-sectional area, ΔL is the change in length of the sample and L_0 is the sample's gauge length. The photogrammetry results revealed that the fabric started to slip in the grips after around the 5% strain point.

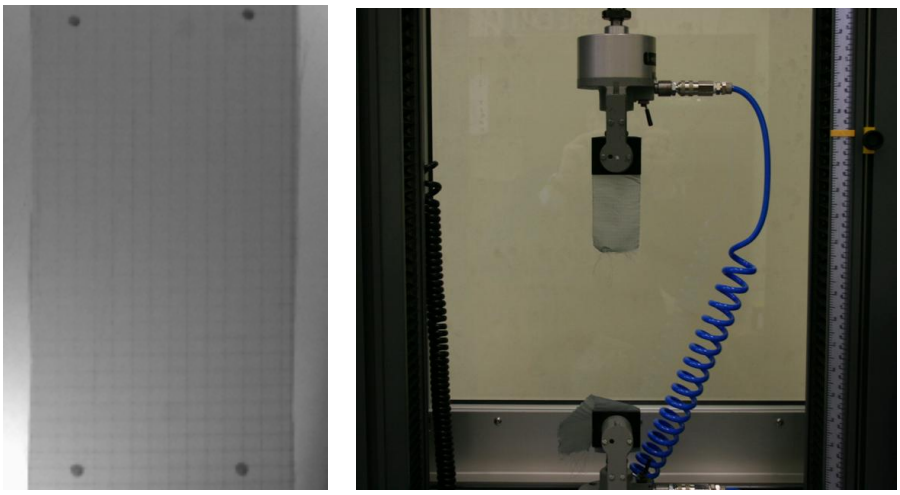
As the stress-strain curves in Figure 4.9 show, there is a large difference in the gradients of the curves for the two directions, showing a large difference in the Young's modulus of the fabric in each direction. The curves also show that for the early part of the stress-strain curves, the relationship is linear. This linear region persists to greater levels of strain in the weaker weft direction.

As it was unlikely that the problem of the fabric slippage would be resolved readily, and given that the design of a typical parafoil canopy is designed to cope with much higher stresses experienced during the deployment and opening phases, than those experienced during steady flight, it is reasonable to conclude that a 5% strain limit is a reasonable window from which to extract the material strength data. Furthermore, for the sake of simplicity, the linear portion of the curve is used to calculate the Young's Modulus, demarcated by the dashed lines on Figure 4.9.

The stress-strain data from the upper surface material is provided in Figure 4.9 as an example, and it is also representative of the behaviour of the lower surface material. The calculated Young's modulus values, based on the region of initial linear behaviour, for the upper and lower surface fabrics in both directions are provided in Table 4.4.

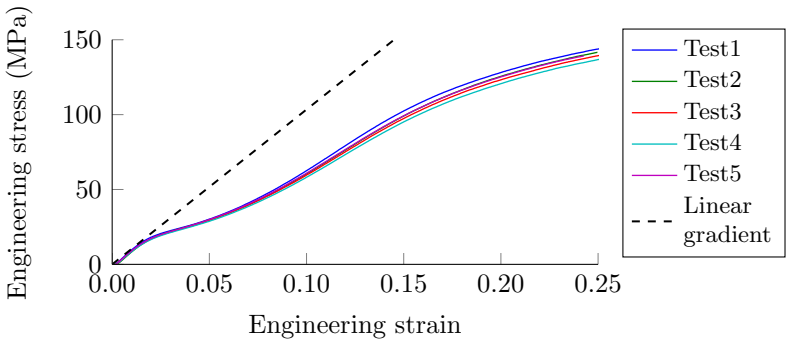
Table 4.4: Results for Young's Moduli of nylon fabric in two principle directions, calculated from tensile test results. Values from tests performed by DGA Toulouse included for comparison.

Material	E Average (MPa)	E - DGA value (MPa)
Lower surface: warp	924	658
Lower surface: weft	425	453
Upper surface: warp	1034	760
Upper surface:weft	442	407

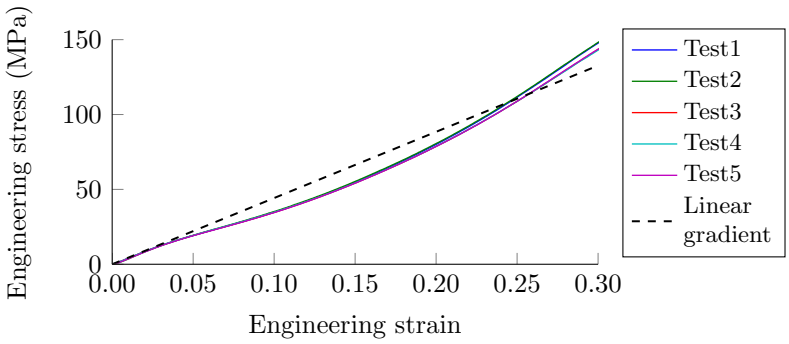


(a) (b)

Figure 4.8: (a) close up of lower surface fabric under tensile load and (b) view of fabric sample within testing apparatus post-failure.



(a) Warp direction



(b) Weft direction

Figure 4.9: Stress-strain results from tensile testing of upper surface fabric, in warp and weft directions.

It is clearly apparent from the values in Table 4.4 that both materials are highly orthotropic. Also included in the table are comparison values of the material properties from tests carried out by the DGA, Toulouse [123]. This shows a degree of variability to the fabric properties, but confirms the significant difference observed between the strength properties in the weft and warp directions.

In order to reduce the complexity of the structural model, the orthotropic nature of the fabric is ignored, and instead the fabric is modelled as being isotropic, with a representative value of Young's modulus of $E = 500$ MPa. Due to the difficulties experienced during the tensile tests, shear tests were not attempted, and so a value for the shear modulus, G, consistent with an isotropic material, was calculated using Equation 4.2 [124], for a Poisson's ratio $\sigma = 0$.

$$G = \frac{E}{2(1+\sigma)} \tag{4.2}$$

This choice of shear modulus is a major approximation for a woven material, but preliminary structural simulations with a shear modulus reduced by an order of magnitude, showed only minor adjustments in the resulting structural shape. Incorporation of a more comprehensive fabric characterisation is easily accomplished within the FSI framework utilised for this study, and so remains a subject for further work.

Observations of the fabric under tensile load via the high-speed camera footage showed a negligible Poisson's ratio in all cases. The same fabric material properties are applied to all of the fabric surfaces within the parafoil structure, summarised in Table 4.5.

Table 4.5: Material properties utilised in half-symmetry, unloaded rib parafoil model, BC2 (Section 5.3).

Parameter	Value
Young's Modulus E_1 (MPa)	500
Young's Modulus E_2 (MPa)	500
Density (kg/m^3)	3100
Poisson's ratio	0
Shear Modulus (MPa)	250
Thickness (μm)	55

Chapter 5

Results

This chapter details the results of a number of different CFD and FSI simulations of parafoils performed using the methods outlined above. Two-dimensional CFD simulations of the parafoil cross-section with varying leading edge configurations are described first, followed by the results produced by three-dimensional FSI simulations of the parafoil cell using two different configurations of structural boundary conditions. Lastly, the influence of the choice of turbulence model on the predicted flow field behaviour is compared for the flow around the inflated shape predicted by the previous three-dimensional FSI simulation.

5.1 Two-dimensional simulations of flow around ram-air parachute

One of the main points to consider when discussing the design and performance of ram-air parachutes is the difference in performance between aerofoils and parafoils, and the reasons behind this. Computational studies of the flow around 2D parafoil and aerofoil cross-sections, such as those of Mittal et al.[50], Balaji et al.[51], and Mohammadi and Johari[52], provide some insight into this key question. However, the results of these studies are seen to be highly dependent on the design and shape of the cross-section being studied. Therefore, a simple 2D study has been performed, to simulate the flow field around the undeformed parafoil cross-section used throughout this study, which is adjusted from cross-section data provided by Ware & Hassell[33],

Section 3.2.2, to observe how this particular design behaves.

The objectives for this study are twofold. Firstly, a 2D study of the flow around the unstressed cross-section used for the manufacture of the parafoil provides a basis from which the 3D flow results can be evaluated. The similarities and differences between the 2D flow field, and the flow that occurs around the predicted 3D shape, will highlight the magnitude of any 3D effects and will clearly show how the change in geometry produced by inflation influences the aerodynamic performance of the canopy. Secondly, as already stated, a comparison of this nature offers an opportunity to investigate the influence of the ram-air inlet on the performance of the parafoil, both in terms of flow field observations and the impact that the observed flow features have on the resulting lift and drag values. Thus, the difference between the aerodynamics of the open inlet parafoil and an equivalent aerofoil with a closed, rounded leading edge will be demonstrated.

The method employed by previous studies has been to take a parafoil cross-section, which is based on an aerofoil cross-section, and compare the flow around the parafoil and base aerofoil. In this study, however, the selected parafoil cross-section is not known to be based on any standard aerofoil design, as discussed in Section 3.2.2. Therefore, to allow for a comparison to be made, a rounded leading edge, adapted from that of a Clark-Y aerofoil, has been attached to the parafoil geometry, to create a comparison aerofoil shape. This can be seen in Figure 5.1, and the coordinates for all three 2D cases are found in Appendix A.

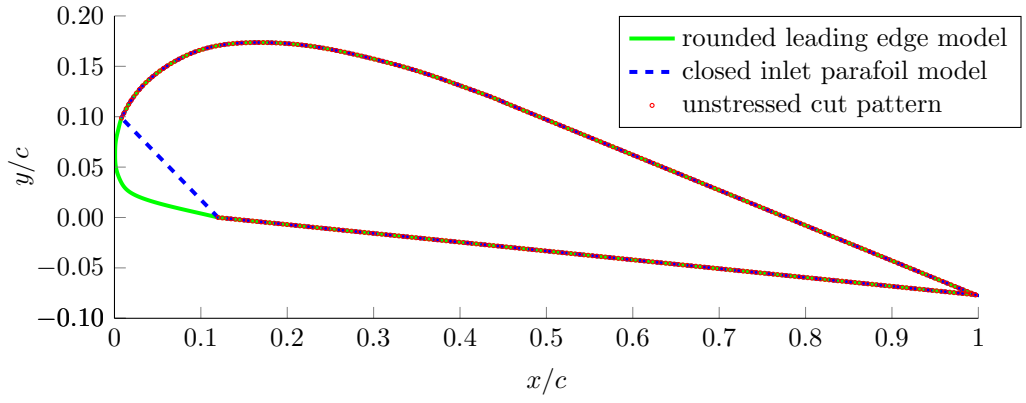


Figure 5.1: Two-dimensional chordwise cross-sections showing geometry of three different leading edge configurations.

As an extension to previous studies of this nature an intermediate geometry is also considered, in the form of the parafoil cross-section with a closed leading edge. The motive behind this is to differentiate between the effects of the parafoil leading edge being open, and those caused solely by the bluff nose profile that results from the creation of the ram-air inlet. The resulting geometry is exactly the same as the open parafoil geometry, but with a solid, blunt nose, cut at an angle of $\theta_{\text{cut}} = 36^\circ$ to the horizontal lower surface. Air is therefore prevented by the solid nose from entering the interior cavity of the parafoil.

All three cases were simulated at an angle of attack of 5°, in line with the 3D cases. The angle of attack for all three shapes is measured against the flat lower surface, rather than the chord line, as is conventional for parafoil analyses. A single angle of attack was chosen in order to allow for an in-depth analysis of the observed flow properties and resultant lift and drag coefficients, in an environment similar to that for the steady glide phase of a ram-air canopy. Further work would of course be required if more detailed characteristics, such as lift and drag behaviour up to stall, were desired. All cases were simulated for a Reynolds number of 2.1×10^6, which corresponds to a velocity of 33ms^{-1} for a reference length based on the chord of 1m. All simulations are performed purely within the fluid domain using steady RANS with the Realizable k-ε turbulence model. All structures are treated as rigid.

5.1.1 Flow field results

Closed, rounded leading edge

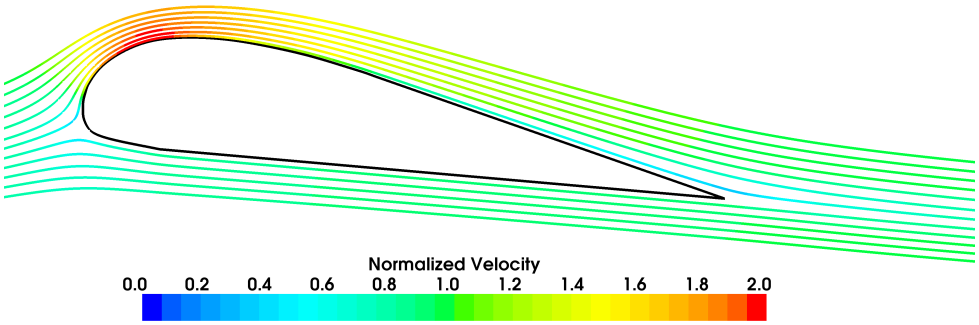


Figure 5.2: Streamlines for the flow around rounded leading edge model. Contours of velocity normalized against freestream velocity magnitude.

The streamlines for the flow around the rounded leading edge geometry are shown in Figure 5.2. The streamlines show a smooth movement of the airflow around the leading edge, with the streamlines closely following the contours of the aerofoil. High flow velocity gradients, indicated by a higher density of streamlines, are present above the upper surface near the point of maximum thickness. This can be seen in greater detail in the velocity contours of the flow field plotted in Figure 5.3. As expected for an aerofoil at a positive angle of attack, an area of accelerated flow is observed on the upper, or suction side. The presence of this "suction peak" is confirmed by the area of peak low pressure on the upper surface observed in Figure 5.4. In turn, the flow around the lower, or pressure side, shows little deviation from the freestream velocity, but does show a slight increase above freestream pressure. This pressure differential between the upper and lower surfaces generates lift.

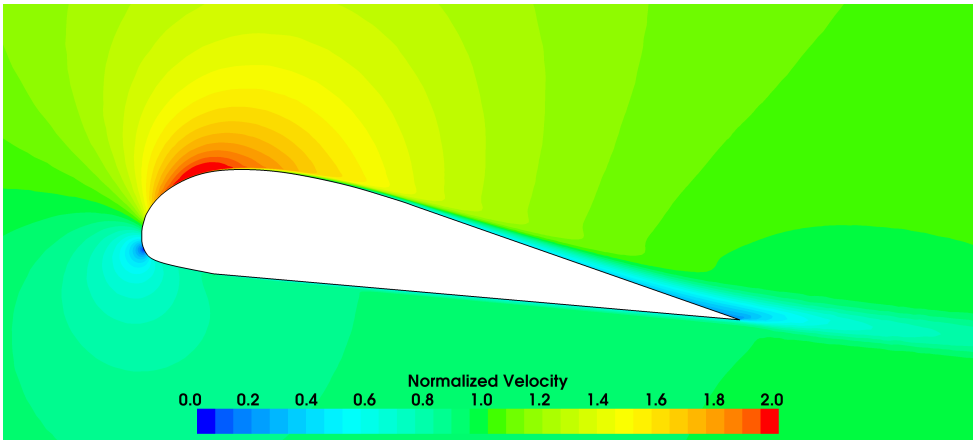


Figure 5.3: Normalized velocity contours for flow around rounded leading edge model. Contours of velocity normalized against freestream velocity magnitude.

The area of very low velocities seen at the tip of the leading edge in Figure 5.3 corresponds to the stagnation point, confirmed by a value of $C_{\mathrm{p}} = 1$ observed in the pressure contours plotted in Figure 5.4, and the surface pressure plot in Figure 5.14.

A further feature of note in Figure 5.3 is the region of low velocity directly adjacent to the upper surface, which appears to extend downstream from the edge of the suction peak, down to the trailing edge. This area of high velocity gradient is the boundary layer, which is seen to noticeably grow in thickness as it moves towards the trailing edge, and its location is also highlighted by the vorticity contours in Figure 5.5. A boundary layer is also observed on the lower side of the section which remains much

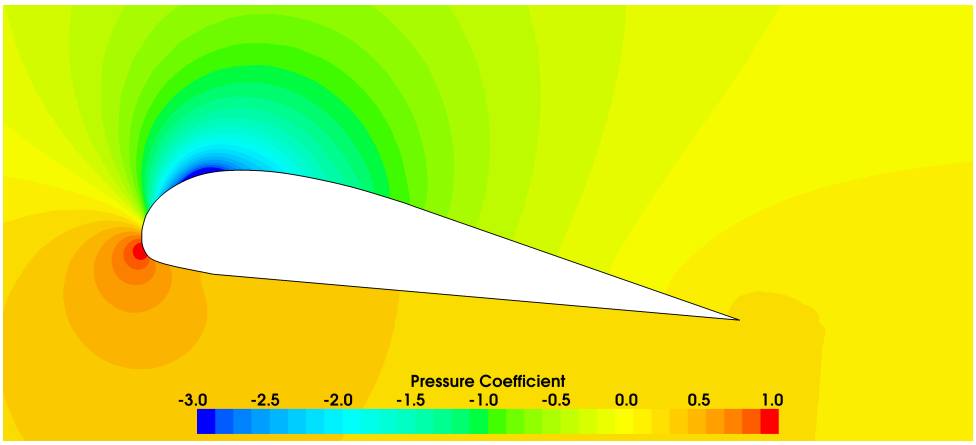


Figure 5.4: C_p contours for flow around rounded leading edge model.

thinner than the one on the upper surface, due to the lack of adverse pressure gradient. Both boundary layers stay attached at all times, confirmed by the wall shear stress values recorded on the upper and lower surface plotted in Figure 5.15. Overall, the results show a very smooth, attached flow for the rounded leading edge case.

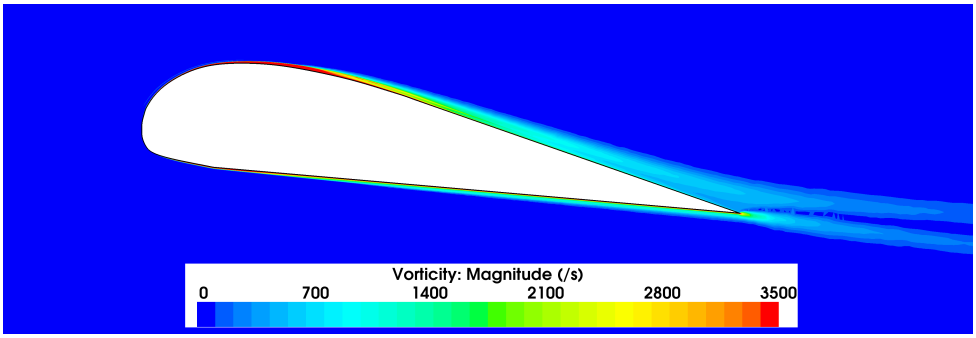


Figure 5.5: Vorticity contours for flow around rounded leading edge model.

Parafoil section with closed leading edge inlet

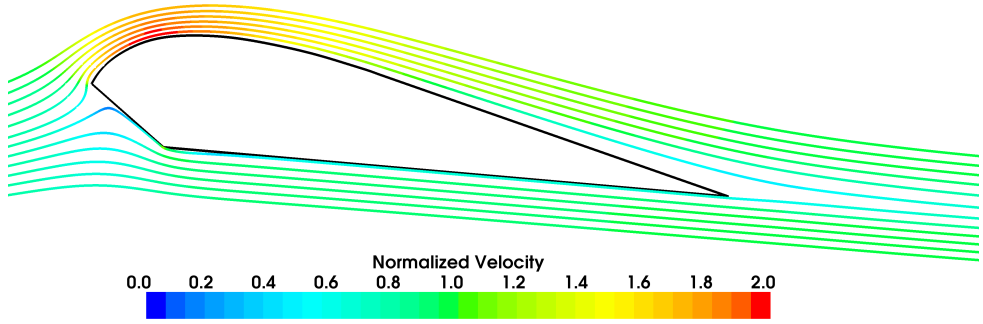


Figure 5.6: Streamlines for flow around parafoil cross-section with closed inlet. Contours of velocity normalized against freestream velocity magnitude.

A plot of the streamlines for the flow around the parafoil cross-section with a closed leading edge inlet is given in Figure 5.6. As a whole, the streamlines appear to show a number of similarities to those observed for the rounded leading edge case, with any differences appearing to be confined to the diversion of the flow around the face of the closed inlet. The velocity and pressure contours, found in Figures 5.7 and 5.8, show a similar region of high velocity and low pressure present over the upper surface, and the near freestream properties of the flow on the pressure side are maintained, confirming that lift is generated by the section as expected.

Notable differences in the flow behaviour are brought about as a consequence of the movement of the flow around the truncated, more angular shape of the nose. The blunt face of the closed inlet appears to cause stronger deviations of the flow than the smooth, rounded leading edge of the previous case. A consequence of this change can be seen in the position of the stagnation point, which is now located further downstream, taking up a position towards the upper end of the inlet face. The pressure contours in Figure 5.8 clearly show the associated area of high pressure, which is present over a much wider section of the nose than for the rounded leading edge model. The figure also shows that in addition to the area of low pressure at the suction peak there is also a small region of low pressure observable at the upper leading edge, on the corner of the inlet and the upper surface. This is also just visible in Figure 5.7 as a region of low velocity.

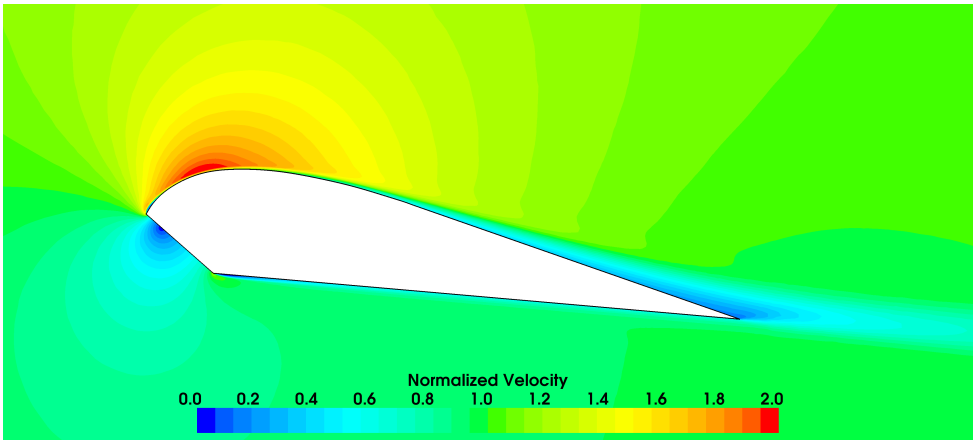


Figure 5.7: Normalized velocity contours for flow around parafoil cross-section with closed inlet. Contours of velocity normalized against freestream velocity magnitude.

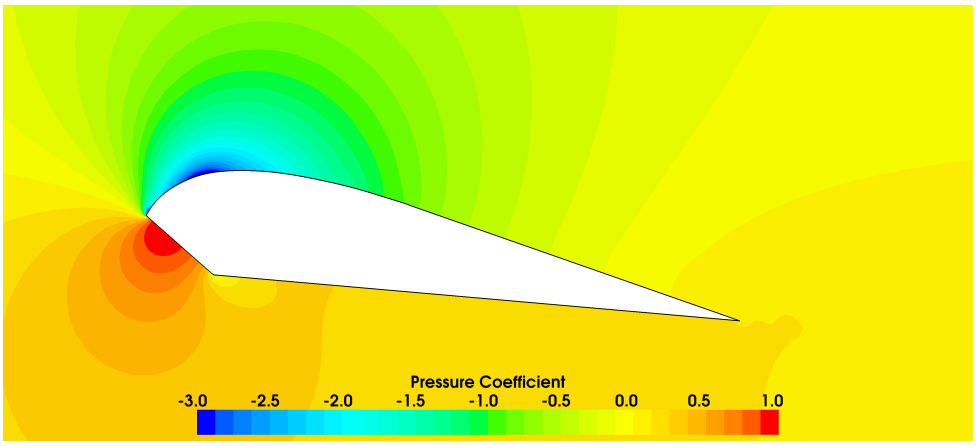


Figure 5.8: C_{p} contours for flow around parafoil cross-section with closed inlet.

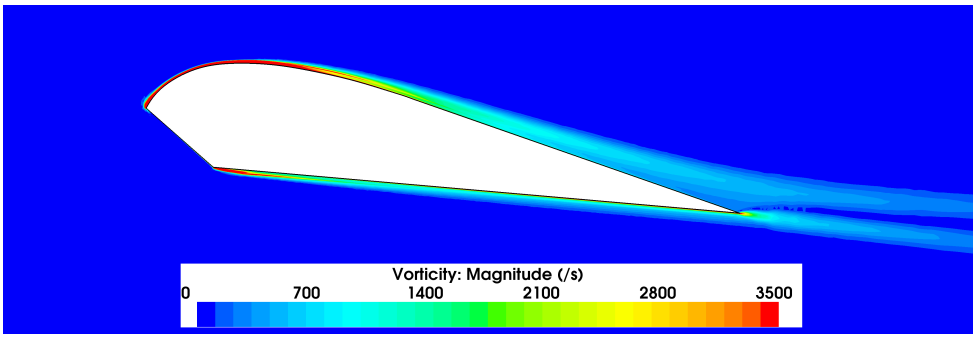


Figure 5.9: Vorticity contours for flow around parafoil cross-section with closed inlet.

Inspection of the vorticity contours in Figure 5.9 provides a clearer view of how the flow develops in this region around the upper leading edge. A boundary layer develops on the upper surface which increases in thickness and reduces in velocity as it reaches the trailing edge, as observed for the rounded nose geometry. Unlike the case of the rounded leading edge however, the suction side boundary layer begins upstream of the suction peak, with a layer of high vorticity starting at the upper leading edge itself. The initial thickness of the upper boundary layer is clearly greater than that of the rounded leading edge case, and is seen to grow in size from the point of maximum section thickness onwards.

The greater initial thickness of the boundary layer appears to originate from flow which has been diverted past the inlet face, intersecting the flow over the upper surface, thereby changing direction to turn round the sharp leading edge. This observation is further supported by the local maximum in the boundary layer thickness immediately behind the leading edge, which then reduces. The flow over the upper leading edge also coincides with the previously described local region of low pressure and reduced velocity.

A similar set of flow features is seen at the junction of the inlet face with the lower surface. The sharp edge at the lip of the lower surface results in an area of high vorticity as the flow is redirected by the inlet face and flows past the corner to rejoin the rest of the flow. The resulting boundary layer is thicker than in the rounded leading edge case, and the flow around the sharp lower leading edge results in a local pressure disturbance observed in Figure 5.8.

The localised disturbances in pressure and velocity created as the flow moves round the sharp edges at the two ends of the inlet face, as well as the high levels of vorticity, suggest that these two regions are likely points of separation. The wall shear stress plot for the upper surface in Figure 5.15(a), which plots the shear stress values at the wall for the upper surface and the inlet face, shows a negative value of wall shear stress corresponding to the location of the upper leading edge. This confirms a small separation bubble on the lip of the upper surface, beyond which the flow reattaches. Figure 5.15(b) confirms a similar occurrence at the lower leading edge where a region of negative shear stress follows an initial sharp negative peak. This area of separation persists for longer before reattachment than the one on the upper surface.

The truncation of the leading edge also has an effect on the pressure values seen at the surface of the parafoil section. Figure 5.14 provides a plot of C_p at the lower surface, and at the inlet and upper surface. A local negative pressure peak is located at the leading edge corresponding to the acceleration of the flow around the corner of the inlet and upper surface, and the resulting flow separation. The pressure recovers to levels similar to that for the rounded leading edge case once reattachment has occurred, but beyond this point the suction pressures are consistently reduced for the majority of the length of the upper surface. Accordingly, the upper surface suction peak is slightly reduced from the value reached in the case of the rounded leading edge.

Surface pressures seen on the inlet exhibit a different pattern to those observed on the rounded nose of the first geometry. In the case of the latter, the pressure peak resulting from the stagnation point is focused at a narrow point on the rounded leading edge, whereas the peak is spread out over a larger proportion of the inlet face on the closed parafoil section. This pressure reduces at the intersection with the lower surface, and then reduces further, due to an area of suction present in the separation bubble just downstream of the lower leading edge. This area of suction gradually recovers, but the pressure on the lower surface does not reach similar levels to the rounded nose case until past the point of reattachment.

Furthermore, the positive pressure seen on the lower surface is slightly lower overall compared to the section with a rounded leading edge. Therefore, due to the combination of reduced peak suction pressures, reduced positive pressure levels on the lower surface, and an area of suction pressure on the lower surface due to separation, it is expected that this geometry will show a reduction in lift over the rounded leading edge geometry

Parafoil section with open leading edge inlet

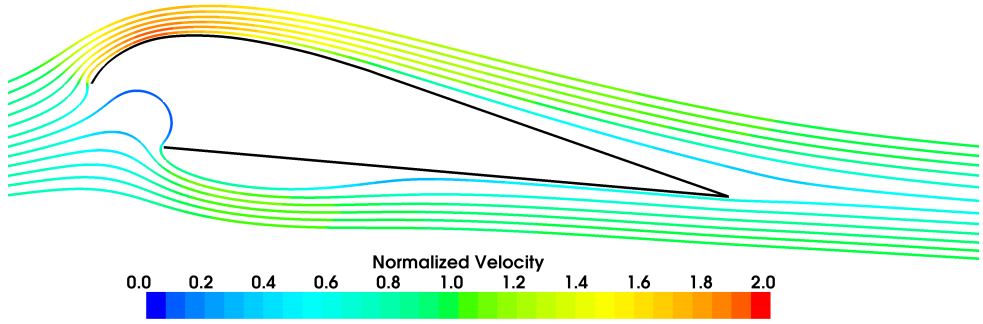


Figure 5.10: Streamlines for flow around parafoil cross-section with open inlet. Contours of velocity normalized against freestream velocity magnitude.

Unlike the preceding two models, a parafoil cross-section has an open leading edge which is required for the canopy to inflate through the build up of ram-air pressure inside the cell. This creates a very sharp leading edge on the upper and lower surfaces and the impact this opening has on the nature of the flow around the cross-section is immediately obvious from the streamlines plotted in Figure 5.10.

The flow around the inlet and lower surface is clearly different to the closed inlet case, but the streamlines show a very similar flow pattern over the upper surface with flow still being diverted up and around the upper leading edge, despite there being no solid face to force the air to take this path. The sharp leading edge again has the effect of producing localised disturbances in the velocity and pressure fields, seen in Figures 5.11 and 5.12. As in the closed inlet case, the vorticity is high in the region around the upper leading edge, and a boundary layer is seen to begin at this point, Figure 5.13. The wall shear stress plot for the upper surface in Figure 5.15(a) confirms that the pressure and velocity disturbances are a result of the flow separating and reattaching, in a similar manner to the closed inlet case.

