Multidisciplinary Design Optimisation of Aero-Engine Fan Blades



Christopher Chahine

Mansfield College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2018

To my family.

"*There is nothing permanent except change.*"

Heraclitus

Acknowledgements

This thesis is the result of many years of work that was performed at the von Karman Institute for Fluid Dynamics and the University of Oxford.

A great many people have guided and supported me throughout my journey leading up to the finalisation of this work, some directly and intentionally, others indirectly and unaware. While it is a virtual impossibility to thank everyone, this section is aimed at thanking those whose efforts and contributions have been central.

First and foremost, I would like to express my deepest appreciation to Professor Tom Verstraete (von Karman Institute) for giving me the opportunity to work at the VKI, and for his supervision, guidance, support and commitment in all matters related to this work and beyond. He has not only been a great motivator, role model and source of inspiration, but also a friend and mentor throughout the many ups and downs that the process of the generation of this work brought about.

I would equally like to extend my greatest gratitude to my supervisor at the University of Oxford, Professor Li He, for enabling me to perform this work in an extraordinary configuration in an out-housed lab. I would like to thank him for the trust he put in me by giving me a lot of freedom in my work, and yet always being available for discussions when it was needed. I consider it a great honour to have been guided by him.

I would like to thank Professor Budimir Rosic and Professor Peter Ireland of the Oxford Thermofluids Institute (Osney Lab) for their assessments during the transfer of status and confirmation of status procedures and their useful feedback provided in the process. Further to that, I also thank Professor Rosic and Dr. Shahrokh Shahpar of Rolls-Royce plc. for acting as my thesis examiners.

I would like to extend my gratitude to all members of the Aerospace Structures and Materials Department of the TU Delft for their warm welcome

during a three-month research secondment on composite material design and optimisation. This time in Delft was not only a great learning experience, but also highly enjoyable. I greatly appreciate the many helpful discussions and highly positive atmosphere in the Department. In particular, I thank Professor Christos Kassapoglou for many helpful suggestions provided on the failure modelling of composite materials and for including me in the evaluation committee for one of his Master's thesis students.

A part of this work was performed under the EU Marie Curie Initial Training Network AMEDEO (Aerospace Multidisciplinarity Enabling DEsign Optimisation). Having been part of this group has been a highly rewarding experience not at least due to the great work of the hosting team at the University of Leeds. Particular thanks to Juliet Jopson, Professor Harvey Thompson and Professor Vassili Toropov (now at Queen Mary University of London) for their excellent efforts in leading the consortium and for their work in making this a very special experience for all who were involved. Many thanks also to my fellow Early Stage Researchers for the ambition, fun and shared learning experiences.

I would like to express my great appreciation to the European Commission and the Engineering and Physical Sciences Research Council of the UK for their contributions in funding this work.

I thank all my colleagues and the staff at the von Karman Institute and the University of Oxford, as well as at Mansfield College for many enjoyable years. With the risk of being incomplete, I particularly thank Mohamed, Delphine, Memo, Domingo, Stefano "Smakelijk" V. and Stefano "Sheldon" D'A., Sebastian, Sara P. and Sara G., Andrea, Tania, Cansev, Orkun, Clara, Imre, Chiara, Vincent, Elissavet and Bogdan for walking parts of the journey with me.

I wish to thank my parents, my brothers and my sister for their continuous support and encouragement throughout my studies and this work. Finally, I would like to thank Iryna for her support, understanding and patience through times of intense work that the completion of this thesis required.

Christopher Chahine

Rhode-Saint-Genèse, Belgium, February 2019

Abstract

This thesis presents a multidisciplinary and multi-objective design optimisation framework for transonic fan blades of modern high bypass ratio turbofan engines. The optimisation strategy applied is based on a two-level approach consisting of a Differential Evolution algorithm coupled to a sequentially updated ordinary Kriging metamodel to accelerate the optimisation process. Aerodynamic as well as structural static and dynamic performance criteria are considered by means of high-fidelity three-dimensional Computational Fluid Dynamics and Computational Structural Mechanics analysis tools. Multiple key operating points are considered in the optimisation process; aerodynamic performance is evaluated at top-of-climb and cruise conditions, while peak stresses are evaluated at take-off operation, taking into account centrifugal and gas loads. Blade vibration is furthermore assessed over the entire operating range. Aerodynamic performance is separately evaluated for core and bypass flows in order to match the requirements specified by the engine cycle design. The successful application of the optimisation system is demonstrated by example of two complex multi-objective, aerostructural fan blade design problems, including finite element structural models for a full titanium and a composite fan blade made from a carbon-fibre reinforced plastic material.

In addition to aerodynamic and structural considerations, aeroelastic effects are becoming of increasing importance for modern lightweight fan blade designs. Flutter, which is a harmful aeroelastic instability, is of particular relevance. An investigation of computational flutter prediction techniques is presented in this thesis with the goal of assessing their potential applicability in an automated optimisation process, showing that today's most commonly applied method, known as the energy method, might lead to mispredictions of flutter stability for modern lightweight, low-stiffness blade designs.

Contents

1 Introduction **1**

1.1 Background and scope 1
1.2 Introduction to the design problem 2
1.3 The industrial aerostructural design process 4
1.4 Multidisciplinary Design Optimisation 5
1.5 Prediction of transonic fan flutter 6
1.6 Engine cycle and baseline design 6
1.7 Thesis objectives 8
1.8 Thesis outline 8

2 Literature Review **11**

2.1 Optimisation problem formulation 11
2.2 Optimisation methods 13
2.2.1 Gradient-free methods 14
2.2.1.1 Evolutionary algorithms 15
2.2.1.2 Particle swarm optimisation 15
2.2.2 Gradient-based methods 16
2.2.3 Metamodeling in numerical optimisation 17
2.3 Blade parametrisation methods 19
2.3.1 Node-based parametrisation 19
2.3.2 Free-form deformation (FFD) 20
2.3.3 CAD-based parametrisation 21
2.4 Performance analyses 22
2.4.1 Steady flow analysis 22

2.4.2 Structural analysis . . . 23
2.4.3 Aeroelasticity . . . 23
2.5 Review of transonic compressor and fan blade design optimisation . . 23
2.5.1 NASA Rotor 37 and Rotor 67 optimisation . . . 24
2.5.2 Highly loaded compact transonic compressor optimisation . . 25
2.5.3 High bypass ratio aero-engine fan blade optimisation . . . 26
2.6 Computational blade flutter prediction . . . 27
2.7 Discussion . . . 29

3 Optimisation Method **33**
3.1 Optimisation system . . . 33
3.2 Description of core components . . . 35
3.2.1 Differential Evolution . . . 35
3.2.2 Database . . . 36
3.2.3 Kriging . . . 37
3.2.4 High-fidelity evaluation chain . . . 38

4 CAD-Based Fan Parametrisation and Automatic Geometry Generation **41**
4.1 Fan assembly overview . . . 41
4.2 Fluid domain definition . . . 43
4.2.1 Meridional contour definition . . . 44
4.2.2 Camber surface definition . . . 44
4.2.3 Lean and sweep . . . 47
4.2.4 Thickness . . . 49
4.3 Solid blade generation . . . 49

5 Aerodynamic Performance Analysis **55**
5.1 An introduction to the aerodynamic fan design problem . . . 55
5.2 Navier-Stokes equations . . . 57
5.3 Reynolds-Averaged Navier-Stokes equations . . . 58
5.4 Turbulence modelling . . . 60
5.5 Automatic mesh generation . . . 60
5.6 Spatial discretisation . . . 62
5.7 Boundary conditions . . . 63
5.8 Solution procedure . . . 63
5.8.1 Steady flow solution . . . 63

5.8.2 Unsteady flow solution 64
5.8.2.1 Time marching of RANS equations 64
5.8.2.2 Non-linear harmonic method (NLH) 64
5.9 Some current challenges in aerodynamic fan design/analysis 65
5.9.1 Stall limit prediction 65
5.9.2 Noise prediction 65

6 Structural Performance Analysis 67
6.1 An introduction to the structural fan design problem 67
6.2 Linear static analysis 67
6.3 Element choice 73
6.4 Automatic mesh generation 73
6.5 Displacement boundary conditions 75
6.6 Material modelling 75
6.6.1 An introduction to composite materials for fan blade structural design 75
6.6.2 Finite element analysis of composite materials 77
6.6.2.1 Micro-scale analysis 77
6.6.2.2 Meso-scale analysis 78
6.6.2.3 Macro-scale analysis 78
6.6.3 Constitutive relations 78
6.7 Automatic element coordinate system generation 81
6.8 Failure criteria 83
6.9 Discussion of non-linear static effects 85
6.9.1 Blade root-disk contact 86
6.9.2 Material non-linearities 86
6.9.3 Geometric non-linearities 86
6.10 Vibration analysis 87

7 Multidisciplinary Design Optimisation of a Titanium Fan Blade 91
7.1 Fan blade parametrisation 91
7.1.1 Meridional contour 91
7.1.2 Camber surface 92
7.1.3 Lean and sweep 94
7.1.4 Thickness 95
7.1.5 Blade root parametrisation 96
7.1.6 Number of blades 96

7.1.7 Summary of optimisation parameters 96
7.2 Aerodynamic performance evaluation 97
7.2.1 Fluid mesh generation 97
7.2.2 Fluid and turbulence modelling 100
7.2.3 Operating points and boundary conditions 100
7.3 Structural performance evaluation 102
7.3.1 Solid mesh generation 102
7.3.2 Displacement boundary conditions and material model 103
7.3.3 Static stress analysis 103
7.3.4 Vibration analysis 104
7.4 Objectives and constraints 105
7.5 Optimisation settings 108
7.6 DOE database 108
7.7 Optimisation results 108
7.7.1 Survey of optimisation process 108
7.7.2 Comparison of optimised and baseline performance 116
7.7.2.1 Aerodynamic performance 117
7.7.2.2 Structural performance 122
7.7.2.3 Further geometry and performance data 124
7.8 Summary 124

8 Multidisciplinary Design Optimisation of a Composite Fan Blade 129
8.1 Fan blade parametrisation 129
8.1.1 Meridional contour 129
8.1.2 Camber surface 129
8.1.3 Lean and sweep 131
8.1.4 Thickness 132
8.1.5 Blade root parametrisation 132
8.1.6 Number of blades 133
8.1.7 Summary of optimisation parameters 133
8.2 Aerodynamic performance evaluation 135
8.2.1 Fluid mesh generation 135
8.2.2 Operating points and boundary conditions 135
8.3 Structural performance evaluation 135
8.3.1 Solid mesh generation 135
8.3.2 Material properties 136

8.3.3 Boundary conditions . . . 136
8.3.4 Static stress analysis . . . 136
8.3.5 Vibration analysis . . . 137
8.4 Objectives and constraints . . . 138
8.5 Optimisation settings . . . 140
8.6 DOE database . . . 140
8.7 Optimisation results . . . 140
8.7.1 Survey of optimisation process . . . 140
8.7.2 Comparison of Pareto optimal solutions . . . 149
8.7.2.1 Aerodynamic performance . . . 149
8.7.2.2 Structural performance . . . 158
8.7.2.3 Further geometry and performance data . . . 160
8.8 Summary . . . 160

9 Analysis of Aerodynamics-Structural Dynamics Coupling for Blade Flutter Prediction 171
9.1 Aeroelastic system and modal equation . . . 173
9.2 Structural model . . . 174
9.3 CFD solver, mesh and operating points . . . 176
9.4 Description of flutter prediction methods . . . 177
9.4.1 Decoupled method . . . 177
9.4.2 Coupled method . . . 179
9.4.3 Static aeroelastic blade untwist . . . 180
9.5 Comparative results and discussion . . . 181
9.5.1 Global aero-damping comparisons and frequency shift results . 181
9.5.1.1 Peak efficiency . . . 181
9.5.1.2 Higher load condition . . . 184
9.5.2 Discussion of frequency shift . . . 185
9.5.3 Low mass ratio, low stiffness case . . . 192
9.6 Summary . . . 193

10 Conclusions and Future Work 195
10.1 Conclusions and Recommendations . . . 195
10.1.1 Multidisciplinary Design Optimisation . . . 195
10.1.2 Flutter prediction . . . 198
10.2 Future work . . . 199

A Validation Study of Unsteady Aerodynamics for Blade Dynamic Aeroelasticity (Flutter) **203**

A.1 Test case description 203

A.2 Computational method description 205

A.3 Comparison of numerical and experimental results 206

Bibliography **212**

List of Figures

1.1 Schematic of a modern high bypass ratio turbofan engine. Adapted from [181]. 2
1.2 Overview of fan performance requirements. 4

2.1 A schematic Pareto front for a two-objective optimisation problem. Design A is dominated by design B, but designs B and C are non-dominated. 14
2.2 Samples along the Pareto-front obtained with different weights for the objectives using a single-objective weighted sum approach. 16

3.1 Flow chart of the optimisation system. 33

4.1 Overview of a typical fan assembly. Adapted from [51]. 42
4.2 Meridional flow path definition. Red squares indicate the B-Spline curve control points. (a) global view. (b) close-up of bypass splitter. . 43
4.3 Parametrisation of the camber surface. 45
4.4 Meridional view showing the initial radial stacking of camber lines through the center of gravity of the associated two-dimensional blade sections. Gravity centers are shown as rectangles. 47
4.5 Selected tip section profiles generated with the implemented parametrisation. 48
4.6 Parametrisation and definition of lean and sweep. 48
4.7 Construction of the final aerofoil shape by adding a thickness distribution to the camber surface. 49

4.8 Step 1 of the dovetail root generation process: Definition of the dovetail cross section. 50
4.9 Step 2 (left) and step 3 (right) of the dovetail root generation process. 51
4.10 Step 4 of the dovetail root generation process: generation of the shank surfaces. 52
4.11 Visualisation of exported fan blade CAD-models in a commercial CAD program. 53

5.1 A structured (a) and unstructured (b) mesh in 2D. 62

6.1 Depiction of the general boundary value problem for an elastic fan blade with volume Ω and boundary Γ, subjected to surface tractions $\bar{t}$ and body forces b. 68
6.2 Depiction of the discretised general boundary value problem. 71
6.3 A linear (a) and quadratic tetrahedral element (b). 73
6.4 Solid mesh generation process. 74
6.5 Schematic illustration of composite laminates. (a) plies with unidirectional fibres and (b) 2D woven fabric. 76
6.6 Reinforced composite plies and their principal material directions. (a) A unidirectionally reinforced ply and (b) two-dimensional woven ply. 79
6.7 Definition of a Cartesian element coordinate system (x',y',z') in dependency of the global coordinates (x,y,z). 82
6.8 Element coordinate systems for the definition of orthotropic material properties for a solid mesh of the baseline fan blade. 84
6.9 Campbell diagram of the baseline titanium fan blade. 88

7.1 Meridional contour of the baseline geometry. 92
7.2 Camber surface parametrisation. Black arrows indicate direct degrees of freedom of control points (optimisation parameters), while red arrows indicate indirect degrees of freedom (quantities that are defined in dependency to other optimisation parameters). For convenience, the figures on the top of the plot point out which quantity is being defined by the respective parametric curve. 92
7.3 Impact of the non-dimensional mid-point beta value β_{mid}^{ND} on the camberline shape for $u_{\beta_{mid}} = 0.5$. λ denotes the camber angle. 95
7.4 Camberline shape for a precompression airfoil with $u_{\beta_{mid}} = 0.5$. λ denotes the camber angle. 95

7.5 Parametrisation of lean and sweep. Note that the baseline values of all control points defining lean and sweep are zero. 96
7.6 Parametrisation of the blade thickness distributions at hub and tip (a) and spanwise interpolated distributions (b). 97
7.7 Parametrisation of the dovetail root cross section. The width of the shape shown is adapted to the thickness of the baseline blade. 98
7.8 Geometry of the parametrised baseline blade with fully constructed root geometry. 98
7.9 Fluid domain and surface mesh point distributions of the baseline design. 100
7.10 Mesh independence study for the baseline blade at the cruise design point. 101
7.11 Solid meshes of the baseline blade. 102
7.12 Solid mesh of the baseline fan blade and its restraint in a typical disk. The nodes on the dovetail flank surfaces shown in red are fixed in all degrees of freedom to simulate the disk attachment (opposite constrained surface not visible). 104
7.13 Definition of objectives and the top-of-climb evaluation point shown in the performance map of the baseline design. 106
7.14 Definition of the non-dimensional frequency margin. 107
7.15 Objective space after 20 iterations. Circles indicate DOE samples, while squares indicate designs generated during the optimisation. . . 109
7.16 Geometries of selected individuals that have been successfully evaluated by the high-fidelity evaluation chain. Numbers correspond to Fig. 7.15 (a). 110
7.17 Overview of constraint spaces after 20 optimisation iterations. Every symbol shown represents a design evaluated by the high-fidelity evaluation chain. Symbols in red quadrants are violating the respective constraint. Circles indicate DOE samples, while squares indicate designs generated during the optimisation. 111
7.18 Overview of the metamodel prediction errors for all performance parameters in dependency of the optimisation progress. The error values are computed as the average prediction error of all candidate individuals that are re-evaluated by the high-fidelity evaluation per iteration. Errors are given in percent. 112

7.19 Evolution of design parameters. Green symbols denote feasible designs, while red symbols indicate infeasible ones. The baseline design is shown as orange diamond. 114
7.20 Impact of the tip and near-hub camber angles on the cruise peak efficiency objective and the stall margin objective. Green symbols denote feasible designs, while red symbols indicate infeasible ones. 115
7.21 Impact of the hub thickness scaling factor on the cruise peak efficiency objective and the stall margin objective. Green symbols denote feasible designs, while red symbols indicate infeasible ones. 116
7.22 Detail view on the feasible region in the objective space and the chosen individual from the Pareto front (IT020_IND002). 117
7.23 Comparison of the meridional geometries of baseline and optimised designs and pitch averaged absolute total pressure distributions in the bypass inlet. 118
7.24 Comparison of the isentropic Mach number distributions at 10, 50 and 90% span of baseline and optimised design. 119
7.25 Comparison of entropy contours for baseline and optimised design at their respective cruise peak efficiency point. 120
7.26 Performance map comparison of baseline and optimised designs. . . . 121
7.27 Comparison of the von Mises stress distributions in the baseline and optimised design. 123
7.28 Campbell diagram of baseline and optimised designs. 124
7.29 Comparison of blade mode shapes and eigenfrequencies at take-off. . . 125
7.30 Blade geometries of baseline and optimised design. 126
7.31 Full fan assemblies of baseline and optimised design. 127
7.32 Relative Mach number contours at 10, 50 and 90 percent span for baseline and optimised design at the cruise design point. 128
7.33 Pitch averaged spanwise distribution of absolute total pressure at the fan exit, upstream of the bypass splitter, at the cruise design point. . 128

8.1 Camber surface parametrisation. Black arrows indicate direct degrees of freedom of control points (optimisation parameters), while red arrows indicate indirect degrees of freedom (quantities that are defined in dependency to other optimisation parameters). 130
8.2 Parametrisation of lean and sweep. Note that the baseline values of all control points defining lean and sweep are zero. 131

8.3 Parametrisation of the blade thickness distributions at hub and tip (a) and spanwise interpolated distributions (b). 132
8.4 Parametrisation of the dovetail root cross section. The width of the shape shown is adapted to the thickness of the baseline blade. 133
8.5 Solid mesh of the baseline fan blade and its restraint in a typical disk. The nodes on the dovetail flank surfaces shown in red are fixed in all degrees of freedom to simulate the disk attachment (opposite constrained surface not visible). 137
8.6 Objective space after 59 iterations. Circles indicate DOE samples, while squares indicate designs generated during the optimisation. . . 141
8.7 Overview of constraint spaces. Every symbol shown represents a design evaluated by the high-fidelity evaluation chain. Circles indicate DOE samples, while squares indicate designs generated during the optimisation. Negative values indicate constraint satisfaction. The baseline design is shown as orange diamond. 142
8.8 Overview of the cruise stall margin constraint. The symbol convention follows that of Fig. 8.7. Negative values indicate constraint satisfaction. 143
8.9 Overview of the metamodel prediction errors for all performance parameters in dependency of the optimisation progress. The error values represent the average error of all candidate designs per iteration. . . . 144
8.10 Evolution of design parameters - part 1. Green symbols denote feasible designs, while red symbols indicate infeasible ones. The baseline design is shown as orange diamond. 145
8.11 Evolution of design parameters - part 2. Green symbols denote feasible designs, while red symbols indicate infeasible ones. The baseline design is shown as orange diamond. 146
8.12 Impact of hub chord scaling factor and the near hub leading edge blade angle on the cruise peak efficiency objective and the total fan mass objective. Green symbols denote feasible designs, while red symbols indicate infeasible ones. 147
8.13 Impact of leading edge blade angle and the rear camber angle at the tip on the cruise peak efficiency objective and the total fan mass objective. Green symbols denote feasible designs, while red symbols indicate infeasible ones. 148
8.14 Detail view on the Pareto front obtained after 59 optimisation iterations and indication of individuals chosen for further examination. . . 149

8.15 Comparison of baseline and optimised blade profiles at 10% span. . . 150
8.16 Comparison of isentropic Mach number distributions at 10% span at the cruise design point and the top-of-climb evaluation point. 151
8.17 Comparison of the entropy contours at 10% span between the highest efficiency and lowest weight design at the cruise design point. 152
8.18 Comparison of baseline and optimised blade profiles at 90% span. . . 153
8.19 Comparison of isentropic Mach number distributions at 10% span at the cruise design point and the top-of-climb evaluation point. 154
8.20 Comparison of the stacking axes of the optimised geometries in the meridional plane. Squares indicate the centre of gravities of the blade sections. The straight dashed lines represent a radial stacking axis. . 155
8.21 Influence of the tip sweep angle on the cruise peak efficiency objective (a) and the Tsai-Wu stress constraint (b). Green symbols denote design that are satisfying all constraints, while red symbols indicate designs with at least one constraint violation. 156
8.22 Impact of different levels of backward sweep at the tip on the shock surface at the cruise design point (red surface shows the iso-surface of unity relative Mach number). 156
8.23 Performance curves of the chosen designs. 157
8.24 Comparison of the Tsai-Wu failure criterion distribution in the front root of the three optimised and the baseline design. 158
8.25 Influence of the leading edge point offset value (front root tilting) on the Tsai-Wu stress constraint. 159
8.26 Stress distributions in a straight vs. tilted front root section in terms of the Tsai-Wu failure criterion. Values greater than one indicate material failure. White edges indicate the deformed shape with tenfold magnification. 159
8.27 Campbell diagram of the baseline design. 162
8.28 Campbell diagram of the optimised design with highest efficiency. . . 162
8.29 Campbell diagram of the optimised design with compromise between efficiency and weight. 163
8.30 fig:campbell-diagram-he 163
8.31 Campbell diagram of the optimised design with lowest weight. 163
8.32 Blade geometries of baseline and optimised design with highest efficiency. 164
8.33 Blade geometries of the compromise and lowest weight design. 165
8.34 Full fan assemblies of the baseline and highest efficiency optimised design. 166

8.35 Full fan assemblies of the compromise and lowest weight design. . . . 167
8.36 Relative Mach number contours at 10, 50 and 90 percent span for baseline and optimised design with highest efficiency at the cruise design point. . . . 168
8.37 Relative Mach number contours at 10, 50 and 90 percent span for the compromise and lowest weight optimised design at the cruise design point. . . . 168
8.38 Pitch averaged spanwise distribution of absolute total pressure at the fan exit, upstream of the bypass splitter, at the cruise design point. . 169

9.1 Solid mesh and mode shape contour of the first bending mode at design speed, calculated with the baseline titanium alloy. . . . 175
9.2 Overview of CFD mesh. Every second mesh point is shown. . . . 176
9.3 Energy conservation test of coupled computation for baseline blade at peak efficiency. . . . 179
9.4 Relative Mach number contours at 90 percent blade span for peak efficiency operation. Left: CFD, right: Experiments [158]. . . . 181
9.5 Coupled time evolution of the tip leading edge displacements at peak efficiency. . . . 182
9.6 Comparison of decoupled and coupled predictions of aeroelastic stability (left) and frequency shift (right) over mass ratio and stiffness changes at peak efficiency. . . . 183
9.7 Relative Mach number contours at 90 percent blade span for operation at the higher load condition. . . . 184
9.8 Coupled time evolution of the tip leading edge displacements at the higher load condition. . . . 185
9.9 Comparison of decoupled and coupled predictions of aeroelastic stability (left) and coupled frequency shift (right) over mass ratio and stiffness changes at the higher load condition. . . . 186
9.10 Argand diagram of the forces in the aeroelastic system. . . . 187
9.11 a) Decoupled frequency shift prediction due to aero added mass as function of the force ratio. b) the same plot with coupled results for the higher load condition included. . . . 191

A.1 Low-speed linear turbine cascade test rig. Adapted from [69]. . . . 204
A.2 Locations of the pressure tappings on one blade section [69]. . . . 205
A.3 Computational domain. . . . 206

A.4 Blade detail of the blade-to-blade mesh at 50% span. 207
A.5 Comparison of the steady pressure distributions at 50% span. 208
A.6 Amplitude and phase of the unsteady pressure at a reduced frequency $\nu = 0.4$, and the least stable IBPA (-60 deg). Lines indicate CFD results, while squares indicate experiments. 209
A.7 Amplitude and phase of the unsteady pressure at a reduced frequency $\nu = 0.4$, and the least stable IBPA (120 deg). Lines indicate CFD results, while squares indicate experiments. 210
A.8 Global comparison of aero-damping versus the interblade-phase angle (IBPA) for a reduced frequency of $\nu = 0.4$. 211

1 Introduction

1.1 Background and scope

Air transport constitutes an integral part of our modern globalised economy. In the year 2016 alone, the number of passengers carried by commercial flights worldwide amounted to 3.7 billion [45]. The market is projected to grow by more than 4% annually in the next two decades [45], making the aircraft industry one of the fastest growing sources of CO_2 emissions worldwide.

Facing increasing environmental concerns, the Advisory Council for Aeronautics Research in Europe (ACARE), has issued ambitious emission goals aiming at a 75% reduction of CO_2 emissions and a 90% reduction of NO_X emissions per passenger kilometre as well as a reduction of the perceived noise level by 65% by the year 2050, in reference to new aircraft entering service in the year 2000 [44]. The achievement of these goals in conjunction with a continuously growing market demand represents an unprecedented challenge for aircraft and engine manufacturers. It is widely acknowledged by industry and policy-makers alike, that a continuation of incremental technology improvements will not suffice to enable the technological advances needed for an environmentally sustainable aircraft industry. Instead, more drastic and revolutionary technology improvements are required.

A design paradigm known as Multidisciplinary Design Optimisation holds promise to be a key technological enabler for the design of novel environmentally friendly aircraft. The subject of this thesis is the application of Multidisciplinary Design Optimisation methods to the design of a central component of modern aero-engines: the fan blades.

1.2 Introduction to the design problem

Modern aircraft are powered by high-bypass ratio turbofan engines. A depiction of a representative example is shown in Fig. 1.1. To enable high propulsive efficiency and low noise, aero-engines for commercial passenger aircraft are designed with a large diameter fan, being located in the front section of the engine. The fan draws in a large mass flow of atmospheric air of which a small fraction is entering the engine core downstream of the fan. The engine core is a gas turbine in which the air flow undergoes a Brayton cycle where it is further compressed, mixed with fuel, burned and expanded in a turbine. The turbine produces shaft power which is used to drive the core compressors and the fan. Any remaining energy of the turbine exit flow is utilized as propulsive jet to contribute to the overall engine thrust. However, most of the fan exit flow bypasses the core engine and is expelled directly into the atmosphere through a bypass nozzle to generate the majority of the thrust. The *bypass ratio* denotes the ratio between the bypass and core mass flows. In modern high-bypass ratio engines, more than 80% of the total engine thrust is generated by the fan alone. The fan design therefore crucially affects the engine's propulsive efficiency and its specific fuel consumption.

The primary purpose of the fan is aerodynamic in nature and can be stated as the generation of a specific amount of thrust whilst producing little aerodynamic losses in

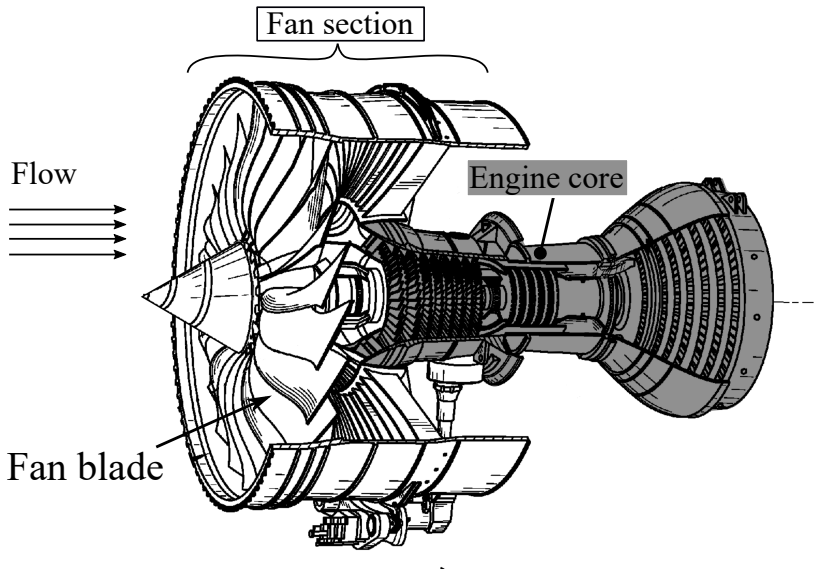


Figure 1.1: Schematic of a modern high bypass ratio turbofan engine. Adapted from [181].

order to attain high efficiencies and correspondingly low engine fuel consumption. The aerodynamic design problem is complicated by the need for the fan to operate under a wide range of operating conditions over the aircraft mission, requiring good design and off-design performance and large aerodynamic stability margins towards stall and surge. Modern fans operate in transonic flow conditions resulting in complex three dimensional shock structures. If the blade is inadequately designed, strong shocks and shock-induced boundary layer separation can lead to considerable performance penalties in the form of increasing losses and reduced operating range.

However, the consideration of aerodynamics alone is insufficient for the design of a fan that is to survive realistic operating conditions and to embody a profitable product for the manufacturer. An illustrative overview of some of the most critical design requirements for modern fan blades originating from various disciplines is given in Fig. 1.2. The origins of those requirements can be manifold and include the need to adhere to regulations (such as for example limitations on noise emissions), favourable operational performance (low maintenance cost), as well as limits imposed by existing manufacturing techniques and materials. It can be argued though, that all requirements can be translated to a primary economic performance metric for the manufacturer.

The weight that is placed on each of those design requirements in the design process is case dependent. For a long range aircraft, for example, high aerodynamic efficiency at cruise is a paramount requirement to ensure good overall fuel efficiency, because the aircraft operates longest in this condition. In contrast, for medium range aircraft which have a shorter cruise segment, it may be more important to trade aerodynamic efficiency for improved maintenance cost.

This argument serves to indicate a critical attribute of the design problem, which is that there typically exists a trade-off between the performances of the disciplines involved. This dependency is schematically illustrated by the grey circular arrow shown in Fig. 1.2.

A good example of this interdisciplinary dependency is given by the fan tip speed, which is typically limited by either structural, or more commonly, aeroacoustic constraints. This tip speed limitation has obviously direct consequences for the overall blade aerodynamics; less obvious is that it strongly affects the obtainable pressure rise in the fan root flow, which feeds the engine core.

Another example is the required stability of the blades against foreign object impact, particularly against large birds. Production blades have to demonstrate that they are able to withstand a standardised bird impact before receiving certification.

In order to achieve sufficient impact stability blades require a certain thickness. For the aerodynamic efficiency of the blades, however, blade thickness is detrimental, and a thinner blade would yield lower losses.

The aforementioned argument outlines that fan design is intrinsically a *multidisciplinary* problem. In addition, modern aerodynamically efficient fan blade shapes are extremely complex and highly three-dimensional. They can be defined by literally hundreds of parameters, whereas each may exhibit complicated coupling relations to other parameters and highly non-linear effects on the individual disciplinary responses.

A central part of the overall multidisciplinary design problem described above is represented by the so-called aerostructural design, which involves, as the name suggests, the aerodynamic and structural design of the fan blade. This is the part of the design problem that this thesis is concerned with.

1.3 The industrial aerostructural design process

Due to the substantial complexity and specialised domain knowledge required for the analysis and prediction of the structural and aerodynamic performances of fan blades, industry is handling the design problem in a de-coupled manner, in which specialist departments are working on each of the disciplines in isolation and in a sequential fashion. When a new design is required, it typically is initiated by the aerodynamicists

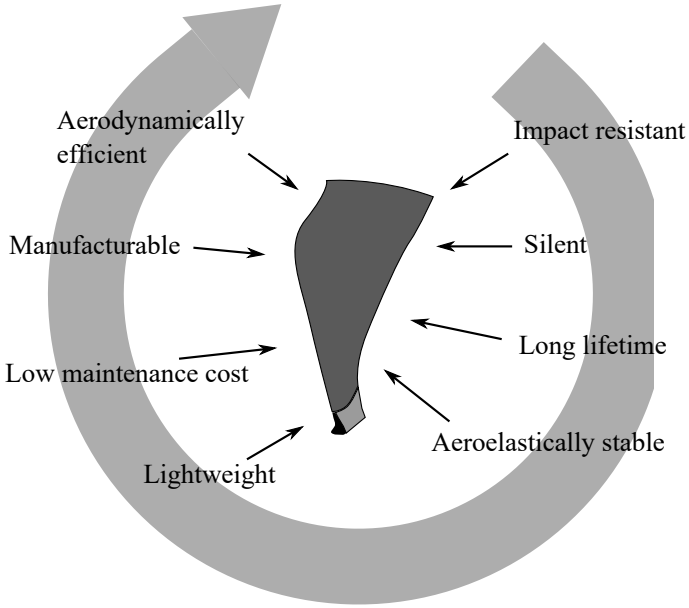


Figure 1.2: Overview of fan performance requirements.

first, who design a blade that meets the set aerodynamic design targets. Subsequently, the design is passed on to the structural engineers in order to ensure that the blade meets the structural integrity and lifetime constraints. Typically, this is initially not the case and the design is modified in order to meet the requirements. As any modification of the blade shape affects its aerodynamic performance, it has to be sent back to the aerodynamicists such that any lost performance may be recovered. This initiates a time consuming iterative loop between aerodynamic and structural departments until the desired design requirements are met.

Typically, the structural and aerodynamics department do not share the same geometry representation due to the different requirements of the respective analysis methods and tools used. This requires a conversion from one representation to another every time the geometry is passed on between the departments, introducing additional sources for error.

The success of this loosely coupled design process hinges largely on the experience of the designers involved and their ability to anticipate implications of design changes in the other discipline. The need for more substantial performance gains of future designs in conjunction with the introduction of new materials, such as composite materials with tailorable mechanical characteristics, greatly increases the complexity of the design problem. An identification of disciplinary dependencies in such design scenarios is virtually an impossibility for a human designer. This is where automatic optimisation methods in conjunction with high-fidelity computational performance prediction tools have the potential to act as an enabler for highly efficient blade designs.

1.4 Multidisciplinary Design Optimisation

Multidisciplinary Design Optimisation (MDO) methods were introduced in the early 1990's [152]. They evolved from the field of structural optimisation, in which numerical optimisation techniques have been used as early as in the 1960's [140] for the efficient and automated solution of complex structural design problems using Computational Structural Mechanics (CSM) analysis methods. With the advancement of computational modelling techniques in other disciplines, such as Computational Fluid Dynamics (CFD), the application of optimisation methods for the solution of multidisciplinary engineering design problems became of considerable interest, and heralded the advent of MDO as a research field.

Great advancements in the field have been made in the past few decades and MDO has been identified as a key enabling technology to reach the stringent environmental goals for future aircraft, set out by the ACARE [44]. The methods found popularity amongst designers of transonic fan blades, and have been applied with success in various works for the design of small fans and compressors (e.g. [91], [92], [123]). However, the application of the methods to the design of large, transonic aero-engine fan blades still presents a generally unsolved problem.

1.5 Prediction of transonic fan flutter

As was indicated in the design problem overview in Fig. 1.2, aeroelastic stability is a crucial requirement for fan blades and needs to be considered in the design process. Blade flutter, which is a self-excited aeroelastic instability, is of particular danger as it can rapidly lead to the destruction of the blade in operation.

Flutter occurs when small blade perturbations in certain structural vibrational modes become amplified by the unsteady flow field surrounding the blades. If the introduced excess vibrational energy cannot be dissipated by some form of structural damping, the vibration amplitude grows rapidly to reach harmful values leading to the destruction of the component. The severity of a flutter incident mainly stems from the fact that it can appear suddenly and its consequences can be disastrous.

Flutter has long been known to be a relevant issue for fan blades. However, in contrast to structural and aerodynamic high-fidelity computational performance prediction tools, which can be considered to be mature, computational flutter prediction still poses a great challenge today. Computational flutter prediction methods vary in the way the fluid-structure interaction is modelled. Most often, coupling effects between aerodynamics and structural dynamics are assumed to be negligible, leading to so-called decoupled techniques, which are computationally efficient. However, whether this assumption holds true for modern lightweight fan blade designs is still a subject of debate. A clear indication for a generally suitable flutter prediction method for the integration into an automated optimisation framework has been so far lacking.

1.6 Engine cycle and baseline design

High bypass ratio aero-engines are typically used for two types of aircraft missions: medium and long range. The targeted aircraft mission influences how much weight is placed on the various key engine performance parameters in the design process,

such as specific fuel consumption (SFC), specific thrust, noise, lifetime targets and maintainability. This choice defines the design boundary conditions in terms of overall engine performance and governs the one-dimensional aero-thermodynamic engine cycle design, which defines the required individual component performances at design and off-design conditions in order to satisfy the overall engine duty.

In this thesis, a long range mission profile similar to that of a Boeing 777 is assumed. A simplified engine cycle design was performed in order to define the boundary conditions needed to initiate the fan design. A detailed explanation of the process can be found in many textbooks, such as [31], [100] and [116]. The end-wall contours of the fan section were chosen to closely resemble those of a GE90-115B engine.

Table 1.1 provides an overview of the resulting one-dimensional design and off-design fan performance requirements in terms of fan speed, mass flow, bypass ratio, bypass and fan root pressure ratio. Correction and referencing of quantities is performed following standard relations as for instance described in [31]. Three key flight conditions are considered; top-of-climb (TOC), cruise and take-off (T/O). Top-of-climb denotes the operating point just before the aircraft initiates the cruise phase. Cruise denotes the cruise average.

In industrial practice, the one-dimensional design is typically succeeded by a two-dimensional component design process. The tools used in this stage of the design process typically seek to model the two-dimensional flow in the meridional plane (i.e. the axial-radial plane) using so-called *through-flow* methods and on several axissymmetric streamsurfaces. Accordingly, a two-dimensional design code was used in this work for the generation of a baseline geometry, which serves as the starting point

Table 1.1: Engine cycle design requirements. Corrected quantities are referenced to standard sea level conditions ($T_{ref} = 288.15$ K, $p_{ref} = 1.01325e5$ Pa).

Flight mode		**TOC**	**Cruise**	**T/O**
Altitude	ft	31000	35000	0
Flight Mach number		0.84	0.84	0
Bypass ratio (BPR)		7.3	8.2	7.1
Corr. engine mass flow	kg/s	1478	1344	1489
Corr. fan rotational speed	rpm	2619	2362	2679
Fan root pressure ratio		1.4	1.3	1.4
Fan bypass pressure ratio		1.7	1.5	1.7

Table 1.2: Geometric parameters of the baseline fan blade.

Hub-to-tip radius ratio	0.35
Aspect ratio (height/mid-height chord)	1.7
Number of blades	19
Tip diameter [m]	3.26

for all optimisation work performed in this thesis. The through-flow computations were performed with a VKI in-house developed tool named ACPreDesign. The initial design process was purely aerodynamic, hence no structural performance was considered. An overview of some key geometrical parameters of the baseline blade, resulting from the preliminary design stage, is given in Table 1.2.

1.7 Thesis objectives

This thesis aims to introduce a multidisciplinary design optimisation framework for the efficient aerostructural design of modern high-bypass ratio aero-engine fan blades. The potential for the integration of fan flutter prediction techniques is further to be investigated.
More specifically, the thesis objectives can be stated as follows:

- Develop, implement and apply a Multidisciplinary Design Optimisation framework for the design of modern high-bypass ratio aero-engine fan blades. This involves the development of a suitable blade parametrisation, the integration of high-fidelity analysis tools for structural and aerodynamic performance analyses (Computational Fluid Dynamics and Computational Structural Mechanics), the derivation of the problem definition, translation to a suitable optimisation setup, and the assessment of the overall method's performance.

- Investigate the aeroelastic coupling strength for flutter prediction of modern lightweight aero-engine fan blades, in order to enable an assessment of common flutter prediction techniques and their potential for an integration into automated optimisation frameworks.

1.8 Thesis outline

The remaining part of this thesis is organised as follows:

Chapter 2 provides a review of the literature. The optimisation problem is introduced and common methods for its solution are presented. This is followed by an overview of common methodologies for three-dimensional turbomachinery blade shape parametrisations and a short review of high-fidelity performance analysis methods. Subsequently, a global overview of existing applications of optimisation methods for the design of transonic fan and compressor blades is given. The chapter also contains a short review of computational flutter prediction methods and concludes with a discussion of the current state-of-the-art.

Chapter 3 introduces the optimisation method used in present work. An explanation of the overall working principle of the method is given and an individual presentation of the core components of the system is provided.

Chapter 4 provides a detailed explanation of the developed fan blade parametrisation and the automated generation of the fluid and solid domains required for the high-fidelity structural and aerodynamic analyses.

Chapters 5 and 6 introduce the fan blade aerodynamic and structural design problems and describe the high-fidelity CFD and CSM analysis tools in detail. Specific modelling requirements for the fan design problem are highlighted.

Chapter 7 and 8 present the multidisciplinary design optimisation of a full titanium and a composite fan blade, respectively.

Chapter 9 presents an investigation of the aeroelastic coupling strength for blade flutter prediction by providing an in-depth comparison of coupled and decoupled flutter prediction methods.

Chapter 10 contains the main conclusions of the thesis and provides some suggested avenues for future work.

2 Literature Review

This chapter provides a review of the current state-of-the-art of single and multidisciplinary transonic fan blade design optimisation. The chapter is divided into seven sections. In the first and second section, the general optimisation problem formulation, related key concepts and terminologies are introduced, and an overview of existing numerical optimisation methods is given. The third section reviews existing methodologies for turbomachinery blade shape parametrisation. In the fourth section, high-fidelity performance analysis methods for the evaluation of the aerodynamic and structural performance of transonic turbomachinery blades are briefly presented and the problem of aeroelastic analyses is introduced. The fifth section reviews existing applications of numerical optimisation methods to transonic compressor and fan blade design. The following section provides a more targeted review of computational prediction methods for fan blade dynamic aeroelastic instabilities, known as flutter. In the final section a brief discussion of the state-of-the-art is given providing a link to the main objectives of the thesis.

2.1 Optimisation problem formulation

A constrained single-objective, non-linear engineering optimisation problem can be formally stated as [169]:

$$\text{minimise} : F(\boldsymbol{x}) \tag{2.1}$$

$$\text{subject to} : g_j(\boldsymbol{x}) \leq 0, j = 1, ..., m \tag{2.2}$$

$$h_k(\boldsymbol{x}) = 0, k = 1, ..., l \tag{2.3}$$

$$x_i^l \leq x_i \leq x_i^u, i = 1, ..., n \tag{2.4}$$

where

$$\boldsymbol{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_n \end{bmatrix} \tag{2.5}$$

denotes the vector of design variables. It contains the parametric variables of the design that can be modified by the optimiser, for instance, blade angles, chord lengths and blade thicknesses. But also boundary conditions such as rotational speeds, or material quantities may be defined. Further, design variables may be *continuous* or *discrete*. Continuous design variables are allowed to change continuously within a given range, defined by upper and lower bounds $x_i^l \leq x_i \leq x_i^u,\ i = 1, ..., n$. Discrete design variables are only allowed to take values from a limited number of permissible values. The number of blades in a turbomachinery blade row or the number of plies in a composite laminate are examples for discrete design variables. In some optimisation problems, both continuous and discrete design variables appear at the same time, such problems are termed mixed discrete-continuous problems.

The vector of design variables in combination with the *side-constraints* Eq. (2.4) defines the so-called *design space*.

An inequality constraint, Eq. (2.2), requires that the value of a performance measure is lower or higher than a specified limit. For example, peak stresses in a component under maximum operational loads are typically required to be well below the material's yield strength in order to ensure structural integrity. A minimum required choking mass flow of a turbomachinery blade row is another example for an inequality constraint. Constraints which are required to assume specific values during the optimisation are so-called *equality constraints*, Eq. (2.3). Some optimisation methods are unable to handle equality constraints directly. In those cases, an equality constraint may be approximated by two inequality constraints specifying a near identical lower and upper limit value.

The region in the design space in which all constraints are satisfied is referred to as the *feasible region* of the design space. Consequently, designs being located in the feasible region of the design space are denoted as feasible designs. Whereas designs that are violating one or more constraints are referred to as infeasible designs.

The objective function $F(\boldsymbol{x})$, Eq. (2.1), is a figure of merit, or measure of goodness, of a design and defines the quantity that needs to be maximised or minimised during the optimisation process. In a turbomachinery design context this might be e.g. the

efficiency, pressure ratio or blade mass, or a combination of those. In the optimisation problem formulation given above, a minimisation problem is considered. However, a maximisation problem can easily be translated to an equivalent minimisation problem by introducing a sign change of the objective function, i.e. the minimisation of -$F(\boldsymbol{x})$ is equivalent to the maximisation of $F(\boldsymbol{x})$.

It is not always adequate to cast a design problem into an optimisation problem with a single objective as defined in Eq. (2.1). Instead, in engineering design it is often desired to maximise or minimise several objectives at the same time. In practice those objectives are frequently conflicting with one another. The design problem is then referred to as a multi-objective optimisation problem, and is defined as

$$\text{minimise} : F_i(\boldsymbol{x}), i = 1, ..., l \tag{2.6}$$

$$\text{subject to} : g_j(\boldsymbol{x}) \leq 0, j = 1, ..., m \tag{2.7}$$

$$h_k(\boldsymbol{x}) = 0, k = 1, ..., n \tag{2.8}$$

$$x_p^l \leq x_p \leq x_p^u, p = 1, ..., q \tag{2.9}$$

In multi-objective problems a single solution does not exist. The concept of optimality therefore changes in comparison to a single-objective problem. To illustrate this, Fig. 2.1 shows the objective space for an optimisation problem with two objectives. The target is to simultaneously minimise the objectives **F1** and **F2**. Comparison of designs A and B shows that B has a better performance than A in both objectives. In this scenario design A is said to be *dominated* by design B, or inversely design B *dominates* design A. On the other hand, design C improves design A only in objective **F2** but worsens **F1**. Here a quantitative statement on which design is better cannot be given. A design for which no design can be found that simultaneously improves all objectives is said to belong to the so-called set of *non-dominated* designs, also known as the *Pareto optimal set*, named after the economist Vilfredo Pareto, who first utilised this concept in economics [119]. The associated curve that describes those designs in the objective space is the so-called *Pareto front*.

2.2 Optimisation methods

Many different numerical optimisation methods exist, and a few of the most prominent will be outlined in this section along with their respective advantages and disadvantages for the solution of engineering design problems and their capabilities to handle multidisciplinary design optimisation (MDO) problems.

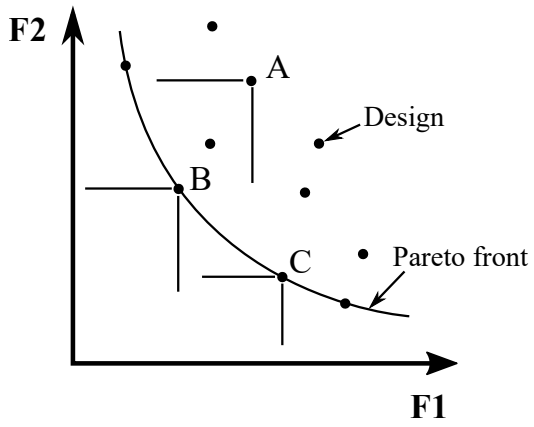


Figure 2.1: A schematic Pareto front for a two-objective optimisation problem. Design A is dominated by design B, but designs B and C are non-dominated.

A convenient way to classify optimisation methods is to consider the order of derivatives of the objective and constraint functions needed by the algorithm. Almost all optimisation methods are based on an iterative process, where a new design is computed based on a previous one by [169]:

$$\boldsymbol{x}_{i+1} = \boldsymbol{x}_i + \alpha_i \cdot \boldsymbol{S}_i \tag{2.10}$$

Where i is an iteration number, $\boldsymbol{S}_i$ is a search direction and α_i is a scaling factor determining the size of the step to be taken in the search direction $\boldsymbol{S}_i$. The optimisation methods generally differ in the way the step size α_i and the search direction $\boldsymbol{S}_i$ are computed.

2.2.1 Gradient-free methods

In gradient-free methods, also known as zero-order methods, only the objective function values and no derivative information is needed during the optimisation process. These methods have a number of advantages which make them suitable for many engineering design problems. They are easy to implement, can handle discontinuous, noisy and multi-modal functions, allow a direct handling of constraints and can be well adapted to multi-objective optimisation problems. Gradient-free methods only require the input-output relations (design variables as input and performance as output) of an evaluation tool and therefore can use any evaluation tool as a black-box function. Thus, established performance evaluation software can be integrated in a gradient-free optimisation process with a relatively small effort.

In addition, gradient-free methods are able to handle discrete optimisation problems, where the design variables are limited to a set of discrete values (e.g. number of blades). These methods do have, however, one major disadvantage, which is that they usually require thousands of objective function evaluations to find an optimum. This drawback prevents the direct application of gradient-free methods to engineering design problems where one function evaluation involves computationally expensive high-fidelity simulation methods like CSM and CFD. However, if the function evaluations are computationally cheap, for instance in the pre-design phase, gradient-free methods provide powerful and robust optimisation algorithms. Two classes of methods that are particularly suitable for engineering design problems will be shortly presented in the following.

2.2.1.1 Evolutionary algorithms

Evolutionary algorithms (EAs) are a class of optimisation methods which make use of the principles of biological evolution, where a population of designs is evolving in the design space by adopting mechanisms like cross-over, mutation and selection. The development of EAs can be traced back to the late 1950s [22], [53], but more recent developments include four strategies that are notable for the application in engineering design: Evolutionary Programming [50], Evolution Strategies [128], Genetic Algorithms [67] and Differential Evolution [157]. Many strategies have been proposed for the extension of EAs to multi-objective optimisation problems, cf. Deb [34]. EAs have been successfully applied in compressor pre-design by e.g. Joly et al. [79], Massardo et al. [98] and Teichel et al. [159].

2.2.1.2 Particle swarm optimisation

Particle swarm optimisation (PSO) is a more recent optimisation method developed by Eberhart and Kennedy [81]. Similar to EAs, PSO is a population based technique. It is inspired by the social behaviour of bird flocks and fish schools. In PSO each design is treated as a particle that is moving in the design space with a certain velocity, which is updated during the optimisation process depending on the best local and global objective function values found. Several adaptations to PSO have been proposed to incorporate constraints, e.g. [68], [122], [127] and for the extension to multi-objective optimisation problems, reviewed in [130]. A recent application of PSO to aerofoil design can be found in [179].

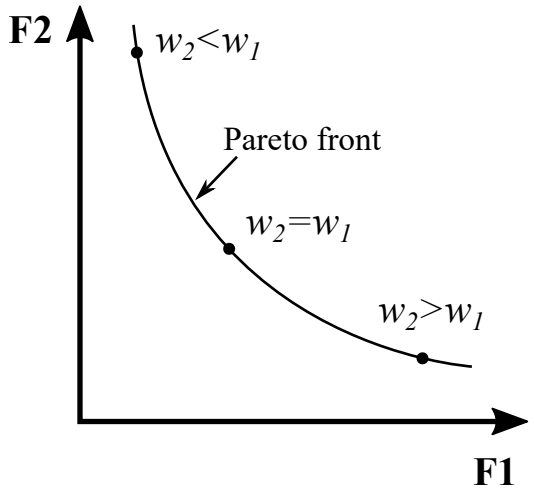


Figure 2.2: Samples along the Pareto-front obtained with different weights for the objectives using a single-objective weighted sum approach.

2.2.2 Gradient-based methods

Gradient-based methods make use of gradient information in order to provide the optimiser with a search direction. These methods are typically based on considerations from differential or variational calculus. Because the methods use gradient information, they are usually more efficient in locating an optimum compared to gradient-free methods considering the number of function evaluations needed.

However, gradient-based methods may perform poorly in the presence of noise and in non-differentiable and discontinuous search spaces. Because of their very nature, gradient-based methods are local optimisers, hence they easily get trapped in a local optimum, if existing. If the design space is multimodal (i.e. it exhibits several local optima), the starting point of the optimisation pre-determines the local optimum that the optimiser will converge to.

Another shortcoming of gradient-based methods is their inability to handle multi-objective optimisation problems in a direct fashion. To make a multi-objective problem tractable for a gradient-based approach, it needs to be translated to a single-objective one. This can be achieved by a weighted sum approach. For example an optimisation problem with two objectives can be translated to a single objective one as

$$F = w_1 \cdot F_1 + w_2 \cdot F2 \tag{2.11}$$

where w_1 and w_2 are the weights associated to the different objectives F_1 and F_2,

respectively. Those weights determine the location of the solution on the Pareto front, which is illustrated in Fig. 2.2.

Whether or not multimodality is a common issue in turbomachinery optimisation problems is a subject of ongoing debate among researchers. With the increasing number of applications of gradient-based optimisation methods in the recent past, there seems to be increasing evidence that multimodality is of far less concern than has previously been assumed. However, to date, this aspect still represents an open question, and more extensive and systematic comparisons of both approaches, gradient-free and gradient-based, might yield some needed insight to put the now rather diffuse understanding of the issue onto stronger footings. Some of the most popular gradient-based methods include steepest descent, conjugate directions and quasi-Newton methods [169].