The boundary layer on the upper surface increases in thickness towards the trailing edge, seen in Figure 5.13, and when compared with Figure 5.9 the increase in boundary layer thickness is greater in the open inlet case than in the two previous cases. The wall shear stress values in this area show that, apart from the separation bubble at the leading edge, the flow stays attached all the way to the trailing edge. The comparison with the other cases in Figure 5.15(a) shows that the shear stress values in the trailing

edge region are lower than the previous two cases, suggesting lower velocity gradients in this area.

Examining the C_p values recorded on the upper surface in Figure 5.14, it is seen that a localised peak of low pressure caused by the upper leading edge separation is recreated for the open inlet case, but the magnitude of this peak is reduced over the closed inlet value. The maximum value of low pressure at the suction peak is also lower than for the closed leading edge case, and therefore is reduced even further when compared to the rounded leading edge case.

As previously stated, the flow around the inlet region and the lower surface is obviously different to the previous two cases. The open inlet allows flow to enter the mouth of the ram-air cell, but the streamlines do not penetrate very far into the interior of the cell. The velocity contours in Figure 5.11 show that the air inside the majority of the interior space is stationary, and this capture of the air means that it equilibrates at the stagnation pressure of the flow. This is displayed very clearly in Figure 5.12. This high pressure region provides structural rigidity in the case of a three-dimensional canopy, and supports the aerodynamic shape of the inflated parafoil.

In addition to contributing to structural rigidity, the stagnation pressure of the air interior to the parafoil also appears to effectively block off the ram-air inlet once inflation is achieved. The velocity contours near the mouth of the parafoil section show that the flow is not completely stationary until just past the lower leading edge, and so the inlet is not blocked off until that point. From the streamline plot it can be seen that this means a portion of the flow is allowed to enter the mouth of the parafoil, and must then change direction to flow back out of the cell, since the high pressure within the cell enforces this condition.

The consequence of these sharp changes in direction, coupled with the influence of the sharp lip of the lower leading edge, is a large deviation in the flow away from the lower surface, seen in Figure 5.10. This pattern of flow produces a large bubble of recirculating air which extends downstream of the lower leading edge to around a third of the way along the lower surface. The area affected by the deviation of the flow is very clear in Figures 5.11 and 5.13, the latter of which also reveals the large region of high vorticity flow that extends downstream of the disturbance.

The wall shear stress data for the lower surface of the open inlet parafoil section in Figure 5.15(b) confirms the flow separation at the lower leading edge, and shows the

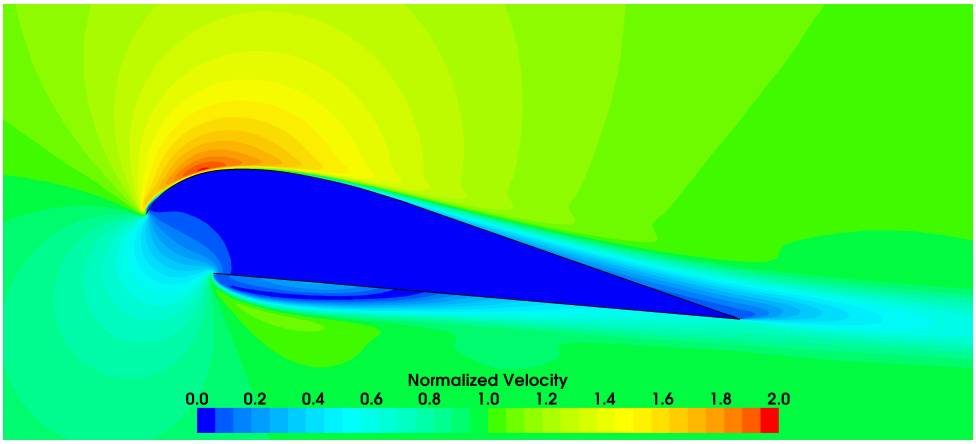


Figure 5.11: Normalized velocity contours for flow around parafoil cross-section with open inlet. Contours of velocity normalized against freestream velocity magnitude.

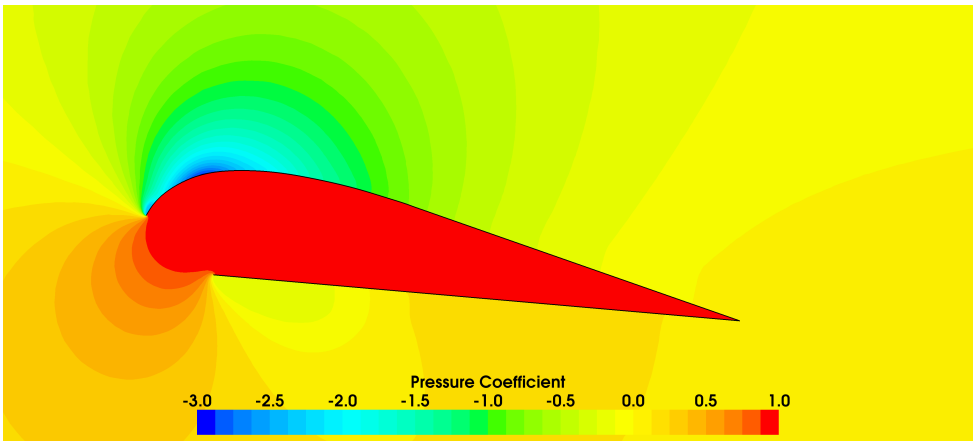


Figure 5.12: C_{p} contours for flow around parafoil cross-section with open inlet.

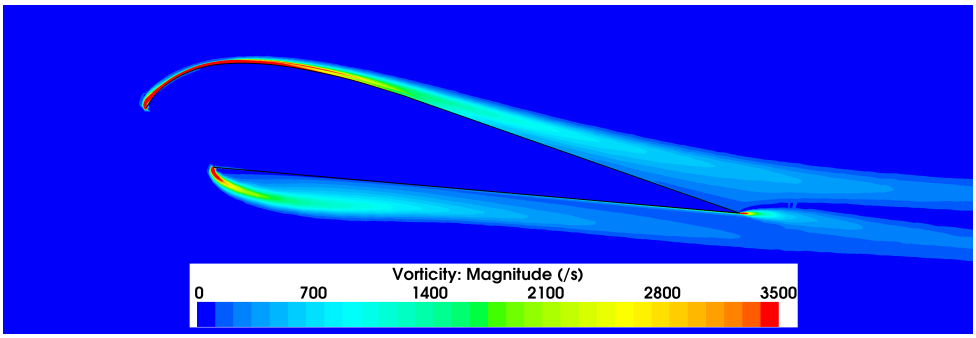


Figure 5.13: Vorticity contours for flow around parafoil cross-section with open inlet.

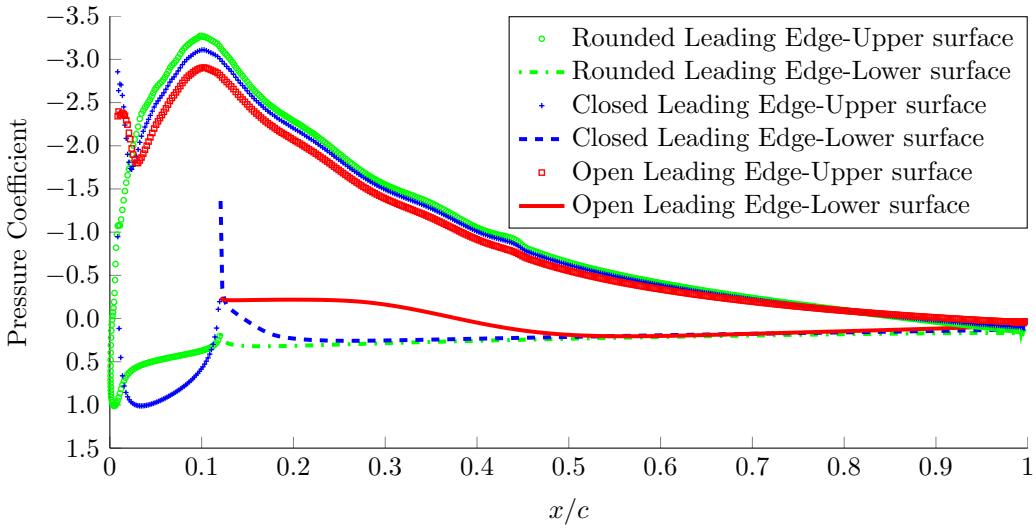
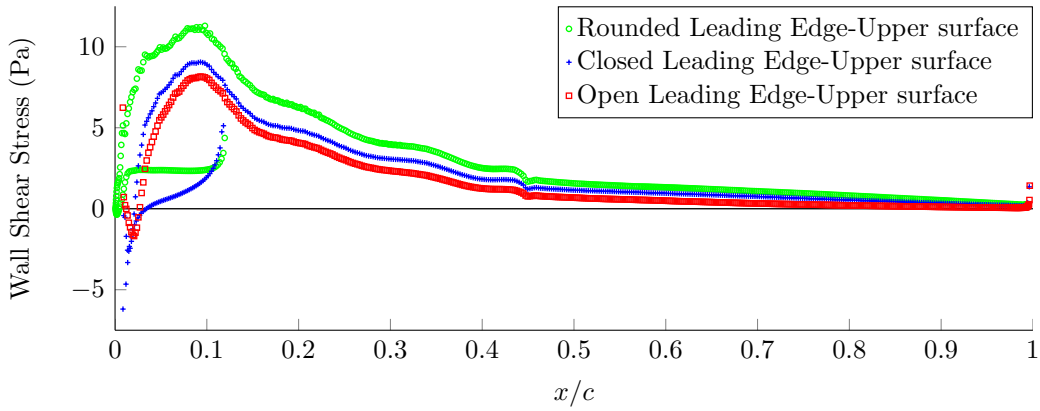


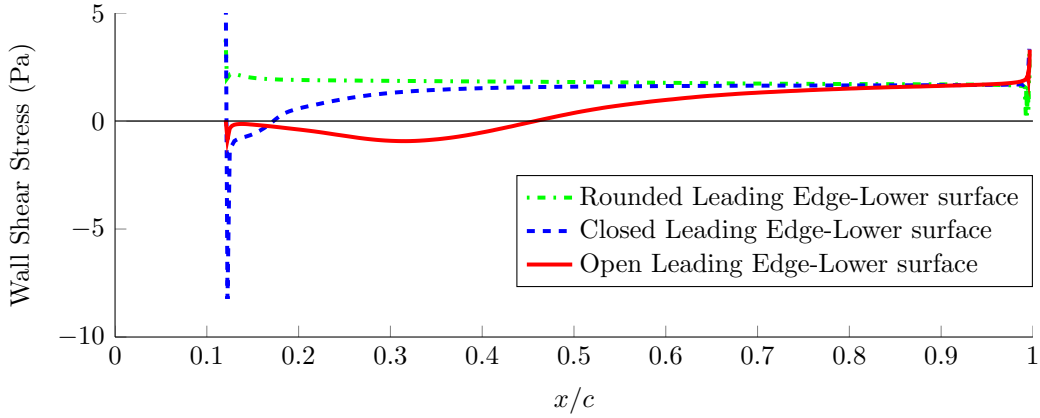
Figure 5.14: Comparison of C_p recorded at the external parafoil surfaces, for three leading edge configurations.

separation bubble to cover just under half the length of the lower surface before the flow reattaches. The wall shear stress levels do not return to expected values until close to the trailing edge. The pressure in this area is reduced to below the freestream pressure and so the separation bubble actually creates a region of suction on the lower surface, which in Figure 5.14 is shown to persist for just under a third of the distance along the lower surface. The pressure does not recover until downstream of the reattachment point. The surface pressure along the remainder of the lower surface is of a similar magnitude to the two previous cases, but shows a slight reduction over the final 15% of the surface.

As in the closed inlet case, the region of suction pressure on the lower surface produced by the flow separation applies an inverse lift force to that area. For the open inlet parafoil geometry, compared to the closed inlet section, this region of suction has a greater magnitude of lower pressure, and is spread over a larger part of the lower surface. Added to the aforementioned reduction in magnitude of the suction peak on the upper surface, it would be anticipated that the open parafoil section would have an even lower lift coefficient than that of the closed leading edge case. The lift and drag coefficients for all three 2D sections are discussed in the following section.



(a) Upper surface

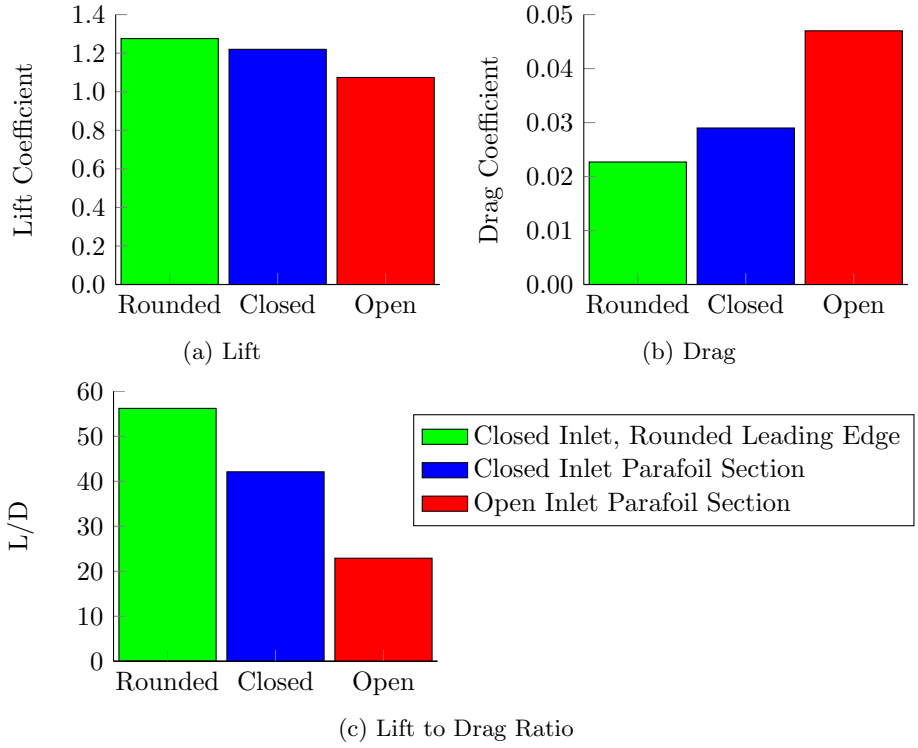


(b) Lower surface

Figure 5.15: Comparison of wall shear stress recorded at the upper and lower external parafoil surfaces, for three leading edge configurations.

5.1.2 Lift and drag comparisons

In the above section, an analysis of various properties of the flow field has enabled the identification of a number of interesting characteristics that arise from the differences in geometry between aerofoils and parafoils. In order to quantify the effect that the introduction of these flow features has on the aerodynamic performance a comparison of the lift and drag values is required.



(a) Lift

(b) Drag

(c) Lift to Drag Ratio

Figure 5.16: Comparison of aerodynamic force coefficient values for three leading edge configurations.

A summary of the lift, drag, and lift to drag ratios for the three two-dimensional cases is found in Figure 5.16. Figure (a) shows a slight reduction in lift as a consequence of the nose being truncated, and a further, more significant reduction in the lift as a consequence of the open inlet. In terms of the drag coefficient, (b) shows an increase in drag for the blunt leading edge when compared to the baseline aerofoil, and a further significant increase in drag is seen in the open inlet case. The addition of an open inlet results in a 2.07 times increase in drag and loss of lift of 16% over

the comparison aerofoil section. Both of these values are extremely similar to the observations from previous computational studies discussed above, which observed a loss of lift of 15%–20% and an increase of drag of around double. The combined influence of the loss in lift and increase in drag produces a considerable drop in the lift to drag ratio, seen in Figure (c), which is a key parameter in parafoil performance. In order to understand the mechanisms behind these losses in aerodynamic performance the force coefficient data can be broken down into the components for the different parts of the geometry.

Lift

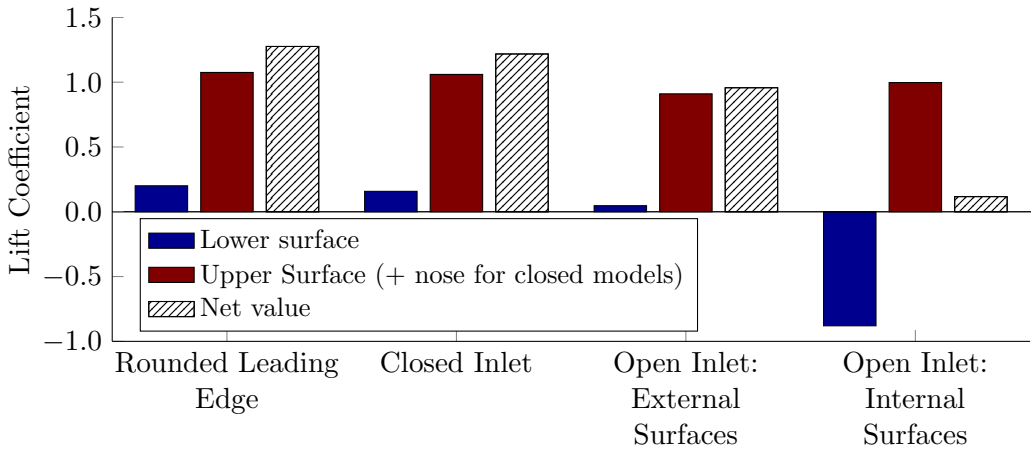


Figure 5.17: Lift coefficient values, for three leading edge configurations, separated in terms of structural components.

A break-down of the lift coefficient generated at the upper and lower surfaces of the three different cross-sections is provided in Figure 5.17. The geometrical dimension of the section assigned to the lower surface component is identical for all three models. The upper surface components for the rounded and closed inlet cases include the rounded nose and inlet face of the respective sections.

As would be expected, the results show that the majority of the lift occurs at the upper surface for all three cases, due to the low pressures created by the acceleration of the flow over the contours of the section. The lift generated at the upper surface of the closed inlet model is very close to that of the leading edge model. Therefore the reduction in total lift generated by the closed inlet model is shown to be a consequence

of the loss in lift at the lower surface. This confirms that the reduction in lift is a direct result of separation on the lower leading edge. The lift reduction at the lower surface is even greater for the open inlet model, again due to the separation at the lower leading edge. The large area of negative pressure created by the separated region reduces the lift on the lower surface by 75% over the equivalent aerofoil. This is compared to a reduction of 21% for the closed inlet case.

The open inlet model also shows a reduction in lift generated at the upper surface, but this is likely to be less severe than it first appears, because the upper surface lift value includes the contribution from the nose section for the first two models, which is obviously non-existent in the open inlet case. Separating this inlet contribution out from the upper surface, as in Table 5.1, verifies that the lift acting on the upper surface of the closed parafoil section is just under 7% higher than for the open parafoil. This reflects the small, but noticeable difference in the magnitude of the negative pressures generated over the length of the upper surface observed in Figure 5.14.

Table 5.1: Lift coefficient for upper surface, for open and closed parafoil model. Lift force contribution from inlet face removed for closed inlet parafoil case.

Model	Lift Coefficient	Difference, %
Closed inlet: upper surface	9.793×10^{-1}	—
Open inlet: upper surface (external)	9.130×10^{-1}	-6.8

A further point of interest for the open inlet case is the fact that aerodynamic forces are applied to the interior surfaces, as well as the exterior ones. This is a departure from the typical consideration of aerodynamic forces over an aerofoil, and has not been considered by the previous parafoil studies found in the literature. The forces are generated by the high pressure air inside the parafoil, which applies a hydrostatic force on the interior walls. As Figure 5.17 shows, the lower interior surface component acts against the positive lift direction, but the larger area of the upper surface ensures that this negative force is counteracted. The net contribution from the interior surfaces to the lift on the structure is therefore positive.

From these observations it is clear that the key mechanism for the loss of lift in the case of the inlet being open is due to the low pressure in the separated region on the lower surface. The effect of this on the total lift far outweighs the reduction of upper surface lift.

Drag

In comparison to the lift coefficient values, it is obvious from Figure 5.16 that the difference in the drag coefficient between the rounded and open leading edge cases is more severe. The contributions in terms of pressure drag and friction drag for the three models can be seen in Figure 5.18. This shows that the addition of a blunt leading edge and an open inlet produces a large increase in the proportion of pressure drag, and a reduction in the proportion of friction drag. This suggests that when the leading edge of the model is truncated, and then opened, the shape is less streamlined, and behaves more like a bluff body, which seems reasonable from the flow field comparisons above.

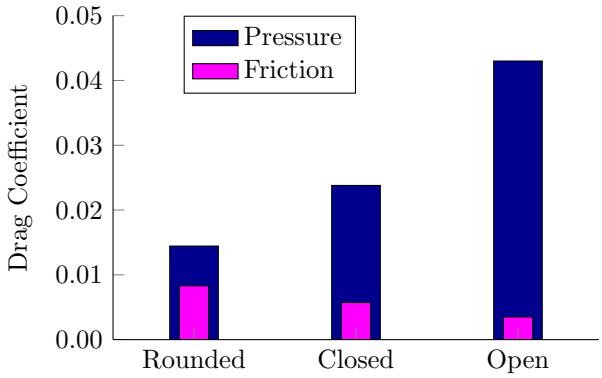


Figure 5.18: Comparison of pressure and friction drag contributions, for three leading edge configurations.

When the drag values are broken down in terms of the different parts of the geometry, displayed in Figure 5.19, the drag produced by the rounded nose model comes predominantly from the lower surface. This is due to the positive pressure seen on the lower surface, which contributes to the drag due to the angle that the lower surface presents to the freestream. The lower surface drag in the closed inlet case is slightly lower, as it is reduced by the negative pressure region within the separation bubble. A larger component of drag is evident on the upper surface and inlet face of the closed model. Given that the flow over the upper surface is very similar for both designs, the increase in drag is likely to be due to the large area of high positive pressure present at the stagnation point on the inlet face.

For the open inlet case, Figure 5.19 shows the drag force on the exterior part of the

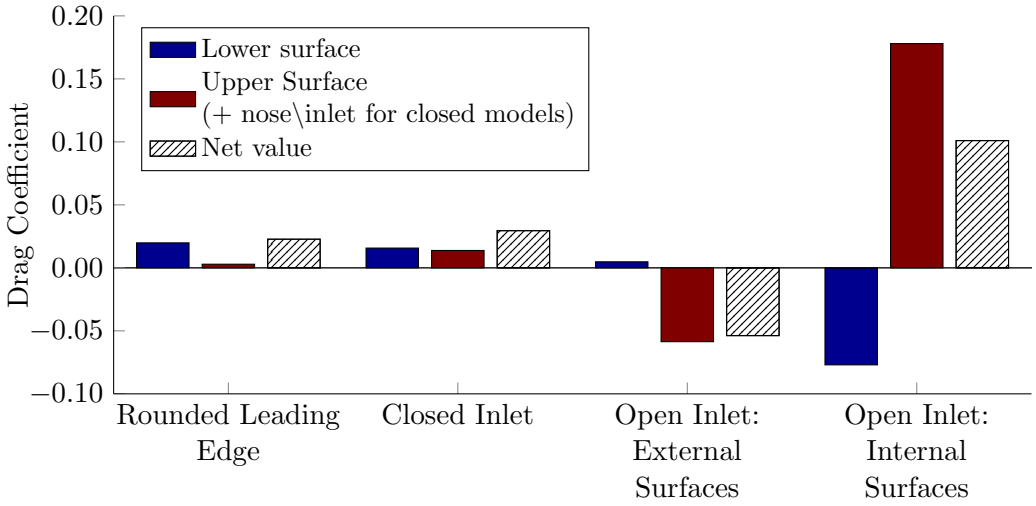


Figure 5.19: Drag coefficient values, for three leading edge configurations, separated in terms of structural components.

Table 5.2: Comparison of drag coefficient for open and closed parafoil model. Contribution from inlet face removed from upper surface value for closed inlet parafoil case.

Model: component	Drag Coefficient
Closed inlet: upper surface	−0.0589
Open inlet: upper surface (external)	−0.0586
Closed inlet: inlet face	0.073
Open inlet: All interior faces	0.101

lower surface to be very low. This is a reduction over the closed inlet case, created by the larger suction region caused by separation of the flow at the lower leading edge. The interior side of the lower surface produces a negative drag contribution, resulting from the pressure forces from the stagnated air inside the cell. The surface is orientated such that a component of the force opposes the positive drag direction. This balance between the presence of interior and exterior pressure forces on the open parafoil section, and the interplay with the angle of attack of the lower surface is visualised in a plot of C_p surface vectors in Figure 5.20.

Table 5.2 provides the drag forces produced at the upper surface of the closed inlet case, with the contribution from the inlet face removed. The drag coefficient on the

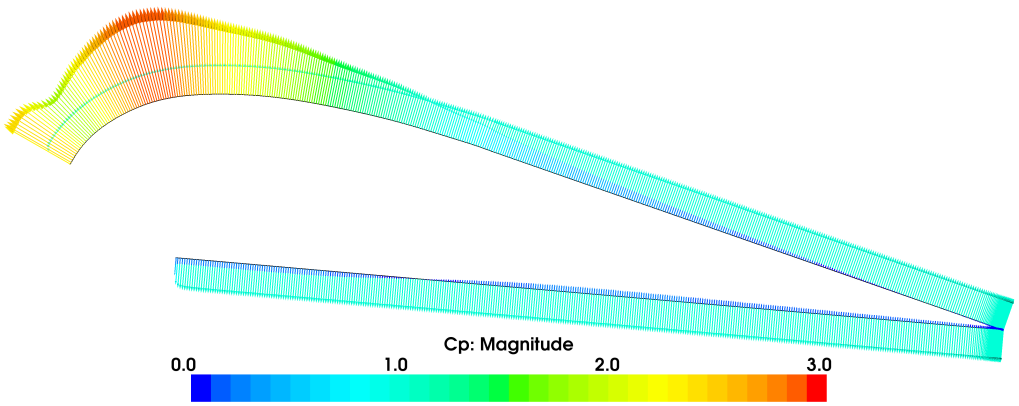


Figure 5.20: Vector plot of pressure coefficient distribution on open ram-air parafoil section. Pressures plotted separately for interior and exterior surfaces.

inlet face is provided as a separate value. The table shows that the exterior upper surface in both cases produce almost identical values of negative drag, driven by the large area of low pressure at the suction peak, most of which can be seen in Figure 5.20 to be acting in a direction that opposes the positive drag direction. The values in Figure 5.19 and Table 5.2 confirm the dominant source of drag on the open parafoil section to be internal stagnation pressure acting on the interior upper surface. Interestingly, Table 5.2 shows that this exceeds the drag resulting from the stagnation pressure on the face of the inlet in the closed parafoil model.

5.1.3 Comparison to other two-dimensional parafoil flow studies

A comparison of the aerodynamic performance of the parafoil design selected by Ware & Hassell with the results from 2D studies in the literature, provided in Table 5.3, reveals some interesting conclusions. The results of this study show marked improvement over the design studied by Mittal et al., based on a Clark-Y aerofoil. The Clark-Y design suffers large amounts of separation from the upper surface, thereby greatly increasing drag and reducing the lift to drag ratio. The parafoil based on the LS(1)0417 section studied by Balaji et al. displayed poor performance for an inlet cut angle of $\theta_{\text{cut}} = 45°$, with strong separation from the upper surface. A change in the cut angle to $\theta_{\text{cut}} = 30°$ offered a vast improvement in the quality of the flow over the upper surface, resulting in an improved lift to drag ratio. The value of θ_{cut} for the

Table 5.3: Comparison of simulation lift and drag results for open-inlet parafoil section to those of other published 2D flow studies. Re=1–6 million in all cases.

Study name and baseline aerofoil (if known)	α (°)	θ_{cut} (°)	C_{L}	C_{D}	$C_{\text{L}}/C_{\text{D}}$
Mittal et al. - Clark-Y	7.5	45	0.99	0.29	3.41
Mittal et al. - Clark-Y	7.5	30	1.01	0.21	4.81
Balaji et al. - LS(1)0417	7.5	45	0.894	0.245	3.65
Balaji et al. - LS(1)0417	7.5	30	0.88	0.034	25.88
Balaji et al. - LS(1)0417	5	30	0.746	0.022	32.75
Mohammadi & Johari	5	–	0.5742	0.04	14.36
Current study	5	36	1.074	0.047	22.9

Ware & Hassell cross-section is 36°, which supports the conclusion that too severe an inlet cut has a negative effect on the upper surface flow.

The parafoil section studied by Mohammadi & Johari, for which the angle of the leading edge cut is not given, also resulted in attached flow over the upper surface, reflected in the values for C_{L} and C_{D}. The performance of the parafoil cross-section analysed within this study compares well to the values found in other studies for the aforementioned cases where the upper surface flow remains attached. The $C_{\text{L}}/C_{\text{D}}$ ratio of the parafoil based on the LS(1)0417 aerofoil of 32.75 is a slight improvement over the result of 23 recorded for the section studied here, due to a lower drag coefficient of 0.022. The cross-section simulated by Mohammadi & Johari however has a lower lift coefficient of ∼0.6, and a similar drag coefficient of 0.04. The glide ratio is therefore noticeably lower than the other comparable sections. All three studies confirm the large reduction in aerodynamic performance in a parafoil over an equivalent aerofoil shape.

5.1.4 Summary and discussion

The results provided by the two-dimensional flow simulations have provided considerable insight into the typical flow phenomena associated with a parafoil section and how this relates to an equivalent aerofoil style section. The flow around the rounded leading edge model was seen to be attached and well-behaved. High pressure was observed beneath the section, and low pressure above the section, as expected. Trun-

cating the leading edge was shown by the closed inlet model to reduce the amount of lift generated by the section, due in some part to a small reduction in suction pressure over the upper surface, and in the most part to the reduction in positive pressure on the lower surface because of flow separation at the sharp lower edge of the inlet. A small increase in drag was also observed, primarily due to the larger stagnation point on the face of the closed inlet. Importantly, the truncation of the leading edge is seen to have less impact on the aerodynamics of the section than might be expected.

A more comprehensive change in the flow field has been observed for the open inlet parafoil model. The presence of the ram-air opening significantly alters the pattern of flow in that region, with flow now able to enter the internal space within the parafoil. High pressure air was seen to fill the vast majority of the interior region, accordant with the concept that ram-air pressure acts to initially support the inflation of the canopy, and then provide rigidity to the structure. Despite the internal volume of stationary, high pressure air, a small proportion of the freestream is still able to enter through the inlet. The air entering the inlet is prevented from travelling further inside the cell by the high pressure and exits the inlet near to the lower leading edge. This generation of flow opposing the freestream, combined with the sharp lower leading edge, results in a large separation bubble under the forward part of the lower surface. The lower surface separation was shown to be greatly exaggerated over the closed inlet model, creating a large area of negative pressure on the lower surface. Separation was also seen at the upper leading edge, to a degree very similar to that of the closed inlet case.