2.2.3 Metamodeling in numerical optimisation

Many zero-order methods, like EAs and PSOs, are powerful optimisation tools, but suffer from one severe drawback: they require a tremendous number of objective function evaluations (often thousands of evaluations are needed) to converge to an optimum. If computationally expensive analysis tools are required which take hours or even days to converge, a direct application of such methods becomes prohibitive even in highly parallelised computational frameworks. An elegant solution to this problem dates back to Schmit and Farshi [141], who recognised about one decade after Schmit's first attempt of structural optimisation [140] that almost all successful applications of the methods were reported for smaller scale component levels. The reason was that the computational expense required for the optimisation of larger problems (with potentially higher relevance) far exceeded the available computational capacities at that time. To overcome this problem, Schmit and Farshi introduced the concept of sequential approximate optimisation. In this concept, the optimisation is not anymore performed using the expensive evaluation directly, but instead a mathematical approximation is constructed from information that is available from previous evaluations of the expensive function. The optimisation is then performed on the approximation, which is computationally cheap. Once the optimum of the approximation is found, it is re-analysed by the high-fidelity tools and the additional information is used to construct a new approximation. This procedure continues until the magnitude of change of the objective function between consecutive iterations drops below a specified tolerance [56]. The main attractiveness of this concept is the reduction of computational cost, but the effectiveness of the method strongly

depends on the accuracy of the constructed approximation model. Over the years, the approach developed into its own field of research, where the terms approximation model, surrogate model and metamodel (i.e. a model of a model) are often interchangeably used. However, in recent years the term metamodel appears to become the most frequently used in the engineering design optimisation community, which is why this term will be used throughout this thesis. Many different metamodels have been proposed in the past decades. A review of their development for structural design can be found in [15], while more recent activities aimed at generic engineering design problems are summarised in e.g. [149], [148]. In general, metamodels can be classified into local, global and mid-range approximations.

Global approximations are spanning over the entire design space. The generation of such a global approximation can be difficult, particularly if the number of design variables is large. However, some sophisticated global metamodels exist, some of the most notable candidates being Response Surface Methods [23], Artificial Neural Networks and Kriging [86]. Hajela and Berke reviewed the application of Neural Networks in structural design [58]. The main drawback of global metamodels is that the number of required function evaluations required for their construction increases dramatically with the number of design variables [56], hence this type of approximation is typically only practical for low dimensional design spaces. An alternative to building an approximation for the entire design space is to use a local metamodel, which only has a limited validity in the design space. Local metamodels are typically built with function and derivative information at the current design point. Different levels of complexity exist, which are normally derived from Taylor series expansion, ranging from linear approximations to more complicated versions like reciprocal and conservative approximations [56]. Their limited validity presents one of the main drawbacks of these methods for complex optimisation problems, particularly for multi-objective cases, because they provide no information about the global characteristic of the problem and are prone to get trapped in local optima. In addition, the availability of gradient information is almost always required. The presence of discontinuities or noise in the original problem is therefore a major limiting factor for the application of local approximations. Mid-range approximations (also known as multipoint approximations) provide a compromise between local and global approximations and have been applied with some success for various engineering problems, e.g. [161], [162] and for transonic blading design in [145]. As the name suggests, these methods use multiple points to generate a model that is valid in a wider region of the design space than a local approximation. Mid-range approximations use the information available from

previous function evaluations to continuously update and refine the existing model as more information becomes available throughout the optimisation process.

Various powerful optimisation strategies have been developed by coupling approximation concepts to optimisation algorithms and high-fidelity evaluation tools. In the context of turbomachinery design this concept found early popularity and was for instance applied by Ahn and Kim (2003) [6], for the optimisation of a transonic compressor rotor using a response surface method. The applications will be more thoroughly discussed in Section 2.5 of this chapter.

2.3 Blade parametrisation methods

For shape optimisation of turbomachinery blades, the optimiser requires a means to change the design. A parametric blade model is thus required, which enables the modification of a blade geometry via a set of design parameters.

Several methods exist for the parametrisation of blade shapes, among which three of the most relevant will be discussed below.

2.3.1 Node-based parametrisation

In a node-based shape parametrisation approach, the coordinates of each individual surface node of a computational finite volume or finite element mesh become a design variable and are therefore allowed to change during the optimisation process, as for example shown by Jameson [73]. This parametrisation strategy represents the richest design space possible with respect to the resolution of a given numerical spatial discretisation. Typically, the nodal displacements are restricted to the direction normal to the element surface to avoid that in-plane node movements lead to a deterioration of the mesh quality or degenerated elements without affecting the shape. Due to the large number of design parameters, node-based parametrisation approaches are only applicable for gradient-based optimisation methods coupled to adjoint-based sensitivity calculations.

A node-based parametrisation method is easy to implement as it avoids the need for a dedicated parametrisation preprocessing step. It further offers a high level of local control.

However, common shortcomings of the approach are that the freedom to move each individual mesh node can result in irregular shapes and potentially non-smooth shape derivatives. Second order finite volume schemes are generally blind to high frequency noise in the surface grid, such that shape perturbations with high frequency

content are undamped and can grow until they destroy the solution. In addition, mesh dependency of the solution can become an issue. Common methods to overcome those problems are the application of filtering techniques e.g. [33], [89].

An additional disadvantage of the method is a missing link to a CAD geometry, which is ultimately needed for manufacturing of the optimised design. This has been a long standing point of criticism for the application of the node-based approach in an industrial framework, because manual reverse engineering of a CAD model from an optimised mesh can be a painstaking and difficult process. In addition, the resulting CAD model would most often only result in an approximation of the optimised mesh. Nevertheless, recent advances have been made to automate the generation of complex CAD surfaces from irregular point clouds using so-called Truncated Hierarchical B-Splines (THB-Splines) [83], which have the potential to reduce the gap between computational mesh and CAD model. Other fitting methods were reported based on NURBS curves or surfaces [16] and a few commercial tools have emerged that can perform geometry transformations between the two domains [1], [2].

In an MDO context, different disciplines involved typically have different mesh requirements with regard to density and node distributions, which results in non-matching meshes. Even within one single discipline, the required mesh densities may differ depending on the desired effect to be resolved, e.g. a finite element modal analysis can be accurately solved with a coarser mesh than is required for an accurate stress analysis, where areas of stress concentrations, like in fillets, require a higher mesh resolution. Those mesh differences lead to inconsistent shape parametrisations and correspondingly different geometry modifications between disciplines, which is an unwanted feature for MDO where one consistent shape parametrisation approach is needed.

2.3.2 Free-form deformation (FFD)

Free-form deformation is a volumetric deformation technique that was introduced in 1986 by Sederberg and Parry [143]. Other volumetric deformation methods exist [10], but FFD is today the most used among those.

The FFD method is based on the definition of a trivariate Bezier volume which encompasses the design, and whose control points represent the design variables. Modifications of the control point coordinates lead to a deformation of the control point net and a corresponding deformation of the geometry inside the volume. The deformation function within the domain is defined by a trivariate tensor product Bernstein polynomial.

FFD approaches are very versatile. They can be applied to virtually any underlying geometric representation (including commonly used boundary representation CAD models), without altering its topology. This is a major advantage when the deformable structure has internal features that need to be preserved, e.g. an aircraft wing with a fuel tank, internal stiffeners, ribs and spars, etc. Further, FFD can be applied locally or globally and can conserve derivative continuity between deformed and undeformed regions of the geometry.

A major disadvantage of the approach is however, that the control points often do not have an apparent physical meaning, which makes their interpretation difficult but also aggravates the setup of a compact and efficient parametrisation. In addition, geometrical constraints are difficult to implement, if at all possible [9].

Examples for the application of free-form deformation methods to problems in aerospace design are reported in [65], [138], and [137]. A comparison of CAD-based (see section 2.3.3) and FFD parametrisation approaches for turbomachinery shape optimisation using an adjoint-method was presented in [9]. A recent application of an FFD approach for the shape optimisation of a transonic compressor was shown in [77].

2.3.3 CAD-based parametrisation

In a CAD-based parametrisation approach the geometry is modified via a change of the geometric representation in form of the CAD-model. Therefore, no domain conversion is necessary for manufacturing. In addition, the approach provides a mesh independent geometry parametrisation and is therefore well suited for the application in a multidisciplinary design optimisation framework. CAD-based parametrisation approaches can be divided into two categories: a) constructive and b) surface-based parametrisation methods.

In a constructive parametrisation approach e.g. [79], [117] the parametrisation is built-in into the geometry construction process, i.e. the parametrisation is defined in the CAD modeller that is used for the generation of the geometry. The approach typically makes use of parametric and feature-based CAD modelling techniques, allowing relevant geometric parameters, such as lengths, angles and thicknesses, to be defined during the construction process. A further level of control for complex geometries can be obtained by defining distributions of geometric parameters using additional parametric curves, such as Bézier or B-Spline curves. As a result, the approach enables very good and in many cases direct control of geometric constraints and facilitates the avoidance of the generation of infeasible geometries. After a change of one or

more parameters, the entire geometry is *re-constructed*, hence its name. The generation of an efficient and robust parametrisation following this approach can be difficult for complex geometries and involves meticulous planning a-priori to the construction of the geometry. The setup is therefore lengthy and labour intensive. However, once a well defined parametric model is existing, the approach provides a powerful means for geometry manipulation. A disadvantage stems from the dependency of the parametrisation from the software tool used for its generation.

In a surface-based CAD parametrisation approach e.g. [94], [65], [178], the control point coordinates of the geometry-bounding surface patches are used as design variables. In contrast to the constructive parametrisation method, the surface parameters are CAD-software independent. This approach enables an increased level of local control and provides the richest design space possible for a given CAD model. However, for the implementation of geometrical constraints it suffers from the same limitations as the node-based and FFD parametrisation methods discussed above. Due to the large number of design variables normally resulting from this parametrisation approach, it lends itself to adjoint-based shape optimisation, but is generally ill-suited for gradient-free techniques.

2.4 Performance analyses

2.4.1 Steady flow analysis

Since the 1990s, steady CFD methods for the solution of the three-dimensional Navier-Stoked equations are routinely applied in the aerodynamic design process of transonic fans [76], [131]. Computational methods and turbulence models have continuously evolved, resulting in an improved modelling fidelity and reduced computational cost. It is widely recognised that steady CFD methods are capable to accurately predict design point performance and performance trends between designs, while accurate off-design performance analysis in terms of absolute numbers still presents a significant challenge. Development efforts are continuing to this day to further improve modelling fidelity and reduce computational cost via software and hardware improvements such as through the use of massive parallelisation with General Purpose Graphics Processing Units (GPGPU).

Steady RANS computations have been successfully integrated in early single-disciplinary transonic fan blade optimisations, e.g. [5], [6], [17] and their application can be considered as state-of-the-art.

2.4.2 Structural analysis

Modelling of the structural static and dynamic blade response using the Finite Element Method (FEM) is established practice in today's industrial design process. The introduction of novel advanced composite materials and manufacturing methods, however, introduces additional modelling challenges due to the material's intrinsic anisotropy and complex failure characteristics on the micromechanical level. Transonic fan blade optimisation cases have to this date been dominated by structural blade models based on isotropic materials, which significantly reduces the structural modelling effort.

2.4.3 Aeroelasticity

Aeroelastic effects are of significance for fan blades of large aero-engines, but have thus far been largely neglected in automated design optimisation frameworks. Flutter, which is a harmful dynamic aeroelastic instability, is of particular concern for modern fan blades. The accurate computational prediction of fan flutter presents to date still a challenging task. Many computational aeroelasticity methods have been developed for that purpose, enabling various levels of modelling fidelity. However, a clear indication on the most suitable method for blade flutter prediction is not existing. Limits of prediction validities of the different available methods need to be known a-priori to an integration into an optimisation framework. A more detailed discussion of this matter will be given separately in Section 2.6 of this Chapter.

2.5 Review of transonic compressor and fan blade design optimisation

This section provides an overview of existing applications of optimisation methods for the design of transonic fan and compressor bladings. Much of the progress in this field has been achieved relying on two publicly available test cases; the transonic compressor NASA rotor 37 and the transonic fan NASA rotor 67. In this literature review it was found that those applications represent a good 75% of all work performed in this research field, and account for 100% of the early applications from the early 2000s until 2008, after which a slight diversification of applications occurs. To account for this development, this section is divided into three application specific parts. The first part solely covers optimisations of the rotor 37 and rotor 67 test cases. This part of the review enables to trace the main milestones in the advancement of

the applied optimisation methods and the increasing complexity of the optimisation problems being solved, ranging from initially single-objective, unconstrained aerodynamic ones, towards multi-objective, constrained, multi-disciplinary and multiple operating point optimisation cases. The remaining applications in the literature are covered in two separate sections discussing the design optimisation of compact highly loaded compressors and finally applications to high-bypass ratio aero-engine fan blade designs.

2.5.1 NASA Rotor 37 and Rotor 67 optimisation

As mentioned above, by far the majority of reported applications of numerical optimisation methods for three-dimensional transonic fan and compressor design are based on the publicly available test cases NASA rotor 37 [129] and 67 [158]. Early investigations started to appear in the early 2000's and were purely aerodynamic in nature. The first applications were characterised by relatively simple, single-objective, unconstrained optimisation problems. Ahn and Kim (2003) [6] utilised a response surface based optimisation using a linear programming method for the maximisation of the isentropic efficiency of rotor 37 using three design variables for the definition of the rotor stacking axis. The lack of constraints, however, leads to trivial solutions with increased efficiency but decreased rotor pressure ratios. Extensions to single-objective constrained problems using a genetic algorithm were soon after proposed by Oyama et al. (2004) [118] for rotor 67. A rather large scale optimisation was performed with a total of 56 design parameters to define the blade shape. Soon after, Benini (2004) [17] successfully extended the approach to multi-objective constrained optimisations of rotor 37 using an in-house developed evolutionary algorithm [160], [18].

While those early applications showed that the direct coupling of evolutionary algorithms to three-dimensional Navier-Stokes solvers proved successful in terms of optimisation outcomes, the actual optimisation run-times turned out to be prohibitively expensive for the implementation of the methods into an industrial design setting. The turnaround time of Benini's rotor 37 optimisation [17] was 2000 hours (or about 83 days) on a four processor workstation. Oyama's rotor 67 optimisation [118] required 1550 hours (about two months) on a 64-core SGI cluster.

To enable more practically relevant design turnaround times while maintaining the advantages offered by evolutionary algorithms, Lian and Liou (2004) [90] presented the optimisation of NASA's rotor 67 with a genetic algorithm coupled to a second-order polynomial response surface metamodel, similar to Ahn and Kim's earlier approach [6], but this time for a more complex optimisation problem. The approach

was shown to successfully reduce the total turnaround time of a single-objective constrained optimisation to roughly 5 days, although no specification on the number of design variables was given. The method was later extended by the same authors in [91] for the same test case. Without mentioning it, Lian and Liou presented the first multidisciplinary optimisation of a transonic fan, including the blade weight into the optimisation problem formulation as a second objective in addition to the aerodynamic performance requirements.

Optimisations using response surface metamodels have subsequently gained great popularity in the community, with further applications being reported by Jang et al. (2005) [74], Yi et al. (2006) [182], Jang et al. (2006) [75], Wang et al. (2008) [173], Samad and Kim (2008) [136] and Kim et al. (2013) [82] for rotor 37 and Zheng et al. (2010) [183] for rotor 67.

The first "true" multidisciplinary optimisation of a transonic fan was introduced by Lian and Liou (2006) [92], where, in addition to the high-fidelity aerodynamic performance evaluation by CFD, a structural finite element solver was included in the optimisation process to evaluate the static stresses in the rotor blade.

Extensions to multiple operating points followed as presented by Pierret (2007) [123]. In contrast to the previously used static metamodels, Pierret made use of a sequentially updated metamodelling approach, which continuously improved metamodel prediction accuracy during the optimisation process. This method will be described in further detail in Chapter 3 of the thesis.

A similar approach using an adaptive Kriging metamodel was later used by Song et al. (2014) [153] for the aerodynamic design-point optimisation of rotor 67. Shahpar et al. (2008) [145] reported on a successful application of a mid-range approximation method for the aerodynamic optimisation of rotor 37.

Gradient-based optimisations using the adjoint method for sensitivity calculations have received increasing attention in the recent past. Single point (Duta et al. (2007) [43]), multi-point (Luo et al. (2013) [94]) and concurrent aerodynamic and aeroelastic optimisations (He and Wang (2011) [62]) have been performed.

2.5.2 Highly loaded compact transonic compressor optimisation

The design of small transonic compressors with very high loading, i.e. exceeding stage pressure ratios of 1.8, presents a particular challenge in terms of both aerodynamic and structural design. A limited amount of past design experience exists for those machines, therefore making them prime candidates for the application of

design optimisation methods. Siller et al. [147] and [146] presented an aerostructural optimisation of a compact highly loaded fan stage with a tandem stator. The applied optimisation method was based on a Kriging-assisted Evolutionary Algorithm. High-fidelity CFD and CSM codes were used for the aerodynamic and structural performance evaluations, respectively. A constrained optimisation problem was solved, while the isentropic efficiency and stall margin were defined as objectives.

Joly et al. [79] presented a purely aerodynamic as well as a multidisciplinary optimisation of a highly loaded compact fan stage aiming at a stage pressure ratio of 2.1. A multi-objective Differential Evolution algorithm [157], [35], [95] was applied without any metamodel. Isentropic efficiency and the maximum von Mises stress in the blade were defined as objectives in the aerostructural optimisation, while constraints on the pressure ratio and the mass flow were imposed. After 20 generations, the chosen geometry from the Pareto optimal set achieved an efficiency gain of 0.4% with a maximum von Mises stress value of 639 MPa, representing a 65% stress reduction with respect to the aero-only optimisation.

2.5.3 High bypass ratio aero-engine fan blade optimisation

Despite the relative long history of successful optimisation applications using the publicly available test cases rotor 37 and rotor 67, only two shape optimisation cases for large high-bypass ratio aero-engine fan blades were found in the literature.

Wilson et al. (2011) [176] used a multidisciplinary optimisation system for the reduction of the tonal noise of a high-bypass ratio fan blade. Aerodynamic, aeroacoustic and structural performance evaluations were included in the optimisation process. Aerodynamic and aeroacoustic performances were quantified using a three-dimensional CFD solver, while a structural finite element code was used to predict the stresses in the blade. The authors report on significant challenges that were faced with regard to the generation of a meaningful blade parametrisation allowing to improve the noise production while not significantly affecting the aerodynamic performance of the baseline design. Ultimately, a blade parametrisation with 5 design parameters was used to define the localised tip leading edge sweep distribution, which led to a successful reduction of the generated tonal noise while satisfying the imposed aerodynamic and structural constraints.

Deng et al. (2013) [38] reported on purely aerodynamic as well as aeromechanical optimisations of a high bypass ratio engine fan blade made integrally from a titanium material. A multi-objective genetic algorithm was applied, connected to a sequentially updated Kriging metamodel. The authors highlight the importance

of the fan blade root geometry for an accurate representation of the structural performance and therefore include it in the solid blade model used in the optimisation process. However, no details on the actual methodology of the root construction nor its parametrisation are provided in the paper. Overall 33 design parameters were used to define the blade geometry. Two objectives were defined for both optimisations, the first being the isentropic efficiency, and the second the slope of the cruise speedline at 95% of the cruise design point mass flow. Two constraints were imposed for the purely aerodynamic optimisation, while in total nine constraints were imposed for the aerostructural optimisation. Both optimisation approaches showed successful outcomes, with a design and off-design efficiency increase of 0.4% and 0.8% respectively, achieved in the aerodynamic optimisation. With structural constraints included, an efficiency gain of 0.2% and 0.67% was achieved, respectively, for the same operating points.

2.6 Computational blade flutter prediction

Flutter is a self-existed aeroelastic instability which can cause rapid structural failure and is of major concern for modern highly-loaded, lightweight aero-engine fan blades. In a multidisciplinary fan blade design optimisation process, it is generally desirable to include a measure of the flutter stability to assess the risk of a potential flutter onset. Various methods are used to this end in current manual design practice, often being based on empirical correlations derived from experience with past designs. In such frameworks, shortcomings of the models are usually offset by the experience of a human designer who can assess the region of the model's validity. However, in an automated design optimisation framework, where designs might be generated that substantially deviate from conventional ones, an accurate computational flutter prediction tool is required to avoid that the optimiser drives the design into a wrong, and potentially unstable, direction.

In the field of computational aeroelasticity, turbomachinery flutter has been treated mainly as a decoupled phenomenon to date, based on an early proposed method by Carta [28], most commonly known as the *energy method*. In this method, aeroelastic coupling effects are assumed to be negligible, which implies that unsteady aerodynamic forces acting on the blade surface are small enough to not alter the vibrational behaviour of the blade. In other words, in the aeroelastic system with the blade subjected to an airflow, the blade is assumed to vibrate purely in one of its structural modes. The implication for the prediction of fan flutter is that the feedback from

the fluid to the structure does not need to be taken into account, and the flutter behaviour can be predicted by imposing the structural free vibration in terms of bladed disk modeshape and eigenfrequency in an unsteady flow computation, hence the term *decoupled*. The application of this method considerably simplifies the prediction of flutter. Further, the unsteady flow field resulting from a forced harmonic blade motion can be computed using computationally efficient frequency domain methods, adding to the method's attractiveness for industrial use as well as for its implementation into automated MDO frameworks.

However, for large, highly loaded aero-engine fan blades with low mass density and low stiffness, the magnitude of unsteady fluid forces may become of significance in relation to the structural ones, potentially leading to non-negligible coupling effects between structure and fluid. Therefore, the validity of fluid-structure decoupling as used in the decoupled approach becomes questionable. Despite the existence of a large body of research work on the development and application of various flutter prediction methods, publications targeting the validity of the energy method for lightweight blade designs are relatively scarce and are often limited to two-dimensional cases or subsonic flows. This is not least because of the significant computational cost associated with fully coupled fluid-structure interaction analyses for three-dimensional transonic cases, that are needed as a reference for the comparison with the decoupled approach. This cost is further increased by the need to compute a large number of vibration cycles for an accurate determination of the aeroelastic response frequency.

In an early comparative study conducted by He and Zhou [63], a coupled eigenvalue method is proposed, based on an unsteady potential flow model and a two degree of freedom structural model in two dimensions. In a comparison with experimental flutter stability boundaries of two transonic test rotors, it was found that mode coupling played a significant role for the aerodynamic damping characteristics. Consequently it was concluded that the single mode energy method can not correctly predict the flutter onset in the cases under investigation.

Srivastava and Reddy [156] carried out a comparative study of three computational flutter prediction methods, including a coupled, decoupled and eigenvalue analysis method for a ducted propfan in subsonic flow. The unsteady flow was modelled by solving the three-dimensional Euler equations in the time domain while the blade structure was modelled using a normal modes method. No detailed account of the blade material or mass ratio is given in the paper, but based on the given reference the blade is a composite blade with a mass ratio of 33. Coupled bending-torsion unstalled flutter was consistently predicted by the coupled and eigenvalue approach. However,

the work-per-cycle method, being intrinsically a single mode analysis method, was naturally unsuited to predict this flutter mode.

Moyroud et al. [109] investigated the effect of material changes on the predicted flutter stability using the energy method and an eigenvalue analysis method for a fan configuration in subsonic flow. A quasi-three dimensional linearised potential flow solver was applied based on a finite element discretisation. A finite element blade structural model constituting a full titanium, full graphite epoxy laminated composite and a sandwich configuration with titanium skins and honeycomb core were investigated. The results indicate significant differences between the stability predicted by the energy method and the eigenvalue analysis method. Mode shape approximation (i.e. one mode only) and frequency approximation (vibration at the blade's natural frequency) were found to introduce roughly similar error magnitudes in the case considered.

In a study on the flutter behaviour of a linear cascade of biconvex uncambered airfoils operating at zero degree incidence in subsonic flow, Sadeghi and Liu [134] investigated the shift of the aeroelastic response frequency from the natural vibration frequency for varying values of mass ratio. By identification of the frequency of neutral stability, i.e. the frequency at which the net work exchanged between fluid and blade is zero, it was found that the potential frequency shift can be as high as 5% for a typical titanium or steel material and significantly higher for lower density materials. It is subsequently assumed, that at low mass ratios the assumption of blades oscillating at their natural frequency as applied in the energy method would lead to significant prediction errors of flutter stability.

2.7 Discussion

– It is clear from the literature review that many successful attempts have been made to date for the optimisation of transonic compressor and fan blades. Optimisation problems that have been solved are ranging from unconstrained single-objective aerodynamic cases to constrained, multi-objective, aeromechanical (multidisciplinary) ones considering multiple aerodynamic operating points. Recently, interest in gradient-based aerodynamic and aeromechanical optimisations using the adjoint approach is gaining increasing momentum. However, for global optimisation, gradient-free Evolutionary Algorithms coupled to global metamodeling techniques are found to provide good results in the reported cases found in the literature.

– The majority of optimisation cases have been reported for small rotors with NASA's rotor 67 and 37 test cases being by far the most popular. A severe lack of work for the multidisciplinary design optimisation of large aero-engine fan blades is identified. This is despite the major practical relevance of the subject. The few existing studies are either a) limited to relatively simple, localised blade shape changes, b) are single disciplinary (structural optimisations), or c) provide only an incomplete, and oversimplified representation of the design problem.

– High-fidelity aerodynamic performance evaluations using Computational Fluid Dynamics and structural static and dynamic computations using Computational Structural Mechanics analysis tools have been successfully integrated into the optimisation workflows. The state-of-the-art is represented by steady RANS methods for the aerodynamic modelling, and Finite Element computations based on solid elements and isotropic material models for the structural computations.

– Modern fan blades are made from carbon-fibre reinforced plastics (CFRP) composite materials. Structural modelling approaches are needed that allow for the integration of those materials into the shape design optimisation framework.

– The bypass splitter, which is responsible for separating the fan exit flow into bypass and core streams, is a central component of the aerodynamic fan design problem, as it influences the required radial fan loading distribution. Not a single optimisation case was found that included the bypass splitter in the optimisation problem.

– The solid root geometry of fan blades is frequently neglected in the literature despite its central influence on the overall structural response of the blade (both static and dynamic). The root geometry of aero-engine fan blades is complex, and in an automated shape optimisation framework it needs to be able to adapt to blade geometry changes. This poses a significant problem for a robust geometry generation process, but nevertheless represents a necessary component if a realistic representation of the structural response is to be obtained.

– Dynamic aeroelastic instabilities (flutter) are known to play a critical role for fan designs, but are mostly disregarded in automated design optimisations to date. Computational cost is likely one of the reasons. Another, more pressing issue, is an uncertainty as to which prediction method (and which fidelity in terms

of modelling of the fluid-structure interaction effects) is needed and sufficiently accurate to qualify for an integration into an automated design framework. Various methods are existing, ranging from so-called decoupled (no feedback from fluid to structure) to coupled methods (two-directional effects between structure and fluid are modelled). Decoupled methods are computationally cheap and frequently applied. However, as it was shown in this review, some evidence in the literature seems to suggest that fluid-structure coupling is non-negligible in large aero-engine fans, although no clear consensus is existing. Further work is needed to clarify this question.

3 Optimisation Method

3.1 Optimisation system

The optimisation system used in present work is the result of two decades of continuous research and development efforts at the von Karman Institute. A flow chart of the optimisation system is shown in Fig. 3.1.