Whilst the observed differences in the flow resulted in reasonably small differences in lift and drag between the rounded and closed inlet cases, the differences observed for the open inlet case contributed to much larger changes in the resultant forces. The large lower surface separation is shown to be responsible for a large reduction in lift due to the associated downward suction pressure. This is coupled with a small reduction in the lift generated over the upper surface, causing a large reduction in the pressure gradient across the upper and lower surfaces. The increase in drag was observed to be much larger than the reduction in lift, related to a large increase in pressure drag. Consideration of the internal drag forces, generated by the hydrostatic action of the interior pressure, has revealed the drag force created by the internal pressure on the upper surface to be the largest component of overall drag, and is seen to be the key factor behind the large drag increase.

Comparison with other studies revealed that the parafoil section used here performs well relative to other designs, and it successfully avoids the severe separation observed for the Clark-Y case of Mittal et al. The demonstrated increase in drag for the open inlet parafoil over an aerofoil equivalent is in line with previous studies found in the literature. These studies conclude that this difference is caused by the blunt shape of the truncated leading edge. It has also been suggested that the effective blocking off of the inlet, caused by the high interior pressure of the stagnated air in the interior space, results in the flow behaving as though the inlet is closed. The results of this study have clearly shown that this is not the case here, and that the flow is able to enter the mouth of the inlet before it exits again, strongly influencing the flow over the lower surface as a consequence. This is further supported by the fact that, interestingly, the results of the closed inlet case show that a blunt shape alone has a surprisingly small impact on aerodynamic performance, when compared with the aerofoil-style section. As a consequence, the assumption that the bluff shape of the inlet region is responsible for the increase in drag seen in a parafoil section over an aerofoil section is seen to be questionable. Results shown here directly contradict this assertion and indicate the situation to be more nuanced than that; the interaction of the aerodynamic forces with the parafoil are revealed to be more complicated than in the case of a structure with purely external flow. Analysis of the drag of the parafoil section demonstrated the need for the internal aerodynamic forces to be considered as part of the analysis of parafoil performance, a factor not previously acknowledged by the literature.

5.2 Initial 3D FSI simulation of infinite span ram-air parachute

This section details the results of the first attempt at the prediction of the inflated shape of a parafoil cell in 3D using a loosely coupled FSI simulation approach. The FSI methodology employed for this has been described in Chapter 3, and it is applied to the single-cell model in Figure 3.1(a). The structural boundary conditions, labelled as BC1, are described in Figure 3.2(a) and in Section 3.2.2, but a brief outline is given here. The cell is modelled in periodic symmetry, as though it is an individual cell in a chain formed of multiple repeating cells, to form an infinitely long canopy of zero anhedral. The half-cell of the parafoil is constructed of two loaded ribs, constrained at an angle of attack of 5°. One of the ribs is able to translate in the spanwise direction, but no forces are applied to the ribs as part of the loading conditions.

5.2.1 Prediction of inflated shape

After six fluid-structure iterations, the differences in the predicted shape of the parafoil cell between iterations are shown in Figure 5.21 to have reduced significantly. Figure 5.21 shows the magnitude of the resultant nodal displacement for three nodes on the centre-line of the upper surface; one at the leading edge, one at the thickest point of the parafoil (in the y-direction), and one near the trailing edge. Whilst the displacements between iterations have not quite reduced to zero, the figure shows that the iterations have reached a reasonable level of convergence, with the differences between the nodal positions less than 5mm, which provides confidence in the results.

The resulting structural shape of the parafoil cell after six FSI iterations is shown from multiple angles in Figure 5.22. It can be seen that the internal pressures inside the cell of the parafoil, and the external pressures created by airflow over the outer surfaces of the parafoil, produce a parafoil shape that is significantly different to the unstressed cut pattern shown in Figure 3.1(a). This difference is also obvious from the cross-section plots provided in Figure 5.23.

The final shape is determined by the balance between the interior and exterior pressure forces acting on the parafoil, and the structural response of the parafoil governed by the fabric behaviour and the canopy design. The underlying mechanism that de-

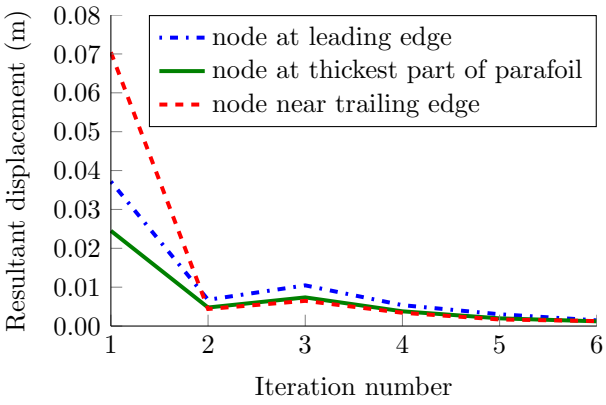


Figure 5.21: Pseudo-residuals showing difference between the predicted structural geometry with each FSI iteration. Values shown are the resultant displacements measured at three nodal points on the spanwise centre-line of the canopy's upper surface.

termines this shape consists of two main actions: the first is the upward movement of the upper and lower surfaces, thus drawing the ribs inwards as a result of membrane tension in the canopy fabric. This subsequently initiates a secondary effect, of introducing a horizontal component to the outward force. This horizontal force opposes the tension force that draws the side walls inwards. The balance of these two forces plays a large part in determining the final shape of the structure.

This force balance within the structure results in a markedly different shape for the parafoil when compared to the cut pattern, which is dominated by the effect of the billowing of the upper and lower surfaces. The cell transitions from the flat, prism-like shape of the uninflated structure to a three-dimensional entity. This is clear from the direct comparison of the centre cross-section of the parafoil before and after inflation in Figure 5.23, which shows that the cross-section of the parachute between the ribs has a greatly increased thickness along the full length of the chord, which is at its peak at the centre-line. Table 5.4 shows that at the centre-line, the distance between the upper and lower surfaces at the thickest point along the chord increases from 18% of chord pre-inflation, to 25% of chord. This compares reasonably well to the increase seen by Ware & Hassell, whose wind tunnel results record an increase in the maximum thickness from 18% of chord to 23% of chord between the ribs. This change in thickness presents a significant deviation from the cross-section of the rib, and is a clear indication that the flow field around the canopy during flight is likely to differ from any predictions made based upon the design of the rib's cross-section. It also

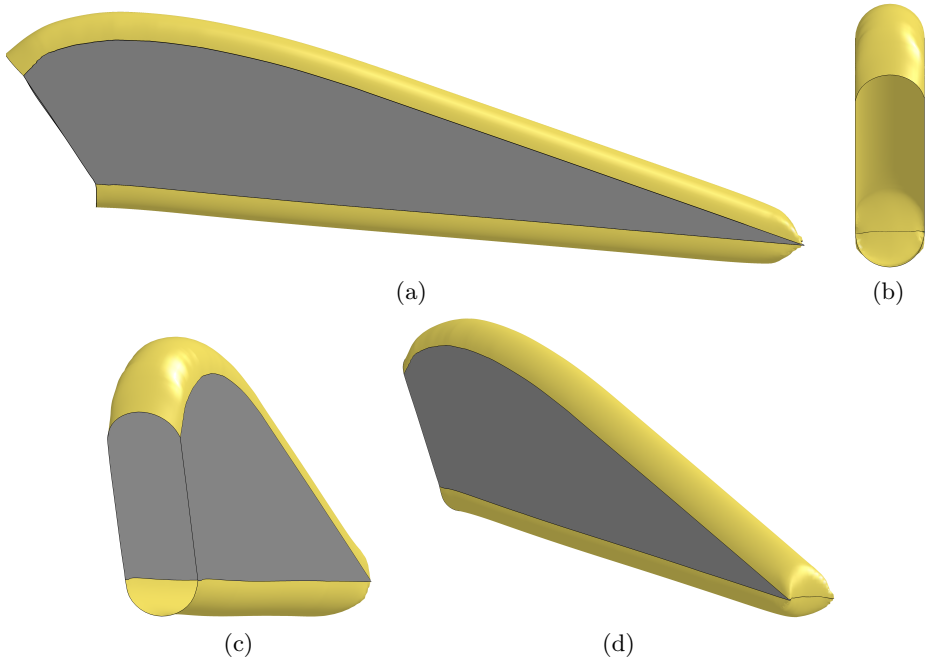


Figure 5.22: Resultant 3D geometry of ram-air parafoil cell, shown from the (a) side, (b) front, (c) front-side, and (d) rear-side. Stoppers not shown for clarity.

suggests that due to the cell billowing the profile of the parafoil varies along the span, and so therefore the flow is likely to exhibit three-dimensional behaviour, which would not be captured in two-dimensional simulations of the unstressed cut-pattern.

It is worth noting that, other than in the region of the leading edge, the shape of the rib is very similar in both the stressed and unstressed cases. This indicates that the expansion of the upper and lower surfaces between the ribs is primarily a consequence of the membrane action of the fabric, and the tension which pulls the ribs together, rather than the fabric in the rib being stretched. The scalloped nature of the canopy is therefore seen to be a result of the way the design of the rib as a whole responds to the inflation pressures, as opposed to an elastic deformation of the ribs.

Apart from the increase in the parafoil profile between the ribs, Figures 5.23 and 5.22 also show other interesting features of the inflated shape. One such feature is seen at the leading edge, where the wrinkling of the fabric in the ribs permits the nose section to move backwards, increasing the angle of the inlet and reducing the distance between the leading edge and the point of maximum thickness. This changes the aerodynamic

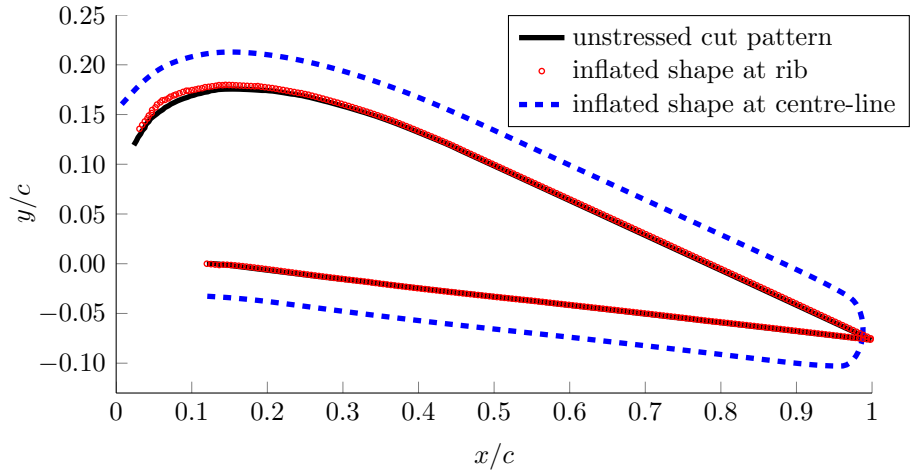


Figure 5.23: Chordwise cross-section of the parafoil in three cases: unstressed shape at the rib (cut pattern), inflated shape at the rib, and inflated shape at the centre-line.

shape of the parafoil, and so will influence the aerodynamic performance. Additionally, Figure 5.23 reveals elastic deformation in the rib in the inlet region, which is not seen in the rest of the rib. The magnitude of the deformation in the rib adjacent to the inlet is obviously dependent on the material properties, which for this case have been approximated, Section 4.4, and so a fabric with a higher elastic modulus is likely to undergo less stretching in this area, and thus a smaller increase in inlet size. It would also be influenced by any stitching or reinforcement that may be present in the parachute design, which have not been included here. These issues are highlighted in the comparison in Table 5.4, which shows that the size of the inlet predicted by the simulation is much larger than the values recorded by Ware & Hassell; with predicted values of 16% at the rib, and 22% of chord between the ribs, differing significantly to the recorded experimental values of 11% and 15% of chord. The discrepancy in the size of the inlet, combined with the closer agreement in the results at the thickest part of the parafoil, suggests that the inlet area is an area of high stress, and that correct material properties and design details are necessary for a more accurate prediction of the shape in this region. The aerodynamic consequences of the size and shape of the leading edge region are discussed in the next section.

A further observation on the inflated shape is that the trailing edge, once inflated, becomes rounded instead of sharp, shown very clearly in the centre-line cross-section plotted in Figure 5.23. This is a major deviation from the original design shape, and

Table 5.4: Comparison of structural dimensions of the simulated ram-air parafoil to wind-tunnel test results.

Dimension (% chord)	Ware & Hassell Wind Tunnel Test	Simulation
Inlet		
Thickness At Rib	11%	16%
Maximum Thickness Between Ribs (% chord)	15%	22%
Thickest part of parafoil		
Thickness At Rib (% chord)	18%	18%
Maximum Thickness Between Ribs (% chord)	23%	25%

the three-dimensional views in Figure 5.22 show that the ribs are the only places along the length of the trailing edge which remain sharp. The two-dimensional CFD simulations around the cut pattern in Section 5.1 demonstrated that the unstressed design was well streamlined over the upper surface, but this would be unlikely to be the case if the trailing edge was more similar to that of the inflated case. The likely consequence of this, in terms of the effect on the aerodynamic performance of the parachute, is that a rounded trailing edge is likely to cause the flow to separate in that region. This leads to an increase in drag, and so parafoil designs which limit this effect may provide improved performance.

In addition to the differences in section thickness and overall shape observed in the inflated canopy, a change in the spanwise dimension of the cell is observed. As already described, the outward pressure forces applied on the upper and lower surfaces create a tension in the fabric, which pulls the ribs towards each other. The reduction in the cell width, as seen in Figure 5.22, is pronounced, and no doubt contributes to the large displacement of the upper and lower leading edges, because the taut fabric membrane is able to displace further upwards thanks to the slack created. The fact that in the structural simulation one of the ribs is completely free to move in the span-wise direction, and that no outward force is applied to either of the ribs, shows that the rib does not require any external structural force to stay open. The width of the cell in this case is purely determined by the effects of the internal and external aerodynamic pressure forces.

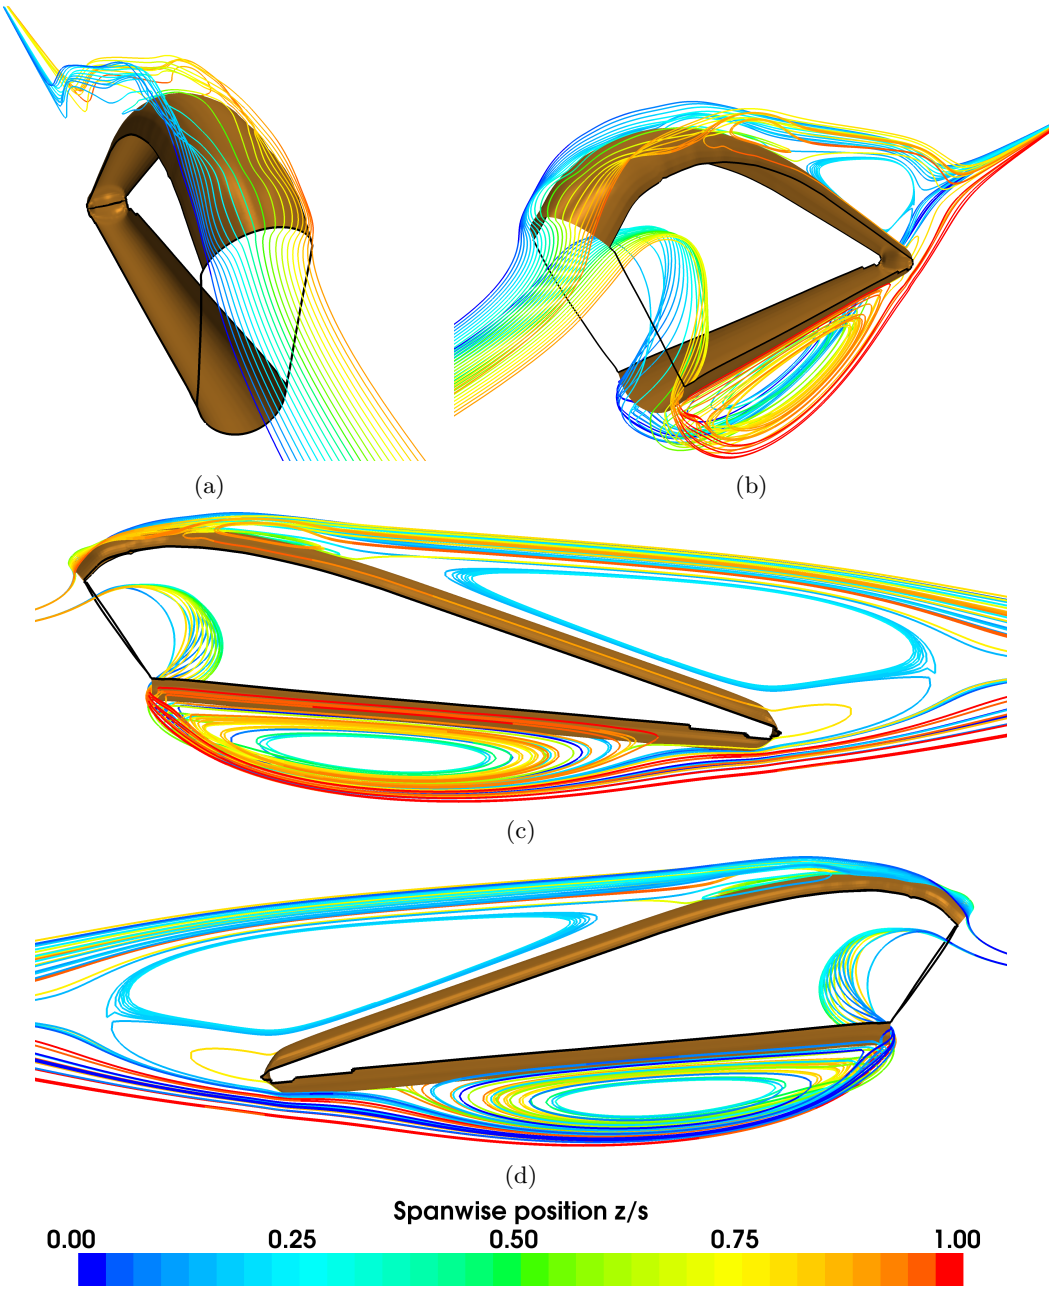
Despite the apparent ability of the cell to remain inflated without any external struc-

tural forces being applied, a comparison of the results to those of Ware & Hassell reveals a large discrepancy. The width of the predicted inflated shape is 63% of the unstressed geometry, whilst Ware & Hassell observed an inflated parafoil span of 95% of the original geometry. This strongly suggests that applying the boundary conditions appropriate for the periodic symmetry of an infinite canopy does not accurately capture the physics of a real-life canopy.

Interestingly, despite the discrepancy in the cell width, and the effect this is likely to have on the potential for upper and lower surface displacement, the values of the displacements for other key dimensions of the parafoil cell discussed still compare well. This observation indicates that there must be a difference between the behaviour of the structure used in the simulation compared to a real-life canopy. A likely explanation is the boundary conditions on the parafoil model, which constrain the ribs on both sides of the cell in the vertical direction. In a modern canopy, this is not the case, and the suspension lines are only present on every other rib, creating unloaded ribs as a means of reducing drag caused by the lines. This difference produces a shape which is more locally scalloped, and this could be the reason for the different relationship between cell width and cell thickness. Despite these differences, investigation of the flow around the structure still provides insight into the behaviour of ram-air parachute canopies.

5.2.2 Flow field results

An overview of the flow around the predicted geometry of the inflated structure is provided by the 3D streamlines presented in Figure 5.24. The different views in the figure show a number of similar features to the results presented for the 2D parafoil geometry. The presence of the open inlet once again results in the flow near the nose region either accelerating over the upper surface, or entering the mouth of the inlet before being diverted out and around the sharp lower leading edge. As before, this diversion of flow around the lower leading edge results in a significant separation bubble on the lower surface. The open inlet leads to a region of stagnated air inside the interior of the cell, creating an area of zero velocity and high pressure displayed in Figures 5.26(a) and (b). This is in agreement with the observations of the flow around the cut pattern in Section 5.1.1.



(a) (b) (c) (d) (e) legend

Figure 5.24: 3D streamlines for flow around inflated shape of ram-air parafoil cell with infinite canopy boundary conditions, shown from (a) above-front, (b) front-side, (c) side 1, and (d) side 2. Streamlines coloured according to spanwise position, (e). Ribs omitted for clarity.

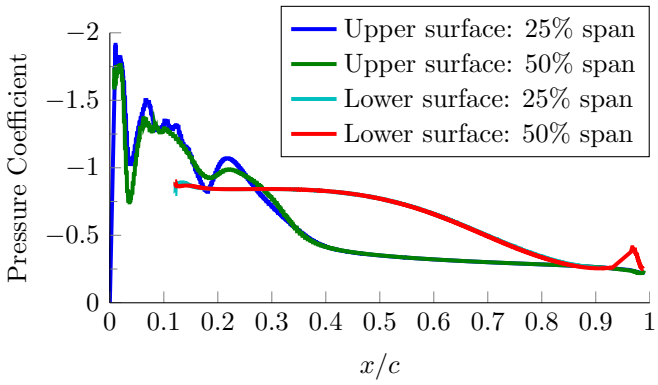
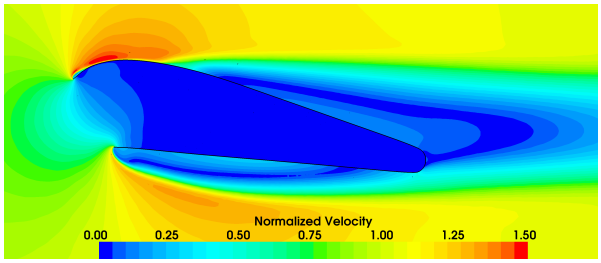


Figure 5.25: Chordwise C_p profiles for external surfaces of parafoil at different spanwise positions.

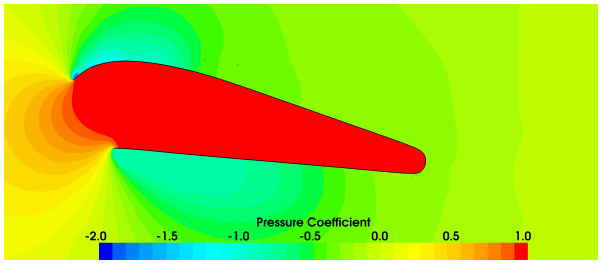
However, beyond these basic similarities it is obvious from Figure 5.24 that the three-dimensionality of the inflated shape produces a corresponding flow field that is also three-dimensional in nature, with obvious differences in the flow across the span of the cell. This signals a departure from the simple 2D treatment of the cut pattern, and requires a more detailed evaluation in order to understand what is happening.

Starting with the upper surface, it can be seen from Figures 5.24 and 5.26(a) that the flow is accelerated over the upper leading edge exactly as would be expected. This acceleration creates an area of low pressure, which acts as the "suction" source that generates lift. Plotting the values of pressure coefficient along the upper surface of the parafoil at the mid-section and quarter section in Figure 5.25, reveals two distinct peaks in the negative pressure. The first is seen in the region right at the upper leading edge, and is clearly seen in Figure 5.27(c) as an area of strong negative pressure. The contours of wall shear stress in (b) show that this corresponds to an area of separated flow, which then subsequently reattaches. This separation bubble was observed in the 2D open inlet case also, and creates a localised region of low pressure, and high vorticity, seen in Figures 5.26(b) and (c).

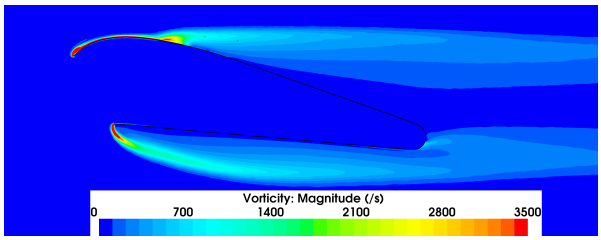
Downstream of this separation bubble the flow quickly reattaches, and accelerates, reaching a peak velocity just before the point of maximum thickness, observed in Figure 5.26(a). After the flow has passed the point of maximum velocity the velocity starts to reduce, and then the flow starts to separate. The wall shear stress contour plot in Figure 5.27(b) shows that separation occurs at different points depending on the spanwise position, and the difference that this makes in the chordwise flow profile is indicated by the constrained streamlines plotted in Figure 5.28.



(a) Normalized Velocity

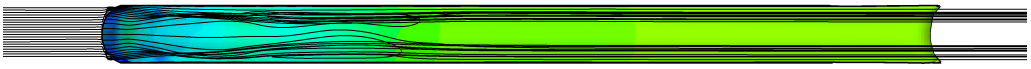


(b) C_p



(c) Vorticity

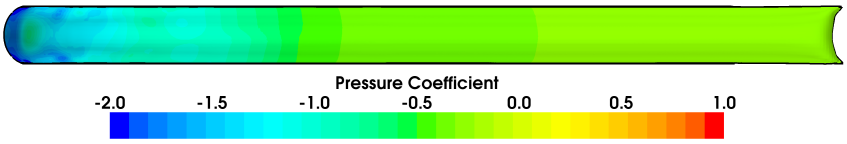
Figure 5.26: Chordwise cross-section contour plots of flow at 50% span.



(a) 3D streamlines and surface contours of C_p

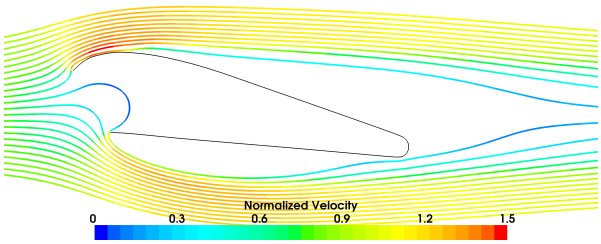


(b) Wall shear stress: positive values red, negative values blue

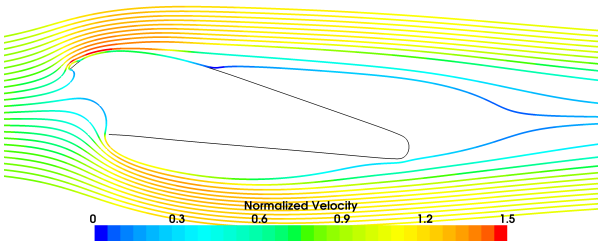


(c) Surface contours of C_p

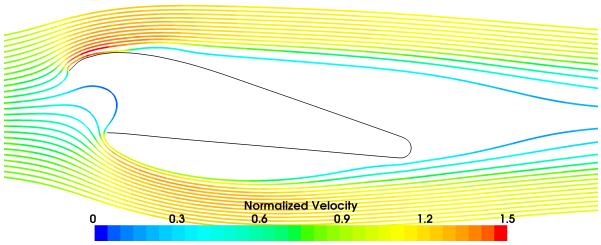
Figure 5.27: Flow contours over the upper surface of canopy. Flow from left to right.



(a) z=25% span



(b) z=50% span



(c) z=75% span

Figure 5.28: Chordwise cross-sections at different spanwise positions showing constrained streamlines. Contours of velocity normalized against freestream velocity magnitude.

Examination of the 3D streamlines in the area where the upper surface separation is initiated, Figures 5.24(a) and (b) indicates a three-dimensional region of flow. Once the flow over the leading edge has reattached the streamlines are seen to spread out, showing a spanwise movement of the air around either side of the dome like shape of the upper surface, near to the point where the parafoil is at its thickest. The fluid moving away from the central area of the flow joins up with the existing flow near the ribs. The combined flow rolls up into two parallel vortical structures, one on each side of the parafoil, Figure 5.24(a). These structures appear to remain intact as they travel downstream.

The vortical structures observed on the upper surface can, in some cases, organise the flow, as the vortices persist downstream due to their stability. This can behave

as a barrier, preventing spanwise flow, thereby reducing drag losses due to wing tip vortex shedding. The separation in the areas to either side of the centre looks likely to be initiated by the spanwise movement of the flow away from the centre-line. The remaining flow over the central portion of the surface however, stays attached for longer, with separation delayed until almost half-way along the chord, shown in Figures 5.27(b) and 5.28(b).

Upon inspection of the pressure and wall shear stress contours for the upper surface in Figure 5.27, a number of asymmetries are obvious, which would not be expected in a symmetrical case such as this. There are also several areas of local disturbance noticeable in the plots, particularly in the case of the wall shear stress, Figure 5.27(b). These may stem from subtle asymmetries introduced into the original structural finite element geometry, as a consequence of subtle rounding errors produced by the pre-processor. The use of an explicit solver for the quasi-static structural simulations also results in slight asymmetries in the final geometries, possibly due to the handling of localised stresses in the fabric. The stress due to changes in the cell width may be absorbed by the compressive stress elimination algorithm, rather than being transmitted through the material by elastic waves, preventing a symmetric equilibrium from being reached. Another possible explanation for this is a build-up of rounding errors due to the large number of iteration cycles performed during a long-duration explicit simulation such as this. The asymmetry introduces subtle differences into the local deformation of the elements, as the localised bumps and lumps in the fabric are not distributed symmetrically. The presence of these asymmetric areas of deformation explains the small patches of localised separation and reattachment, and corresponding pressure fluctuations, observed near the ribs in Figures 5.27(b) and (c).

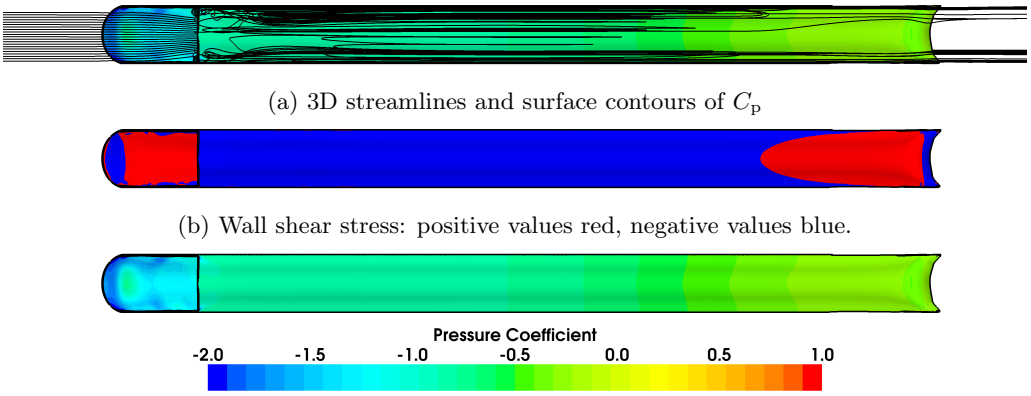
Similarly to the observations of the flow field over the upper surface, the flow around the lower surface also exhibits spanwise variation and is dominated by a large area of separated flow, clearly apparent in Figure 5.24. Similarly to the findings in the two-dimensional case, the stagnated air in the interior region of the parafoil effectively blocks off the flow near the mouth of the inlet, but as is also the case with the 2D model, the blocking effect is not complete until just beyond the leading edge. The streamlines show that a proportion of the flow passes beyond the mouth of the inlet, before being forced to reverse direction and exit the interior region near to the lower leading edge.

This flow reversal was observed in the two-dimensional flow field results, but the three-dimensional results show a dependence in the distance the flow extends into the interior domain on spanwise position. The 2D and 3D streamlines in Figures 5.28 and 5.24 show that the flow encroaches further into the cell at the centre-line than it does at the edges. This suggests a relationship with the inlet size, as perhaps the larger inlet profile at the centre permits the flow to penetrate further into the internal space.

Curiously, from viewing the 3D streamlines, the spanwise variation in the distance to which flow enters the ram-air cavity does not appear to have a strong influence on the path which the flow then takes when it transits around the lower leading edge. One observation that can be made however, is that the streamline at the centre-line can be seen to extend further into the direction of oncoming flow upon exiting the inlet, taking a slightly more pronounced path. This does not appear to have a large effect on the flow, but this flatter profile may be related to the flow reattaching sooner at the centre line. In relation to this, the circulation bubble appears to be flatter towards the centre, with the size of the recirculating region expanding to match the contour of the lower surfaces spanwise curvature. The plan view of the lower surface streamlines in Figure 5.29(a) suggests that once reattachment occurs, there is a tendency for the centre-line flow to move towards either rib.

Upon exit of the flow from the cell's interior it rejoins the rest of the flow diverted around the lower surface, and in a similar manner to that identified in the two-dimensional case, the flow follows a curve as it is forced to change direction once more. In the three-dimensional case the separation is much more severe, and forms a very large area of recirculating flow beneath the lower surface. The scale of this is clear from the 3D streamlines, and a large region of low velocity and reduced pressure is apparent in Figures 5.26(a) and (b), as well as an associated increase in vorticity in (c). The streamline and velocity plots hint at an eventual reattachment of the flow, confirmed by the positive wall shear stress values towards the trailing edge in Figure 5.29(b). The figure shows a symmetry to the reattachment, with the flow reattaching first at the centre-line, at about 80% along the chord.

The large increase in the length of the lower surface separation bubble over the 2D case must be a direct consequence of the change in the ram-air cell geometry brought about by inflation. One factor is the billowing of the upper and lower surfaces, resulting in a scalloped appearance to the cell, which increases the vertical size of the



(a) 3D streamlines and surface contours of C_p

(b) Wall shear stress: positive values red, negative values blue.

(c) Surface contours of C_p

Figure 5.29: Flow contours over lower surface of canopy. Flow from left to right.

parafoil, presenting a larger obstruction to the flow. Naturally, this causes an increased disturbance to the freestream compared to the unstressed cross-section.

Additionally, the analysis of the structural simulation in Section 5.2.1 noted that the height of the parafoil inlet predicted in this study is larger than that recorded by Ware & Hassell, due to stretching of the ribs at the inlet and the large decrease in the cell width. This further increases the size of the obstruction presented to the oncoming flow, and exaggerates the effect on the flow. Coupled with this is the change of the angle of the ram-air inlet relative to the freestream, which is seen above to increase as a result of the upper surface moving rearward as the fabric wrinkles. The steeper angle results in a blunter shape being presented to the flow, and thus air flows around the lower leading edge more obliquely, increasing the length of the separated region. This has been seen to have a negative effect on the aerodynamics in 2D flow studies, as discussed in Section 5.1.4.

Interestingly, the wall shear stresses at the trailing edge show a second separation point on the lower surface, and a late reattachment on the upper surface. This must relate to an interaction of the flow with the rounded trailing edge, suggesting that the flow is not able to follow the sudden deviation of the cross-section at this point. A subsequent recirculation zone is hinted at by the looped streamlines seen in Figure 5.24. This is very unlikely to have a large impact on the aerodynamic performance of the parafoil in this case, but it does allude to the adverse effect a rounded trailing edge would have on a more streamlined parafoil design.

5.2.3 Summary and discussion

The results in this section have shown that the loosely coupled fluid-structure interaction methodology can successfully be implemented, and converges to predict an equilibrium shape for an inflated 3D parafoil cell. The predicted shape for the structure proved to be significantly different to the unstressed geometry. The upper and lower surfaces are seen to translate and expand outwards under the action of the internal and external pressures. The tension created in the fabric surfaces pulls the ribs inwards, and the upper and lower surfaces form a rounded, scalloped shape. The rounded shape creates an outwards spanwise component to the pressure forces, which balances the vertical forces and allows the cell to remain open.

The billowing of the upper and lower surfaces is seen to greatly increase the size of the chordwise cross-section between the ribs, making it much thicker than the unstressed cut pattern. The expansion of the fabric surfaces also caused the trailing edge to become rounded between the ribs, which again is a large and important deviation from the uninflated design. The ribs themselves are shown to remain mostly undeformed, apart from in the vicinity of the inlet, which was shown to be a region of high stress due to the large aerodynamic forces. The edge of the inlet moved backwards due to wrinkling in the fabric, increasing the angle of the inlet to the freestream. The fabric in the rib near to the inlet is also seen to stretch under the aerodynamic forces, further increasing the inlet size.

The associated flow field around the inflated shape shares a number of similarities with the flow around the 2D cut pattern, but is seen to be distinctly three-dimensional due to the spanwise variation in the cross-section of the inflated shape. The flow is therefore significantly different from the 2D case as a consequence. The flow around the inlet and interior of the cell is similar to the 2D case; the internal region was again observed to be full of stationary, high pressure air that prevents further flow entering the space, apart from in the area directly adjacent to the inlet. The localised separation and reattachment at the upper leading edge seen in the 2D case was also present. An area of suction pressure is generated further downstream on the upper surface due to acceleration of the flow. However, unlike the 2D case, the flow is seen to separate from the upper surface near to the point of maximum thickness, and remains separated over the rear half of the structure. The lower surface separation previously seen around the 2D cut pattern is much more severe in this case, and the separation

bubble has been shown to extend along 80% of the length of the chord.

The combination of flow features observed around the inflated parafoil cell, produced very poor aerodynamic performance. Unsurprisingly, the lift is very poor, due to the degradation of the upper surface suction by the onset of separated flow. This is further degraded by the large region of suction in the negative lift direction as a consequence of the massively separated flow over the lower surface. The large degree of separation also resulted in high drag, and therefore a lift to drag ratio of around 1.

Clearly, the level of performance suggested by the aerodynamic coefficients predicted in this model do not reflect the reality of a typical parafoil canopy. This obviously suggests a number of deficiencies in the chosen boundary conditions, as a parafoil with this cross-section was successfully flown by Ware & Hassell, demonstrating more favourable lift and drag coefficients. As discussed above, the separation is very sensitive to the geometry, which a comparison with the observations of Ware & Hassell showed to be enlarged at both the inlet and point of maximum thickness. In terms of encouraging separation, it is likely that the size of the inlet, which was shown to be increased by stretching of the material in that region, is a key aspect to this. The billowing of the inlet also contributed to its size, a large part of which is down to the reduction in cell width due to membrane tension.

As the results have confirmed the need to consider the inflated shape, rather than the cut-pattern when analysing the aerodynamics, and the viability of the loosely coupled FSI framework has been successfully tested, improvement of the parafoil model is shown to be of value. Consequently, the next section of results incorporates a number of changes. Firstly, an outward component to the spanwise force is added, based on the physics of a finite span parafoil, to improve the predicted cell-width. Secondly, improved material properties are included to more accurately account for the forces applied to the ribs. Lastly, the presence of unloaded ribs in modern canopy designs, which have been ignored in this model, is taken into account.

5.3 Improved 3D FSI simulation of infinite span ram-air parachute

This section details the results of fluid-structure interaction simulations of a parafoil that utilise a modified approach over the previous section, in order to improve on the model above and more accurately recreate the behaviour of a ram-air canopy. As discussed in Section 3.2.2, the boundary conditions have been adjusted such that the single parafoil cell is modelled in half-symmetry, which forms a full-symmetry unit-cell consisting of a full-cell, which once again is modelled in periodic symmetry as part of an infinite span canopy. Hence, the simulated half-cell has one loaded rib, which is constrained in the vertical direction, and one unloaded rib, which is free to vertically translate. This more closely recreates the design of modern canopies, which have suspension lines attached only to every other rib, to reduce the aerodynamic drag produced by the lines.

This model also has updated material properties, with a greater Young's modulus, and a zero Poisson's ratio, with full material properties given in Table 4.5. Finally, the single canopy cell has a finite spanwise force applied to the loaded rib. This force is applied as a pressure across the full interior surface of the rib, equal to that of the freestream dynamic pressure. This represents the force associated with the pressure difference present across the interior and exterior surface of the rib at the spanwise tip of a parafoil canopy.

As well as updates to the boundary and loading conditions of the model, there were also alterations made to the FSI iteration procedure. Firstly, an implicit structural solver is used instead of an explicit one, which vastly reduces the structural simulation run-times. Secondly, the Taubin smoothing step is no longer required, because the more accurate value used for the Poisson's ratio, of $\sigma = 0$, removed the issues with element distortion and localised deformation on the upper surface. Lastly, the "stoppers" are no longer required, as the additional spanwise loading prevented the upper and lower surfaces from crossing the symmetry boundaries at the ribs.

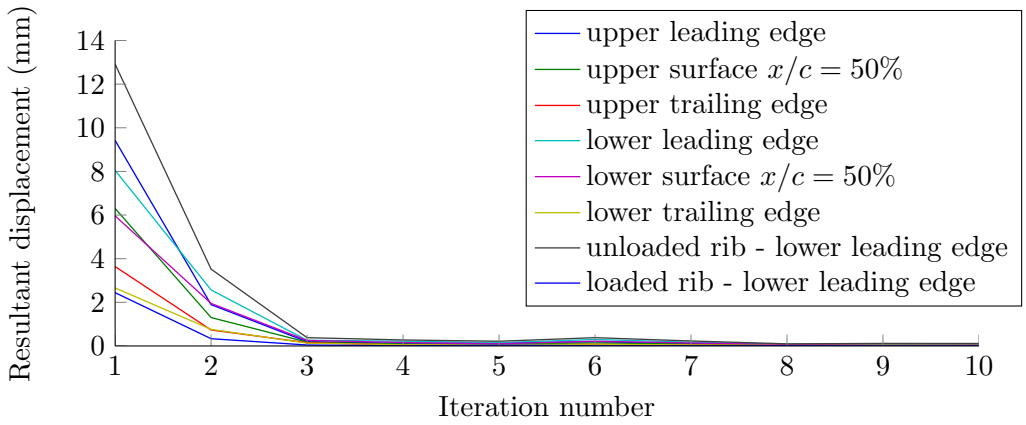


Figure 5.30: Pseudo-residuals showing difference between the predicted structural geometry with each FSI iteration for BC2 model. All measurements taken at points located at the spanwise centre-line unless otherwise stated.

5.3.1 Prediction of inflated shape

The convergence plot in Figure 5.30 shows a very large reduction in the resultant displacements seen in the structure between iterations. By iteration eight, the differences in the nodal positions has reduced to less than 1mm, suggesting a very good level of convergence in the inflated shape. The level of convergence for this model is slightly greater than the previous model, due to the stabilising influence of the applied spanwise force, and the stiffer material properties.

The overall shape of the predicted inflated parafoil cell can be seen from a number of different viewpoints in Figure 5.31. All of the images are reflected about the x-y plane, (the plane of the unloaded rib) to emphasise the unit cell geometry which the boundary conditions of this FSI model represent. The figures clearly show the difference in the shape which the presence of an unloaded rib creates, with the central rib moving in the vertical direction into a higher position relative to the constrained loaded rib. The modelling of the unloaded rib results in a much more realistic overall shape, with a strong curvature to the cell visible in the spanwise direction.

The vertical movement of the unloaded rib is explained by the action of both the lift generating suction pressure on the upper surface and the positive pressure on the lower surface, described in more detail in Section 5.3.2. The difference in vertical position between the loaded and unloaded rib is just over 2% of chord, and the vertical shift of the unloaded rib is displayed clearly in Figure 5.32.

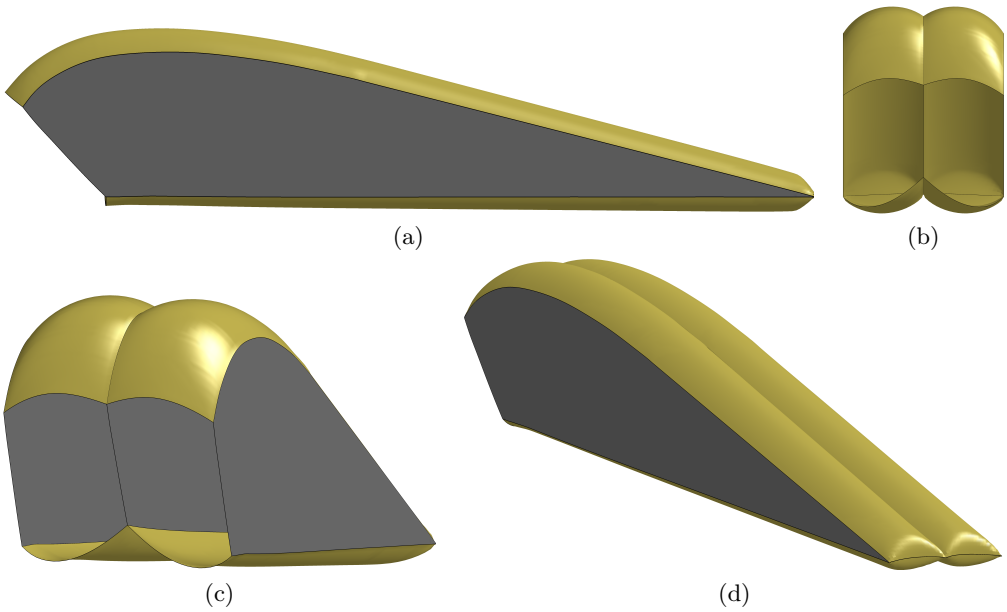


Figure 5.31: Resultant 3D geometry of ram-air parafoil cell using improved boundary conditions, shown from the (a) side, (b) front, (c) front-side, (d) rear-side. Cell reflected about unloaded rib to show full-symmetry case.

This shift in position is seen in the cross-sections of Figure 5.32 to be more pronounced at the leading edge than the trailing edge, where the difference in vertical translation between the two ribs is much smaller. This shows that the tapering of the cell geometry in this region makes the rear part of the structure stiffer, and more resistant to the aerodynamic forces. The pressure differential is likely to be much less pronounced in this region too, resulting in lower forces, which will also contribute.

The views of the structure in Figure 5.31 show a number of features familiar from the previous 3D results; the upper and lower surfaces both billow and expand under the membrane tension in the fabric, creating the typical scalloped shape. Once again, the trailing edge is seen to be rounded across the span, with the sharp trailing edge of the unstressed cut pattern visible solely at the ribs in Figure 5.32. As previously discussed in Section 5.2.2, a rounded trailing edge is likely to have a negative impact on the aerodynamic performance of the canopy.

When observed from the front, the billowing of the canopy observed in Figure 5.31(b) creates a curved leading edge at the upper and lower surfaces, but the presence of the unloaded rib has introduced an asymmetry to the inlet shape between the two

surfaces. The reduction in the cell width due to the tension present in the upper and lower fabric membranes is noticeable in the results, but the effects of this are less than in Section 5.2, due to the prescribed loading in the spanwise direction which acts as a restoring force. The upwards translation of the unloaded rib and the greater cell width causes the size of the inlet between the ribs to reduce over the previous FSI model. A factor in this is the asymmetry now present in the shape of the upper and lower leading edges, which means that the most deformed point on the upper surface edge is no longer in line with the most deformed point on the lower surface edge. As a consequence, the increase in inlet size between the ribs is now less pronounced. This is likely to have a positive effect on the aerodynamics of the inflated shape. The billowing of the cell however still creates a thicker profile in the canopy between the ribs, Figure 5.32, which again is a marked deviation from the design of the unstressed cross-section.

Also of note in the region of the inlet is the shape of the lower surface in the area just behind the leading edge, which displays a noticeable chordwise curvature. This deformation is caused by the movement of the rib, yet no such deformation is seen on the upper surface. The results of the CFD analysis in the following section are likely to shed light on this localised deformation.

Table 5.5: Comparison of predicted structural dimensions of inflated ram-air parafoil to experimental values.

Dimension (% chord)	Simulation	Ware & Hassell Wind Tunnel Test
Inlet		
Thickness At Rib	11%	11%
Maximum Thickness Between Ribs	14%	15%
Width of cell	88%	95%
Thickest part of parafoil		
Thickness At Rib	18%	18%
Maximum Thickness Between Ribs	23%	23%

Comparing the dimensions of the predicted parafoil structure to the experimental values of Ware & Hassell in Table 5.5 shows an excellent agreement between simulation and experiment. The maximum thickness at the inlet, both at the rib and between the

rib, matches very well. This is in contrast to the previous results, where both of these dimensions were much larger in the simulation than the experimental values. This shows that the combination of fabric material properties and boundary conditions used in this simulation ensure a much more realistic response and limit the stretching in the inlet region.

The comparison of the maximum thickness of the parafoil cell also compares well. The previous 3D model also compared well in this area, but the updated FSI model shows further improvement over the previous result for the maximum thickness of the parafoil between the ribs. Due to the low strain seen in the upper and lower surfaces of the cell, Figure 5.33, it is unlikely that any material stretching occurred, thereby eliminating the influence of the updated material model. Therefore, the closer comparison to the observations of Ware & Hassell must be a direct result of the incorporation of the loaded rib, and the application of the spanwise loading. The spanwise loading is likely to be the dominant factor because it forces the cell to be wider, therefore reducing the amount of vertical displacement seen on the upper and lower surfaces.

The resulting cell width is significantly increased over the previous simulation result, but the value of 88% of chord, Table 5.5, is still less than the observed experimental value of 95% chord. A possible explanation for this is the lack of spanwise curvature, or anhedral, considered in the model. As explained in Section 2.1, the curvature of the canopy affects the orientation of the lift force vector on each cell, as illustrated in Figure 2.3. Hence, for a curved canopy such as those tested by Ware & Hassell, a component of the lift generated by the outer cells will act in the spanwise direction, and may contribute to increasing the width of the cell. This could explain the extra cell width observed in the experiments. Overall, the current model shows a reasonable quantitative agreement for the cell width and provides confidence in the validity and necessity of the spanwise loading condition.

Whilst there is noticeable deviation from the cut pattern profile between the ribs, there is little difference in the cross-sectional shape of the ribs themselves. Figure 5.32 shows that the loaded rib maintains the same shape as the cut pattern, apart from a shortening of the area around the nose on the upper surface. A similar rearward movement of the upper leading edge was seen in the earlier 3D model, and is present because of the wrinkling permitted by the CSE algorithm used in the fabric constitutive model. Rearward movement of the lower leading edge is prevented by the suspension line

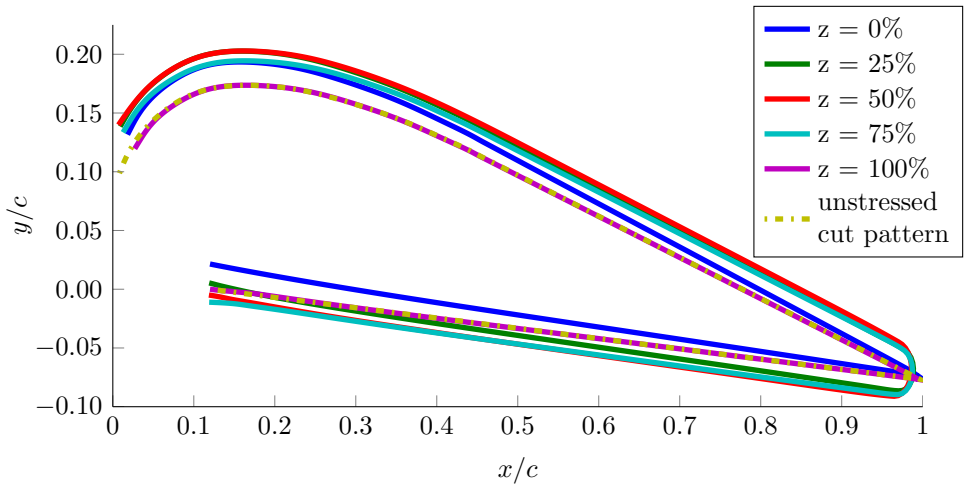


Figure 5.32: Chordwise cross-section of 3D inflated shape taken at various points across the span.

constraints, but rearward movement of the upper leading edge is observed along the full parafoil span, including the unloaded rib, and therefore the angle of the leading edge cut is increased relative to the freestream direction compared to the cut pattern. This angle has been shown to influence the likelihood of separation over the upper surface in 2D simulations, discussed previously in Section 5.1.4.

The cross-section of the unloaded rib in Figure 5.32 maintains its sharp trailing edge, and measurements of the height of the inlet and maximum thickness show that, despite the vertical translation, the profile of the unloaded rib is very close to that of the loaded rib. The comparison presented in Figure 5.32 reveals that as well as the unloaded rib undergoing a translation, it also rotates. As discussed above, the translation at the trailing edge is less than that at the inlet, suggesting an increase in stiffness, and reduction in the applied forces, as x/c increases. This difference in translation between front and back creates a rotation, and the angle of attack of the unloaded rib is seen to increase on both the upper and lower surfaces. The cross-section plots of the structure at different spanwise positions in Figure 5.32 show that the changes in angle of attack vary across the span, and that the angle of attack of the upper and lower surfaces vary independently of each other. This suggests a strong likelihood of a complicated three-dimensional flow field.

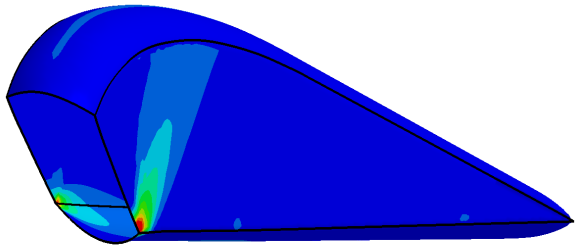


Figure 5.33: Front isometric view of inflated geometry with contours of maximal principal strain. Contours range from blue=0% to red=2.5%.

In terms of the geometrical variations seen in the 2D cross-sections, of particular interest is the gradient of the upper surface, downstream of the point of maximum thickness. As the previous 3D study showed, the inflated shape of the cell is likely to display separation over the upper surface, and the steepness of gradient from the point of maximum thickness to the trailing edge is likely to influence this. If the slope in this region is steep, then the associated adverse pressure gradient in these regions is also likely to be steep, making separation more likely. The results of the associated fluid simulation in Section 5.3.2 will offer an insight into this.

The contour plot of maximal principal strain in the parafoil in Figure 5.33 shows two areas of high strain concentration near the suspension line constraints at the lower leading edge. The areas of high strain also extend upwards into the loaded rib. This is because of the high aerodynamic forces present on the forward section of the upper surface, where the suction peak is located. A corresponding area of high strain concentration is seen at the lower leading edge of the unloaded rib. This is a result of the chordwise constraint applied at that corner to prevent rearward movement. This is unrealistic, but is a small approximation to support the planar constraint on the unloaded rib that maintains the half-symmetry. Realistically, this strain concentration is likely to be present but is unlikely to be as high. It is expected that the removal of the chordwise constraint on the unloaded rib would increase the levels of strain seen at the corresponding point on the loaded rib. This hypothesis is supported by supplementary FE analysis of the structure in this configuration. The remainder of the parafoil cell, comprising of the rearward two-thirds, is shown to be an area of very low strain.

The high levels of strain seen at the lower leading edge on the loaded rib correspond to small, localised points of deformation around the constrained nodes. A similar

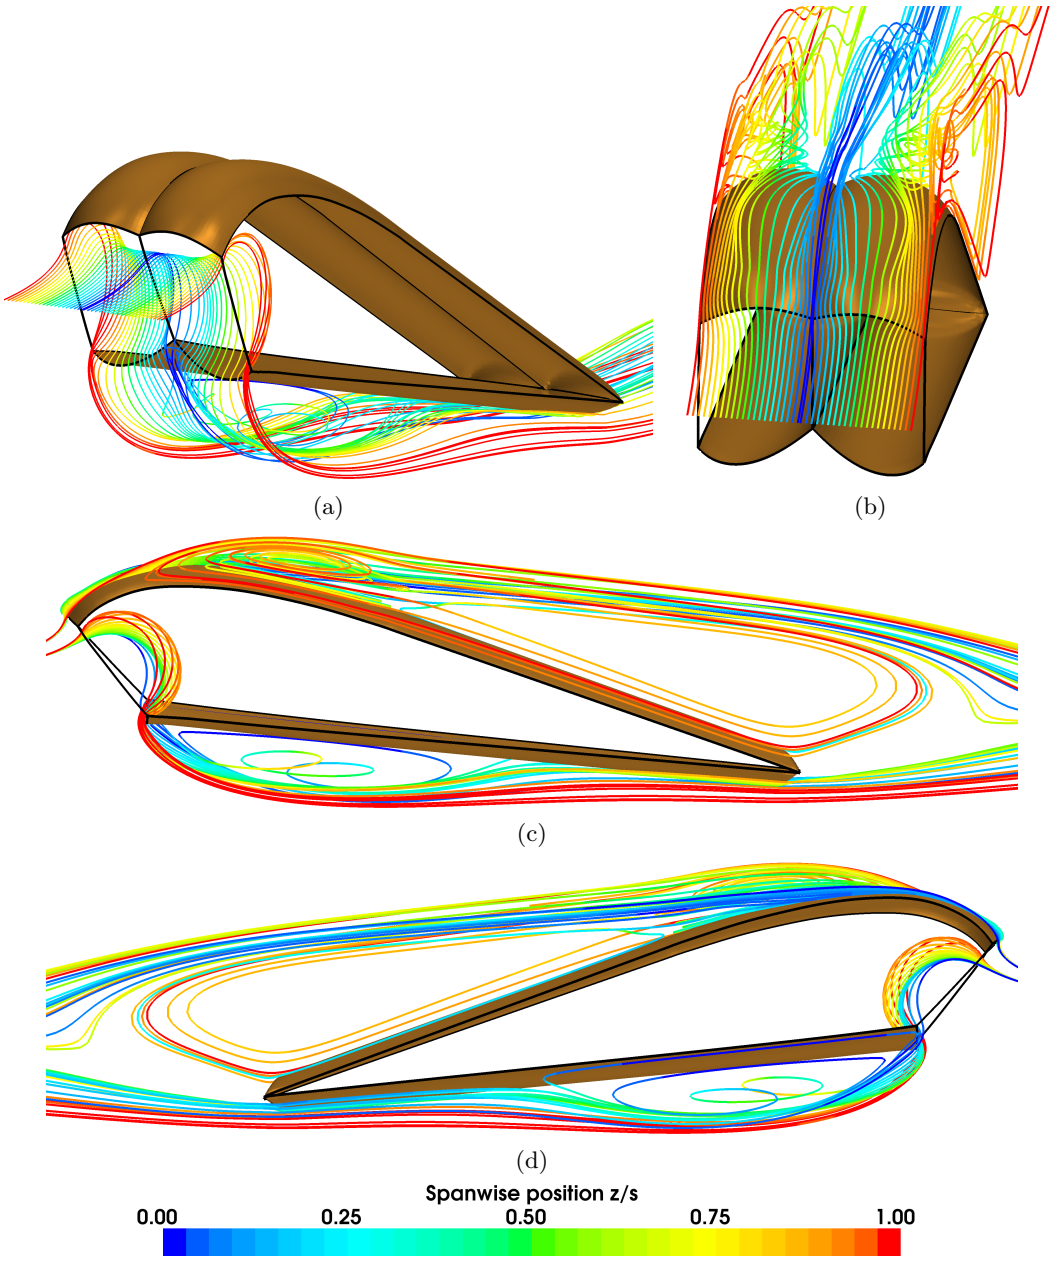
local strain concentration is seen on the other chordwise constraints, where local deformation is also seen, but the strain levels present at these constraint points is much lower. This reflects the high aerodynamic forces caused by the forward location of the suction peak, which are much greater than the forces seen in the rearward half of the parafoil. It is worth noting that the recorded strain values support the fact that the overall levels of strain in the structure are low, and so the approximation of the material properties, which use only the data from the low strain region of the tensile test results, is appropriate and valid. The very highest observed strain value is just over 5%, which is recorded at a very localised, concentrated point at the lower leading edge. This value is just outside the linear region on the stress-strain relation seen in Figure 4.9(b). The remainder, and majority, of the high strain region seen in Figure 5.33 has a stress value of around less than 3%, which is clearly within the linear segment of the material stress-strain relation and so the use of the linear constitutive relationship for the fabric is justified. Of course, this is unlikely to be the case upon opening, but that part of the flight regime is not considered here.

5.3.2 Flow field results

This section details the results of the CFD analysis for the flow around the inflated parafoil shape detailed above. Figure 5.34 shows the three-dimensional streamlines for the flow around the geometry from four different angles. The image of the single cell is reflected about the plane of symmetry at the unloaded rib to show how the streamlines would appear for the full-symmetry case of a full ram-air cell.

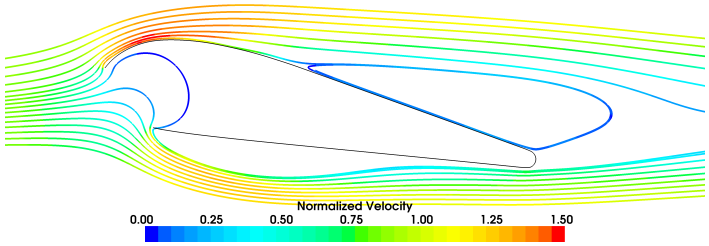
The streamlines show the flow around the structure to have an asymmetrical nature, unlike the flow around the previous structure, due to the asymmetric shape. This is mainly the case over the upper surface, which displays a strong spanwise component to the fluid movement. As well as the difference in symmetry compared with the previous model, the results in Figure 5.34 show a noticeable reduction in the amount of separation seen over both the upper and the lower surfaces, but the flow is still largely separated which is a large deviation from the results for the 2D flow around the uninflated cross-section design.