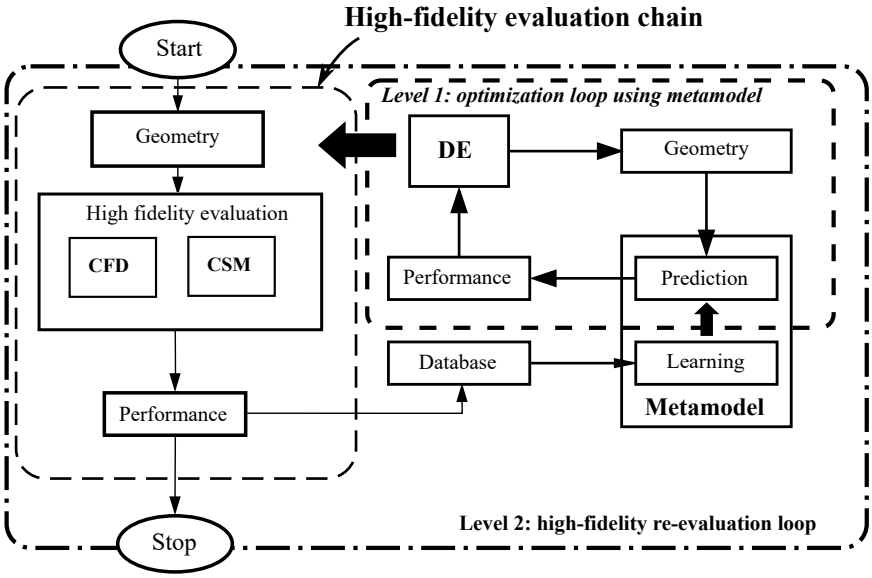


Figure 3.1: Flow chart of the optimisation system.

The core components of the system are a multi-objective Differential Evolution algorithm (DE), a database, several metamodels, including Radial Basis Functions, Artificial Neural Networks and Kriging, and a high-fidelity evaluation chain consisting of a fully automatic geometry generation process, automatic meshing, high-fidelity performance evaluations by Computational Fluid Dynamics (CFD) and Computa-

tional Structural Mechanics (CSM) analysis tools and automatic post-processing and performance parameter extractions. In the context of the present thesis, the fully automated high-fidelity evaluation chain for the fan blade optimisation was developed and implemented. The optimisation strategy, including the DE and metamodel implementations were existing at the outset of this work and have not been altered. Their explanation will thus be kept brief. For further details the reader is referred to [170] and [171]. An explanation of the overall working principle of the system follows. A pseudo-code of the method is provided in Algorithm 1.

Before the optimisation is initialised, an initial database is generated using a suitable design space sampling technique, whereas each sample is evaluated by the high-fidelity evaluation chain. The resulting relationships between design variables and performance parameters are stored in the database.

Algorithm 1 Metamodel-assisted multi-objective DE

```
create an initial database
for iteration=1 to limit do
    generate metamodel on database
    perform multi-objective DE using the metamodel
    select a specified number of candidate individuals from the Pareto front using a
    distance metric
    evaluate the candidate individuals by the high-fidelity analysis tools and add the
    result to the database
end for
```

The optimisation process is based on a two-level strategy. It is initialised by generating the first metamodel based on the initial database samples. On the first level, the evolutionary optimisation is performed using exclusively the metamodel for the performance evaluations. Since metamodel evaluations are computationally cheap (the computational cost for a single evaluation is typically in the order of milliseconds), thousands of generations can be computed by the Differential Evolution algorithm within seconds. However, the metamodel predictions at early stages of the process can be expected to be rather inaccurate. Therefore, a set of non-dominated designs is selected from the predicted Pareto front which are submitted for re-evaluation to the high-fidelity evaluation chain. The results are added to the database which is used to re-train the metamodel, thus increasing its prediction accuracy in the regions where it previously predicted optimal designs. The high-fidelity re-evaluation process constitutes the second level of the optimisation system. A full iteration consists of the first and second level processes (metamodel generation, optimisation, high-fidelity

re-evaluation and addition of the results to the database). Once new designs have been added to the database, a new iteration is started. The process is repeated until a specified number of iterations is reached or the process is manually stopped.

The individual components of the optimisation system will be explained in some further detail below.

3.2 Description of core components

3.2.1 Differential Evolution

Differential Evolution (DE) is an evolutionary algorithm (EA) developed by Price and Storn [157]. Like other EAs, DE is a population based optimisation method. DE was originally developed for global optimisation over continuous search spaces and has shown to outperform Genetic Algorithms for some single-objective [125] and multi-objective problems [164].

In the original single-objective, unconstrained implementation of DE, an individual is represented by the parameter vector in generation t as

$$\boldsymbol{x}_t = (x_1, x_2, ..., x_n) \tag{3.1}$$

where n is the number of design variables. A mutation operator is defined based on the choice of three other randomly chosen parameter vectors from the current population $\boldsymbol{a}_t$, $\boldsymbol{b}_t$ and $\boldsymbol{c}_t$, which are combined to yield a mutant vector

$$\boldsymbol{y}_t = \boldsymbol{a}_t + F \cdot (\boldsymbol{b}_t - \boldsymbol{c}_t) \tag{3.2}$$

where $F \in [0, 2]$ is a real-valued constant controlling the amplification of the differential variation $(\boldsymbol{b}_t - \boldsymbol{c}_t)$. A candidate vector is generated by using a cross-over operation of the form

$$z_i = \begin{cases} y_i, & \text{if } r_i \leq C \\ x_i, & \text{if } r_i > C \end{cases} \tag{3.3}$$

for $i = 1, ..., n$ and where r_i is a uniformly distributed random variable $r_i \in [0, 1[$ and C is the so-called cross-over constant $C \in [0, 1]$.

Lastly, considering a minimisation problem a greedy selection operator is defined as

$$\boldsymbol{x}_{t+1} = \begin{cases} \boldsymbol{z}, & \text{if } f(\boldsymbol{z}) \leq f(\boldsymbol{x}_t) \\ \boldsymbol{x}_t, & \text{if } f(\boldsymbol{z}) > f(\boldsymbol{x}_t) \end{cases} \tag{3.4}$$

Therefore, the individual $\boldsymbol{x}_t$ will remain a member in the next generation if its objective function value is lower than the candidate's vector one, otherwise it is replaced by the candidate vector.

For constrained optimisation problems, a penalty function approach may be used in which the objective function is essentially augmented by the constraint functions, but the selection operator can also be adapted to enable a direct handling of constraints. This can be done in the following way: If the original vector satisfies all imposed constraints it will be replaced by the candidate vector only if it equally satisfies all constraints and has a better objective function value. On the other hand, if the original vector does not satisfy the constraints, it will be replaced by the candidate vector if it satisfies the constraints better even if this leads to a higher objective function value. This places the first priority on the identification of the feasible design space and then switches towards objective function improvement.

Many methods have been proposed to enable the solution of multi-objective problems with DE, e.g. [3], [12], [7]. In the present approach, a Pareto ranking process similar to the method proposed by Deb et al. [35] is used. The methodology is based on an adaptation of the selection operator. In a first step, candidate vectors resulting from mutation and cross-over are added to the population, which results in a population twice as large as the original one. Then, a ranking process takes place in which the original size of the population is recovered. First, all non-dominated solutions are assigned with a rank 0 and are temporarily removed from the population. Next, the rank 1 is assigned to the non-dominated individuals in the remaining population and the process is repeated until all individuals have been assigned with a rank. The population is then restored to its original size by filling it with the individuals of consecutive ranks starting from 0, until the original population size is restored. A distance metric is employed to chose individuals from a rank that cannot be included entirely in the new population in order to retain the original population size.

3.2.2 Database

The database stores the relationship between design variables and performance parameters for designs that have been analysed by the high-fidelity evaluation chain. It serves as the knowledge base to generate the metamodels. As noted above, an initial design space sampling is performed in order to generate the first metamodel

and initialise the optimisation process. In this work a fractional factorial Design of Experiments technique [107] is used for that purpose. During the course of the optimisation, new designs are added to the database in each consecutive iteration, thus adding information at strategically relevant points in the design space in order to increase the prediction accuracy of the metamodel.

3.2.3 Kriging

Kriging is used in this work as global metamodel and computationally cheap performance predictor throughout the optimisation process. Kriging models have their origins in mining and geostatistical applications and were first proposed in the Master's thesis of South African geologist D.G. Krige [86]. A further development of the mathematics of the method was later presented by Matheron [99]. In the late 1980's the usage of Kriging models for the prediction of expensive deterministic computer models became of increasing interest [133]. The application of Kriging models for global optimisation started to gain significant momentum in the early 2000's, as for instance reported in [21], [11] and [139]. A recent review of Kriging models for the approximation of computer simulations is given in [84].

The usual purpose of a Kriging model in engineering design optimisation is to approximate the response of an expensive, physics-based, high-fidelity computational model (commonly CFD and/or CSM) with a mathematical model, that is cheap to evaluate. If the response of a computer code (such as the efficiency, pressure ratio, material stress, etc.) to a set of input design variables $\boldsymbol{x} = [x_1, x_2, ..., x_n]$ is denoted as $f(\boldsymbol{x})$, then an approximation of the computational model is sought as

$$\tilde{f}(\boldsymbol{x}) \approx f(\boldsymbol{x}) \tag{3.5}$$

which is to be constructed based on a finite number of systematically generated observation samples (i.e. designs that have been analysed by the expensive computational tools using, for example, a Design of Experiments sampling technique).

Several types of Kriging exist. Most Kriging models are interpolating, hence the approximation matches the known high-fidelity sample points that are used for the Kriging model generation exactly. However, also approximative Kriging models are existing, which are often used in stochastic optimisation. The model that is used in this thesis is known as *ordinary Kriging*, which presents also the most widely applied method in the engineering design optimisation literature. For the sake of conciseness this method will simply be denoted as Kriging for its further description

in this chapter and in the remaining part the manuscript, unless otherwise stated. In mathematical form, a Kriging model is given as

$$\tilde{f}(\boldsymbol{x}) = \mu + Z(\boldsymbol{x}) \tag{3.6}$$

where μ denotes the simulation output averaged over the design space, and $Z(\boldsymbol{x})$, is used to provide a localised correction to the model in order to match the given data points. A more complex form of Kriging, known as *Universal Kriging*, uses a regression model [84] instead of the constant μ, but ordinary Kriging is often sufficient in practice.

The function $Z(\boldsymbol{x})$ is usually formulated such that $\tilde{f}(\boldsymbol{x})$ becomes the so-called *best linear unbiased predictor*. A linear predictor means that the prediction is formed by a linear combination of the observation samples

$$\tilde{f}(\boldsymbol{x}) = \sum_{i=1}^{N} w_i(\boldsymbol{x}) \cdot f(\boldsymbol{x}_i) \tag{3.7}$$

where the weights $w_i(\boldsymbol{x})$ are determined to minimise the Mean Squared Error (MSE) of the predictor

$$MSE = E\left[(\tilde{f}(\boldsymbol{x}) - f(\boldsymbol{x}))^2\right] \tag{3.8}$$

while the *unbiasedness* condition requires the prediction to satisfy

$$E[\tilde{f}(\boldsymbol{x})] = E[f(\boldsymbol{x})] \tag{3.9}$$

The determination of the optimal weights requires the solution of a constrained, non-linear optimisation problem. For problems with a large number of design variables, the computational cost for the solution of this problem can increase dramatically due to the curse of dimensionality. The details of the solution procedure are out of the scope of this thesis, but have been discussed in detail in the literature. For further reference the reader is referred to [46], [149] and [84].

3.2.4 High-fidelity evaluation chain

The high-fidelity evaluation chain indicated in Fig. 3.1 contains the fully automated high-fidelity fan performance evaluation process, including Computational Fluid Dynamics and Computational Structural Mechanics analysis tools. It receives the vector of design variables $\boldsymbol{x}$ for an arbitrary individual as input and generates the required

performance parameters for the construction of the objective and constraint functions as output.
The process consists of four key components:

- Parametrisation and automatic geometry generation for fluid and solid domains
- Automatic mesh generation, and computational setup
- Automatic CFD and CSM computation
- Automatic post-processing and performance extraction

Each of those components involves significant complexity and therefore will be individually discussed in further detail in the following chapters.

4 CAD-Based Fan Parametrisation and Automatic Geometry Generation

4.1 Fan assembly overview

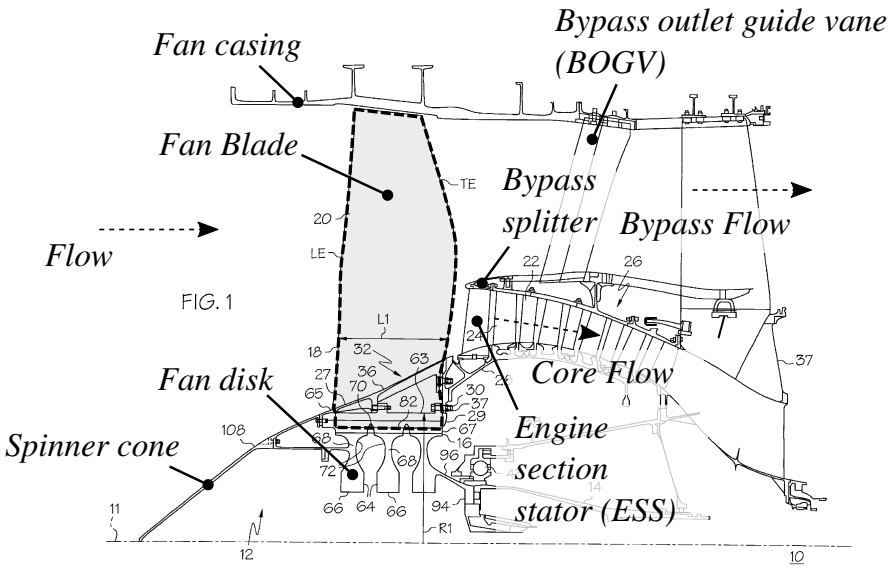
Figure 4.1 provides an overview of a typical fan section assembly. A meridional view is shown in Fig. 4.1 a) while a perspective view is given in Fig. 4.1 b).

The fan section comprises the fan blades, being structurally secured to a fan disk, a spinner cone, a fan casing, a row of bypass outlet guide vanes, a bypass splitter and a row of engine section stator vanes. The fan disk is connected via a drive shaft to the low pressure turbine at the rear of the engine, which provides the power to drive the fan (not shown in the figures).

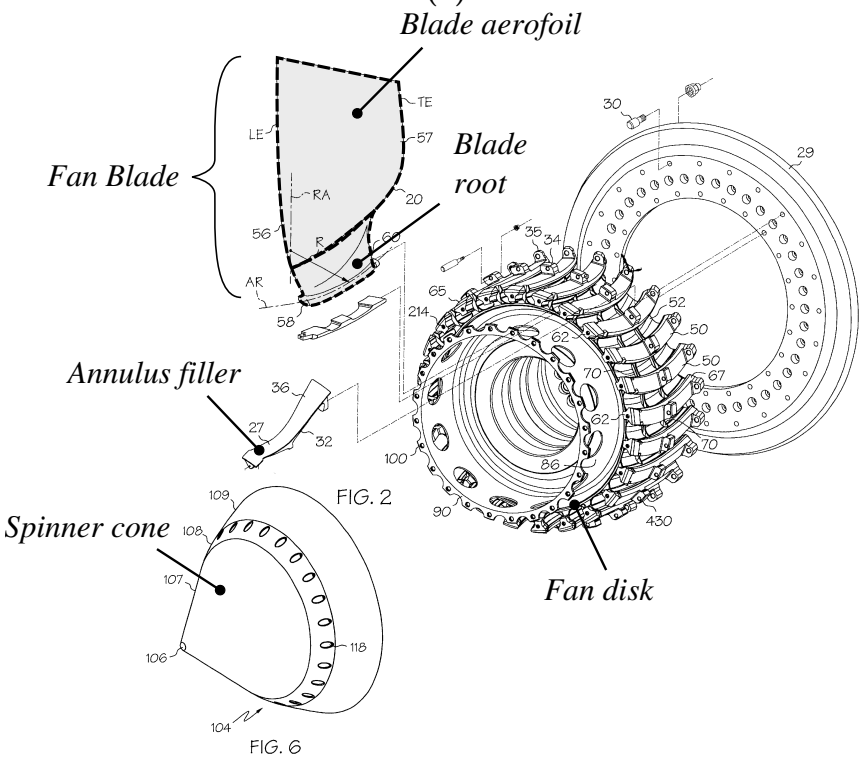
The aerodynamic fan hub contour is typically defined by several subcomponents. Those include the spinner cone located in front of the fan and so-called annulus fillers that are mounted between the fan blades. It is important to note in this context that the annulus fillers and therefore the hub contour do not provide any structural support for the blades; the structural blade restraint is solely realised by the fan disk. Towards the outer diameters, the fan flow is limited by the fan casing which is constructed to yield a minimum gap towards the running blade tips for minimal tip leakage losses. The fan casing is further designed to mitigate fan noise and has the task of containing the fan blades in case of structural blade failure.

As indicated in Fig. 4.1 b) the fan blade comprises an aerofoil region which is submerged in the air flow, and a blade root region, which extends downwards from the aerodynamic hub contour to connect to the disk but has no aerodynamic function.

The fan exit flow is separated into two streams by the so-called *bypass splitter*. The majority of the flow enters the bypass duct, where it is straightened in the bypass outlet guide vanes (BOGV) and exits the engine through the bypass nozzle into the



(a)



(b)

Figure 4.1: Overview of a typical fan assembly. Adapted from [51].

atmosphere to provide typically about 80% of the total engine thrust for current engines. The smaller part of the flow enters the engine core, where it undergoes further compression, is mixed with fuel, burned and expanded in the turbine in order to drive the compressors and the fan. The ratio of the mass flows of air entering the bypass and the core ducts defines the well known bypass ratio; a central performance parameter for civil aero-engines.

For an MDO framework, a definition of a fluid domain and a solid domain is required, which share an identical representation of the blade surface geometry at their interface. The parametrisation and geometry construction process utilizing a constructive CAD-based approach (c.f. Section 2.3.3) as implemented in this work will be discussed next. The implementation is utilizing a VKI in-house CAD kernel, which comprises many standard geometric CAD entities required for the generation of complex Boundary Representation CAD models. All steps are entirely written in C++.

4.2 Fluid domain definition

The fan blade aerofoil and fluid domain are defined by the meridional contour, the camber surface and blade thickness distributions.

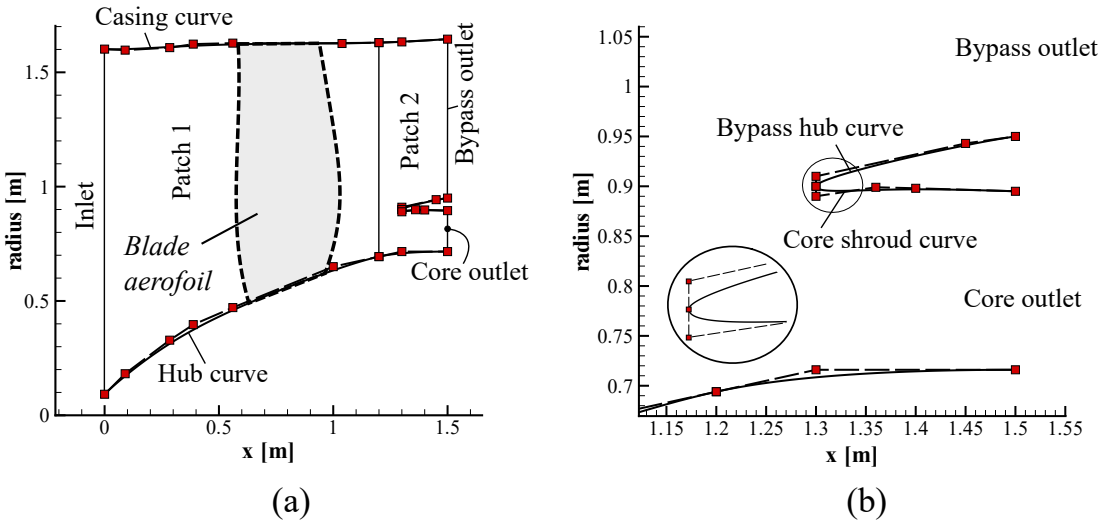


Figure 4.2: Meridional flow path definition. Red squares indicate the B-Spline curve control points. (a) global view. (b) close-up of bypass splitter.

4.2.1 Meridional contour definition

The meridional plane is defined by the axial coordinate x as abscissa and the radius as ordinate, as shown in Fig. 4.2. The end wall contours of hub and casing as well as the bypass splitter contours are defined by parametric B-Spline curves.

The degree of the curves, the control point coordinates as well as the knot vectors can be freely defined by the user to best represent a desired hub and casing contour and enable axis-symmetric end wall contouring, if desired.

The definition of the meridional contour is subdivided into several patches, enabling the definition of separate blade rows, and the introduction of rotor-stator interfaces (sliding-plane, mixing-plane), where required. Figure 4.2 (a) shows an example of a two-patch configuration with a first patch describing the geometry from the inlet to the fan exit right downstream of the trailing edge, whereas a second patch is used for the definition of the bypass splitter, which is defined by two curves that merge at the leading edge as shown in Fig. 4.2 (b). The leading edge control points are here located to enable a G1-continuous bypass splitter leading edge.

4.2.2 Camber surface definition

A schematic summarising the camber surface parametrisation is shown in Fig. 4.3. The parametrisation used here is partly based on previous work [78] and has been further developed in the framework of this DPhil.

The camber surface can be considered as the mean surface of the blade, running from the leading edge to the trailing edge, and having at every point an equal minimum distance to the suction side and the pressure side surfaces. Several *camber lines*, each defined on a specific spanwise position (span $\in$ [0,1], 0 at the hub and 1 at the tip), define the camber surface's shape.

For the construction of a single camber line, a blade metal angle distribution from the leading to the trailing edge $\beta(u)$ (with $u \in [0,1]$ being the non-dimensional meridional length, 0 at the leading edge and 1 at the trailing edge) and a chord length need to be specified (Fig. 4.3). In conventional turbomachinery blade parametrisation methods, the user directly defines the parameters needed for the generation of each camber line. In the present work a different approach is used, in which the blade angle distributions and the chord length are defined as continuous parametric curves over the blade span. In contrast to the conventional method, this parametrisation guarantees a smooth spanwise evolution of the blade angles and furthermore allows

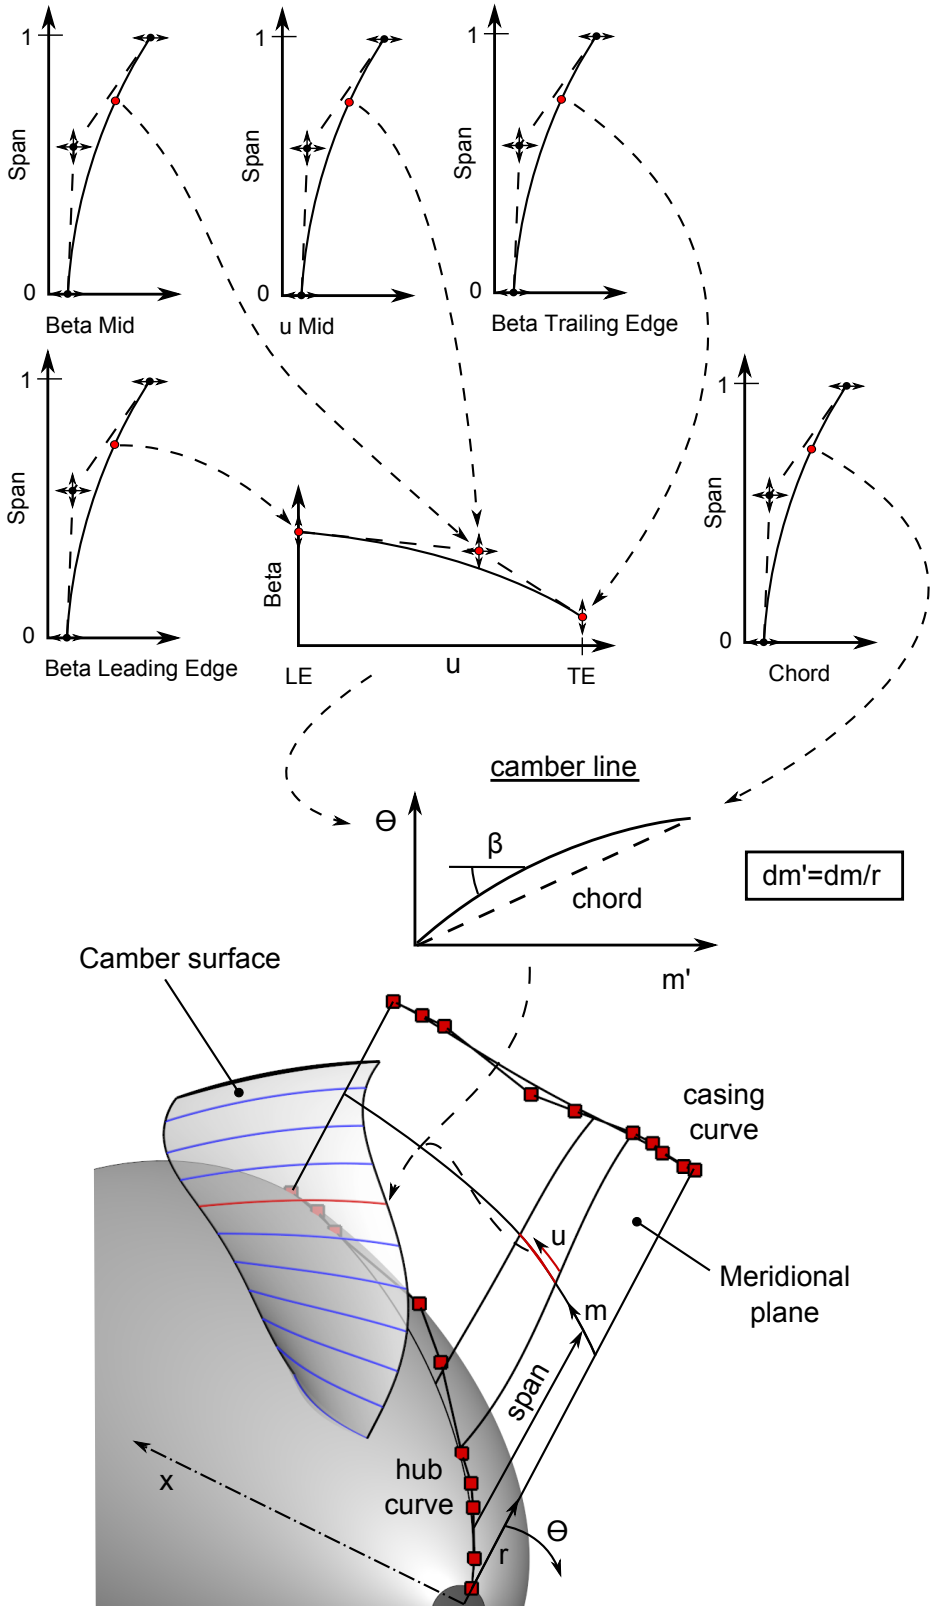


Figure 4.3: Parametrisation of the camber surface.

the parameters to be directly linked to the results of two-dimensional through-flow codes, which are routinely used in the preliminary design process.