Looking more closely, the flow over the upper and the lower surface each display distinct flow behaviour. The upper surface flow starts with a deviation of the flow stream around the leading edge, which can be seen in Figure 5.34(b), and a more

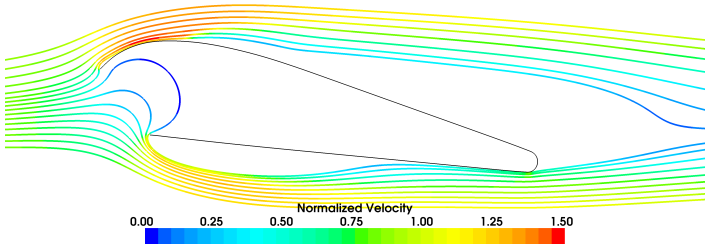


(a) (b) (c) (d) (e) legend

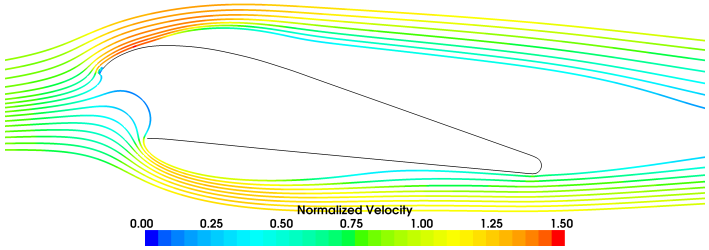
Figure 5.34: 3D streamlines for flow around inflated shape of ram-air parafoil with updated, half-symmetry boundary conditions. Shown from (a) above-front, (b) front-side, (c) side 1, and (d) side 2. Streamlines coloured according to spanwise position, (e). Single cell is reflected to show symmetry, and ribs are omitted for clarity.



(a) z=25%span



(b) z=50%span



(c) z=75%span

Figure 5.35: Chordwise cross-sections at different spanwise positions showing constrained streamlines. Contours of normalized velocity.

obvious bulge is visible in the side view of (c) and (d), as well as the 2D streamline cross-sections in Figure 5.35. As was the case in Section 5.2, the flow over the upper leading edge is associated with an area of high negative pressure, observed in Figures 5.36 and Figure 5.37(c). The wall shear stress contours in Figure 5.37(b) confirm the flow to be separated in this region, reattaching slightly further downstream of the leading edge. The leading edge separation and reattachment occurs along the full width of the upper surface, and the subsequent reattachment is seen to occur at a constant chordwise point across the full span.

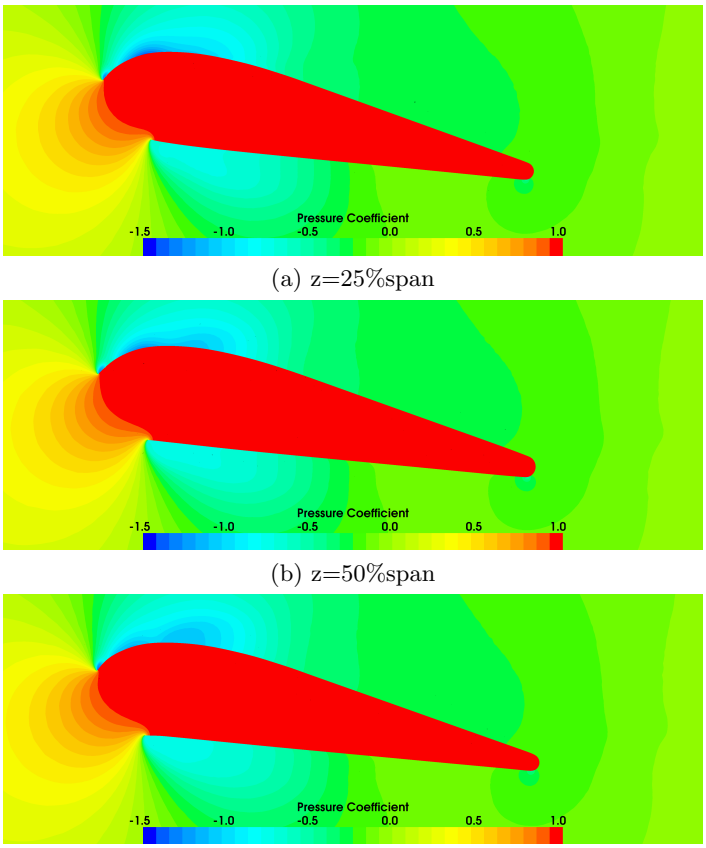
Once the flow past the upper leading edge has reattached it is accelerated over the curvature of the thickest part of the parafoil section, seen clearly in Figure 5.35, which creates an area of high negative pressure, known as the suction peak. The spanwise

cross-sections of C_p in Figure 5.36 show the suction peak to be present just forward of the thickest part of the parafoil, and it is apparent that the magnitude of the negative pressure is greater for the cross-section in (a), nearest the unloaded rib, than for the rest of the upper surface. The plan view of the pressure coefficient contours over the upper surface in Figure 5.37(c) confirm that this is indeed the case, and that the greatest levels of suction are generated in a region near to the unloaded rib. Therefore, the highest proportion of lift is generated in this region too.

In addition to having an impact on the lift generation of the parafoil as a whole, the localised area of high negative pressure over the peak of the unloaded side of the upper surface is also seen to have an influence on the local flow over the rest of the upper surface. Unlike in Section 5.2, the streamlines over the forward section of the upper surface do not move symmetrically across the flow-stream, away from the centre line. Instead the streamlines are seen in Figure 5.34(b) to be drawn spanwise, from right to left, towards the aforementioned area of high suction. This is explained by the flow near to the loaded rib being drawn towards the unloaded rib by the pressure gradient created by the area of high suction. As in Section 5.2, the cross-stream movement of flow affects the fluid already present in the region into which it moves, resulting in separation of the flow directly adjacent to the unloaded rib. The streamlines in this region are shown not to reattach.

The spanwise movement of the flow from right to left is associated with an area of recirculation, which starts around the peak of the upper surface curvature, on the loaded rib side of the span. The recirculating area can be seen to cover around 50% of the span of the parafoil upper surface in Figure 5.37(a). This region is marked out by the contours of wall shear stress plotted in Figure 5.37(b), which shows the region to be the starting point of a large area of separation. The movement of the flow over towards the unloaded rib appears to prolong the distance that the flow remains attached, which at its maximum, on a line towards the unloaded rib, is at roughly halfway along the chord. The large difference in the flow profiles across the span is demonstrated in Figure 5.35, and the reduced severity of the separation between 25% and 50% span is apparent in Figure 5.35(a).

Downstream of the suction peak and the upper surface recirculation region, the streamlines along the centre-line move back towards the direction of the loaded rib. By this point, the upper surface flow has completely separated, and it organises



(a) z=25%span

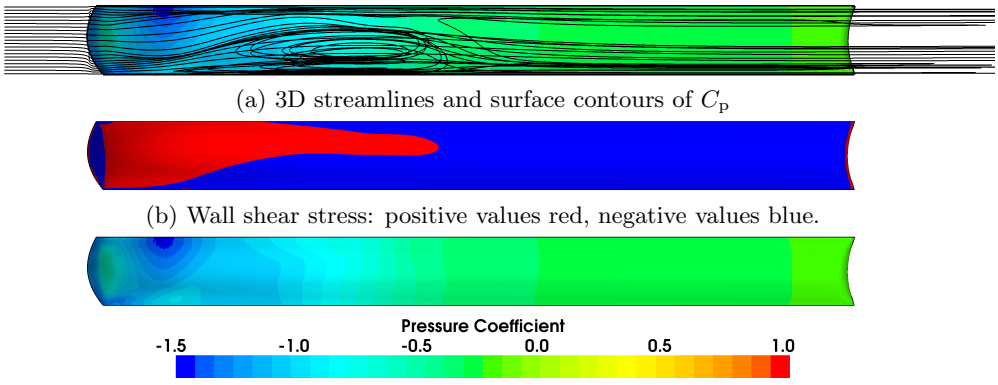
(b) z=50%span

(c) z=75%span

Figure 5.36: Chordwise cross-sections at different spanwise positions showing contours of C_{p}.

into two areas of recirculation either side of the centre-line, seen clearly in Figures 5.34 and 5.37(a).

By comparison, the flow around the lower surface is more homogenous. As has been observed for both the previous 2D and 3D studies described above, the streamlines in Figure 5.34 are seen to enter the interior of the canopy near to the upper surface, and proceed to curve round and reverse direction in order to exit again near the lower leading edge. As observed in the previous 2D and 3D parafoil flow results, the pressure contours within the interior of the parafoil, plotted in Figure 5.36, reveal an area of high ram-air pressure created inside the parafoil, due to the volume filling with stagnated air. This high pressure blocks off the internal cavity to the external flow, but the pressure is not significant enough to seal off the interior completely.



(a) 3D streamlines and surface contours of C_p

(b) Wall shear stress: positive values red, negative values blue.

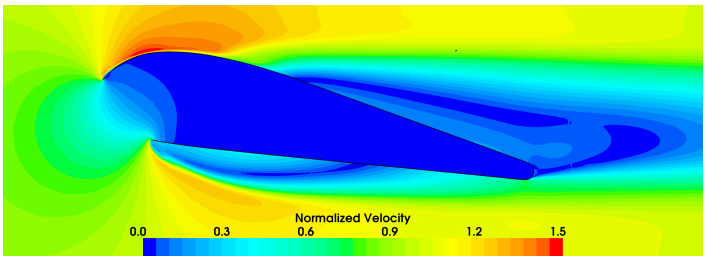
(c) Surface contours of C_p

Figure 5.37: Flow over upper surface of canopy. Flow from left to right.

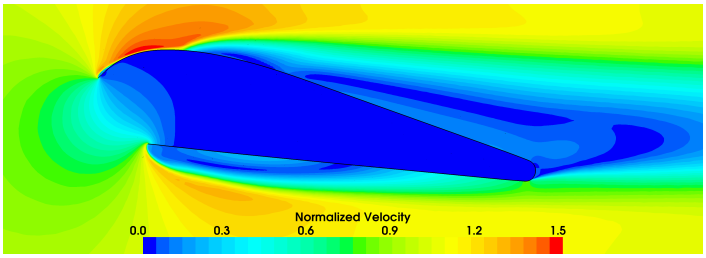
The streamlines in Figure 5.34 show a spanwise variation in the path of the flow within the interior. The streamlines near to the unloaded rib are seen to take a lower trajectory, and to not encroach as far into the interior domain. The streamlines within the half-span nearest the loaded rib take a higher trajectory upon crossing the threshold of the inlet, and travel further into the interior space. The distance travelled past the upper leading edge peaks around 75% of the span from the unloaded rib.

As previously observed, the reversal of the flow in the ram-air intake results in the flow exiting the inlet with a velocity opposing that of the freestream. The flow exiting the canopy interior also has to manoeuvre around the sharp leading edge of the lower surface. These two factors result in the flow separating from the lower leading edge, and a large separation bubble under the forward half of the lower surface occurs, Figure 5.34.

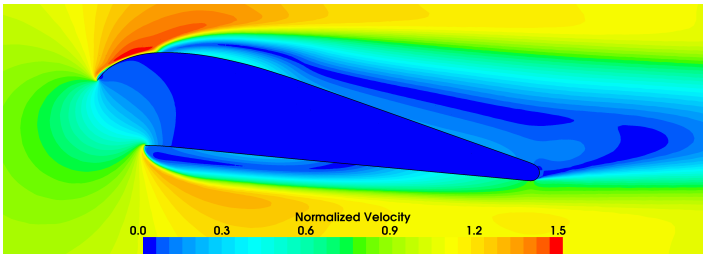
The spanwise variation in the path which the flow takes within the inlet can be seen to influence the curve of the streamlines upon exiting the cell interior, thereby influencing the size and shape of the separation bubble. This is suggested by the observation of variations in the outlines of the separation bubble displayed in the streamlines in Figures 5.34(a), (c) and (d). This variation is further displayed by comparison of the normalized velocity contours for cross-sections at different points along the span, provided in Figure 5.38. The low velocity region below the lower surface which delineates the separated region is seen to be thicker and more rounded in Figure 5.38(a), compared to (b) and (c) where the separation bubble becomes progressively thinner and flatter. This variation appears to be a consequence of both



(a) z=25%span



(b) z=50%span

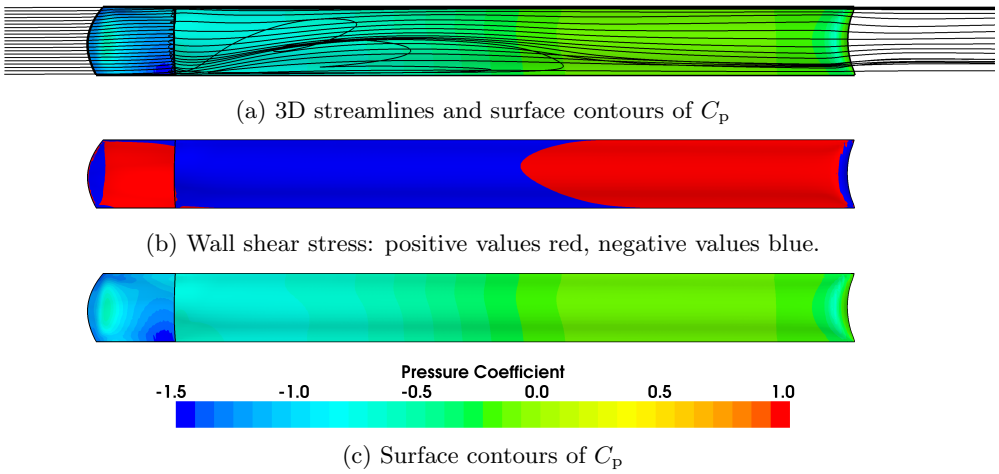


(c) z=75%span

Figure 5.38: Chordwise cross-sections at different spanwise positions showing contours of normalized velocity.

the spanwise curvature of the lower leading edge caused by the upwards movement of the unloaded rib, as well as the localised chordwise curvature of the lip of the lower surface, obvious in Figure 5.31(b). The associated pressure contours plotted in Figure 5.39(c), however, do little to explain the origin of the observed chordwise curvature. A plan view of the lower surface streamlines, Figure 5.39(a) shows a spanwise component to the recirculating fluid, a further consequence of the expansion of the fabric in this area.

Despite the observed differences in the separation bubble, no discernable spanwise change is seen in the pressure fields in the fluid adjacent to that region, Figure 5.36, nor is there any obvious spanwise differential in the pressure on the external surface of the lower half of the canopy, Figure 5.39(c). There is, however, a clear spanwise variation



(a) 3D streamlines and surface contours of C_p

(b) Wall shear stress: positive values red, negative values blue.

(c) Surface contours of C_p

Figure 5.39: Flow over lower surface of canopy. Flow from left to right.

in the reattachment point on the lower surface. The reattachment line outlined by the contrast in the wall shear stress contours in Figure 5.39(b) is asymmetric, and has a clear curve to it. The contours show that the separation bubble is shortest along a line just off-centre, towards the loaded rib. At this spanwise position the reattachment point is located at $x/c = 0.55$, a definite improvement over the results from Section 5.2. The point of reattachment shifts gradually rearward, when moving spanwise towards either rib, with the contours showing the separation bubble to be at its longest on the half closest to the unloaded rib. This coincides with the vertical movement of the unloaded rib and the subsequent cross-flow curvature in the lower surface; reattachment occurs sooner at the lowermost points of the lower surface.

The lower surface wall shear stress contours also highlight a final point of separation just prior to the trailing edge, which corresponds to the beginning of the curvature in the trailing edge region, a feature highlighted in Section 5.3.1. The separation at the trailing edge is likely to contribute to an increase in pressure drag, and is consistent with the observations of the flow around the rounded trailing edge in Section 5.2.2.

5.3.3 Lift and drag results

In order to assess and understand the consequences of the flow features identified above on the aerodynamic performance of the parafoil canopy, the lift and drag coefficients are analysed.

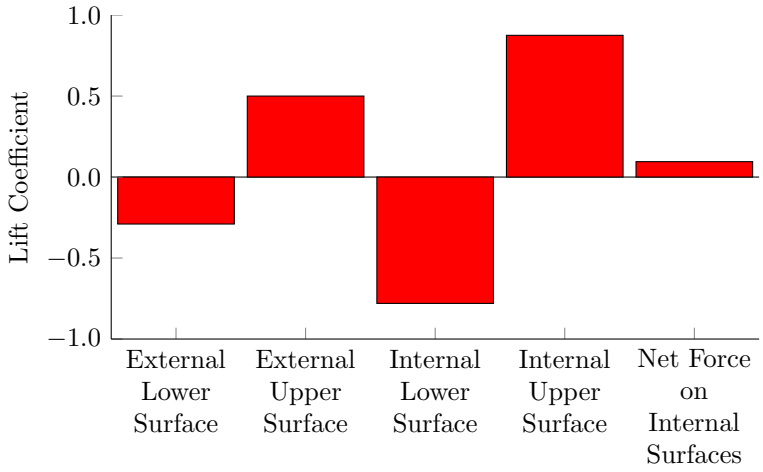


Figure 5.40: Lift coefficient values, for both internal and external surfaces, for inflated parafoil cell, separated by structural component.

The total lift coefficient value of 0.306 is a large reduction over the lift coefficient of the results from the flow around the 2D cut pattern of the parafoil, which has a lift coefficient of ∼1.1. This is not surprising, as the results in Section 5.3.1 show a large deviation of the shape from the initial design, and Section 5.3.2 revealed large differences in the flow behaviour as a result. The results also show a large increase in the drag coefficient; the 3D model has a coefficient of 0.132, much higher than the 2D cut pattern result of ∼0.05.

Examining the breakdown of these coefficient values in terms of the lift and drag generated at each different part provides some explanation for the levels of aerodynamic performance exhibited by the parafoil cell. Firstly, the breakdown of the lift by part is displayed visually in Figure 5.40. As first identified in Section 5.1.2, the presence of an open inlet and the corresponding area of high pressure internal to the structure, means consideration of the lift and drag forces on a parafoil canopy must include those seen on both the interior and exterior surfaces. The interior lift forces are seen to have large components in Figure 5.40 which, due to the difference in surface area between the upper and lower surfaces, sum to provide a small net lift force.

The majority of the lift force, as would be expected from the aerofoil-style form of the parafoil, is generated by the acceleration of flow over the curved upper surface, creating the suction peak identified in Section 5.3.2. Typically, at this angle of attack the upper surface lift would be supplemented by an additional contribution from the

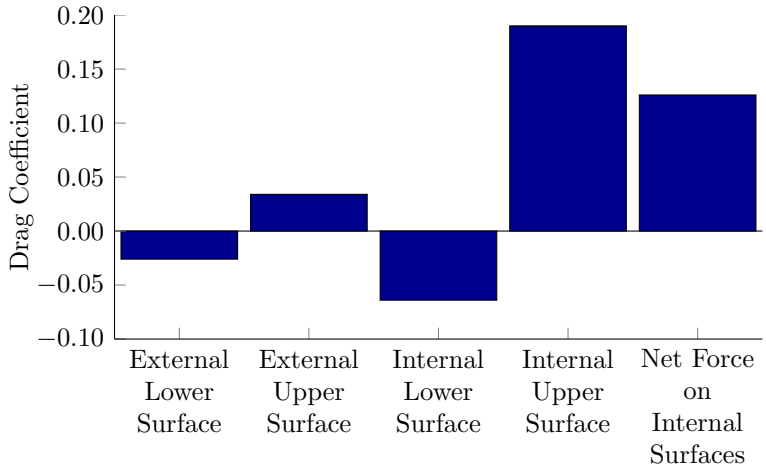


Figure 5.41: Drag coefficient values, for both internal and external surfaces, for inflated parafoil cell, separated by structural component.

lower surface, due to a build up of high pressure under the structure; such was the case in the 2D rounded leading edge design in Section 5.1.2, Figure 5.17. Analysis of the forces generated by the flow around the unstressed cut pattern demonstrated that separation from the lower leading edge generated an area of suction pressure, and therefore reduced the positive lift force, thereby reducing the total lift generated by the parafoil. In a similar manner, the large area of separation downstream of the lower leading edge in this 3D model creates a large area of negative pressure, Figure 5.39(c), that creates a large negative lift contribution on the lower surface. In contrast to the 2D cut pattern case, the 3D model has a much larger separated region, and hence the negative pressure is large enough for the net lift contribution at the lower surface to be negative. The negative lift resulting from the lower surface suction pressure counteracts over 50% of the lift generated at the upper surface. This shows how lower surface separation has a very strong impact on the lift generating capability of the parafoil canopy.

The aerodynamic forces in the drag direction on each part are plotted in Figure 5.41, and a breakdown of the drag forces by component, separated into their pressure and friction contributions, is provided in Table 5.6. Table 5.6 shows that the dominant drag mechanism for the canopy is pressure drag, as opposed to friction drag.

Consideration of the forces must include those acting on the interior surfaces, which arise from the hydrostatic pressure from the interior region of stagnated air. The

Table 5.6: Drag coefficient results for inflated canopy: pressure and shear components separated by component.

Part	Pressure	Shear	Net
Lower surface: Interior	-6.4×10^{-2}	-3.7×10^{-5}	-6.4×10^{-2}
Upper surface: Interior	1.9×10^{-1}	-5.3×10^{-6}	1.9×10^{-1}
Lower surface: Exterior	-2.6×10^{-2}	2.3×10^{-4}	-2.6×10^{-2}
Upper surface: Exterior	3.4×10^{-2}	4.6×10^{-4}	3.4×10^{-2}
Total	1.31×10^{-1}	6.5×10^{-4}	1.32×10^{-1}

Table 5.7: Comparison of simulation results to wind tunnel results, Re=2.1million. Corrected C_d values pertain to canopy only, line drag component has been removed.

Source	C_l	C_d	$C_\mathrm{d} - corrected$	$C_\mathrm{l}/C_\mathrm{d} - canopy$
Ware & Hassell M1.	0.58	0.19	0.135	4.15
Nicolaides	0.59	–	0.16	3.7
Simulation	0.306	–	0.132	2.32

lower interior surface in fact shows a drag gain, that arises because the angle of attack creates a streamwise component to the internal pressure force. However, the force from the stagnation pressure on the upper surface greatly outweighs this and the net interior contribution is a large positive drag value. The internal drag forces are shown to provide the largest contribution to the total drag on the structure by a large margin.

Examining the external forces reveals that whilst the low pressure region on the lower surface caused by separation reduces the net lift, it has a positive effect on the drag, and due to the angle of attack it contributes a drag reduction. The separation on the rearward half of the upper surface creates a large pressure drag force, but 75% of this is counteracted by the driving force seen on the external lower surface.

A comparison of the Lift and Drag results to those recorded in the wind tunnel by Ware & Hassell and Nicolaides [40] is found in Table 5.7. The results for the lift coefficient show a large difference to the experimental values, with the simulation results being around half that of experiment. The drag on the other hand shows very good agreement with Ware & Hassell, to within 2% of the experimental value. The low lift prediction results in a lower than expected value for the lift to drag ratio, by a factor of around 2. Interestingly, the results of Nicolaides for the same canopy and

the same Reynolds number disagree slightly with those of Ware & Hassell in terms of the drag, but the lift values are extremely close. This highlights the difficulties of finding reliable experimental data.

5.3.4 Summary and discussion

The single-cell parafoil model considered previously has been adjusted and improved to provide a more realistic simulation of a ram-air canopy. Once more the loosely coupled FSI methodology has been shown to successfully converge to an equilibrium shape, and the resulting prediction for the inflated structure was shown to be a closer representation to real-life. With the boundary conditions updated to account for the influence of half-cells within the canopy, the unloaded rib was seen to move vertically under the force of the suction present on the upper surface. The combination of unloaded and loaded ribs is seen to cause an asymmetric spanwise curvature in the cell, echoing the shape seen in a typical ram-air canopy. The difference in vertical movement along the length of the unloaded rib has also highlighted a possible difference in stiffness between the front and rear of the parafoil structure, as well as the concentration of high aerodynamic forces near to the nose region. Once more, the inflated shape is significantly different to the uninflated design, and changes first observed in Section 5.2, such as the rounding of the trailing edge, are also observed here.

Comparing the dimensions of the new inflated shape with experimental measurements revealed a close match, a distinct improvement over the original model. The inclusion of a spanwise loading results in a more realistic cell width, which in turn ensures a very close match to the experimental value of the maximum thickness of the parafoil. The improvement in material properties displayed a positive impact on the shape of the inlet, by preventing the unwanted elastic deformation previously recorded in that area. This has ensured a similarly good match to experiment for the size of the inlet both at and between the ribs.

The flow field results were shown to be considerably different to previous results, due to the asymmetry of the structure introduced by the presence of the unloaded rib. The suction peak on the upper surface is concentrated towards the side of the unloaded rib, which is likely a consequence of the greater vertical deflection at that side of the canopy. A distinct spanwise component to the flow was also observed over the upper surface, and both upper and lower surface flows deviated significantly from the 2D

cross-section case. Flow separation was reduced over the upper surface, but was still present over the rear-half of the upper surface. Previously observed lower surface separation was present in this case, resulting in a large region of suction pressure, but the length of separation was reduced over the previous results. This is a consequence of the reduced size of the inlet, but it is not immediately clear from these results to what extent the altered inlet shape, caused by the movement of the unloaded rib, contributed to this.

In terms of aerodynamic coefficients, the net lift has been shown to be greatly reduced by the suction pressure acting on the forward lower surface. The drag was shown to be formed predominantly of pressure drag on the internal surfaces, further emphasizing the importance of considering the internal forces. Comparison of the force coefficients to the experimental data of Ware & Hassell has shown a good agreement for drag, but the lift coefficient is under-predicted by about 50%. This descrepancy may be due to the ability of the chosen RANS turbulence model to predict the separation over the leading edge, and therefore the impact of this will be investigated in the next section. It is acknowledged that a more comprehensive comparison, made to experimental measurements of the flow field, such as surface pressures, rather than integrated forces such as lift and drag, would permit a more detailed and informed assessment of the accuracy of the results. Unfortunately, however, due to the previously discussed difficulties in making these measurements, they are not included in the work of Ware & Hassell, nor are they reported elsewhere in the literature, and so this has not been possible.

Furthermore, it is also important to recognize that it was necessary to make a number of assumptions in relation to the selected design of the canopy, such as an estimation of the width of the parafoil cell, and the lack of specification for the fabric material used for the Ware & Hassell canopy design. For these reasons, it is unlikely that the structural result, and associated fluid dynamics predictions, would identically replicate the experimental values. Therefore, the correlation between the results and the reported values is extremely encouraging.

A final point worthy of consideration, is the effect on the aerodynamics of neglecting end-cell effects, due to the symmetry conditions on the two side-walls of the fluid domain. In other words the simulated flow is two-dimensional, and therefore the effects of induced drag that will be seen in reality in a finite-span canopy are not included in

the analysis. This should mean a higher predicted value of lift, and a lower predicted level of drag, for the current case. However, it is interesting to note the previously discussed observations of Speelman et al.[47], who observed experimentally the occurrence of trailing edge separation on the upper surface that progressed from the centre of the canopy towards the tips of the span as the angle of attack increased. This would suggest that separation would be more likely to occur in this type of simulation, because it models the region away from the edges of the canopy. Following this line of reasoning, the ability for this style of simulation to represent the performance of a full canopy may be influenced by a balance between the effects of a greater chance of separation, versus those resulting from neglecting the effects of induced drag. Further work on this topic is required, to enable a calculation of the likely errors produced by these two factors.

5.4 Use of Standard k-ε model for simulation of 3D flow field around inflated parafoil

The previous section described the prediction of the inflated shape of a single half-cell of a ram-air parafoil and the simulation of the flow field around the parafoil structure. The steady RANS simulations used to compute both the fluid portion of the FSI iterations and the final fluid dynamic analysis of the flow, all utilised the Realizable k-ε (R.k-ε) turbulence model. It is well known that turbulence models are empirical relations, derived from experimental observations of the flow for particular geometries and flow types. Consequently, one turbulence model may give excellent results for one case, but be completely incorrect in another. Therefore, the influence of the choice of turbulence model on the results must be considered.

Due to the lack of detailed experimental data available for the flow around parafoils, a validation study for the turbulence model was not possible. As explained in Section 3.2.3, the turbulence model was chosen with the type of flow in mind, specifically one that was likely to feature separation and reattachment, and for demonstration of robustness and convergence. Validation studies against experimental observations of flow over a backward facing step also suggested that the R.k-ε model was better able to predict separation and flow reattachment distances than other eddy-viscosity models, including the Standard k-ε (S.k-ε) model [93, 100]. The results of the numerical analysis described here have shown that the flow around a parafoil canopy is characterised by a three dimensional flow field, incorporating strong separation caused by both adverse pressure gradients and flow around sharp edges, and subsequent reattachment. This pushes the capability of RANS modelling to its limit, and so the choice of turbulence model is very likely to have an influence on the results.

In an effort to quantify how much of an impact the choice of turbulence model is likely to have, results are presented here that detail the simulation of the flow around the previously predicted canopy shape from Section 5.3, using the Standard k-ε turbulence model in place of the previously employed Realizable k-ε model. This is not an exhaustive analysis, but the choice of turbulence models available for this geometry has been shown by experience of the author to be limited to the k-ε family, due to difficulties in achieving convergence when applying other models such as the Spallart-Allmaras and k-ω SST models. Therefore, the following observations will provide

an insight into the limit of the applicability of the steady RANS method for this case and the consistency of the acquired results. This will then inform future studies and ascertain the need for experimental data to support validation of the numerical methods.

5.4.1 Flow field results

The three dimensional streamlines for the flow field predicted using the Standard k-ε turbulence model in Figure 5.42 provide an overall view of the flow around the parafoil structure. It is immediately obvious from the figure that the flow pattern is very similar to that predicted using the Realizable k-ε model. The flow over the upper leading edge is seen to separate and reattach, as predicted in the R.k-ε case, Figure 5.42. After the reattachment downstream of the leading edge, the flow over the upper surface is seen in 5.42(b) to be drawn spanwise from the side of the loaded rib towards the unloaded rib, and an area of recirculation is generated over the peak of the upper surface on the loaded rib side. Downstream of this the flow is seen to separate, generating two obvious regions of recirculation over the rear upper surface of the cell.

A further similarity in the flow is seen at the inlet region of the parafoil, Figure 5.42(a), where flow is observed to enter the inlet towards the upper surface, and curve round to exit near the lower leading edge. This is then seen to cause a large separation bubble downstream of the lower leading edge, observed in 5.42(c) and 5.42(d), in exactly the same manner as the previous fluid case.

Whilst it is clear that the observed flow features are very similar in their type and their location, closer comparison between the two sets of results reveals a number of differences in the magnitudes and dimensions associated with the observed features. An example of this is found in Figure 5.43(c) where the low pressure region associated with the separation at the leading edge has a much larger magnitude of negative pressure than that predicted by the R.k-ε model. A stronger level of negative pressure is also recorded at the suction peak, located adjacent to the unloaded rib near to the point of maximum thickness. The location and size of the suction region is very similar to the R.k-ε predictions.

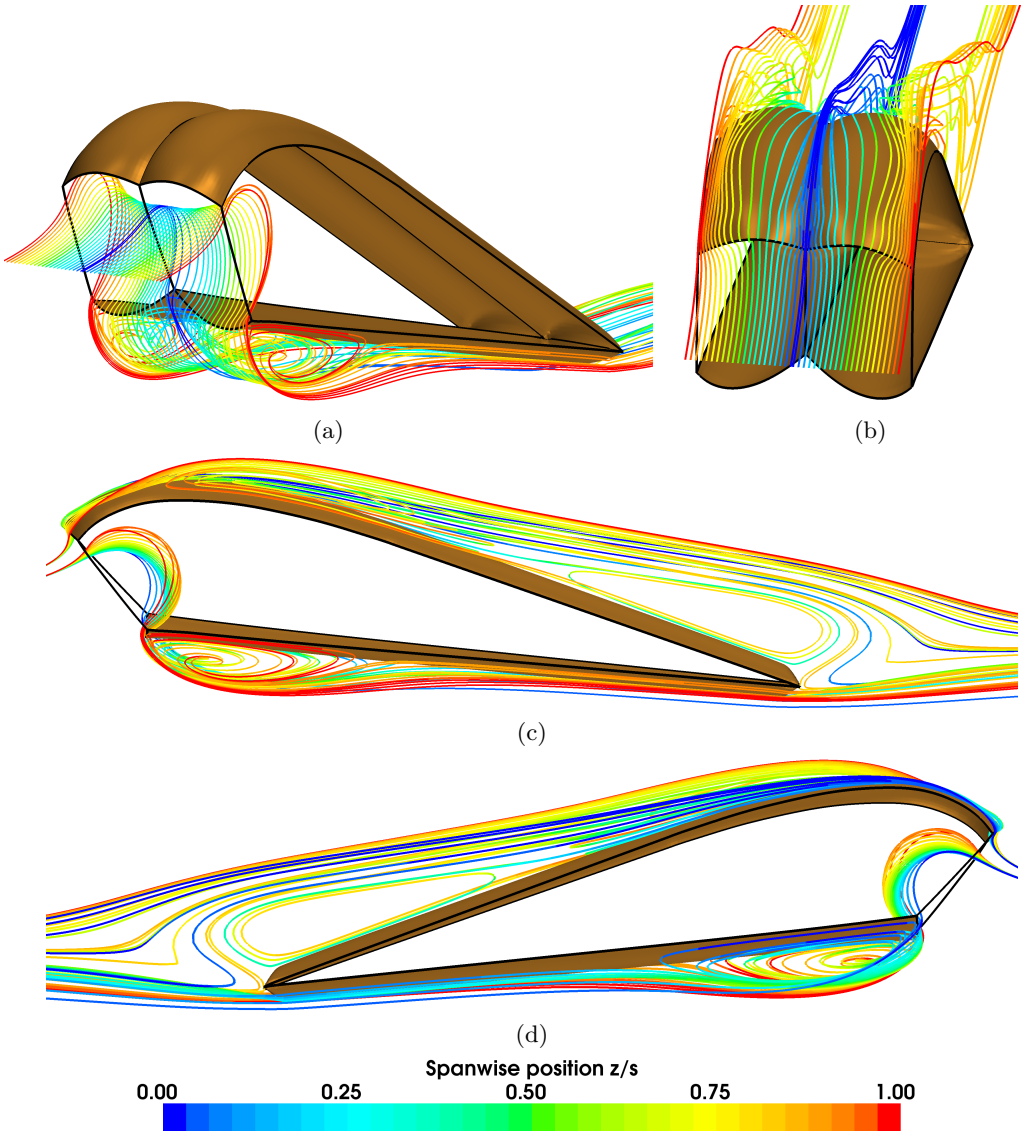
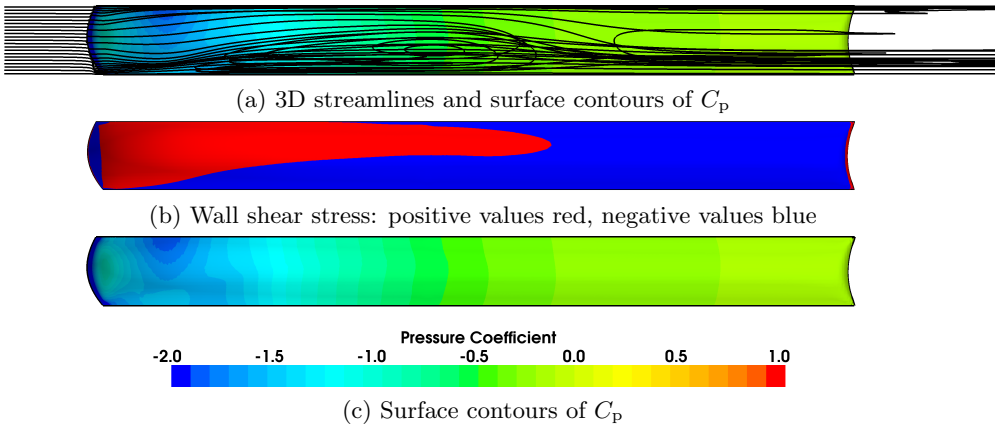


Figure 5.42: 3D streamlines for flow around inflated shape of ram-air parafoil cell, simulated using Standard k-ε turbulence model. Shown from (a) above-front, (b) front-side, (c) side 1, and (d) side 2. Streamlines coloured according to spanwise position, (e). Single cell is reflected to show symmetry, and ribs are omitted for clarity.



(a) 3D streamlines and surface contours of C_p

(b) Wall shear stress: positive values red, negative values blue

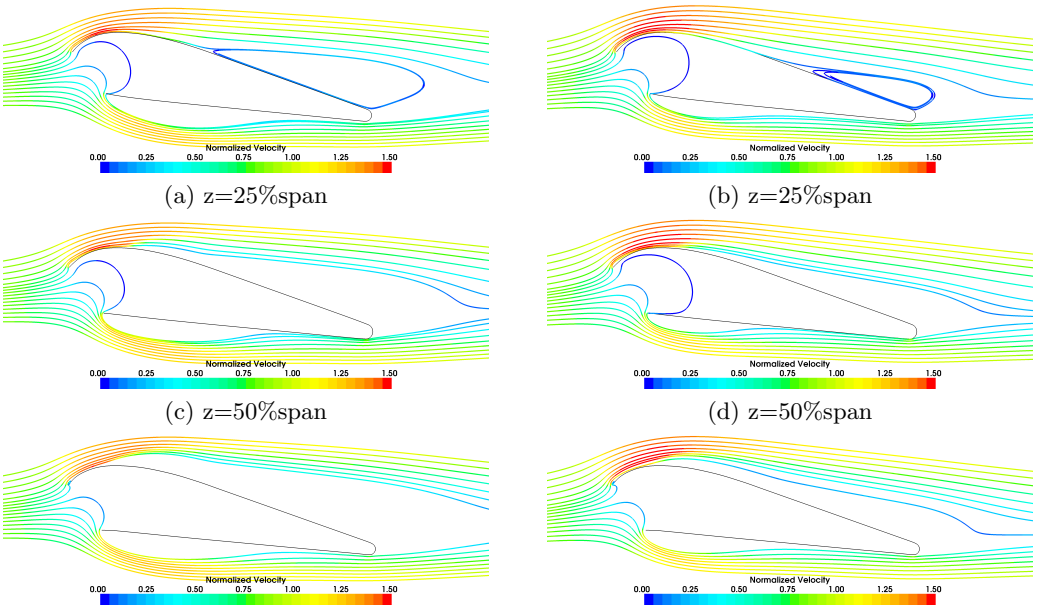
(c) Surface contours of C_p

Figure 5.43: Flow over upper surface of canopy predicted using S.k-ε model. Flow from left to right

Figure 5.44 provides a comparison of the constrained streamlines at different spanwise locations for the two cases. From this it is clear that the degree of separation predicted by the Standard k-ε model is greatly reduced compared to the Realizable model. This is the case on both the upper and lower surfaces. The spanwise variation in the flow is present in both cases, but the upper surface separation occurs further downstream for the S.k-ε model, resulting in a smaller recirculating area. Both models predict a large area of separation at the lower leading edge, but the separation bubble is smaller and reattachment occurs sooner when the S.k-ε model is employed.

The delay in separation from the upper surface in the results of the S.k-ε case is clearly visible in the wall shear stress contour plot in Figure 5.43(b), confirming the observations made from the streamline data. The point of separation varies across the span of the upper surface, and the flow in the S.k-ε case remains attached for much longer on the half of the span closest to the loaded rib than in the R.k-ε case. The observed separation at the upper leading edge is much the same in both cases, other than the aforestated increase in the negative pressure produced as a consequence. The area of recirculation present near to the point of maximum thickness of the cell, on the loaded rib side, extends further in the chordwise direction than previously, but is still aligned with the latest point of separation.

A further point of note regarding the flow over the upper surface is found in Figure 5.44, in which the normalized velocity contours show the acceleration of the flow to be greater over the peak of the upper surface for the S.k-ε case, indicated by higher



(a) z=25%span (b) z=25%span

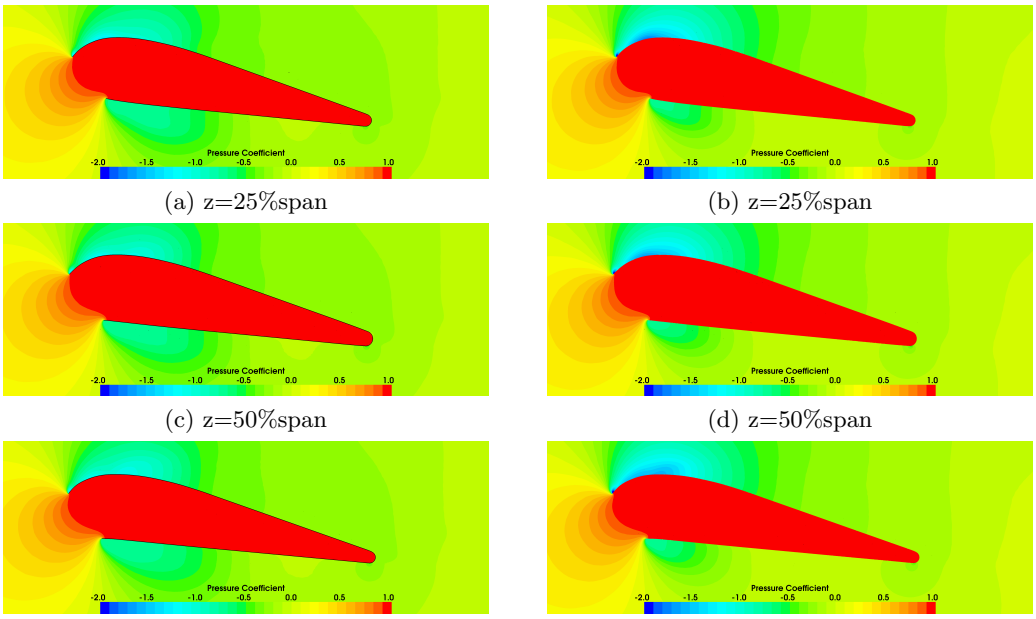
(c) z=50%span (d) z=50%span

(e) z=75%span (f) z=75%span

Figure 5.44: Chordwise cross-sections at different points across the span, showing constrained streamlines with contours of normalized velocity, predicted using (left) Realizable k-ε and (right) Standard k-ε turbulence models.

velocities over the thickest region of the parafoil. This explains the aforementioned increase in the magnitude of the negative pressure regions on the upper surface, that is also visible in the cross-section plots of C_{p} in Figure 5.45.

In addition to the lesser degree of separation seen on the upper surface, 3D streamlines of Figure 5.42 show reduced separation on the lower surface. The separation bubble downstream of the lower leading edge is reduced in size, in both the vertical and chordwise directions. The wall shear stress contours of Figure 5.46(b) show the reattachment line to be noticeably closer to the lower leading edge. The asymmetry of the reattachment point, in terms of the spanwise variation, is in agreement with the previous results. The influence of the reduced scale of the lower surface separation is evident in the flow field pressure values in Figure 5.45, where the corresponding low pressure region is noticeably smaller. Interestingly, the smaller recirculation region beneath the lower surface is shown by the streamlines in Figure 5.46(a) to have greater spanwise variation to it. Also of note is the separation just prior to the lower trailing edge, revealed by the wall shear stress values; a detail present in both sets of results.

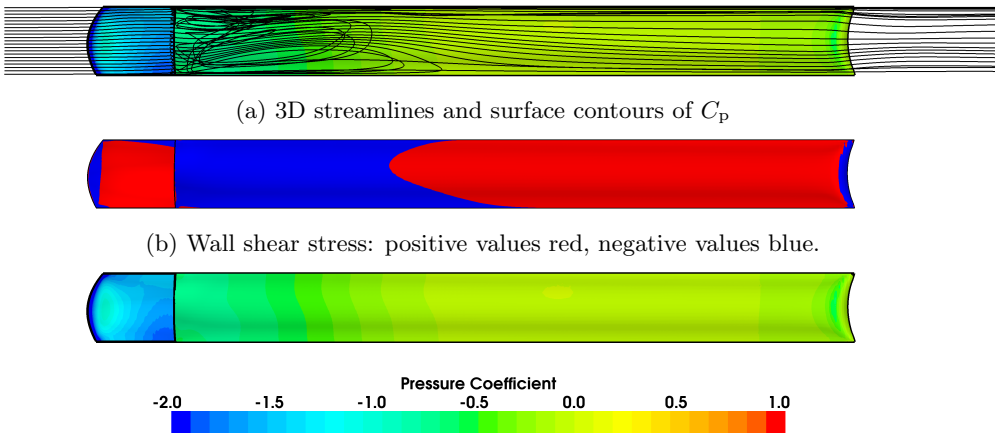


(a) z=25%span (b) z=25%span

(c) z=50%span (d) z=50%span

(e) z=75%span (f) z=75%span

Figure 5.45: Chordwise cross-sections at different points across the span, showing C_{p} predicted using (left) R.k-ε and (right) S.k-ε turbulence models.



(a) 3D streamlines and surface contours of C_{p}

(b) Wall shear stress: positive values red, negative values blue.

(c) Surface contours of C_{p}

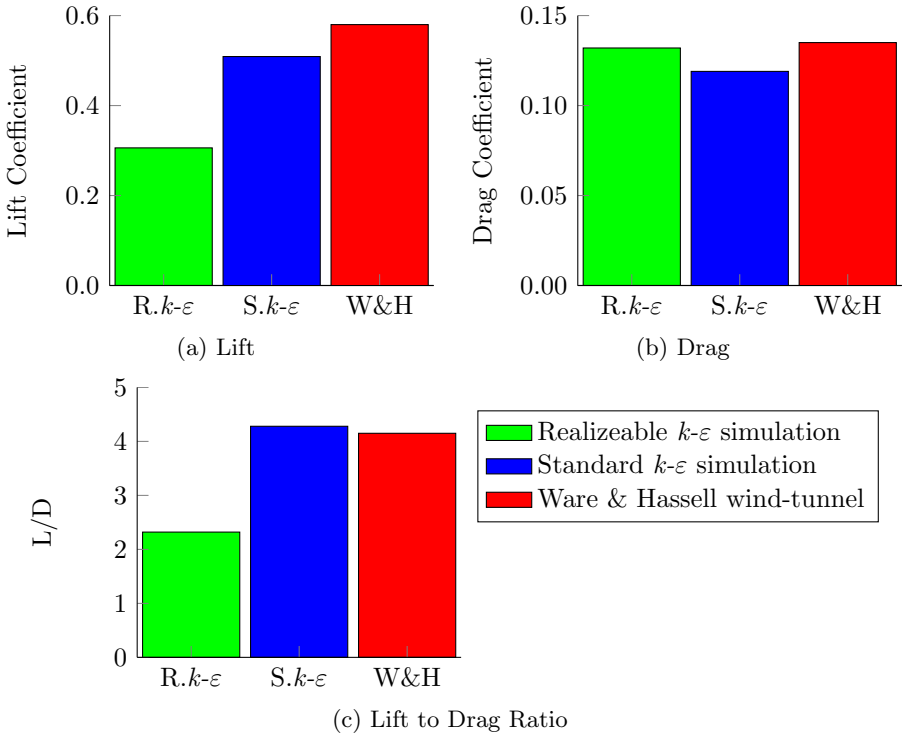
Figure 5.46: Flow over lower surface of canopy predicted using S.k-ε model. Flow from left to right

5.4.2 Lift and drag coefficient comparison

As a consequence of the differences detailed above between the results produced by the Standard and Realizable k-ε models, the resulting lift and drag values for the parafoil differ also. A comparison of the lift, drag and lift-to-drag ratio for the two models is plotted in Figure 5.47, along with the related experimental values from Ware & Hassell. The comparison shows an increase in the predicted lift to $C_{\mathrm{L}} = 0.509$ when using the S.k-ε model, which is closer to the experimental value of $C_{\mathrm{L}} = 0.58$ recorded by Ware & Hassell. In fact it is 12% less than the experimental value. The drag is predicted to be $C_{\mathrm{D}} = 0.119$ by the S.k-ε model, which is also 12% below the value measured in the wind tunnel, which is actually further from the value of $C_{\mathrm{D}} = 0.135$ observed by Ware & Hassell than that predicted by the Realizable model. The difference in the lift is much more substantial than the difference in the drag however, and with these changes combined they result in a lift to drag ratio of 4.28, around double that of the R.k-ε model, slightly more than the experimental value of 4.15.

Separating out the drag contributions for the different parafoil components in Figure 5.48 provides an insight into where the difference in force coefficients come from, and how they relate to the observed features in the flow. The drag produced on the external lower surface was actually seen to be negative in the R.k-ε case, and the same is true for the S.k-ε case too, but the magnitude is lower. The reduction in the negative drag value is due to the smaller separation bubble observed beneath the lower surface in the S.k-ε result, and therefore the reduced suction means less of a negative drag force.

The drag produced at the upper surface is also less in the S.k-ε result, which is due to the delayed separation. The reduction in pressure drag was seen to be significant, and it far outweighed the increase in skin friction drag. As would be expected, the contribution of the drag from the force on the interior surfaces of the parafoil is the same for both cases. Figure 5.48 shows this to be the most significant proportion of the drag by far, and it remains equal in both cases because it is a product of the pressure exerted in the internal surfaces by the high pressure air inside the parafoil cavity. As this air is stationary, the turbulence model will not have an effect on how its behaviour is predicted, and so no difference is observed. The net reduction in drag seen in the S.k-ε case is therefore a consequence of the reduction in pressure drag on the upper surface permitted by the later onset of separation. This reduction outweighs the increase seen on the lower surface.



(a) Lift

(b) Drag

(c) Lift to Drag Ratio

Figure 5.47: Comparison of predicted aerodynamic force coefficient values to to experimental measurements for two different RANS turbulence models.

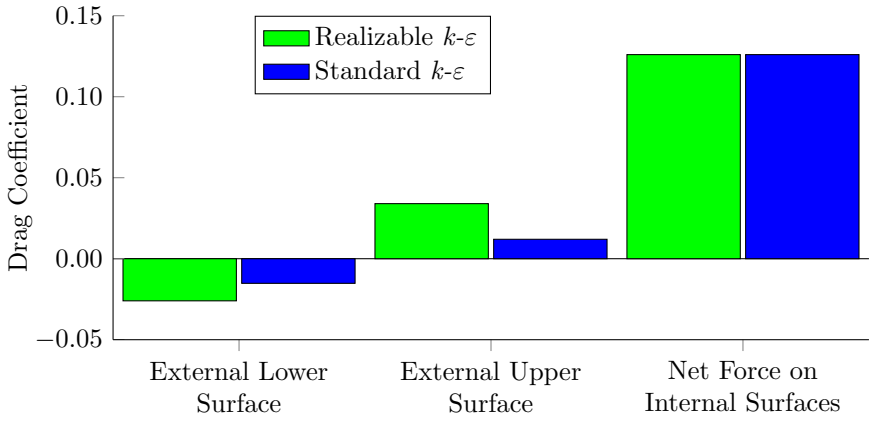


Figure 5.48: Comparison of drag coefficient values, by structural component, for inflated parafoil shape simulated using R.k-ε and S.k-ε models.

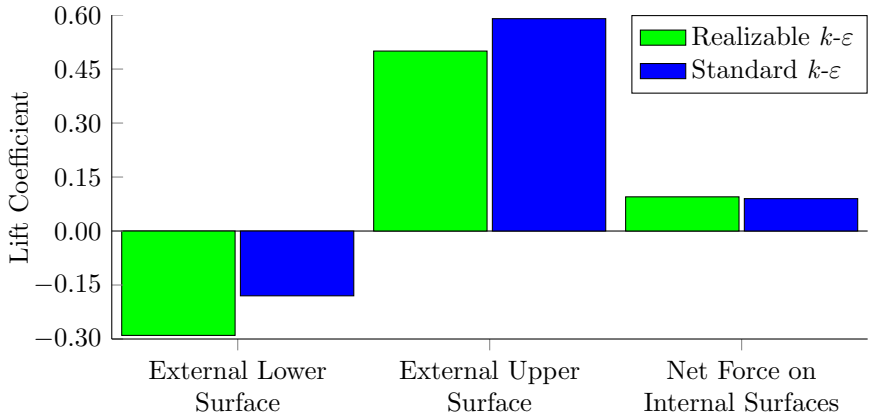


Figure 5.49: Comparison of drag coefficient values, by component, for inflated parafoil shape using R.k-ε and S.k-ε models.

Assessing the lift forces by component in Figure 5.49 provides similar insight. The difference in force at the exterior of the lower surface is due to the smaller separation bubble predicted by the S.k-ε model, lowering the associated suction force exerted on the lower canopy. This reduction in negative lift thereby reduces the degradation effect that lower surface separation has on the net value. Improvement of the lift performance is also noted at the upper surface, which shows a greater generated lift force predicted by the S.k-ε model. This is explained by the greater acceleration of the flow over the curvature of the upper surface and the corresponding increase in the maximum negative pressure exhibited at the suction peak. This is also likely to be positively influenced by the reduction in separation on the upper surface. Analogously to the drag forces, the interior lift forces are very similar for both models. Thus, the improvements at the two exterior surfaces culminate in a significant improvement in the predicted lift for the parafoil cell.

5.4.3 Summary and discussion

In this section, the influence of the choice of RANS turbulence model on the fluid dynamic results was assessed. This was achieved by performing a purely fluid dynamic simulation of the flow around the predicted inflated parafoil geometry generated in the previous section. The previously used turbulence model, the Realizable k-ε model, was replaced by the Standard k-ε model. The results of the simulation showed a very

similar flow pattern, with the same key flow features developing over both the upper and lower surfaces, in very similar positions relative to each other. The key differences were in the dimension or magnitude of these features.

The suction peak over the upper surface was, for example, seen to have a greater value of peak negative pressure, located in a similar position to the results obtained using the R.k-ε model. Separation of the flow from the upper surface was also shown to occur, similarly to the R.k-ε case. However, the separation point was delayed, occurring further downstream than in the previous results. Flow in the inlet region was shown to be unaffected by the change in turbulence model but this was not seen to be the case for the lower surface. The size of the separation bubble at the lower leading edge is reduced, and the flow reattaches sooner than for the R.k-ε case.

Due to the observed modifications in the flow, changes were seen in the aerodynamic coefficients when using the S.k-ε, with a small decrease in the drag coefficient, and a large increase in the lift coefficient. The decrease in the drag coefficient was shown to be due to the reduction in pressure drag on the upper surface, thanks to a reduction in the degree of separation over the rear of the upper part of the parafoil. The increase in lift was due to a reduction in negative lift on the lower surface due to reduced separation in that area, and an increase in lift over the upper surface due to a stronger suction peak in that region. Overall, the S.k-ε model has been shown to produce a greatly improved prediction of the lift coefficient, to within 12% of the experimental value, but the reduced drag value compares less favourably to experiment than previous results.

It is perhaps not surprising that the flow pattern in this case is similar, given that both models come from the same family of eddy-viscosity models. However, the comparison does confirm the importance of careful validation of the fluid dynamic simulation methods used, by highlighting the importance of accurately capturing the degree of separation that is observed. Given that lower surface separation is a consistent feature of all of the simulations within this work, including both 2D and 3D simulations, as well as cases in the literature such as Eslambolchi & Johari [53], the accurate prediction of separated flow is a significant issue for parafoil simulations. It is fair to say that this section has not provided an exhaustive assessment of the various different RANS closure models, but such a study would only be worthwhile if compared against a comprehensive experimental dataset, including flow field measurements such as a

map of surface pressures for an accurately measured parafoil shape. Needless to say, this is not currently available. It is also left to future work for the investigation of more sophisticated simulation strategies, with greater predictive capabilities, such as Detached Eddy Simulation (DES) or Large Eddy Simulation (LES) techniques, that rely less on the modelling of turbulence and more on simulation. Now that the benefits of CFD and FSI simulations have been demonstrated, the additional investment in research effort can be justified.

Chapter 6

Conclusions

This thesis has presented the development of a loosely coupled fluid-structure interaction methodology applied to the study of a ram-air parafoil during steady glide. The loosely coupled FSI framework focused on the modelling of a single-cell of a ram-air parachute such that it behaved as though it was part of a complete canopy, of infinite span. Both the coupled FSI methodology and standalone computational fluid dynamics have been used to analyse the typical flow behaviour and aerodynamic performance of a ram-air parafoil canopy, as well as to investigate the nature of the interaction between the flow and the canopy structure.

A summary of the results and associated conclusions relating to the three research objectives stated in Section 1.2 is presented below, followed by a list of the primary contributions of the thesis in Section 6.2 and a discussion of possible avenues for future work in Section 6.3.

6.1 Summary

Research Objective 1: Development of loosely coupled fluid-structure interaction simulation methodology

The challenges surrounding the computational simulation of the flow around a three-dimensional ram-air parachute canopy in flight, due to the difficulties present in acquiring an accurate representation of the inflated shape of the parafoil, were emphasized

in the background to this study as a key motivation for the work. To this end, a loosely coupled fluid-structure interaction approach was selected, and a methodology successfully developed in order to apply it to a ram-air parachute. Due to the computational requirements of coupled simulations a novel concept of representing the canopy as a chain of cells was introduced, so that the model could be reduced down to the representation of a single-cell, which responded as if it were part of a full canopy.

Several pivotal modelling features were required for the loosely coupled FSI methodology and single-cell approach to be successful. Inclusion of a wrinkling model for the parafoil fabric was integral, particularly in relation to accurately capturing the behaviour at the trailing edge, which was observed to buckle and wrinkle due to the reduction seen in the cell width. The necessity of using realistic material properties for the fabric material was demonstrated, particularly the values of the Young's modulus and Poisson's ratio. The incorporation of an under-relaxation step in the fluid-structure iteration cycle was also an essential inclusion to the methodology.

The single-cell approach was successfully implemented, and was shown to be a beneficial tool to reduce the complexity associated with the analysis of ram-air parafoils. As well as successfully producing a prediction of the inflated shape of a ram-air parafoil, the derivation and improvement of the single-cell model, particularly in terms of determining the appropriate structural boundary conditions, yielded a number of important conclusions regarding the manner in which different forces act to create the inflated shape. This included the importance of the spanwise loading on the cell width, and its subsequent impact on the section thickness, which was shown to greatly improve the accuracy of the inflated shape. The single-cell approach also revealed a tendency for the cell to expand spanwise near the trailing edge upon internal inflation, which may not otherwise have been observed. This may have negative repercussions on the full canopy shape and indicates that the planform may need to be thinned in the aft region of the cell, which may offer an additional benefit of saving weight. The use of the single-cell model permitted an in detail analysis of the key features of parafoil behaviour, both structurally and aerodynamically, which was more focused as a consequence of removing the complications surrounding simulating a full canopy.

Research Objective 2: Predict inflated shape of a parafoil and its associated flow field

The loosely coupled FSI framework discussed above was applied to the geometry

provided in a report by Ware & Hassell and a converged result for the inflated shape of the parafoil during steady glide was successfully produced. Two different sets of boundary conditions were utilised; the content of the second set was informed by the results from the first. In both cases, the three-dimensional shape that was produced varied considerably from the uninflated cut-pattern, reinforcing the need for the use of 3D simulation tools. The results of the simulations showed two main mechanisms behind the formation of the inflated shape. Firstly, the pressure forces cause the upper and lower surfaces to billow outwards, increasing the height of the cross-section, and creating a three-dimensional curvature to both surfaces. Secondly, the membrane tension in the fabric was seen to pull the ribs inwards, reducing the width of the cell, allowing further increases in the height of the cell and inlet.

Comparison of the cell height at both the inlet and thickest part of the parafoil with the observations of Ware & Hassell, showed the inflated shape of the cell to be higher than the experiment for the initial set of boundary conditions. This led to the addition of an outward spanwise loading to the improved set of boundary conditions, BC2. The effect of this was seen to be significant, and the value of the cell height compared exactly to the experimental value as a consequence, both at the inlet and at the point of maximum thickness. The cell width of the inflated model for the BC2 case was still lower than the experimental value despite the addition of a spanwise loading. This suggests that in the experiment by Ware & Hassell the spanwise components of lift generated in the outer cells due to the arc-anhedral of the canopies may act to further increase the canopy span. This additional loading was not included in the models, which were designed to represent a flat canopy. In both cases, the cell width reduces upon inflation which, in combination with the billowing results in a rounding of the trailing edge; a significant deviation from the uninflated shape that could have a negative influence on the aerodynamics.

The results produced using the first set of boundary conditions also highlighted the high amount of stress present in the ribs near the inlet, due to the high aerodynamic forces in that area, which caused the rib material in that region to stretch. This caused the inlet size to increase to a size that was much greater than the experimental measurement. Use of more accurate material data in the BC2 case prevented the stretching in this area and, along with the introduction of a spanwise loading force, improved the predicted shape such that the inlet size matched that of experiment. The forces on the canopy in this area were observed to be much higher than the rest

of the cell, but overall the stresses in the canopy structure were very low, justifying the use of the linear stress-strain data used to derive the values of the fabric's Young's modulus. This also highlighted the fact that the expansion of the upper and lower surfaces was due to membrane action but not material stretching.

In addition to improved material properties and the use of a spanwise loading, the adaptations made for boundary conditions BC2 to represent an unloaded rib were seen to further improve the prediction of the shape, and contributed to the closer comparison to experiment. The removal of the vertical constraint to create an unloaded rib meant that the rib was able to move upwards under the force of the suction peak on the upper surface. This is a key example of the interplay between the fluid forces and the structure. The upwards movement of the rib was seen to be greatest at the front of the canopy, where the aerodynamic forces were highest. The inclusion of the unloaded rib was shown to produce a more realistic shape to the canopy, particularly regarding the curved shape of the inlet.