In total four B-Spline curves are used for the definition of the spanwise blade angle distributions; three curves are needed for the specification of blade angles at the leading edge, trailing edge and a mid point while one additional curve defines the local position of the mid point in non-dimensional meridional length u. A fifth B-Spline curve is introduced for the specification of the spanwise chord length distribution C' in (m',θ) space. The number of control points of these B-Spline curves, their coordinates and the knot vectors are parameters which can be specified by the user in order to define and modify the geometry. By extracting the curve values from the angle distributions at an arbitrary spanwise position, the control point coordinates of a two-dimensional second order Bézier curve are given defining the blade metal angle distribution from the leading to the trailing edge, see Fig. 4.3. The fact that the meridional stream line with arc length m can be arbitrarily curved introduces the requirement of generating the camber line in different mathematical spaces in order to accurately conserve angles and distances during the camber line construction process. By integration of the blade angle distribution a two-dimensional camber line is constructed in (m',θ) space, with dm' being the differential arc length dm at span s normalized with respect to the radius r ($dm' = dm/r$), see Fig. 4.3. The space defined in (m',θ) coordinates is angle preserving with respect to physical space, therefore allowing the definition of critical blade angle properties, and the conformal generation of the camber lines.

With knowledge of the meridional curve at a specific spanwise position, the camber line can be transformed to a three-dimensional curve in Cartesian coordinates. In order to avoid bending moments in the blade resulting from centrifugal forces, the camber lines are positioned relative to each other in such a way that the center of gravity of each two-dimensional spanwise blade section is lying on a radial stacking axis, whereas the axial coordinate of this axis is an additional parameter to be defined by the user, see Fig. 4.4. This radial stacking of two-dimensional cross sections involves operations like the interpolation of spanwise thickness distributions and multiple coordinate transformations. A detailed explanation of this procedure will therefore not be given here. The camber surface is finally created by a surface loft through all previously generated three-dimensional camber lines.

It is well known that the tip section camber line profile has an important influence on the aerodynamic performance of supersonic compressor sections, and the optimal tip profile shape depends on the relative in-flow Mach number e.g. [175]. It is shown

here, that the implemented parametrisation strategy can incorporate all relevant blade profiles that are currently found in practice. This includes flat, C-shaped or J-shaped profiles and S-shaped profiles (also known as precompression profiles) as shown in Fig. 4.5.

4.2.3 Lean and sweep

In transonic fans it can prove beneficial for the aerodynamic performance to deviate from the purely radial stacking of blade profiles as presented above. This is usually done by the introduction of three-dimensional design features known as *lean* and *sweep*. However, although widely applied, lean and sweep are not uniquely defined in the literature. One frequently found convention is to define sweep as a translation of a blade section in the axial direction, while defining lean as a translation in the circumferential direction. This definition is therefore independent of the blade geometry. An alternative strategy is to define lean and sweep in reference to the local blade chord, such that sweep represents a translation of a blade section along the blade chord, while lean describes a translation perpendicular to it [39]. This description is advantageous as it enables a clearer decoupling of the physical implications

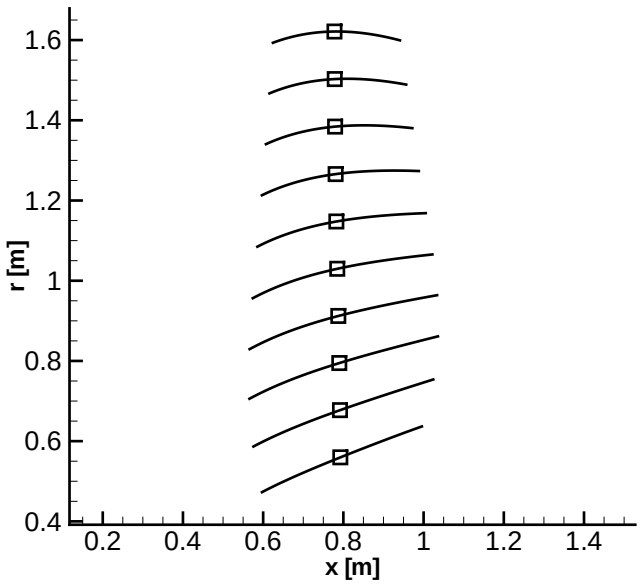


Figure 4.4: Meridional view showing the initial radial stacking of camber lines through the center of gravity of the associated two-dimensional blade sections. Gravity centers are shown as rectangles.

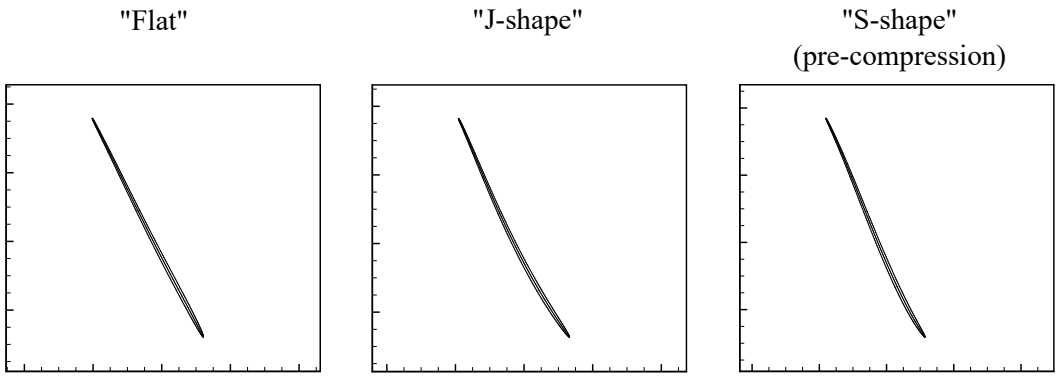


Figure 4.5: Selected tip section profiles generated with the implemented parametrisation.

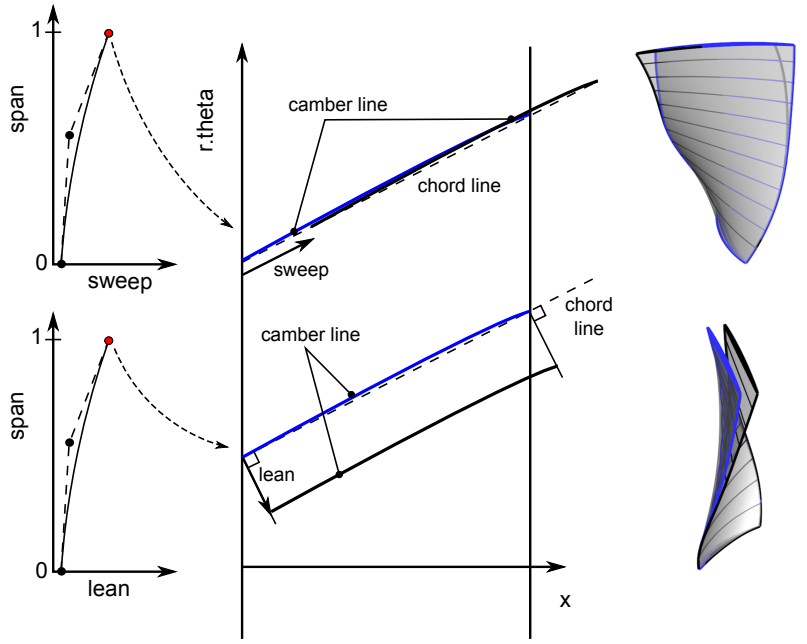


Figure 4.6: Parametrisation and definition of lean and sweep.

of both design features in terms of aerodynamic and structural performance. Hence, this definition is adopted in the present work and is shown in Fig. 4.6.

The effects of blade lean and sweep on the flow through blade rows have been described in detail in [40]. The main incentive behind the introduction of sweep is to influence the relative position of the leading edge shock with respect to the incident flow so that the component of Mach number perpendicular to the shock, and hence the shock loss, is reduced [39]. Through the application of lean, radial oriented forces can be imposed on the flow inside the blade passage which affects the meridional

streamline curvature and consequently the spanwise blade loading distribution.

As can be seen on the left hand side of Fig. 4.6, lean and sweep are both parametrised by two-dimensional B-Spline curves over the blade span.

4.2.4 Thickness

The last step required to generate the suction and pressure side surfaces of the blade is the application of a thickness to the camber surface. The blade thickness is parametrized by B-Spline curves at several spanwise positions, usually at hub and tip, see Fig. 4.7, but an arbitrary number of intermediate distributions can be defined. By adding the thickness perpendicular to the camber surface, the final blade shape is obtained, as shown schematically on the right hand side of Fig. 4.7.

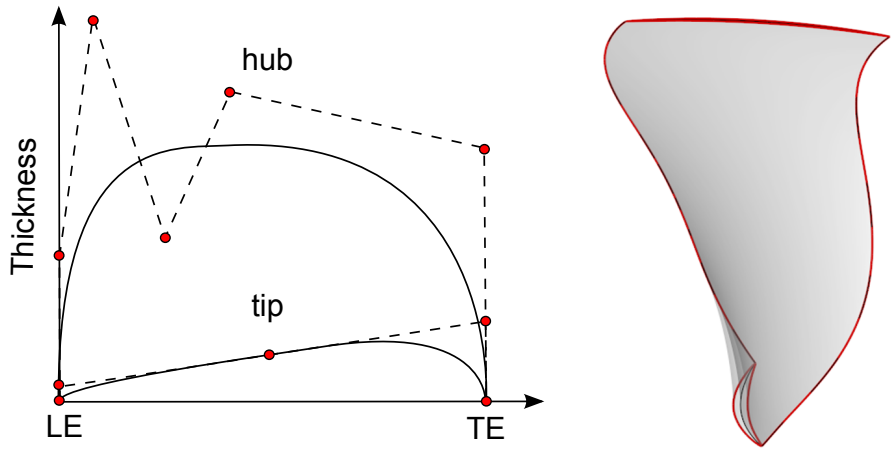


Figure 4.7: Construction of the final aerofoil shape by adding a thickness distribution to the camber surface.

4.3 Solid blade generation

The generation of the solid blade model is initiated from the constructed aerofoil geometry, resulting from the fluid domain generation described in the previous section. The process consists of two steps: a) the generation of a dovetail, which serves as structural connection of the fan blade to a disk, and b) the generation of a shank, smoothly connecting the dovetail to the aerofoil. A smooth connection is required to avoid the occurrence of stress concentrations at the interface between dovetail, shank and aerofoil. To the author's best knowledge, a methodology for a fully automated CAD-based generation of a complex dovetail root geometry has not been reported elsewhere in the literature to date. The following presents one possible approach in

greater detail, which has been developed in the framework of this DPhil. The solid blade generation process can be described in four distinctive steps.

Step 1 In the first step, a two-dimensional dovetail cross sectional shape is defined, as shown in Fig. 4.8. Two first order Bézier curves (straight lines 1, 2), and two second order Bézier curves (3, 4) are used to define one side of the cross section. Both sides of the cross section can be defined independently, therefore enabling non-symmetric shapes to be generated. However, in the present work only symmetric cross sections are considered.

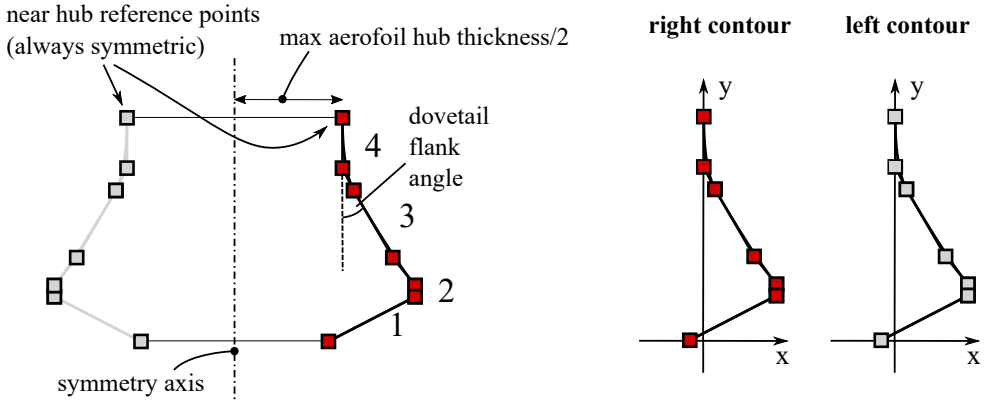


Figure 4.8: Step 1 of the dovetail root generation process: Definition of the dovetail cross section.

The curves are defined in two-dimensional physical space. The y-axis defines the dovetail height, while the x-axis defines its width. However, the definition of the dovetail width as an independent parameter is not desirable in an optimisation process in which the thickness of the blade may change. To allow the dovetail thickness to adapt to the blade thickness, the width coordinates of the control points are not defined in an absolute sense, but are given in dependency of the control points at the top of the dovetail, referred to as near hub reference points in Fig. 4.8.

The control point coordinates can be added as optimisation parameters in case an optimisation of the cross section would be desired. Via the definition of control point dependencies, other relevant geometric quantities such as angles and distances can be specified, enabling for instance additional control of the dovetail flank angles.

Step 2 In the second step, a rotational axis is defined which serves to generate the rotated surfaces of the dovetail contours, as depicted on the left hand side of Fig. 4.9. The rotation axis is positioned at the centre of a circle, which is defined by in-plane projections of the leading edge, trailing edge and mid points of the aerofoil hub section camber line. The horizontal plane on which the points are projected to is defined by a user defined vertical offset of the leading edge point p_{LE}. This offset similarly defines the shank height at the leading edge, as indicated in Fig. 4.9. Additional control of the front root geometry is provided by enabling an additional in-plane offset of the projected leading edge point along the curve radius r, which will be denoted as Δp_{LE}. A non-zero value of Δp_{LE} introduces a tilted shank surface in the front root, in contrast to a straight geometry for $\Delta p_{LE} = 0$.

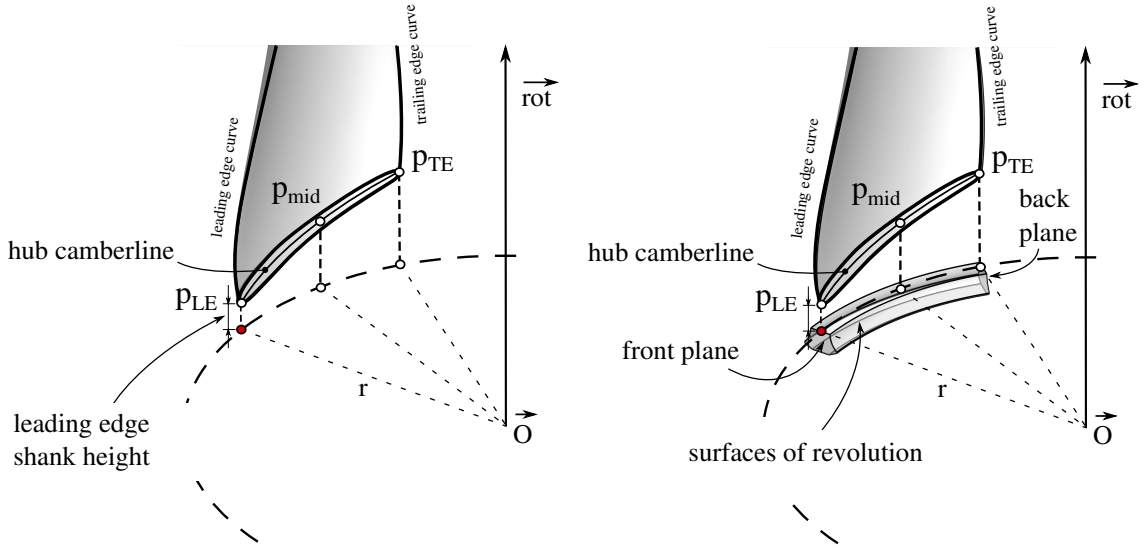


Figure 4.9: Step 2 (left) and step 3 (right) of the dovetail root generation process.

Step 3 In the third step, the curves defining the previously generated root cross section (now automatically adapted to the thickness of the blade aerofoil section at the hub), are rotated around the rotation axis for the generation of the surfaces of revolution defining the dovetail side and bottom surfaces, as shown on the right hand side of Fig. 4.9. The dovetail geometry is subsequently completed by the generation of front and back planes positioned at the projected leading edge and trailing edge points, with surface normals tangent to the circle. Other orientations are easily implemented to enable different plane orientations to match specific disk installation requirements, if necessary.

Step 4 The right hand side of Fig. 4.9 shows that the dovetail geometry and the aerofoil are still disconnected. The methodology used to create a smooth connection between the two disjoint geometries will be explained in the following for the example of the left shank surface (extending down from the suction side surface), see Fig. 4.10: First, march along the top dovetail curve and the lower suction surface curve with an identical non-dimensional stepsize from the leading edge to the trailing edge. Second, for each point on the respective curve, compute the local directional derivative in the direction corresponding to the parametric u-direction of the surface. The definition of the parametric u and v directions of the respective surfaces is shown in Fig. 4.10. Third, construct a control point in the direction of the tangent with a given distance from the origin point on the respective curve. In the current implementation the distance is computed as 20% of the distance between the two respective points that are to be connected. Fourth, create a third order Bézier curve based on the four control points; the tangent control points ensure G1-continuity of the curve with the adjacent surface curves. Fifth, generate the shank surface by creation of a surface loft through all previously generated curves, as depicted in the middle of Fig. 4.10. Note, that the first point on the suction surface curve shown in Fig. 4.10 does not coincide with the leading edge point, but a small distance is kept for the generation of a front shank surface.

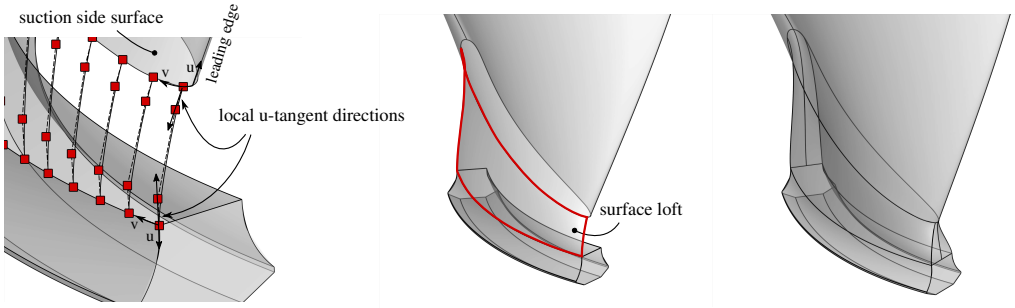


Figure 4.10: Step 4 of the dovetail root generation process: generation of the shank surfaces.

The shank surfaces to connect front, back and right blade surfaces are constructed using a similar approach. The full set of surfaces is shown on the right hand side of Fig. 4.10. All blade surfaces now represent a closed shell to finalise the definition of a Boundary Representation (B-Rep) solid model. The model can be exported in any common CAD format (e.g. STEP) for viewing and/or further manipulation in an external CAD program or for manufacturing purposes. A rendered example is shown in Fig. 4.11.

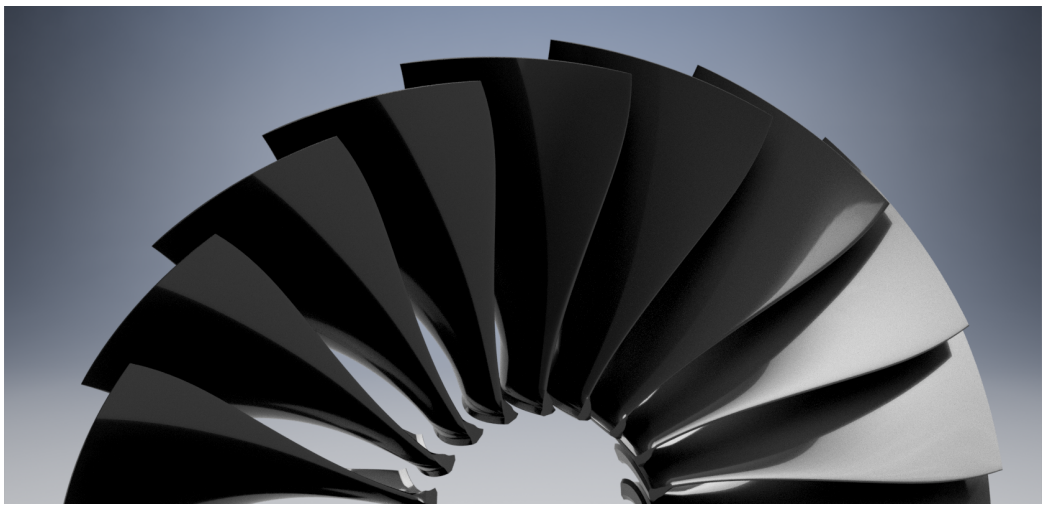


Figure 4.11: Visualisation of exported fan blade CAD-models in a commercial CAD program.

5 Aerodynamic Performance Analysis

5.1 An introduction to the aerodynamic fan design problem

Simply stated, the primary purpose of the fan is the generation of a specific amount of thrust with the lowest possible amount of aerodynamic losses in order to attain high efficiencies and correspondingly low engine fuel consumption. A consequent definition of an aerodynamic loss may be given as any flow mechanism that extracts energy from the rotating blade, without contributing to the forward thrust. The aerodynamic design problem is complicated by the need for the fan to operate under a wide range of operating conditions, requiring good design and off-design performance and large stability margins. Further, fans operate in transonic flow conditions resulting in complex three dimensional shock structures. If the blade is inadequately designed, strong shocks and shock-induced boundary layer separation can lead to considerable performance penalties in the form of increasing losses and a reduced operating range. The fan tip speed is typically limited by either structural, or more commonly, aeroacoustic constraints, which results in low circumferential blade velocities in the aerofoil root sections due to typically low hub-to-tip radius ratios (long blades). In order to attain a sufficient fan root flow pressure rise, fan root sections need to be highly loaded, which increases the risk for flow separation.

Further, fan blades do not operate in isolation, but their design and off-design performance is dictated by the requirements of the engine cycle design. This implies that the bypass outlet guide vane (BOGV) in the bypass duct, and the engine section stator (ESS) in the engine core duct downstream of the fan (c.f. Fig. 4.1) need to receive a healthy inlet flow in order to operate efficiently. Hence, rotor alone performance is not sufficient if it leads to stage performance penalties. The presence of the

bypass splitter downstream of the fan adds further three-dimensionality to the flow and additionally influences the operating point dependent spanwise blade loading distribution, with a corresponding impact on the bypass ratio. Secondary flow effects in the end-wall regions, such as induced by the tip-leakage flow, add further to the complexity of the design problem.

In industrial practice, a combination of two- and three-dimensional methods is applied in the aerodynamic fan design process. Two-dimensional methods solving the flow in the blade-to-blade (S1) and hub-to-tip (S2) streamsurfaces and combinations of those for quasi-three dimensional flow solutions have been used in the turbomachinery design community since the 1950's [177]. So-called "through-flow" methods which solve the axissymmetric flow in the hub-to-tip (S2) streamsurface are used extensively in the preliminary fan design process, after the required design and off-design stage performance, annulus shapes, rotational speeds and mean blade angles have been defined in the one-dimensional engine cycle design. Through-flow calculations provide the spanwise flow angle variations at inlet and outlet to the blade rows and enable the generation of spanwise blade angle distributions for the generation of an initial blade geometry. Blade-to-blade methods in design (inverse) or analysis mode can be used for the definition of the detailed two-dimensional blade shape. Iterations between the two approaches are often performed to enable a quasi-three dimensional design approach.

However, in transonic aero-engine fans the three-dimensionality of the flow field is of such extent that the usefulness of two-dimensional and quasi-three dimensional methods for design is limited, and three-dimensional CFD is required for an accurate resolution of the flow physics. CFD methods for the solution of the steady three-dimensional Reynolds-Averaged Navier-Stokes (RANS) equations have been applied in the engine fan design process since the early 1990's (e.g. [76], [131]) and have since matured to become a routine tool in today's design process. The application of steady RANS methods has enabled significant improvements in the fan design, and played a substantial role in the development of the highly swept, three-dimensional fan blades that are common today.

The long standing history of CFD development and application strengthens the confidence in the methods' applicability in an automated design optimisation process. However, some relevant issues remain. CFD is still not able to accurately predict losses, particularly in off-design conditions and under adverse pressure gradients. Limitations of the turbulence models applied play a major role in this context,

and the turbulence model chosen will greatly affect the tendency of the boundary layers to separate. Effects of increasing flow unsteadiness may further become dominant for the overall component performance. It is important to take those limitations into consideration in the application and interpretation of CFD computations. However, a perhaps more important requirement for the application of CFD in an automated design optimisation process is the ability to correctly rank the performances of different designs. In this context, the prediction of absolute performance is not as crucial as the ability of the code to accurately capture the trends of performance changes between differing geometries. It is today widely accepted that well validated modern CFD codes are generally capable to do so [87], although a healthy amount of scepticism as well as proper error quantifications should always be part of the evaluation process.

In this work a commercial CFD code is applied for the solution of the three-dimensional fan flow. The solver applied is the 3D Reynolds-Averaged Navier-Stokes code FINE$^{\text{TM}}$/Turbo v11.1 from Numeca International. In the sections following below, an introduction to the theory and the solution procedure of the CFD code will be given. The description will be kept brief to provide a general, basic understanding of the applied method to a level deemed reasonable for the understanding of the aspects involved in the numerical setup and the analysis of the fan flow. A detailed coverage of the CFD implementation is, however, out of the scope of the present thesis.

5.2 Navier-Stokes equations

The unsteady, three-dimensional, compressible, viscous flow through a fan stage is described by the Navier-Stokes equations. The equations are derived by the application of the conservation laws of mass, momentum and energy to the fluid flow. While the Navier-Stokes equations in their original form contain solely the momentum equations, it has become common practice to unite the full set of conservation equations under the terminology.

The conservation equations in Cartesian coordinates in a rotating frame of reference read

$$\frac{\partial \boldsymbol{U}}{\partial t} + \frac{\partial \boldsymbol{F}_i}{\partial x_i} = \boldsymbol{S} \tag{5.1}$$

where $\boldsymbol{U}$ represents the so-called state vector

$$\boldsymbol{U} = \begin{bmatrix} \rho \\ \rho w_1 \\ \rho w_2 \\ \rho w_3 \\ \rho E \end{bmatrix} \tag{5.2}$$

and w_i are the Cartesian components of the relative velocity, ρ is the fluid density and E is the total energy given by

$$E = e + \frac{1}{2} w_i w_i \tag{5.3}$$

with e being the specific internal energy. The column vectors $\boldsymbol{F}_i$, $i = 1, 2, 3$ contain the inviscid and viscous flux vectors

$$\boldsymbol{F}_i = \begin{bmatrix} \rho w_i \\ \rho w_1 w_i + p\delta_{1i} - \tau_{i1} \\ \rho w_2 w_i + p\delta_{2i} - \tau_{i2} \\ \rho w_3 w_i + p\delta_{3i} - \tau_{i3} \\ (\rho E + p)\, w_i - q_i - w_i \tau_{ij} \end{bmatrix} \tag{5.4}$$

where the heat flux components are defined by

$$q_i = k \frac{\partial T}{\partial x_i} \tag{5.5}$$

with k being the thermal conductivity. τ_{ij} defines the shear stress tensor and δ_{ij} is the Kronecker delta. p denotes the static pressure. $\boldsymbol{S}$ is the vector of source terms, including the centrifugal and Coriolis forces.