The prediction of the inflated shape of the ram-air cell enabled the simulation of their associated flow fields. This was achieved for both the BC1 and BC2 cases. In both cases, the internal region was shown to contain a region of high-pressure stagnated air, which provided rigidity to the parafoil structure and prevented further air entering the cell. Some of the freestream flow was still able to enter the mouth of the ram-air inlet, but was forced by the high pressure within the cell to re-emerge near to the lower leading edge. In both cases, this led to large areas of separation beneath the lower surface of the parafoil.

Due to the inclusion of the unloaded rib in the BC2 case, the flow field over the upper surface of the structure was noticeably different to the BC1 case, as the upwards displacement of the unloaded rib introduced a spanwise component to the flow on the upper surface, resulting in the suction peak now being concentrated on the side of the cell nearest the unloaded rib. Separation was observed from the upper surface in both cases, which was a departure from the simulation of the flow around the uninflated cross-section of the parafoil, where the upper surface flow remained attached. Despite the changes introduced to the upper surface flow by the unloaded rib and the reduced maximum thickness of the cell, the flow in the BC2 case still detached from the upper surface around half way along the chord. However, the onset of separation from the upper surface was delayed compared to the BC1 model.

The flow around the lower surface was seen to separate in the 2D simulations around the unstressed parafoil cut pattern due to the combination of flow being unable to enter further inside the parafoil, and the presence of a sharp leading edge on the lower surface. In both 3D models lower separation was also seen but the improved shape of the inlet and cross-section thickness, due to the more realistic material properties and boundary conditions, significantly reduced the reattachment length in the BC2 model. The presence of the unloaded rib was shown to influence the reattachment on the lower surface, which occured further downstream in the area adjacent to the unloaded rib due to the curvature of the lower surface.

Due to the updated boundary and loading conditions in the BC2 model the lift and drag performance values were more realistic than for the BC1 model. Comparing the lift and drag values to those of Ware & Hassell revealed a close match to the predicted drag coefficient, which was within 2%, whilst the predicted lift was 50% lower than experiment. These are encouraging results given the assumptions and approximations that were made, including an estimation of the uninflated cell-width which was not given in the report, and the use of fabric properties from a completely different model of parafoil.

Finally, owing to the importance of the separated regions in the resultant aerodynamic coefficients, and because of the known difficulties with predicting separation using RANS, the influence of the chosen turbulence model on the results was investigated. This was achieved by performing a purely CFD simulation around the inflated parafoil shape, using the Standard k-ε model instead of the Realizable model. The use of the S.k-ε model was shown to produce a very similar flow field as for the Realizable k-ε model, but with differing magnitudes and dimensions to the key flow features. The altered flow field produced more lift, due to a shorter reattachment length on the lower surface, which was closer to the experimental value than the R.k-ε results. The lift value was closer to the value recorded by Ware & Hassell than the R.k-ε model, as it was 12% lower than experiment. The predicted drag value was slightly lower, however, due to a delayed separation from the upper surface, and therefore the drag coefficient was actually not as close to the recorded value as the previous result, but was again within 12%. The result has thus established a need for the fluid dynamics model to be carefully validated before CFD and FSI simulations can be relied upon for exact prediction of ram-air canopy aerodynamic coefficients. It may also be necessary to use more sophisticated simulation methods, a point discussed further in Section 6.3,

due to the previously stated difficulties of predicting separation.

Research Objective 3: Use simulation results to improve understanding of parafoil performance

In addition to the three-dimensional FSI simulations, two-dimensional CFD simulations were performed to analyse the flow around the uninflated cross-section, or cut pattern, of the chosen parafoil. Also, the difference in flow caused by varying the design of the leading edge was explored by simulating the parafoil with a closed inlet and a rounded leading edge. Analysis of the results revealed that the leading edge cut, typical of a parafoil, was by itself responsible for only small negative changes in lift and drag. The further addition of the open inlet, however, had a much more severe impact, strongly suggesting that the ram-air inlet rather than the blunt leading edge, is a significant contributor to the observed discrepancy between parafoil and aerofoil performance. The observed reduction in performance of the parafoil section over the aerofoil-style rounded leading edge section was significant; the lift coefficient was 16% lower and the drag was 2.07 times greater. This was extremely consistent with the observations from previous studies by Mohammadi & Johari and Balaji et. al.

The use of CFD simulations provided further insight into the underlying causes of the negative effects related to the presence of the ram-air inlet. The primary cause of the reduction in lift is demonstrated to be the separation from the lower leading edge, creating an area of suction under the lower surface, resulting in a negative component of lift force. The presence of an internal region of ram-air stagnation pressure was also confirmed in the simulation, and a strong relevance of this to the aforementioned increase in drag force was revealed; this was due to the action of the hydrostatic pressure acting on the internal surfaces. Consequently, the importance of the behaviour of the flow in and around the interior of the parafoil on its aerodynamics has been made clear. The simulation of the flow around this parafoil cross-section has also added to the small number of flow predictions for other designs available in the literature, and the aerodynamic performance is seen to compare favourably.

Similar conclusions can be drawn from the three-dimensional results, with lower surface separation seen to occur in both the BC1 and BC2 models. The BC2 lift and drag forces were seen to be strongly influenced by this, particularly the lift, which was degraded by the aforementioned lower surface suction. The results strongly suggest that lower surface separation is a typical flow feature for a ram-air canopy, in agreement

with Eslambolchi & Johari [53]. It appears that the size and angle of the inlet has a large effect on the magnitude of the separated region, which controls the magnitude of the lift loss due to the suction pressure on the lower surface. The influence of the size of the lower surface separation bubble on the generation of lift was further illustrated by the comparison of the Standard and Realizable k-ε models.

Consideration of the flow around the predicted 3D inflated shape revealed separation from the upper surface too, unlike the 2D models, demonstrating that the increased thickness of the canopy section due to inflation may have a negative impact on the aerodynamics, and for this reason three-dimensional simulations are required to accurately assess the flight performance of canopies. The upper surface separation in the BC2 model was shown to occur near to the loaded rib first, with the flow being completely separated at a distance of 50% chord. The flow does not reattach and the separation of the flow from the upper surface contributes significantly to the total drag, which was seen to be governed by pressure drag rather than friction drag, in contrast to a typical aerofoil.

Overall, a loosely coupled FSI simulation method has been proven to be a viable technique for the prediction of the inflated shape of a ram-air parafoil. Numerous benefits and advantages of applying computational analysis to the design of ram-air parachutes have been demonstrated, underlined by several hitherto unknown or undocumented features of parafoil behaviour, many of which would be difficult to observe through the use of experiment alone. Features such as the spanwise force supplied by the end cell ribs and the strength of the rib near the inlet have been identified to be crucial in the determination of the inflated shape, and should therefore always be considered in such an analysis. A single-cell has been successfully modelled to behave as though it was a component in a full canopy, reducing the computational overhead associated with FSI simulation. Insight into the relation between the parafoil shape and its associated flow field has been gained, and the incorporation of independent and coupled computational fluid-structure analysis into the design process for ram-air parachutes has been brought a step closer.

The results therefore show the loosely coupled framework to be a valid approach, and the single-cell representation is very promising for reducing the computational burden associated with the FSI simulations. It must be acknowledged, however, that the effect of the 2D nature of the flow on the flow field predictions, i.e. the neglect of

induced drag, needs further investigation, particularly with regards to the likelihood of separation in light of the observations of Speelman et al. [47]. The research has also highlighted the need for validation of the fluid dynamics models being used, especially the treatment of turbulence.

6.2 Primary contributions

The primary contributions of this thesis are as follows:

- Demonstrating that the presence of an open inlet has a greater impact on parafoil aerodynamics than that of the blunt nose arising from the associated leading edge cut.
- Development of a viable methodology for performing loosely coupled FSI simulations of ram-air parachutes.
- Presentation of a novel approach for performing the aforementioned FSI analysis in a computationally efficient manner, namely modelling a single-cell of a parafoil such that it behaves as though it was part of a full canopy.
- Successfully utilised FSI techniques to predict the inflated shape of a ram-air parachute and the surrounding flow field.
- Discovery of the importance of a spanwise force in determining the shape of the canopy.
- Improvement of derived boundary conditions for single-cell model, to include representation of unloaded ribs and spanwise loading.
- A computational investigation into the main sources of fluid-structure interaction and their influence on the resulting canopy shape or flow field.
- Established the potential benefits of including FSI simulation techniques in parafoil design.
- Highlighted the need for high quality, comprehensive experimental data for model validation, particularly in terms of the fluid dynamics, due to the detached nature of the flow field.

6.3 Recommendations for future work

The thesis presented above has shown the application of FSI simulation approaches, in this case a loosely coupled one, to be a valuable and a worthwhile endeavour for increasing the understanding of ram-air parafoil flight behaviour and performance. The multidisciplinary aspect of the work has meant that a number of aspects could not be explored to their fullest within the time available. Additionally, a number of potential improvements that could be made to the methodology, in the greatest scientific tradition, revealed themselves as a consequence of this first attempt. In this manner, the work provides a foundation for the commencement of a number of potential future studies. Several suggestions of what this future work could include are provided here.

Firstly, the most significant point raised by this work is one of validation. This is particularly pertinent for the fluid dynamics simulations, which are shown in Section 5.4 to be dependent on the choice of turbulence model. Therefore, if the current choice of CFD approach, i.e. the use of RANS, is continued, then a validation exercise is required to assist in the evaluation and selection of the most appropriate turbulence closure model. Given the results of Section 5.1, that demonstrate a clear difference between the flow field around an aerofoil and that of a parafoil, the data used to validate against needs to be specific to an open-inlet parafoil model, so that key flow features such as the upper surface separation and lower surface reattachment are captured. Given that such data does not currently exist in the literature, a detailed experimental characterisation of the flow around a parafoil shape would be extremely valuable.

Given the difficulties repeatedly highlighted in this work that surround the fulfilment of such a task, it is expected that the validation process would require a number of stages. Simulations could be validated initially for a 2D case, using rigidly constructed prism-style cross-sections in the wind-tunnel. The process could start with the unstressed cut-pattern of the parafoil, following which the variation in the 3D inflated cross-section observed in this study could be taken into account using prismatic cross-sections, based on the shapes observed at different points along the span in FSI simulations. Proceeding in this systematic way would permit a steady increase in the complexity of the flow field, whilst also providing a steady increase in confidence in the CFD results, which would in-turn benefit the later stages of the validation pro-

cess. Ultimately, for a validation of the 3D canopy geometry, a reliable method for replicating the predicted inflated shape would be required. It is speculated that the incorporation of 3D printing technology could provide a solution to this, enabling the 3D flow simulations to be formally validated.

In addition to the validation of the simulation of the fluid domain, the structural simulations performed in this study have used a number of assumptions which need investigating. Principally, this involves a more comprehensive characterisation of the fabric material, requiring a solution to the problem of slippage for the thin parachute fabric, and inclusion of the data into the FEM constitutive model. The results of the fabric tests in Section 4.4 displayed strong anisotropy, which was not included within the current analysis. It would also be of interest to investigate the influence of this on the resultant shape of the canopy, and determine if it could be used for optimisation of the shape. The same is true of the shear modulus, which would be expected to be much lower than the isotropic value used in this study. Once the requisite material testing is accomplished, implementation of more accurate material data is easily achieved within the current methodology, allowing the study of the effect of material properties to be analysed.

Closely related to the improvement of the material properties is the use of a complete and more detailed design for the canopy. A number of the details and dimensions of the parafoil have been approximated, such as the fabric used which was based on a completely different canopy model, and there is no knowledge of the placement of any stitching or structural reinforcements. Furthermore, the cross-section design used within this study is approximately 40 years old. If a more modern, up-to-date canopy design was used for which a full knowledge of the structural design features and types of material used is assured, then the fidelity of the simulations would be improved. The use of a more modern, high-performance parafoil design would also be of interest in terms of the flow field, as it would be presumed that a more-modern canopy would be of superior performance and therefore the flow would be better behaved, and the amount of separation would likely be reduced. This is likely to place less demand on the simulation method required for the fluid dynamics, and would provide an idea of the typical flow behaviour around more modern parafoil shapes.

A final aspect of the approach to be validated is the shape prediction produced by the FSI iterations. To make this practicable a wind-tunnel experiment would likely

have to be performed where a fabric model of a single parafoil cell was tested, but without an attempt to replicate the infinite-span canopy boundary conditions. If the experiment was designed such that it could be recreated exactly by using the loosely coupled FSI framework, then photogrammetry measurements on the experimental result could provide a reliable means by which the predicted shape could be compared to the physical result. Further to this, there is potential for 3D printing to be used to provide a rigid recreation of the measured or simulated shape, allowing for the associated fluid dynamic simulation results to be compared to experiment conjointly.

As well as validation, further investigation into the choice of simulation techniques applied within this work are of interest. In particular, the selection of the steady RANS method for CFD simulation warrants further consideration. As the flow results in Section 5.3 are shown to involve large areas of separated flow, the steady RANS approach is pushed to the limit of its validity. Therefore, a simulation method with a greater predictive capability may be required, such as the Detached Eddy Simulation (DES) or Large Eddy Simulation (LES) techniques, which rely less on empirical models for their treatment of turbulence. These methods are also unsteady, and so applying them to the simulation of flow around a parafoil is likely to give a more complete picture of the flow field. This would inform the decision as to the suitability of both the RANS method and the assumption of steady flow to the problem.

Further extensions to the study could involve utilising the loosely coupled FSI simulations to simulate canopies at different angles of attack. This could, for example, inform the development of novel control methodologies that take advantage of glide-slope control to improve system accuracy. It is also hoped that a validated methodology could be used as part of the development of new actuation concepts such as spoilers and flaps.

Lastly, the simulation techniques used here could be used to investigate the influence of leading edge cuts and angles, as well as cross-section thicknesses, on the aerodynamics. A number of potentially advantageous adaptations to the parafoil design, such as non-rectangular planforms, the internal placement of ties to prevent the rounded trailing edge, or a curved lower surface could also be researched using FSI to assess the likely benefit to performance.

References

[1] T. W. Knacke. Technical-Historical Development of Parachutes and Their Applications since World War I. In *9th AIAA Aerodynamic Decelerator and Balloon Technology Conference*, Albuquerque, USA, 7-9th October 1986.

[2] W. B. Pepper and R. C. Maydew. Aerodynamic Decelerators - An Engineering Review. *Journal of Aircraft*, 8(1):3–19, 1971.

[3] C. P. Knapp and W. R. Barton. Controlled Recovery of Payloads at Large Glide Distances, Using the Para-Foil. *Journal of Aircraft*, 5(2):112–118, 1968.

[4] D. Jalbert. Multi-cell Glide Canopy Parachute, 1964.

[5] M. Seeger, L. Muller, P. K. Carlsson, Y. Bury, Y. Pressigny, G. Lallemand, M. Vallance, R. Wheeler, R. Benney, R. Charles, S. Morton, and S. Tavan. Four-Powers Long Term Technology Projects: "Airflow Influence on Airdrop" and "2nd Precision Airdrop Improvements". In *18th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Munich, Germany, 23-26th May 2005.

[6] J. D. Nicolaides, R. J. Speelman, and G. L. C. Menard. A Review of Para-Foil Applications. *Journal of Aircraft*, 7(5):423–431, 1970.

[7] J. A. H. Potter and L. Giles. The United Kingdom's Air Drop Capability. Technical Report RTO-MP-AVT-133, NATO RTO, Neuilly-Sur-Seine, 2006.

[8] R. J. Benney, W. J. Krainski, P. Onckelinz, C. Delwarde, and L. Muller. NATO Precision airdrop initiatives and modeling and simulation needs. In *Fluid Dynamics of Personnel and Equipment Precision Delivery from Military Platforms*, number RTO-MP-AVT-133, pages KN2–1 – KN2–22. NATO RTO, Neuilly-sur-Seine, France, 2006.

[9] K. Bergeron, A. Fejzic, and S. Tavan. AccuGlide 100: Precision airdrop guidance and control via glide slope control. In *21st AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Dublin, Ireland, 23-26th May 2011.

[10] A. P. Taylor. An investigation of the apparent mass of parachutes under post inflation dynamic loading through the use of fluid structure interaction simulations. In *17th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Monterey, USA, 19-22th May 2003.

[11] J-C. Berland, S. Dunker, S. George, and J. Barber. Development of a Low Cost 10,000lb Capacity Ram-Air Parachute, DRAGONFLY Program. In *18th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Munich, Germany, 23-26th May 2005.

[12] R. Benney, K. Stein, W. Zhang, M. Accorsi, and J. Leonard. Controllable Airdrop Simulations utilizing a 3-D Structural Dynamic Model. In *15th CEAS/AIAA Aerodynamic Decelerator Systems Technology Conference*, Toulouse, France, 8-11th June 1999.

[13] J. W. Wegereef, F. Doejaaren, S. Benolol, and F. Zapirain. FASTWing CL flight tests with a high-glide ram-air parachute for 6000kg payloads. In *21st AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Dublin, Ireland, 23-26th May 2011.

[14] J. Potvin, K. Bergeron, G. Brown, R. Charles, K. Desabrais, H. Johari, V. Kumar, M. McQuilling, A. Morris, G. Noetscher, and B. Tutt. The Road Ahead: A White Paper on the Development, Testing and Use of Advanced Numerical Modeling for Aerodynamic Decelerator Systems Design and Analysis. In *21st Aerodynamic Decelerator Systems Technology Conference and Seminar*, Dublin, Ireland, 23-26th May 2011.

[15] B. Tutt, R. Charles, S. Roland, and G. Noetscher. Development of Parachute Simulation Techniques in LS-DYNA. In *11th International LS-DYNA Users Conference*, Detroit, USA, 2010.

[16] A. P. Taylor, B. Tutt, and J. Sanders. On the application of explicit finite element analysis and coupled fluid/structure simulations as they apply to escape and recovery systems. In *41st Annual SAFE symposium*, Jacksonville, USA, 2003.

[17] K. Stein, T. Tezduyar, S. Sathe, M. Senga, C. Ozcan, T. Soltys, R. Benney, R. Charles, and V. Kumar. Simulation of Parachute Dynamics During Control Line Input Operations. In *17th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Monterey, USA, 19-22th May 2003.

[18] J. H. Strickland and H. Higuchi. Parachute Aerodynamics: An Assessment of Prediction Capability. *Journal of Aircraft*, 33(2):241–252, 1996.

[19] P. R. Löhner. *Applied Computational Fluid Dynamics Techniques: An Introduction Based on Finite Element Methods*. Wiley, 2008.

[20] V. Kalro and T. E. Tezduyar. A parallel 3D computational method for fluid-structure interactions in parachute systems. *Computer Methods in Applied Mechanics and Engineering*, 190(34):321–332, 2000.

[21] H. Johari and K. J. Desebrais. Stiffness Scaling for Solid-Cloth Parachutes. *Journal of Aircraft*, 40(4):631–638, 2003.

[22] K. R. Stein. *Simulation and modeling techniques for parachute fluid-structure interactions*. PhD thesis, University of Minnesota, December 1999.

[23] B. A. Tutt, A. P. Taylor, J. Berland, and B. Gargano. The use of LS-DYNA to assess the performance of Airborne Systems North America Candidate ATPS Main Parachutes. In *18th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Munich, Germany, 23-26th May 2005.

[24] J. S. Lingard. The Performance and Design of Ram-Air Parachutes. Technical Report 81103, Royal Aerospace Establishment, Ministry of Defence, Farnborough, UK, 1981.

[25] J. S. Lingard. Ram-Air Parachute Design. In *13th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Clearwater Beach, FL, USA, 15-18th May 1995.

[26] T. W. Knacke. Parachute Recovery Systems Design Manual. Technical Report NWC TP 6575, Naval Weapons Centre, China Lake, USA, 1991.

[27] D. J. Cockrell. The Aerodynamics of Parachutes. Technical Report AGARD-AG-295, NATO Advisory Group for Aerospace Research and Development, 1987.

[28] E. Puskas. Ram air parachute design considerations and applications. In *8th AIAA Aerodynamic Decelerator Systems Technology Conference*, Hyannis, MA, USA, 2-4th April 1984.

[29] D. Poynter. *The Parachute Manual: A Technical Treatise on Aerodynamic Decelerators*. Para Pub., 1984.

[30] R. Charles. Personal Communication, March 2011.

[31] Spectra Fibers description and physical properties. `http://www.honeywell-advancedfibersandcomposites.com/products/fibers/`. Accessed: 27-12-2013.

[32] K. Bergeron, D. Ecklebe, K. McClure, H. Johari, T. Curlett, and B. Pitman. Parachute Suspension Line Drag Analysis. In *20th AIAA Aerodynamic Decelerator Systems Technology Conferences*, Seattle, WA, USA, 4-7th May 2009.

[33] G. M. Ware and J. L. Hassell. Wind-Tunnel Investigation of Ram-Air-Inflated All-Flexible Wings of Aspect Ratios 1.0 to 3.0. Technical report, NASA, 1969. TM SX-1923.

[34] J. Sobieski. The aerodynamics and piloting of high performance ram-air parachutes. Online, May 1994. Accessed 20/06/2013.

[35] F. W. Martin. Parafoil Aerodynamic Characteristics Derived from Flight Measured Suspension System Loads. In *15th AIAA/CEAS Aerodynamic Decelerator Systems Technology Conference*, Toulouse, France, 8-11th June 1999.

[36] A. Gavrilovski, M. Ward, and M. Costello. Parafoil Control Authority with Upper-Surface Canopy Spoilers. *Journal of Aircraft*, 49(5):1391–1397, 2012.

[37] M. Ward, A. Gavrilovski, and M. Costello. Flight Test Results for Glide Slope Control of Parafoil Canopies of Various Aspect Ratios. In *21st AIAA Aerodynamic Decelerator Systems Technology Conference*, Dublin, Ireland, 23-26th May 2011.

[38] M. Ward, A. Gavrilovski, and M. Costello. Glide Slope Control Authority for Parafoil Canopies with Variable Incidence Angle. *Journal of Aircraft*, 50(5):1504–1513, 2013.

[39] G. Brown. Tethered parafoil test technique. In *10th AIAA Aerodynamic Decelerator Systems Technology Conference*, Cocoa Beach, USA, 18-20th April 1989.

[40] J. D. Nicolaides. Parafoil Wind Tunnel Tests. Technical Report AFFDL-TR-70-146, Air Force Flight Dynamics Laboratory, June 1971.

[41] R. H. Geiger and W. K. Wailes. Advanced Recovery Systems Wind Tunnel Test Report. Technical Report NASA CR-177563, Pioneer Aerospace Corporation, August 1990.

[42] S. M. Burk and G. M. Ware. Static Aerodynamic Characteristics of three Ram-Air-Inflated Low-Aspect-Ratio Fabric Wings. Technical Report NASA TN D-4182, NASA Langley, September 1967.

[43] T. Tezduyar, V. Kalro, and W. Garrard. Advanced computational methods for 3D simulation of parafoils. In *15th CEAS/AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Toulouse, France, 8-11th June 1999.

[44] P. Askins, P. Zell, and J. Ross. Aerodynamic decelerator testing in the National Full-Scale Aerodynamics Complex. In *14th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, San Francisco, USA, 3-5th June 1997.

[45] C. Matos, R. Mahalingam, G. Ottinger, J. Klapper, R. Funk, and N. Komerath. Wind tunnel measurements of parafoil geometry and aerodynamics. In *36th AIAA Aerospace Sciences Meeting & Exhibit*, Reno, USA, 12-15th January 1998.

[46] C. K. Lee and P. Li. Geometric properties of parachutes using 3-D laser scanning. *Journal of aircraft*, 44(2):377–385, 2007.

[47] R. Speelman, C. A. Babish, R. J. Berndt, J. H. DeWeese, S. R. Metres, W. R. Pinnell, and L. L. Watson. Para-Foil Steerable Parachute, Exploratory Development for Airdrop System Application. Technical Report AFFDL-TR-71-37, US Air Force Flight Dynamics Laboratory, April 1972.

[48] J. Heeg, C. Spain, and J. Rivera. Wind Tunnel to Atmospheric Mapping for Static Aeroelastic Scaling. In *45th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Palm Springs, CA, USA, 19-22th April 2004.

[49] J. Furfaro, C. Hailey, V. Behr, J. Furfaro, C. Hailey, and V. Behr. An experimental study of parafoil sections. In *14th Aerodynamic Decelerator Systems Technology Conference and Seminar*, San Francisco, USA, 3-5th June 1997.

[50] S. Mittal, P. Saxena, and A. Singh. Computation of two-dimensional flows past ram-air parachutes. *International Journal for Numerical Methods in Fluids*, 35(6):643–667, 2001.

[51] R. Balaji, S. Mittal, and A. K. Rai. Effect of leading edge cut on the aerodynamics of ram-air parachutes. *International Journal for Numerical Methods in Fluids*, 47(1):1–17, 2005.

[52] M. A. Mohammadi and H. Johari. Computation of Flow over a High-Performance Parafoil Canopy. *Journal of Aircraft*, 47(4):1338–1345, 2010.

[53] A. Eslambolchi and H. Johari. Simulation of Flowfield Around a Ram-Air Personnel Parachute Canopy. *Journal of Aircraft*, 50(5):1628–1636, 2013.

[54] V. Kalro, S. Aliabadi, W. Garrard, T. Tezduyar, S. Mittal, and K. Stein. Parallel finite element simulation of large ram-air parachutes. *International Journal for Numerical Methods in Fluids*, 24(12):1353–1369, 1997.

[55] T. Tezduyar, V. Kalro, and W. Garrard. Parallel computational methods for 3D simulation of a parafoil with prescribed shape changes. *Parallel Computing*, 23(9):1349 – 1363, 1997.

[56] C. Ibos, C. Lacroix, A. Goy, and P. Bordenave. Fluid-Structure Simulation of a 3D Ram Air Parachute with SINPA software. In *15th CEAS/AIAA Aerodynamic Decelerator Systems Technology Conference*, Toulouse, France, 8-11th June 1999.

[57] C. Ibos, C. Lacroix, L. Chuzet, D. Granville, C. Ibos, C. Lacroix, L. Chuzet, and D. Granville. SINPA, a full 3D fluid-structure software package for parachute simulation. In *14th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, San Francisco, USA, 3-5th June 1997.

[58] Y. Coquet, P. Bordenave, G. Capmas, and C. Espinosa. Improvements in Fluid Structure Interaction simulations of parachutes using LS-DYNA. In *21st AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Dublin, Ireland, 23-26th May 2011.

[59] I. J. Oye. Development of Aerodynamic Analysis Software Tools within the FASTWing CL Project. In *20th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Seattle, WA, USA, 4-7th May 2009.

[60] T. Chatzikonstantinou and E. Puskas. Numerical analysis of three-dimensional non rigid wings. In *10th Aerodynamic Decelerator Systems Technology Conference*, Cocoa Beach, USA, 18-20th April 1989.

[61] T. Chatzikonstantinou. Problems in ram air wing modeling and their solution in the three dimensional simulation code PARA3D. In *15th CEAS/AIAA Aerodynamic Decelerator Systems Technology Conference*, Toulouse, France, 8-11th June 1999.

[62] H. Altmann. Numerical Simulation of Parafoil Aerodynamics and Dynamic Behaviour. In *20th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Seattle, WA, USA, 4-7th May 2009.

[63] K. Lu, M. Accorsi, and J. Leonard. Finite element analysis of membrane wrinkling. *International Journal for Numerical Methods in Engineering.*, 50(5):1017–1038, 2001.

[64] X. Liu, C. H. Jenkins, and W. W. Schur. Large deflection analysis of pneumatic envelopes using a penalty parameter modified material model. *Finite Elements in Analysis and Design*, 37(3):233–251, March 2001.

[65] M. Pruett, M. Accorsi, and R. D. Charles. Validation of Computational Structural Dynamics Models for Parachute Systems. In *20th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Seattle, WA, USA, 4-7th May 2009.

[66] S. M. Murman and S. A. Suresh. Modeling Effective Stiffness Properties of IAD Fabrics. In *21st AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Dublin, Ireland, 23-26th May 2011.

[67] A. Jarasjarungkiat, R. Wchner, and K.-U. Bletzinger. A wrinkling model based on material modification for isotropic and orthotropic membranes. *Computer Methods in Applied Mechanics and Engineering*, 197(68):773–788, January 2008.

[68] W. Wong and S. Pellegrino. Wrinkled membranes I: experiments. *Journal of Mechanics of Materials and Structures*, 1(1):3–25, 2006.

[69] R. Rossi, V. Renato, and E. Onate. FE Analysis of Membrane Systems Including Wrinkling and Coupling. In Eugenio Onate and Bern Kroplin, editors, *Computational Methods in Applied Sciences*, volume 3, pages 89–108. Springer Netherlands, 2005.

[70] A. Tessler, D. Sleight, and J. Wang. Nonlinear Shell Modeling of Thin Membranes with Emphasis on Structural Wrinkling. In *Structures, Structural Dynamics, and Materials and Co-located Conferences*. American Institute of Aeronautics and Astronautics, April 2003.

[71] F. G. Flores and E. Onate. Wrinkling and folding analysis of elastic membranes using an enhanced rotation-free thin shell triangular element. *Finite Elements in Analysis and Design*, 47(9):982–990, September 2011.

[72] W. Wong and S. Pellegrino. Wrinkled membranes III: numerical simulations. *Journal of Mechanics of Materials and Structures*, 1(1):63–95, 2006.

[73] S. Kang and S. Im. Finite Element Analysis of Wrinkling Membranes. *Journal of Applied Mechanics*, 64(2):263–269, June 1997.

[74] T. Raible, K. Tegeler, S. Lhnert, and P. Wriggers. Development of a wrinkling algorithm for orthotropic membrane materials. *Computer Methods in Applied Mechanics and Engineering*, 194(2124):2550–2568, June 2005.

[75] D. G. Roddeman. *Force transmission in wrinkled membranes*. PhD thesis, Eindhoven University of Technology, 1987.

[76] R. Rossi, M. Lazzari, R. Vitaliani, and E. Onate. Simulation of light-weight membrane structures by wrinkling model. *Int. J. Numer. Meth. Engng.*, 62(15):2127–2153, 2005.

[77] LSTC. LS-DYNA Keyword User's Manual, Version 971, R6.0.0, Vol.II. pp. 180-187, February 2012.

[78] A. Carney, C. Niezrecki, E. E. Niemi, and J. Chen. Parachute Strain and Deformation Measurements using Imaging and Polymer Strain Sensors. In *19th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Williamsburg, VA, USA, 21-24th May 2007.

[79] V. Gidzak. *A Method of Simulating Fluid Structure Interactions for Deformable Decelerators*. Phd thesis, University of Minnesota, 2010.

[80] W. Dettmer and D. Perić. A computational framework for fluid structure interaction: Finite element formulation and applications. *Computer Methods in Applied Mechanics and Engineering*, 195(4143):5754–5779, 2006.

[81] S. Rugonyi and K. J. Bathe. On finite element analysis of fluid flows fully coupled with structural interactions. *CMES- Computer Modeling in Engineering and Sciences*, 2(2):195–212, 2001.

[82] R. Charles, M. Accorsi, S. Morton, R. Tomaro, K. Stein, S. Sathe, and T. Tezduyar. Overview of the Airdrop Systems Modeling Project within the Collaborative Simulation and Test (CST) Common High Performance Computing Software Support Initiative (CHSSI) Portfolio. In *18th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Munich, Germany, 23-26th May 2005.

[83] H. G. Matthies and J. Steindorf. Partitioned strong coupling algorithms for fluid-structure interaction. *Computers & Structures*, 81(8–11):805–812, 2003.

[84] R. van Loon, P. D. Anderson, F. N. van de Vosse, and S. J. Sherwin. Comparison of various fluid-structure interaction methods for deformable bodies. *Computers & Structures*, 85(1114):833–843, 2007. Fourth MIT Conference on Computational Fluid and Solid Mechanics.

[85] T. Belytschko, W. K. Liu, B. Moran, and K. Elkhodary. *Nonlinear Finite Elements for Continua and Structures*. CourseSmart. Wiley, 2013.

[86] LSTC. *LS-DYNA*, 2011. Version R5.1.1 Revision 65543.

[87] B. Tutt, S. Roland, R. Charles, and G. Noetscher. Finite Mass Simulation Techniques in LS-DYNA. In *21st AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Dublin, Ireland, 23-26th May 2011.

[88] C. L. Tanner, J. R. Cruz, and R. D. Braun. Structural Verification and Modeling of a Tension Cone Inflatable Aerodynamic Decelerator. In *51st AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Orlando, FL, USA, 12-15th April 2010.

[89] J. O. Hallquist. LS-DYNA Theory Manual. 2006.

[90] PIA Technical Standard TS-100: Standardized Nomenclature for Ram-Air Inflated Gliding Parachutes. `http://www.pia.com/piapubs/TSDocuments/ts-100.pdf`, January 1984. Accessed: 27-12-2013.

[91] T. J. Chung. *Computational Fluid Dynamics*. Cambridge University Press, 2002.

[92] E. L. Houghton, P. W. Carpenter, S. Collicott, and D. Valentine. *Aerodynamics for Engineering Students*. Elsevier Science, 2012.

[93] J. Tu, G. H. Yeoh, and C. Liu. *Computational Fluid Dynamics: A Practical Approach*. Elsevier Science, 2007.

[94] S. B. Pope. *Turbulent Flows*. Cambridge University Press, 2000.

[95] M. A. Leschziner. *Encyclopedia of Aerospace Engineering*, chapter Reynolds-Averaged Navier-Stokes Methods. Wiley, 2010.

[96] M. A. Leschziner and D. Drikakis. Turbulence modelling and turbulent-flow computation in aeronautics. *The Aeronautical Journal*, 106(1061):349–383, 2002.

[97] C. T. Shaw. *Using Computational Fluid Dynamics*. Prentice Hall, 1992.

[98] T. H. Shih, W. W. Liou, A. Shabbir, Z. Yang, and J. Zhu. A new k-epsilon Eddy Viscosity Model for High Reynolds Number Turbulent Flows - Model Development and Validation. Technical Report TM-106721, NASA, 1994.

[99] D. C. Wilcox. *Turbulence Modeling for CFD*. DCW Industries, Inc, 3rd edition edition, 2006.

[100] S. J. Collie, M. G. Gerritsen, and M. J. O'Sullivan. Numerical simulation of the turbulent flow past upwind yacht sails. Accessed: 3/11/2013, Sep 2002.

[101] W. C. Lasher and P. J. Richards. Validation of Reynolds-Averaged Navier-Stokes Simulations for International America's Cup Class Spinnaker Force Coefficients in an Atmospheric Boundary Layer. *Journal of Ship Research*, 51(1):22–38, 2007.

[102] W. C. Lasher and J. R. Sonnenmeier. An analysis of practical RANS simulations for spinnaker aerodynamics. *Journal of Wind Engineering and Industrial Aerodynamics*, 96(2):143–165, 2008.

[103] B. Andersson, R. Andersson, L. Hakansson, M. Mortensen, R. Sudiyo, and B. van Wachem. *Computational Fluid Dynamics for Engineers*. Computational Fluid Dynamics for Engineers. Cambridge University Press, 2012.

[104] W. Rodi. Experience with Two-Layer Models Combining the k-epsilon Model with a One-Equation Model Near the Wall. In *29th Aerospace Sciences Meeting*, Reno, NV, 7-10th January 1991.

[105] J. H. Ferziger and M. Perić. *Computational Methods for Fluid Dynamics*. Springer London, Limited, 2002.

[106] S. Patankar. *Numerical Heat Transfer and Fluid Flow*. Series in computational methods in mechanics and thermal sciences. Taylor & Francis, 1980.

[107] M. Schäfer, D. C. Sternel, G. Becker, and P. Pironkov. Efficient Numerical Simulation and Optimization of Fluid-Structure Interaction. In Hans-Joachim Bungartz, Miriam Mehl, and Michael Schäfer, editors, *Fluid Structure Interaction II*, volume 73 of *Lecture Notes in Computational Science and Engineering*, pages 131–158. Springer Berlin Heidelberg, 2010.

[108] M. Schäfer and I. Teschauer. Numerical simulation of coupled fluid-solid problems. *Computer Methods in Applied Mechanics and Engineering*, 190(28):3645–3667, 2001.

[109] LSTC. *LS-PrePost*, 2012. Version 3.2 64-bit.

[110] G. P. Guruswamy. A review of numerical fluids/structures interface methods for computations using high-fidelity equations. *Computers & Structures*, 80(1):31–41, 2002.

[111] The MathWorks Inc. *MATLAB Release 2012b*. Natick, MA, USA.

[112] Altair Engineering Ltd. *HyperMesh*. Royal Leamington Spa, UK, 2013. Version 11.

[113] CD-adapco. *STAR-CCM+*. London, UK. Version 7.04.011 double precision.

[114] Oracle Corporation. *Java programming language*. Redwood Shores, CA, USA. Version 1.6.0_27.

[115] M. Accorsi, K. Lu, J. Leonard, R. Benney, and K. Stein. Issues in parachute structural modelling: Damping and Wrinkling. In *15th CEAS/AIAA Aerodynamic Decelerator Systems Technology Conference*, Toulouse, France, 8-11th June 1999.

[116] P. Wriggers. *Nonlinear Finite Element Methods*. Springer, 2010.

[117] G. Taubin. A signal processing approach to fair surface design. In *Computer Graphics Proceedings*, pages 351–358, 1995.

[118] R. van Loon. *A 3D method for modelling the fluid-structure interaction of heart valves*. PhD thesis, Eindhoven University of Technology, 2005.

[119] Pisa Italy Visual Computing Lab, CNR-ISTI. MeshLab, Ver.1.3.2. Available at `http://www.meshlab.sourceforge.net/`, August 2012.

[120] AerAzur, 2013. Accessed: 1/12/2013.

[121] E. J. Barbero. *Introduction to Composite Materials Design, Second Edition*. Composite Materials. Taylor & Francis, 2010.

[122] ISO 13934-1:2013 Part 1: Determination of maximum force and elongation at maximum force using the strip method, 2013.

[123] B. Perrine. Personal Communication, April 2013.

[124] R. Turton. *The Physics of Solids*. Oxford University Press, 2000.

Appendix A

Parafoil cross-section co-ordinates

The co-ordinates for the parafoil cross-section provided by Ware & Hassell in [33] are repeated in Table A.1. The adjusted co-ordinates used throughout this study are found in Table A.2. The coordinates for the rounded leading edge geometry created for the 2D flow simulations in Section 5.1 are provided in Table A.3.

Table A.1: Original Ware & Hassell parafoil cross-section coordinates.

x/c co-ordinate	y co-ordinate	
	Upper	Lower
0.0000	0.0886	0.0886
0.0200	0.1206	0.0722
0.0400	0.1392	0.0578
0.0600	0.1518	0.0434
0.0800	0.1616	0.0290
0.1000	0.1688	0.0145
0.1056	0.1702	0
0.1200	0.1730	0
0.1400	0.1756	0
0.1600	0.1780	0
0.1800	0.1784	0
0.2000	0.1790	0
0.2200	0.1790	0
0.2400	0.1768	0
0.2600	0.1752	0
0.2800	0.1730	0
0.3000	0.1708	0
0.3200	0.1676	0
0.3400	0.1646	0
0.3600	0.1606	0
0.3800	0.1562	0
0.4000	0.1524	0
0.4200	0.1476	0
0.4400	0.1426	0
1.0000	0	0

Table A.2: Adjusted parafoil cross-section coordinates used throughout this thesis.

x/c co-ordinate	y co-ordinate	
	Upper	Lower
0.0000	0.0886	0.0886
0.0200	0.1206	0.0722
0.0400	0.1392	0.0578
0.0600	0.1518	0.0434
0.0800	0.1616	0.0290
0.1000	0.1688	0.0145
0.1056	0.1702	0.0120
0.1200	0.1730	0
0.1400	0.1756	0
0.1600	0.1780	0
0.1800	0.1784	0
0.2000	0.1790	0
0.2200	0.1790	0
0.2400	0.1768	0
0.2600	0.1752	0
0.2800	0.1730	0
0.3000	0.1708	0
0.3200	0.1676	0
0.3400	0.1646	0
0.3600	0.1606	0
0.3800	0.1562	0
0.4000	0.1524	0
0.4200	0.1476	0
0.4400	0.1426	0
1.0000	0	0

Table A.3: Rounded leading edge geometry cross-section coordinates used within 2D study, Section 5.1.

Upper surface		Lower surface	
x/c co-ordinate	y co-ordinate	x/c co-ordinate	y co-ordinate
0.0000	0.0886	0.0000	0.0886
0.0200	0.1206	-0.0011	0.0848
0.0400	0.1392	-0.0020	0.0810
0.0600	0.1518	-0.0031	0.0752
0.0800	0.1616	-0.0038	0.0714
0.1000	0.1688	-0.0044	0.0675
0.1056	0.1702	-0.0048	0.0616
0.1200	0.1730	-0.0044	0.0538
0.1400	0.1756	-0.0039	0.0479
0.1600	0.1780	-0.0029	0.0420
0.1800	0.1784	-0.0019	0.0382
0.2000	0.1790	0.0000	0.0326
0.2200	0.1790	0.0014	0.0308
0.2400	0.1768	0.0024	0.0291
0.2600	0.1752	0.0035	0.0274
0.2800	0.1730	0.0046	0.0257
0.3000	0.1708	0.0058	0.0241
0.3200	0.1676	0.0085	0.0211
0.3400	0.1646	0.0100	0.0198
0.3600	0.1606	0.0207	0.0143
0.3800	0.1562	0.0300	0.0116
0.4000	0.1524	0.0402	0.0096
0.4200	0.1476	0.0501	0.0080
0.4200	0.1476	0.0601	0.0067
0.4400	0.1426	0.0701	0.0054
|	|	0.0800	0.0043
|	|	0.0900	0.0033
|	|	0.1000	0.0022
|	|	0.1060	0.0016
|	|	0.1100	0.0011
|	|	0.1200	0.0000
1.0000	0.0000	1.0000	0.0000

Appendix B

Examples of scripts and macros employed in loose coupling framework

A number of scripts and macros are used to automate the process of performing the loosely coupled FSI simulations, and to ensure that the exchange of data between the fluid and structure domains is carried out correctly. A description of how these combine to complete a fluid-structure iteration is provided here, followed by some code examples covering some of the more individual approaches that have been used to achieve the coupling.

B.1 Workflow example for fluid-structure iteration

A typical workflow example for a fluid-structure iteration is as follows:

1. **Structural simulation is run in LS-DYNA.**
   Simulation started by a script run in MATLAB.
2. **Output of deformed geometry and under-relaxation**
   (a) Deformed geometry is output using a command file for LS-PrePost, executed by MATLAB.

(b) Under-relaxation is performed in MATLAB by reading in the nodal coordinates of the new and old versions of the structural geometry, and then using Equation (3.19) to create a biased combination of the two.

3. **Create fluid boundaries around under-relaxed geometry**
   A MATLAB script executes a command file for LS-PrePost. This combines the under-relaxed geometry with a pre-prepared geometry for the fluid boundary, creating one geometry file.

4. **Output centroid coordinates for elements on upper and lower surface of under-relaxed geometry**
   MATLAB executes a command file in LS-PrePost that outputs the centroid coordinate values as a text file, ready to be used to extract the surface pressures from the results of the fluid simulation.

5. **Align side walls of fluid boundary with parafoil rib, and convert to .stl format**
   This step is performed using Altair HyperMesh, which is controlled by a TCL script executed by MATLAB, and consists of three main actions:

   (a) Position of rib in structural geometry is measured and fluid boundaries are adjusted to match.

   (b) All four-sided elements are split into three-sided triangular elements.

   (c) Geometry consisting of structure shape and fluid boundaries is exported as an .stl file.

6. **Pre-processing of fluid model**
   This is controlled via a Java macro, that does the following:

   (a) Import .stl file of structural geometry and fluid boundaries into STAR-CCM+.

   (b) Create surface and volume meshes of fluid domain using pre-determined mesh parameters.

   (c) Assign appropriate boundary conditions.

7. **Run fluid simulation in STAR-CCM+**

8. **Post-process fluid model to extract surface pressures**
   A Java macro is executed in STAR-CCM+ upon completion of the fluid simulation. This carries out the following tasks:
   (a) Read in coordinates of the structural geometry element centroids
   (b) Create point probes at each centroid position
   (c) Extract surface pressure value at each point
   (d) Write surface pressures directly to a loading file for use in the next structural simulation
9. **The next structural simulation is run using the newly created pressure loading, and the iteration loop starts again**

Several examples of the code used are given in the following sections to illustrate how the coupling was achieved.

B.2 Scripts for post-processing structural geometry

The post-processing procedure for the structural simulations contain several different steps, as described above. A number of these are performed using simple, standard matrix or array operations, acting on an array of values that have been imported from a text file, such as the under-relaxation step. Further details of these are not covered here, as they are commonplace programming operations.

Several of the steps, however, involve using MATLAB to execute files in other programmes, namely LS-PrePost or HyperMesh, which are slightly more unusual steps and so some examples of how this is done are given in this section.

Firstly, all instances of calling other programmes from MATLAB are done in the same way, using the `!` command to execute the following syntax in the shell. An example of this would be

```
! lspp32 c=export_deformed_geom.cfile
```

which executes the command `lspp32 c=export_deformed_geom.cfile` in the command line, which starts LSPrePost and executes the command file called "`export_deformed_geom.cfile`".

For consistency, the command files that are executed are written by MATLAB prior to their execution, to ensure that all of the destination paths are correct for the current iteration. The command file that controls the export of the deformed geometry, Φ_{new}, is written and executed within MATLAB using the following code.

```
%variable destination is created by the script to provide a path to
    write the command file to.
%variable structureNo is also created by the script according to the
    current iteration number.

%create the file
fid=fopen('export_deformed_geom.cfile','w');

%Write the command file
fprintf(fid, '$# cfile created by matlab, to export deformed geometry\n
    ');
fprintf(fid,'open d3plot "d3plotaa"\n');
fprintf(fid,'ac\n');
fprintf(fid,'state -1\n');
fprintf(fid,'selectpart on S1/0 S2/0 S3/0 S4/0\n');
fprintf(fid,'output "%s/deformed-geom-%s.k" 253 1 0 1 1 1 0 0 0 0 0 0 0
     0 0 1.000000\n',destination,int2str(structureNo));
fprintf(fid,'exit\n');

%close all files
closeresult=fclose('all');

%use shell prompt to run lsprepost with the newly created cfile
! lspp32 c=export_deformed_geom.cfile
%
```

Similarly, the command file controlling the output of the element centroids from the under-relaxed geometry, is written and executed in MATLAB as follows:

```
%variable targetfile is the path to the under-relaxed structural
    geometry file, created earlier in
% the script.
% variable destination is created earlier in the script to provide a
    path for the command file and
% deformed geometry to be written to.

%create the file
fid=fopen('deformed_centroid_output.cfile','w');
```

```
%write the file
fprintf(fid,'$# cfile created by matlab, to export deformed geometry
    centroids\n');
fprintf(fid,'open keyword "%s"\n',targetfile);
fprintf(fid,'selectpart on S1/0\n');
fprintf(fid,'output "%s/deformed_upper_surface_centroids" 1 1 0 1 0 0 0
     0 0 0 0 0 0 1 0 1.000000\n',destination);
fprintf(fid,'genselect clear\n');
fprintf(fid,'selectpart on S2/0\n');
fprintf(fid,'output "%s/deformed_lower_surface_centroids" 1 1 0 1 0 0 0
     0 0 0 0 0 0 1 0 1.000000\n',destination);
fprintf(fid,'exit\n');

%close all files
closeresult=fclose('all');

%
%use shell prompt to run lsprepost with the newly created cfile
! lspp32 c=deformed_centroid_output.cfile
%
```

Finally, the last stage of the structural post-processing procedure requires running a TCL script in HyperMesh. The script is executed by executnig a command file in the same manner as for LS-PrePost. The command file contains the command `*evaltclscript()` to execute the TCL script. Some key elements of the TCL script are provided below.

```
#measure distance between position of rib, and position of side wall
    fluid boundary

#Select the two nodes to measure
*createmark node 1 300188
set node1 [hm_getmark node 1]
*createmark node 2 90466
set node2 [hm_getmark node 2]

#find z coordinates of the two nodes
set node1_Z [hm_getentityvalue nodes $node1 z 0]
set node2_Z [hm_getentityvalue nodes $node2 z 0]
tk_messageBox -message "Node $node1 = $node1_Z ; $node2 = $node2_Z"
```

```
#Calculate the z-distance between the two nodes
set distz [expr ($node2_Z - $node1_Z)]
tk_messageBox -message "Difference between $node1 and $node2 is: $distz
    "

#Select the nodes on the side wall fluid boundary that need translating
*createmark node 2 300002 300004 300006 300008 300010 300012 300014

#translate the nodes by the required distance in the z-direction
*createvector 1 0.0000 0.0000 1.0000
*translatemark nodes 2 1 $distz

#Convert to triangles and export as .stl

#split shell elements into triangles on the ram-air cell
*createmark elements 1 "displayed"
*splitelements 2 1

#export structural geometry and fluid boundary to .stl file
*displaycollectorwithfilter components "off" "6" 1 0
*retainmarkselections 0
*entityhighlighting 1
*createstringarray 1 "CONNECTORS_SKIP"
*feoutputwithdata "/home/nick/altair/templates/feoutput/misc/
    triasto.stl" "$structural_dir_path/URgeom${iteration}
    _cell_unSmoothed.stl" 0 0 0 1 1
```

B.3 Post-processing macro for extracting pressure results from Star-CCM+

The post-processing of the fluid simulation results, and the extraction of the surface pressure data is an important part of the coupling process, and so an example of how this was achieved is presented here.

The first step was to create a number of points, known as “Point Probes”, at the positions of the centroid of each element in the structural FEM model. Each probe is then used to sample the pressure at that point, to be used as a loading in the subsequent structural simulation. The locations of each probe are output to a text file in an earlier step in the process. These locations are read in and the probes

created, as in the code example below.

```
//Declare and initialize variables
double[][] matrix = new double [7741][5];

//Use a scanner to read the point probe coordinate data from the file,
//into the array, called matrix
Scanner scanner = new Scanner (datafile);
scanner.useDelimiter("[,\r\n]");
for (int i=0;i<7740;i++)
{
  for (int j=0;j<4;j++)
  {
    matrix[i][j]=scanner.nextDouble();
  }
}

//Get the region data so that it can be referenced.
Region region_0 = simulation.getRegionManager().getRegion("Region 1");

//Use region data to get boundary object for upper surface, and do the
    same for the lower surface
DirectBoundaryInterfaceBoundary directBoundaryInterfaceBoundary_0 = ((
    DirectBoundaryInterfaceBoundary)region_0.getBoundaryManager().
    getBoundary("deformed-geom-parts.Upper Surface[deformed-geom-parts.
    Upper Surface]"));
DirectBoundaryInterfaceBoundary directBoundaryInterfaceBoundary_1 = ((
    DirectBoundaryInterfaceBoundary) region_0.getBoundaryManager().
    getBoundary("deformed-geom-parts.Lower Surface [deformed-geom-parts.
    Lower Surface]"));

//Get the coordinate system and units that will be applied to the point
     parts upon creation
LabCoordinateSystem labCoordinateSystem_0 = simulation.
    getCoordinateSystemManager().getLabCoordinateSystem();
Units units_0 = simulation.getUnitsManager().getPreferredUnits(new
    IntVector(new int[] {0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0,
     0, 0, 0, 0, 0, 0, 0, 0}));

//Use a loop to create all of the point probes at the appropriate
    location.  Apply correct units and coordinate system.
//Rename the points according to their id so that they can be easily
    accessed in future
```

```
//
//For the upper surface
for (int a=0;a<4220;a++)
{
  PointPart pointPart_16 = simulation.getPartManager().createPointPart(
    new NeoObjectVector(new Object[] {directBoundaryInterfaceBoundary_0
    }), new DoubleVector(new double[] {matrix[a][1], matrix[a][2],
    matrix[a][3]}));
  pointPart_16.setCoordinateSystem(labCoordinateSystem_0);
  Coordinate coordinate_11 = pointPart_16.getPointCoordinate();
  coordinate_11.setCoordinate(units_0, units_0, units_0, new
    DoubleVector(new double[] {matrix[a][1], matrix[a][2], matrix[a
    ][3]}));
  coordinate_11.setValue(new DoubleVector(new double[] {matrix[a][1],
    matrix[a][2], matrix[a][3]}));
  String strPointName = String.format("%.0f",matrix[a][0]);
  pointPart_16.setPresentationName(strPointName);
}
//
//For the lower surface
for (int a=4220;a<7740;a++)
{
  PointPart pointPart_16 = simulation.getPartManager().createPointPart(
    new NeoObjectVector(new Object[] {directBoundaryInterfaceBoundary_1
    }), new DoubleVector(new double[] {matrix[a][1], matrix[a][2],
    matrix[a][3]}));
  pointPart_16.setCoordinateSystem(labCoordinateSystem_0);
  Coordinate coordinate_11 = pointPart_16.getPointCoordinate();
  coordinate_11.setCoordinate(units_0, units_0, units_0, new
    DoubleVector(new double[] {matrix[a][1], matrix[a][2], matrix[a
    ][3]}));
  coordinate_11.setValue(new DoubleVector(new double[] {matrix[a][1],
    matrix[a][2], matrix[a][3]}));
  String strPointName = String.format("%.0f",matrix[a][0]);
  pointPart_16.setPresentationName(strPointName);
}
```

The above code produces a set of "Point Probes" on the external upper and lower surfaces of the structural boundary within the CFD model in STAR-CCM+, which correspond to the locations of the centroids of the elements in the structural FEM model. These point probes can then be used to extract the pressure values at that location.

Performing the extraction can be done in two ways. Initially, an approach was taken where the pressures were added to an XY-plot within the model, that was grouped in terms of chordwise strips. The pressures were then extracted from the plots, allowing them to be written to a file that was used directly as an input (.k) file for the structural simulation.

However, this method was quite slow, and a quicker method was found. In this case, the pressures were extracted from the PointProbes directly using a "Report" that extracted the maximum pressure variable given at that point. The pressures could then be written to a loading file (in .k format for LS-DYNA) and used as an input in the structural simulation. This significantly sped up the process of extracting the surface pressures. An example of the code required for this is provided below.

```
//Declare and initialise the array used to get the point probes
//as PointParts and use it to add the PointPart to the report
//in order to extract the pressure
PointPart[] pressureValue = new PointPart[7741];

//Declare and initialise the array to hold all of the pressure values,
//so that they can be looped over for writing to file.
//This holds the values read into yDataLoop below,
//which can be used as a temporary variable in the loop.
double[] pressureMatrix = new double[7741];

//Create a buffered writer object for when the pressure data is written
    to file
BufferedWriter outputPressuresLoop = null;

//Create and initialise the matrix to hold the probe point id numbers
//and co-ordinates when read in from the file.
double[][] matrix = new double [7741][5];

//Use a getter method to get the file path of the point probe data file
    , and the iteration number
String pointProbeFilePath = ReadProbeDataAndCreatePoints.
    getPointProbeFilePath();
int iterationNo = ReadProbeDataAndCreatePoints.getIterationNo();
int nextIterationNo = iterationNo + 1;

//Use a scanner to read in the probe point id numbers into an array
//(and coordinates but they're not used)
try
```

```
{
  Scanner scanner = new Scanner (new File(pointProbeFilePath));
  scanner.useDelimiter("[,\r\n]");

  for (int i=0;i<7740;i++)
  {
    for (int j=0;j<4;j++)
    {
      matrix[i][j]=scanner.nextDouble();
    }
  }
}
catch(FileNotFoundException e)
{
}

//Create a new directory for the pressure file to be loaded into,
//and create a path for this to be referenced from
String pressuresDirPath = "/media/Models/sim1/structure"+
    nextIterationNo;
String pressuresFilePath = pressuresDirPath + "/surfacePressuresLoop.k"
    ;
File loadingFile = new File(pressuresFilePath);
File dir = new File(pressuresDirPath);
dir.mkdir();

//Define the load curve idea for the ls-dyna model that will be used to
     apply the loads.
String LCID = "20";

//Create the report used in the loop to get the pressures at the probe
    locations
MaxReport maxReport = simulation.getReportManager().createReport(
    MaxReport.class);
//Assign pressure to the report
maxReport.setScalar(primitiveFieldFunction_0);

//Use a buffered writer to create and write a loading file containing
    the pressures for the
//structural analysis
//The writer loops over each point probe.
try
{
```

```
  outputPressuresLoop = new BufferedWriter(new FileWriter(loadingFile))
    ;
  outputPressuresLoop.write("*KEYWORD\n");
  outputPressuresLoop.write("*LOAD_SEGMENT_SET\n");

  //loop over first XYPlot
  for(int i=1;i<7741;i++)
  {
    pressureValue[i] = ((PointPart) simulation.getPartManager().
    getObject(String.format("%.0f",matrix[i-1][0])));
    maxReport.getParts().setObjects(pressureValue[i]);
    pressureMatrix[i]=maxReport.getValue()*-1;
    outputPressuresLoop.write(pressureValue[i] + "," + LCID + "," +
    String.format("%.3f",pressureMatrix[i]));
    outputPressuresLoop.newLine();
  }
  outputPressuresLoop.write("*END");
  outputPressuresLoop.close();
}
catch(IOException e)
{
}
finally
{
  try
  {
    outputPressuresLoop.close();
  }
  catch(IOException e)
  {
  }
}
```

Appendix C

Header files for structural simulation

C.1 Example LS-DYNA input files

C.1.1 Example LS-DYNA main header file

Below is an example of an LS-DYNA main header file for the structural simulation, using the implicit non-linear solver.

```
*KEYWORD 20000000
*TITLE
$# Structural iteration
*STRESS_INITIALIZATION
*CONTROL_TERMINATION
$#  endtim    endcyc     dtmin    endeng    endmas
160.0000000         0     0.000     0.000     0.000
*CONTROL_IMPLICIT_GENERAL
1,0.04
*CONTROL_IMPLICIT_AUTO
1,
*CONSTRAINED_NODE_SET
4,3,
*DATABASE_BINARY_D3PLOT
$#      dt      lcdt      beam     npltc    psetid
     0.250         0         0         0         0
```

```
*INCLUDE
Structural_header_file.k
geometry_file.k
spanwise_pressure_loading_file.k
Surface_pressure_loading_file.k
Pressure_segment_sets_file.k
*END
```

C.1.2 Example LS-DYNA structural header file

```
*KEYWORD
$------------DEFINE-PARTS---------------------------------
*PART
Upper-Surface
$#     pid     secid       mid     eosid      hgid      grav    adpopt      tmid
         1         1         1                              0         0         0
*PART
Lower-Surface
$#     pid     secid       mid     eosid      hgid      grav    adpopt      tmid
         2         1         1                              0         0         0
*PART
Rib1
$#     pid     secid       mid     eosid      hgid      grav    adpopt      tmid
         3         1         1                              0         0         0
*PART
Rib2
$#     pid     secid       mid     eosid      hgid      grav    adpopt      tmid
         4         1         1                              0         0         0
*PART
Stopper1
$#     pid     secid       mid     eosid      hgid      grav    adpopt      tmid
         5         1         1                              0         0         0
*PART
Stopper2
$#     pid     secid       mid     eosid      hgid      grav    adpopt      tmid
         6         1         1                              0         0         0
*CONTROL_HOURGLASS
4
*DAMPING_PART_STIFFNESS
1,0.2
2,0.2
3,0.2
```

```
4,0.2
*SECTION_SHELL
1,16,1,1
0.055e-3,0.055e-3,0.055e-3,0.055e-3
*MAT_FABRIC
$      mid        ro        ea        eb        ec      prba      prca      prcb
         1      3100     500E6                              0
$      gab       gbc       gca       cse        el       prl    lratio      damp
                                       1                                    0.75
$     aopt    flc/x2    FAC/x3       ela      lnrc      form     fvopt    tsrfac



$------------CONTACT-----------------------------------------
*CONTACT_AUTOMATIC_SINGLE_SURFACE



$------------BOUNDARY/INITIAL CONDITIONS---------------------------------
*BOUNDARY_SPC_SET
2,0,1,0,1
3,0,0,0,1
$4,0,0,0,0,1,1,1
5,0,0,1,0
6,0,0,0,1
7,0,1,1,0
$------------Load Curve for Spanwise Loading------------
*DEFINE_CURVE
2
0,0
40,1
1E6,1
$------------Load Curve for External Surface Pressure Loading------
*DEFINE_CURVE
20
0,0
60,0
80,0
120,1
1E6,1
$------------Load Curve for Internal Loading--------------
*DEFINE_CURVE
200
```

```
0,0
60,0
61,1
1E6,1
*END
```

C.1.3 Example LS-DYNA loading file

```
*KEYWORD
*LOAD_SEGMENT_SET
1,2,656
2,200,656
1001,20,830.364
*END
```

C.1.4 Example LS-DYNA loading segment set file

```
*KEYWORD
*set_segment
1001
308,4439,4436,309
*END
```