The problem set now consists of five equations in six unknowns (ρ, w_1, w_2, w_3, E and p). In order to close the system of equations, one additional relation is necessary, which is provided in the form of the equation of state for the fluid. For air modelled as a perfect gas, the relation is

$$p = \rho R T \tag{5.6}$$

where R is the gas constant.

5.3 Reynolds-Averaged Navier-Stokes equations

The Navier-Stokes equations describe both laminar and turbulent flows. The flow in aero-engine fans is characterised by high Reynolds numbers leading to almost fully

turbulent flow fields over the entire range of operation. At those Reynolds numbers, length and time scales of the turbulent eddies in the flow vary widely. To capture the smallest scales, required mesh resolutions and time-steps would be intractable with currently available computational resources. However, a full resolution of the turbulent spectrum is not needed to resolve the relevant flow physics in the fan flow.

A solution to this problem is obtained by Reynolds-averaging the Navier-Stokes equations, which involves the assumption that the unsteady flow quantities that change randomly in time can be decomposed into a time mean and a fluctuating component. The resulting *Reynolds-Averaged Navier-Stokes (RANS)* equations have the same form as the Navier-Stokes equations, Eq. (5.1), but the vectors now read

$$\boldsymbol{U} = \begin{bmatrix} \bar{\rho} \\ \bar{\rho}\tilde{w}_1 \\ \bar{\rho}\tilde{w}_2 \\ \bar{\rho}\tilde{w}_3 \\ \bar{\rho}\tilde{E} \end{bmatrix} \tag{5.7}$$

$$\boldsymbol{F}_i = \begin{bmatrix} \bar{\rho}\tilde{w}_i \\ \bar{\rho}\tilde{w}_1\tilde{w}_i + \overline{\rho w_1' w_i'} + \bar{p}\delta_{1i} - \bar{\tau}_{i1} \\ \bar{\rho}\tilde{w}_2\tilde{w}_i + \overline{\rho w_2' w_i'} + \bar{p}\delta_{2i} - \bar{\tau}_{i2} \\ \bar{\rho}\tilde{w}_3\tilde{w}_i + \overline{\rho w_3' w_i'} + \bar{p}\delta_{3i} - \bar{\tau}_{i3} \\ \left(\bar{\rho}\tilde{E} + \bar{p}\right)\tilde{w}_i + \overline{\rho E' w_i'} + \overline{p w_i'} - \bar{q}_i - \overline{w_i \tau_{ij}} \end{bmatrix} \tag{5.8}$$

where the overbar denotes Reynolds-averaged quantities while the tilde denotes mass averaged quantities. The mass averaged total energy is given by

$$\tilde{E} = \tilde{e} + 0.5\tilde{w}_i\tilde{w}_i + \tilde{k} \tag{5.9}$$

with $\tilde{e}$ being the mass averaged specific internal energy and $\tilde{k}$ representing the turbulent kinetic energy defined by

$$\tilde{k} = \frac{1}{2}\left(\frac{\overline{\rho w_i' w_i'}}{\bar{\rho}}\right) \tag{5.10}$$

Note that the Reynolds-averaging process has introduced additional terms in the equations which are of the form

$$\rho\overline{w_i' w_j'} \tag{5.11}$$

Those terms are called Reynolds stresses. Hence, there are now again more unknowns existing than equations. Additional relations are needed to close the system of equations, which are provided in the form of turbulence models.

5.4 Turbulence modelling

Turbulence models seek to model the Reynolds stress terms in the RANS equations. A common assumption is that the effects of turbulence can be represented by an increased viscosity over the purely laminar case, leading to so-called eddy-viscosity models for the Reynolds stresses

$$-\bar{\rho}\widetilde{w_i''w_j''} = \mu_t \left(\frac{\partial \tilde{w}_i}{\partial x_j} + \frac{\partial \tilde{w}_j}{\partial x_i} \right) - \frac{2}{3}\bar{\rho}\delta_{ij}\tilde{k} \tag{5.12}$$

Here, μ_t is the eddy-viscosity and $\tilde{k}$ is the turbulent kinetic energy introduced above. Although the eddy-viscosity assumption is not correct in detail, it does provide reasonably good results for many flows if applied with care and is further relatively easy to implement.

Eddy-viscosity turbulence models have been proposed with various levels of complexity. The simplest ones are so-called algebraic models (also known as mixing length models). In those models the turbulent kinetic energy and eddy-viscosity are determined based on representative length scales of the time mean flow. Those are typically applicable for relatively simple flows, but not for separated or highly three-dimensional flows [48]. So-called one-equation models use a partial differential equation, most commonly to model the turbulent kinetic energy. Consequently, so-called two-equation models utilise two partial differential equations; in the most common forms one to model the turbulent kinetic energy and one to model either the turbulent dissipation per unit mass ϵ (k-ϵ models) or the turbulent dissipation per unit turbulence kinetic energy ω (k-ω models).

For many aerospace and turbomachinery flows under adverse pressure gradients, a one-equation model that has found to deliver good results is the model of Spalart and Allmaras [154]. Due to being a one-equation model, it further adds robustness to the flow solution over the usage of two-equation models, which is a desired characteristic for the optimisation process. The Spalart-Allmaras model is frequently used for transonic compressor and fan computations, e.g. [62], [144], [94], [93]. This model has been chosen for the computations carried out in this thesis.

5.5 Automatic mesh generation

A necessary first step for the numerical prediction of the flow field is the discretisation of the geometrical fluid domain into a finite number of subdomains, also called elements, cells or control volumes. The so-called mesh (commonly also referred to

as grid), enables the continuous flow field, which is physically described by a set of Partial Differential Equations (PDE), to be approximated by a system of algebraic equations which can be solved by a computer.

There are generally two different types of meshes; structured and unstructured ones, which differ in the way the data structures are organised in the computer implementation. Structured meshes utilise an ordered data structure, in which the position of each grid point is uniquely identified by a set of indices (two in two dimensions and three in three dimensions), see Fig. 5.1 (a). This has positive implications for the solution procedure. For structured meshes, the matrix of the algebraic system of equations has a regular structure, which enables the development of highly efficient solution methods [48]. The disadvantage of structured meshes is that they are difficult to generate for complex solution domains, such as given in three-dimensional turbomachinery cases. A subdivision of the domain into separate structured mesh blocks resulting in so-called multi-block structured meshes increases meshing flexibility, but there are as of today no fully automatic (multi-block) structured mesh generation methods available for reasonably complex geometries.

On the other hand, unstructured meshes are not constrained by a regular data structure and are therefore much easier to generate in an automated fashion even for very complex three-dimensional geometries. In such meshes, the elements may have any shape, and there is no restriction on the number of neighbouring elements or nodes. They are defined by a node list, which contains the nodal coordinates in two or three dimensions and an element connectivity list, which contains the node indices each element is connected to. A 2D example made of triangular elements is shown in Fig. 5.1 (b). Unstructured meshes are widely used in structural analysis using the finite element method (as will be discussed in the following chapter), however, they are also applicable in fluid analyses. While they can in principle be used with any discretisation scheme, they are best suited for finite volume and finite element methods. In fluid analyses, the price to pay for the gain in meshing flexibility is a loss of solution efficiency due to an irregular data structure. Further, the application of tetrahedral or triangular elements in the boundary layers in viscous flow analyses also places a penalty on solution accuracy, which can be alleviated using quadrilateral or hexahedral elements in the near wall regions. This approach is sometimes also referred to as hybrid meshing. However, the overall data structure of such meshes remains unstructured, with the corresponding penalties in solution performance.

In view of solution performance and accuracy it is overall desirable to use a structured meshing approach if a structured solver is available (which is the case in present

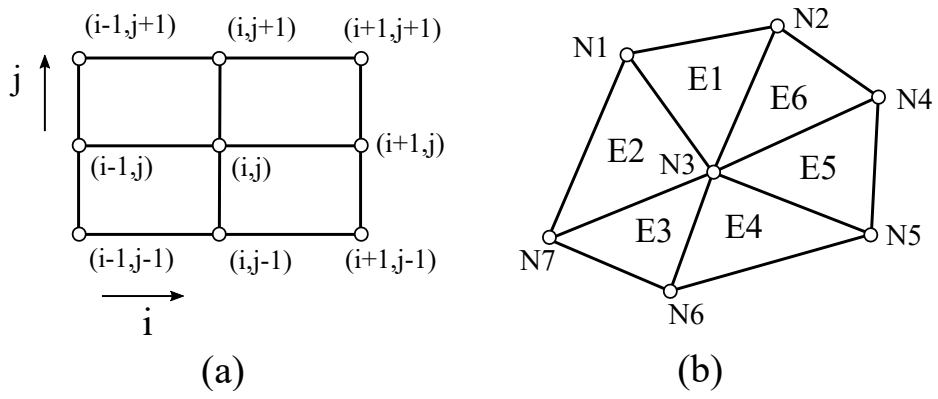


Figure 5.1: A structured (a) and unstructured (b) mesh in 2D.

work). However, as noted above, fully automated methods do not exist. Nevertheless, in a shape optimisation framework, generated geometries are typically quite similar and therefore a fully automatic mesh generation from scratch is not necessary. Instead, a semi-automatic meshing approach can be used, in which the topology and mesh point distributions from a manually generated multi-block structured template mesh are mapped to a different geometry with similar topology. The meshing tool AUTOGRID from Numeca enables such a templated mesh generation for turbomachinery geometries. The approach was found to be sufficiently flexible and robust for the purposes of this work and has therefore been chosen as fluid meshing tool for the optimisation process.

It can be intuitively appreciated that, as the size of the individual elements decreases, the approximation of the continuous (real) problem should improve, and in the limit of infinitely small elements, the solution should coincide with the "exact" solution of the equations. The difference between the numerical solution of the flow field with a given mesh and the exact solution of the equations, is known as the *discretisation error*. The estimation of the discretisation error is an important part of a CFD computation and is typically performed by a so-called mesh independence study, which involves consecutive and systematic mesh refinements and comparison of the results.

5.6 Spatial discretisation

The problem now being defined in terms of a set of Partial Differential Equations, needs to be solved using a numerical solution technique, for which the PDEs needs to be "translated" to a set of non-linear algebraic equations. For this purpose, most

CFD methods, including FINE/Turbo, use a finite volume discretisation, in which the mesh (described in Section 5.5) defines the control volumes in the fluid domain in which the flow is to be solved.

The translation of the continuous problem to the discrete one is performed by approximation of the surface and volume integrals in the RANS equations for each respective control volume. A large number of different methodologies exist, ranging from first order, to second order and higher order schemes. This is a large topic, and the reader is referred to the literature for a more in-depth coverage, e.g. [48], [66].

Summation of all the flux and source terms for the elements results in the global system of non-linear algebraic equations to be solved, which contains as many equations and unknowns as there are control volumes in the mesh.

5.7 Boundary conditions

To enable a unique solution of the approximated problem inside the computational domain, information of the solution at the domain boundaries needs to be given. In finite volume approximations, this involves the definition of fluxes for the control volume faces that coincide with the domain boundaries, including inlet, outlet, periodic and solid boundaries. For inlet and outlet boundaries, different treatments are necessary depending on whether the flow is subsonic or supersonic. In unsteady flow computations, non-reflective inlet and outlet boundaries are important in order to avoid that disturbances propagating out of the domain are numerically reflected back into it, which would lead to an incorrect solution.

5.8 Solution procedure

The resulting system of non-linear algebraic equations needs to be solved using a suitable solution technique, of which many different variants are existing. Most frequently, explicit or implicit iterative techniques are applied, while direct methods are seldom used. The solution methods also differ depending on whether the steady or unsteady flow is to be computed.

5.8.1 Steady flow solution

For civil fan blade design purposes, the primary interest is usually in the steady flow field, while unsteady flow effects are initially considered to be negligible. FINE/Turbo solves the steady flow problem using an explicit pseudo time-marching scheme in

which the system of non-linear equations is successively linearised and solved using an iterative method.

5.8.2 Unsteady flow solution

Turbomachinery flows are inherently unsteady. While the computation of the steady flow field often gives an acceptable representation of the flow physics for design purposes, the solution of the unsteady flow field is a relevant problem in fan design for several reasons: first, at off-design conditions unsteady flow effects such as flow separations may quickly gain central importance for the overall aerodynamic fan stability and time-averaged performance. Increasing blade loading for purposes of blade count reduction also increases the possibility of significant unsteady effects to appear at design conditions. Further, unsteady component interactions, e.g. wake-blade interactions, can influence the aerodynamic performance and stability as well as the aeroelastic stability in terms of forced response. In addition, self-excited unsteady flow effects govern the dynamic aeroelastic stability of fans in terms of flutter. Various methods can be used for the prediction of the unsteady flow field in turbomachinery. A brief outline of two central strategies will be given below.

5.8.2.1 Time marching of RANS equations

The conceptually perhaps most straightforward way of resolving the unsteady flow is by advancing the solution of the RANS equations forward in time from an initial solution (i.e. the steady flow solution) using so-called time-marching methods. These naturally enable to resolve all non-linearities in the flow and therefore provide a high-level of solution accuracy if the time dependent terms of the equations are properly discretised (e.g. second order accurate). However, time-marching of the RANS equations is computationally too expensive to be used for design purposes.

5.8.2.2 Non-linear harmonic method (NLH)

Based on the assumption that the unsteady flow behaves linearly in time, some early attempts have been made for its solution in the frequency domain, which is computationally more efficient than the corresponding solution in the time domain, introduced above. However, flow non-linearities are not always negligible, and thus their consideration is important in many instances to accurately represent the flow physics.

Non-linearities manifest themselves in the form of higher flow harmonics which alter the time-averaged unsteady flow from its steady counterpart. The non-linear

harmonic method (NLH) ([61], [115]), has been developed as an extension to linear frequency domain methods, allowing for the interaction of higher harmonics to be taken into account and therefore providing a computational framework that combines the accuracy of non-linear time-marching methods with the computational efficiency of time-linearised frequency domain methods. A validation of the NLH method implemented in FINE/Turbo for the prediction of turbine flutter is presented in Appendix A.

5.9 Some current challenges in aerodynamic fan design/analysis

Steady RANS CFD methods are today considered as being standard in the design process of fan blades. However, many relevant design and modelling challenges remain for the aerodynamic design process and some of the most prevalent ones will be briefly highlighted below.

5.9.1 Stall limit prediction

Fans are required to operate stably under largely varying flow conditions. A large stability margin with respect to the onset of stall (stall margin) is thus a critical design requirement. The prediction of the stall limit using steady RANS in the design process is, however, challenging, as stall is an intrinsically unsteady effect. The assumption is often made - and will be similarly applied in this thesis - that the occurrence of physical stall coincides with the breakdown of the numerical flow (i.e. "numerical stall"), hence the last converged computation towards reduced mass flows on a speedline is typically considered to be the point where stall occurs. With a well calibrated CFD code, this typically enables a consistent ranking of designs in terms of their stall margin, although the absolute stall margin values might be, and most likely are, different from the real, physical ones.

5.9.2 Noise prediction

Noise has become of major concern for civil engine and aircraft manufacturers over the past decades with increasingly stringent certification regulations having been put in place in the recent past. It was noted by Cumpsty [31] that the specification of the engines for the Airbus A380 required a higher fuel burn to meet noise regulations than would otherwise be possible if noise could have been ignored.

The fan is a dominant contributor to the total engine noise levels. Fan noise can be broken down into the following components (e.g. [87]):

- Tonal noise generated by the interaction of the fan blade wakes with downstream vanes and other stationary components.
- Rotor-alone and so-called buzz-saw noise due to the interaction of the fan rotor shock systems.
- Broadband noise which is a complicated turbulence related phenomenon, the resolution of which requires higher fidelity modelling of turbulence such as e.g. Large Eddy Simulations.

While noise prediction in the fan design process is highly desirable, the accurate numerical prediction of noise is difficult and expensive. The reasons for this are numerous, but a central issue is related to the typical scales that define the aeroacoustic behaviour which are much smaller than corresponding aerodynamic ones, although both are governed by the Navier-Stokes equations: Acoustic wavelengths are much shorter than aerodynamic wavelengths, therefore requiring a considerably finer mesh for their resolution and propagation without excessive numerical dissipation.

6 Structural Performance Analysis

6.1 An introduction to the structural fan design problem

The problem of structural analysis in the framework of the aerostructural fan design loop typically involves two different types of problems. The first is the computation of the displacement, stress and strain response of the blade under static or quasi-static loading. The second is the computation of the blade's free vibration response in order to avoid the occurrence of harmful resonant vibrations in the blade's operating range, arising from matching of the blade's natural frequencies with harmonics of periodic excitations from e.g. in-flow non-uniformities, rotor-vane interaction resonance, structural unbalance, etc. The prediction of those effects is today routinely performed using the finite element method (FEM). In this chapter, the theoretical background of the finite element method will be briefly introduced, and its application for the prediction of the structural static and vibration response of a composite fan blade will be explained. The structural computations in this work are performed using the open-source finite element solver CALCULIX [41].

6.2 Linear static analysis

Consider the general boundary value problem of solid mechanics depicted in Fig. 6.1, where the fan blade with volume Ω and boundary Γ is subjected to body forces b (such as centrifugal forces) and surface tractions $\bar{t}$ (such as pressure forces) and is clamped at the connection to the disk. In order for the body to be in a static equilibrium, the internal three-dimensional stresses at any point within the blade must be in balance

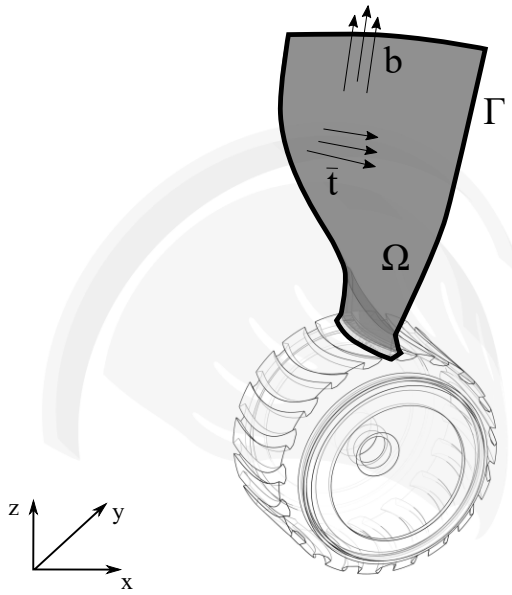


Figure 6.1: Depiction of the general boundary value problem for an elastic fan blade with volume Ω and boundary Γ, subjected to surface tractions $\bar{t}$ and body forces b.

with the internally acting forces in that same point. Thus, for any given point the following equilibrium relations must hold

$$\frac{\partial \sigma_{xx}}{\partial x} + \frac{\partial \sigma_{xy}}{\partial y} + \frac{\partial \sigma_{zx}}{\partial z} + b_x = 0 \tag{6.1}$$

$$\frac{\partial \sigma_{xy}}{\partial x} + \frac{\partial \sigma_{yy}}{\partial y} + \frac{\partial \sigma_{yz}}{\partial z} + b_y = 0 \tag{6.2}$$

$$\frac{\partial \sigma_{zx}}{\partial x} + \frac{\partial \sigma_{yz}}{\partial y} + \frac{\partial \sigma_{zz}}{\partial z} + b_z = 0 \tag{6.3}$$

where σ_{xx}, σ_{yy} and σ_{zz} are the normal stresses, σ_{xy}, σ_{yz} and σ_{zx} denote the shear stresses and b_x, b_y, b_z are the body forces per unit volume acting in the coordinate directions x, y and z, respectively.

A more compact notation of the equilibrium equations may be obtained by expressing them in matrix form

$$[\mathbf{S}]^T \boldsymbol{\sigma} + \boldsymbol{b} = \mathbf{0} \tag{6.4}$$

where the six independent stress components are given as

$$\boldsymbol{\sigma} = [\sigma_{xx} \;\; \sigma_{yy} \;\; \sigma_{zz} \;\; \sigma_{xy} \;\; \sigma_{yz} \;\; \sigma_{zx}]^T \tag{6.5}$$

and the strain operator is

$$[\mathbf{S}]^T = \begin{bmatrix} \dfrac{\partial}{\partial x} & 0 & 0 & \dfrac{\partial}{\partial y} & 0 & \dfrac{\partial}{\partial z} \\ 0 & \dfrac{\partial}{\partial y} & 0 & \dfrac{\partial}{\partial x} & \dfrac{\partial}{\partial z} & 0 \\ 0 & 0 & \dfrac{\partial}{\partial z} & 0 & \dfrac{\partial}{\partial y} & \dfrac{\partial}{\partial x} \end{bmatrix} \tag{6.6}$$

Note that $\boldsymbol{\sigma}$ given in Eq. (6.5) is not a vector, but merely an ordered approach to represent the independent components of the full stress tensor in a column array. The vector of body forces is given consistently as

$$\boldsymbol{b} = [b_x \;\; b_y \;\; b_z]^T \tag{6.7}$$

Strains are a non-dimensional measure of the local displacement of a point within the body. Normal strains are defined as the change of the displacement components u, v, w in their respective coordinate directions x, y and z as

$$\varepsilon_{xx} = \frac{\partial u}{\partial x}, \;\; \varepsilon_{yy} = \frac{\partial v}{\partial y}, \;\; \varepsilon_{zz} = \frac{\partial w}{\partial z} \tag{6.8}$$

For small displacements, the shear strains acting respectively in the xy, yz and zx planes are defined as

$$\varepsilon_{xy} = \frac{\partial u}{\partial y} + \frac{\partial v}{\partial x}, \;\; \varepsilon_{yz} = \frac{\partial v}{\partial z} + \frac{\partial w}{\partial y}, \;\; \varepsilon_{zx} = \frac{\partial w}{\partial x} + \frac{\partial u}{\partial z} \tag{6.9}$$

Using the strain operator Eq. (6.6), a more explicit strain-displacement relation is obtained in matrix form as

$$\boldsymbol{\varepsilon} = [\mathbf{S}]\boldsymbol{u} \tag{6.10}$$

where the strains are defined in column array form as

$$\boldsymbol{\varepsilon} = [\varepsilon_{xx} \;\; \varepsilon_{yy} \;\; \varepsilon_{zz} \;\; \gamma_{xy} \;\; \gamma_{yz} \;\; \gamma_{zx}]^T \tag{6.11}$$

with

$$\gamma_{ij} = 2\varepsilon_{ij}, \;\; i \neq j \tag{6.12}$$

being the engineering shear strains and

$$\boldsymbol{u} = [u \;\; v \;\; w] \tag{6.13}$$

is the vector of displacements.

The relationship between stresses and strains is described by the constitutive relations, which are material dependent. For a linear elastic material the constitutive relations can be given in matrix form as

$$\boldsymbol{\sigma} = [\mathbf{D}](\boldsymbol{\varepsilon} - \boldsymbol{\varepsilon}_0) \tag{6.14}$$

where the 6 x 6 matrix $[\mathbf{D}]$ is the so-called elasticity matrix which contains the elastic moduli of the material. $\boldsymbol{\varepsilon}_0$ in Eq. (6.14) denote initial strains which are originating from effects other than displacements, as may be due to, for example, thermal expansion or compression.

Further, forces acting on the domain boundaries, which are commonly denoted as traction forces, originating for instance from gas loads, may be described as [185]

$$\bar{\boldsymbol{t}} = -p\boldsymbol{n} \tag{6.15}$$

where p is the pressure magnitude and $\boldsymbol{n}$ is the outward pointing unit surface normal.

The above equations in combination with appropriate boundary conditions describe the strong form of the boundary value problem, as the equations need to be satisfied at each point within the domain. An analytical solution of those equations cannot be obtained for the majority of practically relevant problems in engineering.

A weak form of the equations, which only requires the equations to be satisfied in an integral sense, can be derived based on the principle of virtual work, which provides the basis for the finite element discretisation.

Introducing an arbitrary kinematically admissible virtual displacement $\delta\boldsymbol{u}$ into the strong form of the equilibrium equations Eq. (6.4) and using Gauss' divergence theorem relating a surface integral to a volume integral, in addition to the virtual strain

$$\delta\varepsilon = [\mathbf{S}]\delta\boldsymbol{u} \tag{6.16}$$

the governing equations in weak form are obtained as

$$\int_\Omega \delta\varepsilon^T\boldsymbol{\sigma} \; d\Omega - \int_\Omega \delta\boldsymbol{u}^T\boldsymbol{b} \; d\Omega - \int_{\Gamma_t} \delta\boldsymbol{u}^T\bar{\boldsymbol{t}} \; d\Gamma = 0 \tag{6.17}$$

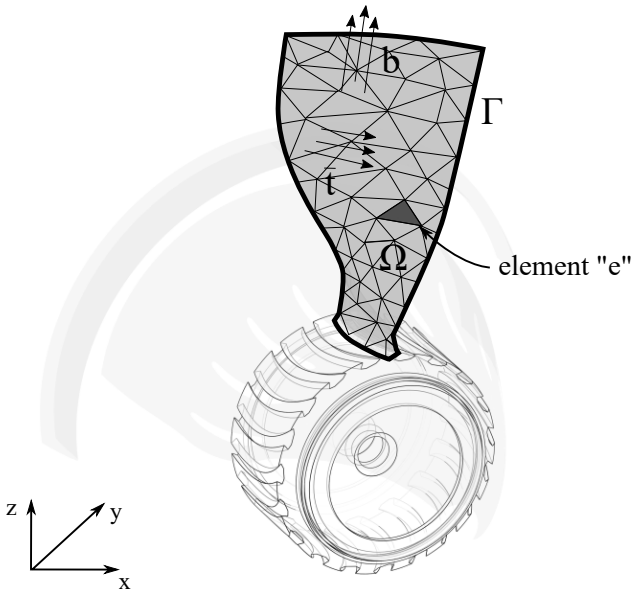


Figure 6.2: Depiction of the discretised general boundary value problem.

To obtain the equilibrium equations in finite element form, the domain is discretised into a number of elements or cells. This enables the required integrations to be performed element-wise, therefore providing an approximation of Eq. (6.17), see Fig. 6.2. The problem then translates to finding the displacements at discrete nodal points of the solid domain, while appropriate interpolation functions are applied to recover the internal distributions of displacements, stress and strain. The final problem then becomes the solution of a sparse system of linear equations, for which effective algorithms exist. The approximated governing equations read

$$\sum_e \int_{\Omega_e} \delta\varepsilon^T [[\mathbf{D}](\boldsymbol{\varepsilon} - \boldsymbol{\varepsilon}_0)] \; d\Omega - \sum_e \int_{\Omega_e} \delta \boldsymbol{u}^T \boldsymbol{b} \; d\Omega - \sum_e \int_{\Gamma_{t_e}} \delta \boldsymbol{u}^T \bar{\boldsymbol{t}} \; d\Gamma = 0 \qquad (6.18)$$

where the constitutive relations Eq. (6.14) have been introduced to replace $\boldsymbol{\sigma}$. An approximation of the element-internal displacements in dependency of the nodal displacements is obtained as

$$\boldsymbol{u}_e \approx \hat{\boldsymbol{u}}_e = [\mathbf{N}] \begin{bmatrix} \tilde{u} \\ \tilde{v} \\ \tilde{w} \end{bmatrix}_e = [\mathbf{N}] \; \tilde{\boldsymbol{u}}_e \qquad (6.19)$$

where $[\mathbf{N}]$ is the matrix of shape functions and $\tilde{\boldsymbol{u}}_e$ is the vector of nodal displacements.

Introducing Eq. (6.19) into Eq. (6.10) yields the strain approximation as

$$\boldsymbol{\varepsilon} = \begin{bmatrix} \varepsilon_{xx} \\ \varepsilon_{yy} \\ \varepsilon_{zz} \\ \gamma_{xy} \\ \gamma_{yz} \\ \gamma_{zx} \end{bmatrix} \approx \hat{\boldsymbol{\varepsilon}} = \begin{bmatrix} \dfrac{\partial}{\partial x} & 0 & 0 \\ 0 & \dfrac{\partial}{\partial y} & 0 \\ 0 & 0 & \dfrac{\partial}{\partial z} \\ \dfrac{\partial}{\partial y} & \dfrac{\partial}{\partial x} & 0 \\ 0 & \dfrac{\partial}{\partial z} & \dfrac{\partial}{\partial y} \\ \dfrac{\partial}{\partial z} & 0 & \dfrac{\partial}{\partial x} \end{bmatrix} [\mathbf{N}]\tilde{\boldsymbol{u}}_e = [\mathbf{B}]\tilde{\boldsymbol{u}}_e \tag{6.20}$$

where the matrix [**B**] is the so-called strain-displacement matrix which describes the derivatives of the shape functions with respect to global coordinates. The approximated virtual displacements and strains can be derived in a similar fashion. Introducing the displacement and strain approximations into Eq. (6.18) gives, after rearrangement, the approximated equilibrium equations of the body as

$$[\mathbf{K}]\tilde{\boldsymbol{u}} = \boldsymbol{f} \tag{6.21}$$

where [**K**] is the global stiffness matrix resulting from the summation over all element stiffness matrices [$\mathbf{K_e}$]

$$[\mathbf{K}] = \sum_e [\mathbf{K_e}] = \sum_e \int_{\Omega_e} [\mathbf{B}]^T [\mathbf{D}][\mathbf{B}] d\Omega \tag{6.22}$$

and $\boldsymbol{f}$ is the global load vector

$$\boldsymbol{f} = \sum_e \left(\boldsymbol{f}_{i_e} + \boldsymbol{f}_{t_e} + \boldsymbol{f}_{b_e} \right) \tag{6.23}$$

containing the element contributions of the forces due to initial strains $\boldsymbol{f}_{i_e}$, surface forces $\boldsymbol{f}_{t_e}$ and body forces $\boldsymbol{f}_{b_e}$ as

$$\boldsymbol{f}_{i_e} = \int_{\Omega_e} [\mathbf{B}]^T [\mathbf{D}] \boldsymbol{\varepsilon}_0 \, d\Omega \tag{6.24}$$

$$\boldsymbol{f}_{t_e} = \int_{\Gamma_{t_e}} [\mathbf{N}]^T \bar{\boldsymbol{t}} \, d\Gamma \tag{6.25}$$

$$\boldsymbol{f}_{b_e} = \int_{\Omega_e} [\mathbf{N}]^T \boldsymbol{b} \, d\Omega \tag{6.26}$$

The surface and volume integrals for the determination of the element stiffness matrices and element forces are typically evaluated numerically using Gauss integration. Together with appropriate boundary conditions, the system of algebraic linear

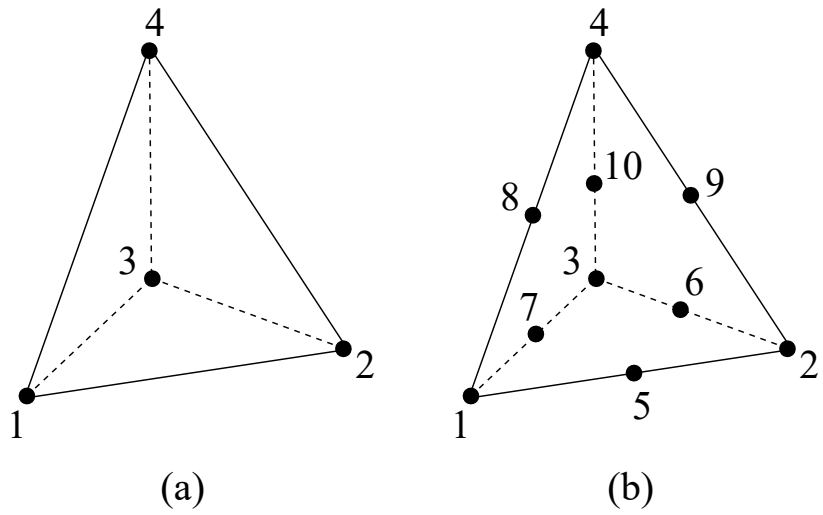


Figure 6.3: A linear (a) and quadratic tetrahedral element (b).

equations Eq. (6.21) can be solved using a direct or iterative solution method to yield the nodal displacements of the solid. Using the strain-displacement and constitutive relations, strains and stresses can be recovered.

6.3 Element choice

It was chosen in this work to discretise the solid domain using quadratic tetrahedral elements. Fully automated mesh generation methods exist for the generation of tetrahedral meshes, even for complex geometries. While four-node linear tetrahedral elements, shown in Fig. 6.3 (a), are known to be too stiff, quadratic tetrahedral elements with 10 nodes (four corner nodes and six mid-side nodes) provide a good accuracy and are thus widely used general purpose elements. The shape functions of the elements, which are contained in [**N**], Eq. 6.19, can be found in the literature (e.g. [184]). An example of a quadratic tetrahedral element is depicted in Fig. 6.3 (b).

6.4 Automatic mesh generation

An automated unstructured mesh generation tool is used for the discretisation of the solid blade geometry. Starting from the blade geometry description in the form of a boundary representation (B-rep) CAD model (resulting from the geometry generation process, Chapter 4), the automatic mesh generation of the solid domain follows a two step process, as illustrated in Fig. 6.4.

First, the blade surfaces are discretised using an unstructured mesh consisting only of triangular elements. Local target surface element sizes can be defined for the different faces of the blade, thus enabling control over the mesh resolution in the domain. For instance, areas of high-stress concentrations (i.e. in the root region) typically require a finer mesh than is required for the blade aerofoil.

In the next step, the volume mesh is generated starting from the existing surface mesh from the previous step, via the creation of solid quadratic tetrahedral elements for the interior part of the domain.

The meshing approach using solid finite elements presented here and used in this thesis assumes that the blade is an integral blade construction. If a hollow blade or a blade with internal cavities would need to be modelled, the present approach can in principle be applied as long as a geometrical representation of the blade in a CAD format is existing. However, such an approach would require the usage of extremely small solid elements in order to resolve thin walled features such as the blade face sheets and internal structures with an adequate mesh quality to avoid ill-conditioning of the problem. Consequently, the computational cost of such a modelling strategy would likely not be economical. A computationally more efficient approach would be to introduce two-dimensional shell elements for the modelling of thin-walled blade features, which intrinsically contain kinematic assumptions on their deformation behaviour under in-plane and out-of-plane loading. Many different shell theories are existing. Two common theories to mention are the Kirchhoff-Love plate model and

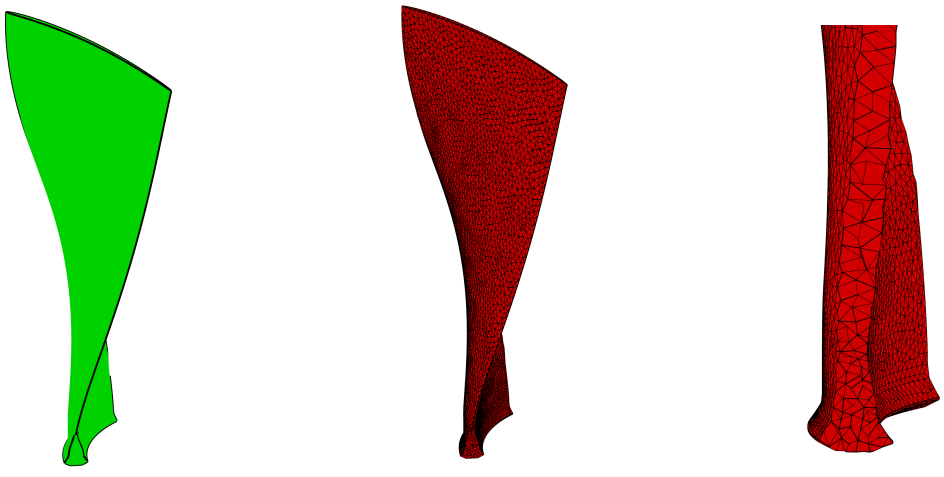


Figure 6.4: Solid mesh generation process.

Reissner-Mindlin shell theory, also known as First-Order Shear Deformation theory. A discussion of those is however out of the scope of the present text. The interested reader is referred to the finite element modelling literature, such as for instance [184] and [126].

6.5 Displacement boundary conditions

The element and global stiffness matrices are singular, therefore Eq. (6.21) cannot be solved directly. The reason is that the body is not restrained and can undergo unlimited rigid body motion. Appropriate boundary conditions in the form of limitations to the degrees of freedom of a set of nodes need to be imposed in order to make the problem solvable.

A common approach to simulate the disk attachment of the fan blade is to prevent displacements in all degrees of freedom for the nodes lying on the dovetail flank surfaces (i.e. the surfaces that are in contact with the disk in the assembled configuration). Hence, if the dovetail flank surfaces are denoted as Γ_u, then the displacement degrees of freedom of all nodes lying on those surfaces are

$$\boldsymbol{u}_{\Gamma_u} = 0. \tag{6.27}$$

6.6 Material modelling

6.6.1 An introduction to composite materials for fan blade structural design

While early aero-engine fan blades were fabricated as integral metallic blades, modern blades are constructed as either hollow titanium or carbon-fibre reinforced plastics (CFRP) composite structures for reduced weight. Due to significant recent advances in composite manufacturing, and computational design and analysis methods, all major aero engine manufacturers have recently transitioned to the development of fan blades made of composite materials. However, a wide variety of design and fabrication techniques exists resulting in a large number of possible configurations. This aspect is perhaps best highlighted by the number of existing patents in the field, e.g. [108], [54], [132], [4], [124], [59], [88], [111], [20], [110], [29], [172], [101], [49], [105], [13], [103], [102], [112], [85], [113], [121], [72], [142], [120], [114], [8], [104].

A composite material is made of two or more materials, which are known as constituents. For engineering materials, those are typically strong and stiff fibres

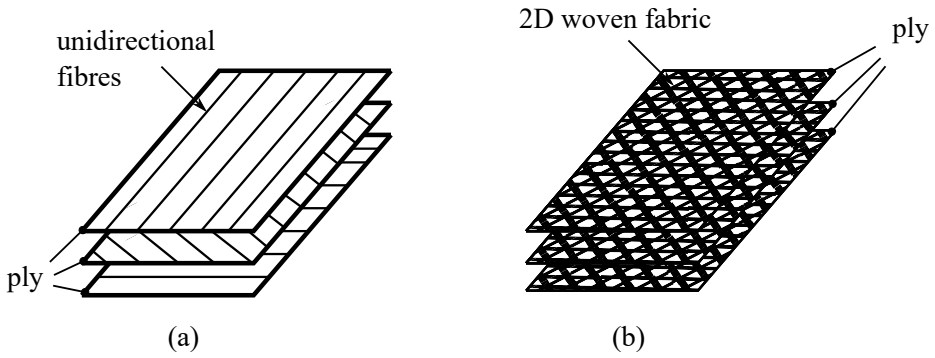


Figure 6.5: Schematic illustration of composite laminates. (a) plies with unidirectional fibres and (b) 2D woven fabric.

embedded in a polymer matrix material. The fibres act as the primary load carrying constituent, while the matrix holds the fibres in place and distributes the load between them. The constituent materials of choice for large aero-engine fan blades are typically carbon-fibres embedded in a high toughness epoxy matrix. The resulting material is widely known as carbon-fibre reinforced plastic (CFRP). Carbon fibres and matrix may be combined in a number of different ways. Generally, fibres may be chopped, continuous, straight, curved or woven and oriented in two or three dimensions. Continuous fibres are typically applied when strength and stiffness are critical, as is the case for aero-engine blading designs.

In their most common form, composite materials are made by stacking a number of layers or plies on top of each other to form a so-called *laminate*, see Fig. 6.5. Each ply represents itself a composite material with typically unidirectional continuous fibres embedded in them (Fig. 6.5 (a)), but fibres in a ply may also be woven, leading to so-called two-dimensional woven fabric composites (Fig. 6.5 (b)). The design of each individual ply (e.g. in terms of fibre volume fractions, layer thickness etc.) and the layer orientations and ply stacking sequence, define the overall mechanical properties of the resulting composite laminate. They can be designed and optimised to meet specific structural performance targets [55] (although many practical laminates are stacked in such a way, that their resulting global in-plane properties are isotropic, thus avoiding unexpected warping in the manufacturing process, and facilitating design and quality control. Such composites are often termed as black-metal, or black-aluminium). Composite laminates are commonly applied for many aerospace structures. They have further been successfully used for the construction of the first composite fan blades to enter commercial service in a high-bypass ratio aero-engine in the mid 1990s and many of those are in service to this day.

A major challenge associated with the application of laminates for fan blade constructions is the risk of delamination upon impact. Delamination describes the process of de-bonding of adjacent plies in the laminate, which can rapidly lead to ultimate failure of the material. So-called three-dimensional woven fabric composite materials, circumvent this issue by adding fibres in the through thickness direction. This avoids the possibility of delamination, however, at the cost of a typically lower in-plane stiffness compared to laminates. Different weave patters can be used to tailor the mechanical behaviour of the material. Those aspects make three-dimensional woven fabrics attractive for the application in aerospace structures. They have been recently introduced in production fan blades for high-bypass ratio turbofan engines.

6.6.2 Finite element analysis of composite materials

The discussion above has outlined two common design and manufacturing strategies for composite fan blades - laminated and three-dimensional woven fabric materials. The structural response of either of those materials can be predicted using the finite element method.

In principle, composites can be numerically analysed on three different levels of fidelity, known as micro-scale, meso-scale and macro-scale levels [14]. The most appropriate analysis approach is the one that yields the most economical resolution of the relevant physics. As usual, a trade-off exists between the accuracy and the computational cost of the modelling approach. All three approaches will be briefly discussed below.

6.6.2.1 Micro-scale analysis

In a micro-scale analysis, the fibres and matrix are directly resolved. Very fine meshes are needed for that purpose, hence only very small parts (e.g. unit cells) of the material can be economically analysed with currently available computational resources. Micro-scale analyses are useful to gain an improved understanding of the micromechanical material behaviour in terms elastic and inelastic matrix-fibre interactions, the mechanics of material failure onset and damage progression. For those reasons, micromechanical analyses are often useful to complement experimental tests, but the methods are too computationally expensive for design purposes of a composite material of any practical size.

6.6.2.2 Meso-scale analysis

In a meso-scale approach, each layer of a composite laminate is modelled as a homogeneous anisotropic material, thus fibres and matrix of a layer are not individually resolved. The layers themselves are however resolved in the computation. This method is less applicable to three-dimensional woven composites. Each layer needs to be numerically discretised, which is most economically performed by using so-called layered shell elements. Interlaminar stresses and deformations can be expressed with this approach, and therefore it does lend itself to design and optimisation of composite structures, see e.g. [55].

6.6.2.3 Macro-scale analysis

In a macro-scale approach, the three-dimensional composite material is modelled as a smeared unit, where only the global resultant mechanical properties of the material are considered. This analysis method does not allow for the resolution of layer stresses or strains, but displacements and mode shapes can be accurately predicted, and an indication of the overall stress levels can be obtained, if the corresponding global anisotropic material properties are known and imposed in the computation. This method is well suited for fan blade shape optimisation, as it allows the desired structural responses to be computed in an economic fashion. For those reasons, this is the chosen approach for the composite blade model used in this work.

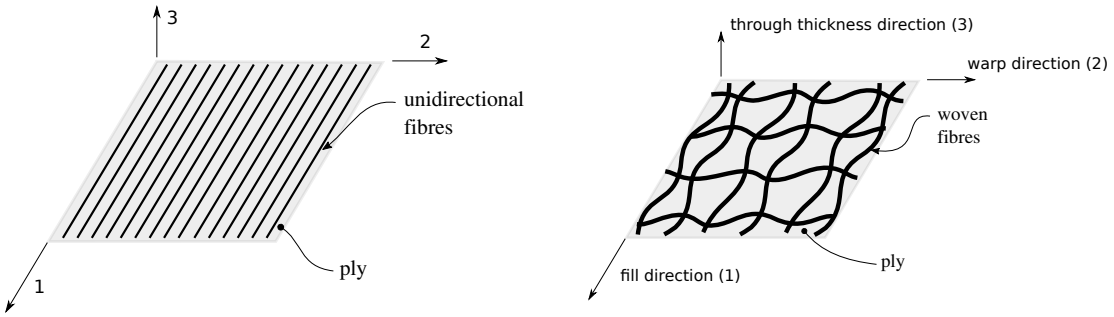
6.6.3 Constitutive relations

Section 6.2 laid out the general derivation of the finite element problem statement for the linear static analysis of an elastic solid. No distinction was made with regard to the material model used. The constitutive relations, which were introduced for a linear elastic material in Eq. (6.14), provide the link between the independent stress and strain components within the material. The relations will be recalled here for convenience

$$\boldsymbol{\sigma} = [\mathbf{D}](\boldsymbol{\varepsilon} - \boldsymbol{\varepsilon}_0) \tag{6.28}$$

As noted earlier, the elasticity matrix $[\mathbf{D}]$ contains the elastic moduli of the material.

The three-dimensional constitutive relation for a linear, fully anisotropic material is given by



(a) **(b)**

Figure 6.6: Reinforced composite plies and their principal material directions. (a) A unidirectionally reinforced ply and (b) two-dimensional woven ply.

$$\begin{bmatrix} \sigma_1 \\ \sigma_2 \\ \sigma_3 \\ \sigma_4 \\ \sigma_5 \\ \sigma_6 \end{bmatrix} = \begin{bmatrix} D_{11} & D_{12} & D_{13} & D_{14} & D_{15} & D_{16} \\ D_{12} & D_{22} & D_{23} & D_{24} & D_{25} & D_{26} \\ D_{13} & D_{23} & D_{33} & D_{34} & D_{35} & D_{36} \\ D_{14} & D_{24} & D_{34} & D_{44} & D_{45} & D_{46} \\ D_{15} & D_{25} & D_{35} & D_{45} & D_{55} & D_{56} \\ D_{16} & D_{26} & D_{36} & D_{46} & D_{56} & D_{66} \end{bmatrix} \begin{bmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \varepsilon_3 \\ \varepsilon_4 \\ \varepsilon_5 \\ \varepsilon_6 \end{bmatrix} \tag{6.29}$$

For an elastic material the elasticity matrix must be symmetric [185], and therefore $D_{ij} = D_{ji}$. Consequently, the linear elastic material characteristic of an anisotropic material is fully defined by 21 independent elastic constants.

A material is said to have *orthotropic* material properties, if it has three mutually orthogonal planes of symmetry, or in other words, the material properties differ in three mutually orthogonal axes. An example that presents this relation in an intuitive fashion is a unidirectional fibre reinforced composite ply shown schematically in Fig. 6.6 (a). Due to the directional nature of the fibres, the ply has its highest stiffness along the fibre direction, while the stiffnesses transversely to the fibres in directions 2 and 3 are much lower. However, the transverse in-plane (2) and out-of-plane (3) stiffness are not identical, which can be explained by fibre and matrix acting effectively "in-series" when loaded along the 2-direction versus "in-parallel" when loaded along the 3-direction, respectively.

While the microscopic material structure of a woven fabric layer is more complex than that of a unidirectionally reinforced one, on a macroscopic level it can still be considered as a quasi-homogeneous orthotropic material with the warp, fill and through thickness directions as the principal material axes [32], as depicted schematically in Fig. 6.6 (b). The coordinate system defined by the axes that coincide with

the principle material directions (in this example 1, 2 and 3), is the so-called *principal material coordinate system.*

Global laminate properties can be approximated from its individual layer properties (e.g. [14]). Alternatively, laminate or three-dimensional woven-fabric properties can be determined experimentally. In many cases, the global elastic behaviour can be reasonably well characterised as an orthotropic material (e.g. [32]). We assume for the further discussion that the global properties have been determined by suitable means and are thus available.

For an orthotropic material, the constitutive relation reduces to [14]

$$\begin{bmatrix}\sigma_1\\ \sigma_2\\ \sigma_3\\ \sigma_4\\ \sigma_5\\ \sigma_6\end{bmatrix} = \begin{bmatrix}D_{11} & D_{12} & D_{13} & 0 & 0 & 0\\ D_{12} & D_{22} & D_{23} & 0 & 0 & 0\\ D_{13} & D_{23} & D_{33} & 0 & 0 & 0\\ 0 & 0 & 0 & D_{44} & 0 & 0\\ 0 & 0 & 0 & 0 & D_{55} & 0\\ 0 & 0 & 0 & 0 & 0 & D_{66}\end{bmatrix} \begin{bmatrix}\varepsilon_1\\ \varepsilon_2\\ \varepsilon_3\\ \varepsilon_4\\ \varepsilon_5\\ \varepsilon_6\end{bmatrix} \tag{6.30}$$

Hence, only nine elastic constants are needed for the full definition of its behaviour. The components of $[\mathbf{D}]$ can be expressed in terms of engineering constants as [14]

$$\begin{aligned}
D_{11} &= \frac{1-\nu_{23}\nu_{32}}{E_2E_3\Delta}\\
D_{12} &= \frac{\nu_{12}+\nu_{32}\nu_{13}}{E_1E3\Delta}\\
D_{13} &= \frac{\nu_{13}+\nu_{12}\nu_{23}}{E_1E_2\Delta}\\
D_{22} &= \frac{1-\nu_{13}\nu_{31}}{E_1E_3\Delta}\\
D_{23} &= \frac{\nu_{23}+\nu_{21}\nu_{13}}{E_1E_2\Delta}\\
D_{33} &= \frac{1-\nu_{12}\nu_{21}}{E_1E_2\Delta}\\
D_{44} &= G_{23}\\
D_{55} &= G_{13}\\
D_{66} &= G_{12}\\
\Delta &= \frac{1-\nu_{12}\nu_{21}-\nu_{23}\nu_{32}-\nu_{31}\nu_{13}-2\nu_{21}\nu_{32}\nu_{13}}{E_1E_2E_3}
\end{aligned} \tag{6.31}$$

With $\nu_{i,j}$, E_i and G_{ij} being the Poisson's ratios, Young's moduli and shear moduli with respect to the global *material coordinates*, respectively. In most FEM software

packages, including CALCULIX, the relations in Eq. (6.31) are hard-coded, enabling the definition of an orthotropic material via a direct input of the engineering constants. A transformation of the elastic constants from material to element coordinates is required in order to correctly determine the individual element stiffness matrices in dependency of the blade geometry. The method implemented in this work for the definition of the local element coordinate systems will be presented in Section 6.7.

For the sake of completeness, it should be noted that a further simplification of the constitutive relations is obtained for so-called *transversely isotropic* materials, which exhibit only one axis of symmetry and for which the number of independent material constants reduces to five. Important representatives of this class of materials are so-called quasi-isotropic laminates, which make up a large part of composite structures presently in use in the aerospace industry. Quasi-isotropic laminates exhibit isotropic in-plane properties, while only the out-of-plane properties are different. For a transversely isotropic material with the isotropic plane being the 2-3 plane, $G_{13} = G_{12}$, $\nu_{13} = \nu_{12}$, and $E_3 = E_2$.

6.7 Automatic element coordinate system generation

For the computation of the element stiffness matrices in global coordinates ($[\mathbf{K_e}]$ in Eq. 6.22), the constitutive relations need to be transformed from material coordinates, in which they are most conveniently measured and defined, to a suitable local element coordinate system, such that the directional properties of the material are correctly taken into account in the global stiffness matrix assembly.

To enable this transformation, element coordinate systems need to be defined for each element in the solid mesh, as schematically depicted in Fig. 6.7. The directions chosen for the element coordinate systems should have a physical meaning in the sense that they represent the principle material directions consistently and in dependency of the blade geometry. This is of particular importance for the out-of-plane properties, which typically differ substantially from the in-plane ones. But if the in-plane properties of the composite are anisotropic, then also the in-plane orientations need to be adequately and consistently defined between the elements. For this purpose, an automated methodology has been implemented which generates the element coordinate systems in reference to the blade geometry, such as to enable the correct assignment of the orthotropic constitutive properties for different blade shapes,

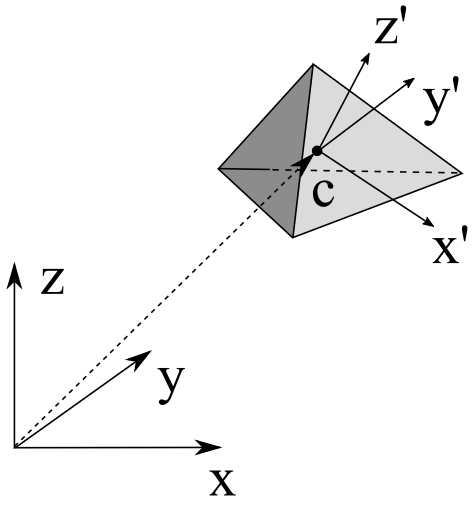


Figure 6.7: Definition of a Cartesian element coordinate system (x',y',z') in dependency of the global coordinates (x,y,z).

which is a central requirement for the application of the structural model in a shape optimisation framework.

The overall process is described in brief in the pseudo-algorithm listed in Algorithm 2. It is initiated after the volume mesh generation has been completed, and utilises the elements and their associated nodal coordinate values available from the meshing process, as well as the geometric CAD model.

As a first step, a reference geometry needs to be defined, which serves to represent the overall blade shape. For this purpose, a parametric mid-surface is generated. The part of the mid-surface that extends over the aerofoil portion of the blade coincides with the camber surface. For the root section, the mid-surface is oriented such that it intersects the geometry in its geometric centre.

Once the midsurface is generated, a loop is initiated over all elements contained in the solid mesh. For each element, the centroid is computed from knowledge of the nodal coordinate points. Subsequently, the point lying on the midsurface which is closest to the element centroid is calculated. On the parametric surface, two unit length tangent vectors, with directions respectively in the u- and v-directions are determined. Those vectors are subsequently translated to have their origin in the element centroid. The vector pointing in the spanwise direction is by convention defined to represent the first coordinate axis of the element coordinate system x', see Fig. 6.7.

The second (out-of-plane or z') coordinate axis is defined next by using the vector cross-product of the two in-plane vectors. The third orthogonal in-plane coordinate

Algorithm 2 Element coordinate system generation

Generate a parametric midsurface of the blade, which includes the aerofoil camber surface and extends down towards the bottom of the dovetail root
for element=1 to number of elements **do**
  Compute the centroid of the tetrahedron.
  Compute the point on the midsurface which is closest to the centroid.
  For this point, compute the unit tangent vectors in the parametric u and v directions of the mid-surface. By convention, the u-direction coincides with the spanwise direction.
  Translate the tangent vectors to the element centroid and define one of them as x'.
  Compute the out-of-plane coordinate z' from the cross-product of the in-plane vectors.
  Compute the second in-plane coordinate axis (y') using the cross-product of x' and z'.
end for

axis (y') then results again from a cross product of the out-of-plane z' and the x' axes. This process results in the definition of a unique Cartesian coordinate system for each element, being in accordance with the blade geometry.

With knowledge of the coordinate systems, the transformation of the constitutive relations can be carried out using an appropriate definition of transformation matrices, which contain as components the direction cosines between the global and the local element coordinate systems. Detailed explanations of those can be found in many standard textbooks on finite element modelling, e.g. [184], [14]. In most finite element packages, the relations are readily implemented, as is the case in Calculix. Hence, no further user intervention is needed once the coordinate systems have been suitably defined in the finite element program's input deck.

An example of the resulting element coordinate systems for a solid mesh of the baseline blade is shown in Fig. 6.8.

6.8 Failure criteria

Failure criteria present curve fits to experimentally determined material failure under uniaxial stress, attempting to predict failure under multiaxial stress. Various failure criteria exist for composite materials. The ones most commonly applied for design purposes are so-called *first-ply failure criteria*, which attempt to predict an initial localised material failure. Analyses of damage propagation leading to ultimate material failure requires higher fidelity modelling that is typically not feasible within

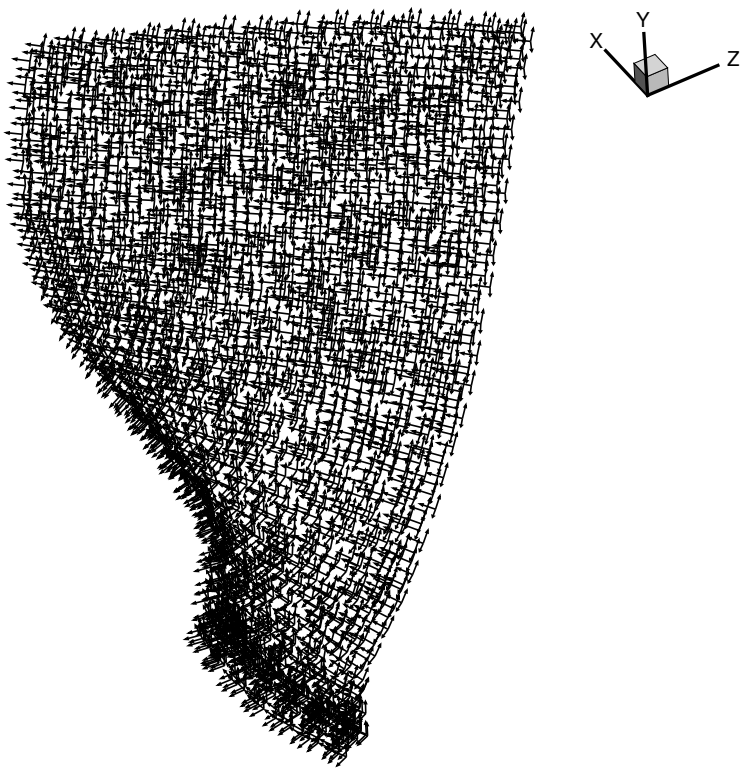


Figure 6.8: Element coordinate systems for the definition of orthotropic material properties for a solid mesh of the baseline fan blade.

preliminary design iterations, but may be used at a detailed design stage. First-ply failure criteria provide a conservative estimate of the composite failure load, since first failure in the form of matrix cracks may not lead to laminate failure [55].

A general overview of various existing first-ply failure criteria can be found in standard mechanics of composite materials textbooks, such as e.g. [80]. Among those are the so-called maximum stress, maximum strain, Hoffmann, Tsai-Hill and Tsai-Wu criteria. In this work the Tsai-Wu failure criterion is applied due to its wide application in the literature.

The Tsai-Wu failure criterion for three-dimensional orthotropic materials is defined as [163]:

$$\begin{aligned}\zeta_{TW} = F_1\sigma_1 + F_2\sigma_2 + F_3\sigma_3 + F_{11}\sigma_1^2 + F_{22}\sigma_2^2 + F_{33}\sigma_3^2 + F_{44}\sigma_4^2 + F_{55}\sigma_5^2 + F_{66}\sigma_6^2 \\ +2F_{12}\sigma_1\sigma_2 + 2F_{13}\sigma_1\sigma_3 + 2F_{23}\sigma_2\sigma_3 \leq 1\end{aligned} \tag{6.32}$$

With F_i and F_{ij} being the non-zero components of the strength tensor. The

strength tensor components relate to engineering strengths in the same way that elastic compliance tensor components relate to elastic engineering constants. This important relationship allows to link experimentally measured engineering strengths (that are not components of a tensor) to the components of the strength tensor. With the uniaxial engineering tension and compression strengths in the principal material directions given as X_t, X_c, Y_t, Y_c, Z_t, Z_c and the shear strengths in the 2-3, 3-1 and 1-2 planes given as Q, R and S, respectively, the strength tensor components in dependency of the engineering strengths become

$$\begin{aligned} F_1 &= \frac{1}{X_t} - \frac{1}{X_c},\ F_2 = \frac{1}{Y_t} - \frac{1}{Y_c},\ F_3 = \frac{1}{Z_t} - \frac{1}{Z_c}, \\ F_{11} &= \frac{1}{X_t X_c},\ F_{22} = \frac{1}{Y_t Y_c},\ F_{33} = \frac{1}{Z_t Z_c}, \\ F_{44} &= \frac{1}{Q^2},\ F_{55} = \frac{1}{R^2},\ F_{66} = \frac{1}{S^2} \end{aligned} \tag{6.33}$$

and the strength components F_{12}, F_{13} and F_{23}, which relate to the interaction of two normal stress components are defined as

$$\begin{aligned} F_{12} &= -\frac{1}{2\sqrt{X_t X_c Y_t Y_c}}, \\ F_{23} &= -\frac{1}{2\sqrt{Y_t Y_c Z_t Z_c}}, \\ F_{13} &= -\frac{1}{2\sqrt{X_t X_c Z_t Z_c}} \end{aligned} \tag{6.34}$$

6.9 Discussion of non-linear static effects

The structural response of fan blades may be subject to non-linear effects, arising from

- Surface contact between the blade root (dovetail) and disk
- Material non-linearities
- Large deformation (geometric non-linearities)

Each of those points will be discussed separately in the sections below.

6.9.1 Blade root-disk contact

The surface contact between the dovetail flanks and the disk presents a non-linear problem [150]. To accurately account for the dovetail root surface loads, this contact problem needs to be resolved in the finite element analysis [184]. However, contact simulations add significant complexity to the structural analysis and degrade the robustness of the solution process. In the shape optimisation framework presented here, where a large number of different blade and root geometries are to be analysed, it is therefore assumed that the disk attachment can be simulated by the application of zero displacement boundary conditions on the dovetail root flank surfaces.

6.9.2 Material non-linearities

Material non-linearities have been observed in composite materials, including two and three-dimensional woven composites [71], [174] under small and large strains. Those are mainly attributed to the complex interactions between the tow and matrix materials, resulting in e.g. transverse matrix cracking and plastic straightening of misaligned tows with the loading direction. Modelling these micromechanical effects is complicated and while some non-linear constitutive models have been derived for e.g. unidirectional plies [57], generality of those cannot be assumed. This is simply due to the sheer variability of material configurations that can be achieved with composite materials. It is therefore not attempted in this work to include material non-linearities in the constitutive relations and the linear model as introduced in Eq. (6.14) is used.

6.9.3 Geometric non-linearities

Geometric non-linearity arises when the deformation of the structure becomes too large to retain the assumption of constant load and stiffness. Fan blades are subject to large deformations particularly at the tip section. Under the small strain assumption applied for geometric linear analyses, the strains are referenced to the undeformed geometry of the body. If the deformation becomes large, the full Green-Lagrange strain tensor needs to be used instead [184]

$$E_{IJ} = \frac{1}{2}\left(F_{iI}F_{iJ} - \delta_{IJ}\right) \tag{6.35}$$

which is expressed in terms of the so-called deformation gradient

$$F_{iI} = \frac{\partial x_i}{\partial X_I} \tag{6.36}$$

relating the Cartesian coordinates of a material point in a current (deformed) configuration x_i, to the initial reference configuration X_I. δ_{IJ} denotes the Kronecker delta in the reference configuration.

Stresses are correspondingly expressed using the second Piola-Kirchhoff stress S_{IJ} which is a symmetric stress measure with respect to the reference configuration. The Kirchhoff stress τ_{ij} results from the second Piola-Kirchhoff stress as

$$\tau_{ij} = F_{iI} S_{IJ} F_{jJ} \tag{6.37}$$

with the Kirchhoff stress being linked to the Cauchy stress σ_{ij} via the determinant of the deformation gradient by

$$\tau_{ij} = J\sigma_{ij}, \;\; J = \det \; F_{iI} \tag{6.38}$$

6.10 Vibration analysis

Fan blades are subjected to various sources of periodic disturbances during operation, which may arise from structural (unbalance) as well as aerodynamic origins. On the aerodynamics side, periodic excitations result from circumferential flow non-uniformities which most commonly originate from non-symmetric intakes, pressure variations due to downstream blade rows, pylons or struts, or wakes from upstream obstacles such as probes. For safe operation, it is critical to ensure that the blades do not enter into a resonant condition with the periodic forcing, in which excessive vibrational stresses can quickly lead to structural failure. Resonance occurs when the forcing frequency, particularly its low-order harmonics, coincides with one of the natural frequencies of the blade.

Assuming a rigid disk (i.e. a disk with a significantly higher stiffness than the blades) resonance will be driven by the blade only vibration behaviour. The existence of possible resonant conditions within the operating range is typically assessed in the Campbell diagram, shown in Fig. 6.9, where the natural frequencies of a blade are plotted in combination with the excitation frequencies (engine orders, EO) over the rotational speed. The engine order lines are independent of the blade geometry. A first engine order excitation line (1EO), for example, represents a disturbance that appears once per revolution. Because, the forcing functions arising in the engine are normally not purely sinusoidal, they contain higher order harmonics. The second harmonic of a one-per-revolution disturbance is represented by the second engine order line (2EO). Similarly, this line represents a disturbance appearing twice per

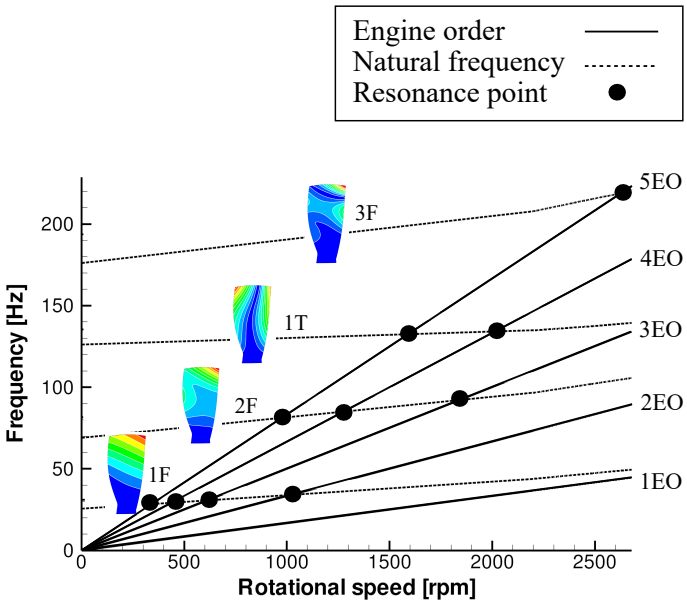


Figure 6.9: Campbell diagram of the baseline titanium fan blade.

revolution with each disturbance appearing 180 degrees phase shifted. The relation between the engine order frequency and the rotational speed is given by the simple expression

$$f = \mathrm{EO} \cdot n/60 \tag{6.39}$$

with f being the frequency in Hz, n is the rotational speed in rpm and EO represents the engine order number.

Blade natural frequencies required for the construction of the Campbell diagram can be predicted using a finite element modal analysis. In addition to the stiffness matrix $[\mathbf{K}]$ derived in the previous sections, the mass matrix $[\mathbf{M}]$ is now required to solve the dynamic problem, which is not particularly difficult to construct, see e.g. [126]. Neglecting forcing and damping, the free vibration response of the blade can be found by solving the following generalised eigenvalue problem

$$([\mathbf{K}] - \omega_n^2[\mathbf{M}])\boldsymbol{\phi} = \mathbf{0} \tag{6.40}$$

which yields the natural frequencies $\omega_{n,j}$, $j = 1, 2, ..., N$ and corresponding mode shape vectors $\boldsymbol{\phi}_j$, $j = 1, 2, ..., N$. The number of natural frequencies and mode shapes N is identical to the number of degrees of freedom in the system. To avoid the occurrence of rigid body modes, displacement boundary conditions need to be applied, similarly to the static analysis, Section 6.5.

Natural frequencies of fan blades typically rise with increasing rotational speed; an effect known as centrifugal stiffening. In the finite element computation, this effect can be considered by an augmentation of the global stiffness matrix [$\mathbf{K}$] with a stress stiffening matrix [$\mathbf{K}_{\boldsymbol{\sigma}}$], which results from a previous static analysis in which the centrifugal loads are imposed.

In the Campbell diagram, resonance is indicated by the crossing of an engine order excitation line with one of the natural frequency curves, thus representing a matching of excitation and natural frequencies. The task in the design process is to ensure that the most dangerous crossings are avoided at the engine's main operating points, such as for example at cruise and take-off. However, in practice, small variations of the actual blade natural frequencies in comparison to the numerically predicted ones are virtually unavoidable due to numerical and modelling errors, manufacturing tolerances, small variations of the actual operating points from the cycle design and changes of the blade geometry due to static aeroelastic untwist, as well as the impact of structural and aerodynamic damping on the resonance frequency. To account for those effects, a minimum frequency margin is typically required between excitation and natural frequencies at critical operating points.

An important note to make is that, while the Campbell diagram presents a means for the investigation of temporal matching of blade natural frequencies and potential periodic excitations, it does not include any information on the vibration amplitudes and vibratory stress levels at a potential resonance point, because the levels of forcing and damping remain unknown at this stage of the analysis. To gain a more detailed insight into the dangerousness of a particular resonance point, a computation of the fluid-structure interaction problem in terms of a forced response analysis is required, which presents a significant increase in computational model detail, complexity and solution cost.

The Campbell diagram presents a useful design tool for the manual assessment of the vibrational behaviour of fan blades, but also presents a suitable conceptual framework for the implementation into an automated optimisation system. This aspect will be further explained in Chapters 7 and 8.

7 Multidisciplinary Design Optimisation of a Titanium Fan Blade

The optimisation framework described in detail in the previous chapters will now be applied to the aeromechanical design of a large titanium high-bypass ratio aero-engine fan blade. The fan blade parametrisation will first be defined, followed by an explanation of the high-fidelity performance analysis setup and the optimisation settings chosen. Next, the derivation of the optimisation problem formulation will be presented, followed by a discussion of the results. The chapter is finalised with a concluding summary.

7.1 Fan blade parametrisation

The parametrisation of the fan blade is now discussed based on the constructive CAD-based parametrisation method developed in Chapter 4. In the following subsections, a detailed description of the chosen parameters is given, which is followed by a summary of all optimisation parameters in table form in subsection 7.1.7.

7.1.1 Meridional contour

The curves defining the hub, casing and bypass splitter end-wall contours of the baseline design are shown in Fig. 7.1 in the meridional plane. Although those curves are fully parametrised, they are not altered throughout the optimisation process presented in this chapter.

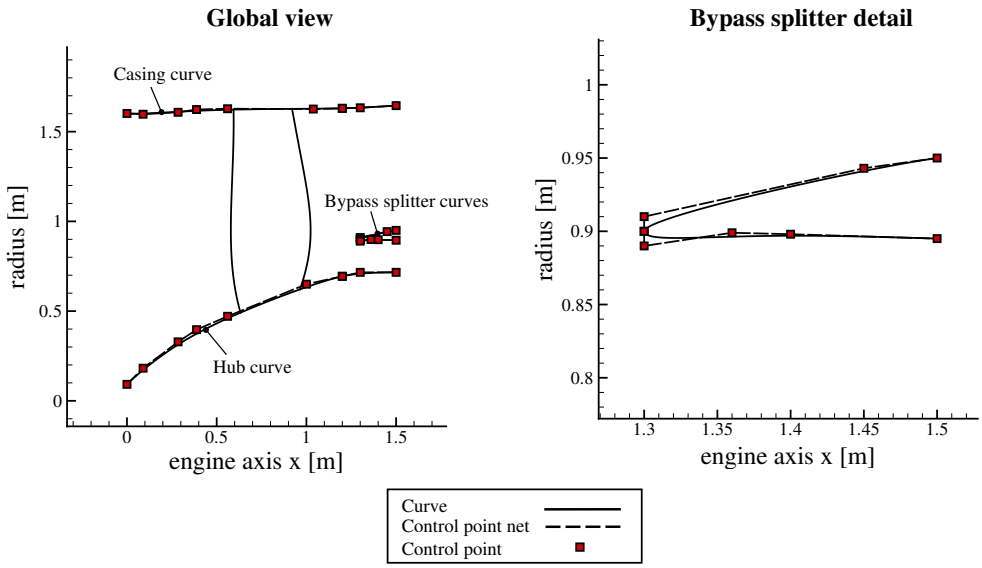


Figure 7.1: Meridional contour of the baseline geometry.

7.1.2 Camber surface

The parametric curves defining the spanwise blade angle and chord length distributions for the definition of the blade's camber surface are shown in Fig. 7.2. Arrows in the plots indicate how the control points are allowed to change throughout the optimisation process. Black arrows indicate direct degrees of freedom, i.e. parameters that can be directly modified by the optimiser, while red arrows indicate indirect de-

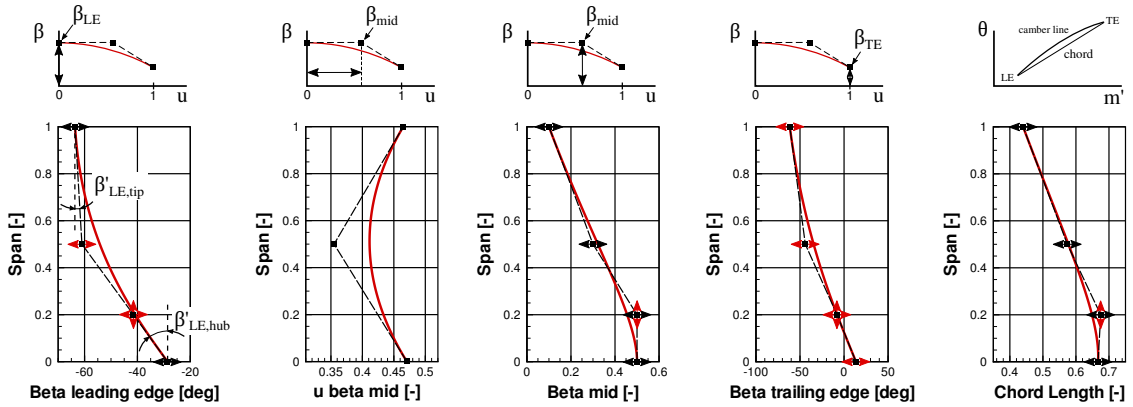


Figure 7.2: Camber surface parametrisation. Black arrows indicate direct degrees of freedom of control points (optimisation parameters), while red arrows indicate indirect degrees of freedom (quantities that are defined in dependency to other optimisation parameters). For convenience, the figures on the top of the plot point out which quantity is being defined by the respective parametric curve.

grees of freedom, i.e. parameters that are defined in dependency to other optimisation parameters. All spanwise distributions are allowed to change with the exception of the axial position of the mid-point beta control point $u_{\beta_{mid}}$, for which the baseline distribution is kept unchanged in the present optimisation. For the remaining curves, control points are defined respectively at 0, 20, 50 and 100% span. In order to allow additional control of the camber surface close to the bypass splitter, the spanwise coordinate of the control point initially defined at 20% span is allowed to change. The value of the optimisation parameter simultaneously affects all distributions in order to retain consistency of the control points.

All four control points defining the leading edge curve are allowed to be changed by the optimiser. The hub and tip values can be modified directly, while the two intermediate control points are defined via angle dependencies to the hub and tip control points. The angles are defined as optimisation parameters, which allows to retain information on the relative position of the control points to each other, which facilitates the optimisation process.

A pseudo camber angle is defined for the determination of the control point values defining the spanwise trailing edge blade angle distribution in dependency of the corresponding leading edge blade angle control point values. The pseudo camber angle is defined as

$$\lambda' = \beta_{TE} - \beta_{LE} \tag{7.1}$$

where β_{LE} denotes the blade angle coordinate value of the leading edge control point and β_{TE} denotes the corresponding control point value of the spanwise trailing edge blade angle curve. At the hub and the tip, where the spanwise curve values coincide with the control point values, the pseudo camber angle is equivalent to the actual camber angle.

The pseudo camber angles at the four dedicated spanwise positions are defined as optimisation parameters, therefore defining the control point values of the spanwise trailing edge blade angle curve in dependency of the leading edge blade angle control points and the pseudo camber angle. The camber angle has a direct physical implication as it controls the blade loading distribution, which makes it easier for the optimiser to identify the impact of the optimisation parameter changes on the blade performance.

The mid-point blade angle is defined in a non-dimensional form in dependency of the leading edge and trailing edge blade angle control point values as $\beta_{mid}^{ND} \in (0, 1)$, resulting in a dimensional mid-point beta value as

$$\beta_{mid} = \beta_{LE} \cdot (1 - \beta_{mid}^{ND}) + \beta_{TE} \cdot \beta_{mid}^{ND} \tag{7.2}$$

Hence, for $\beta_{mid}^{ND} = 1$, the mid-point blade angle is identical to the trailing edge blade angle, while it equals the leading edge blade angle for a value of $\beta_{mid}^{ND} = 0$. An example of how this parameter affects the camberline shape is presented in Fig. 7.3 for the definition of the mid-point beta value at $u_{\beta_{mid}} = 0.5$. As shown in in Fig. 7.3 (a), a value of $\beta_{mid}^{ND} = 1$ results in a camberline shape with the curvature being concentrated at the front. In contrast, $\beta_{mid}^{ND} = 0$, shown in Fig. 7.3 (c) shifts the curvature and hence the flow turning towards the rear of the camberline. A value of $\beta_{mid}^{ND} = 0$ consequently leads to a uniformly distributed curvature, as can be observed in Fig. 7.3 (b).

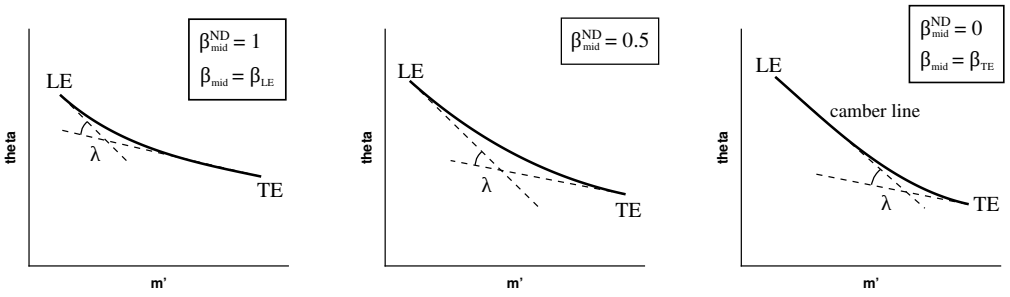
A negative value of the non-dimensional mid-point beta value can be used to obtain a precompression (S-shaped) camberline, as shown in Fig. 7.4 for a value of $\beta_{mid}^{ND} = -0.3$.

In the parametrised spanwise non-dimensional mid-point blade angle distribution shown in Fig. 7.2, all control points are allowed to be changed independently by the optimiser. Similarly, the control point values defining the chord length can be changed independently, whereas a chord length scaling factor is defined for each of the control point values.

7.1.3 Lean and sweep

The parametrisation of the spanwise lean-to-chord and sweep-to-chord ratio distributions is shown in Fig. 7.5. The distribution of the lean-to-chord ratio is defined with a second order Bézier curve with control points being specified at 0, 50 and 100% span. The control points at 50 and 100% span are allowed to be directly modified by the optimiser.

A fourth order Bézier curve is used for the definition of the sweep-to-chord distribution. Control points are defined at 0, 20, 70, 80 and 100% span, enabling an increased control of sweep at the higher span region where the relative inflow is supersonic. The spanwise position of the near hub control point at 20% span is changed according to the corresponding camber surface control points described previously. The impact of blade sweep on the flow field mainly stems from the spanwise variation of sweep and not from the absolute value of sweep at any one blade section. Three sweep angles are therefore defined as optimisation parameters, indicated as



(a) **(b)** **(c)**

Figure 7.3: Impact of the non-dimensional mid-point beta value β_{mid}^{ND} on the camberline shape for $u_{\beta_{mid}} = 0.5$. λ denotes the camber angle.

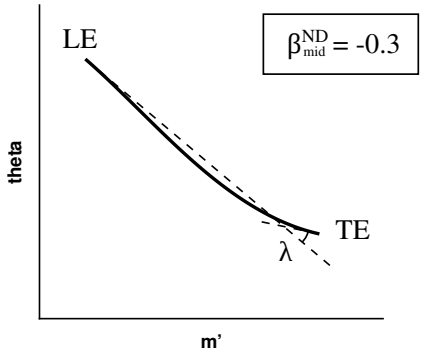


Figure 7.4: Camberline shape for a precompression airfoil with $u_{\beta_{mid}} = 0.5$.
λ denotes the camber angle.

σ_{hub}, $\sigma_{neartip}$ and σ_{tip} in Fig. 7.5. Solely the control point defining the tip value of sweep is allowed to be varied independently.

7.1.4 Thickness

Figure 7.6 (a) shows the parametrisation of the blade thickness distributions at hub and tip. In the present optimisation the thickness distributions are allowed to be uniformly scaled with an independent thickness scaling factor for the hub and tip section, respectively. An example of the spanwise interpolation of the parametrised thickness distributions as resulting from the baseline distributions shown here is presented in Fig. 7.6 (b).

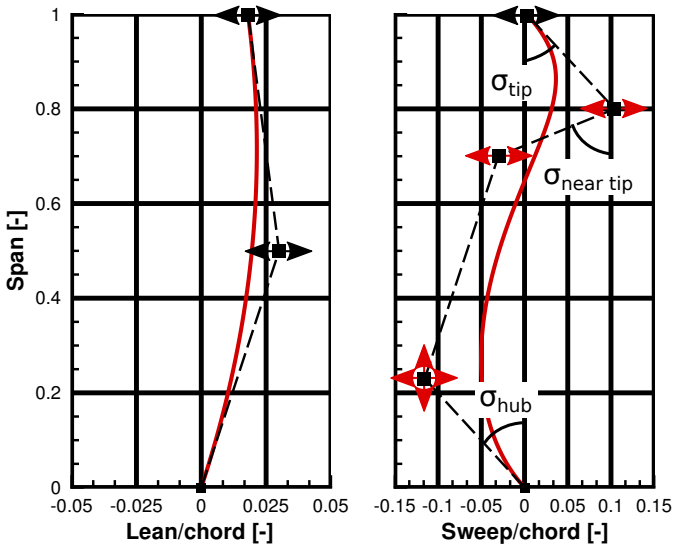


Figure 7.5: Parametrisation of lean and sweep. Note that the baseline values of all control points defining lean and sweep are zero.

7.1.5 Blade root parametrisation

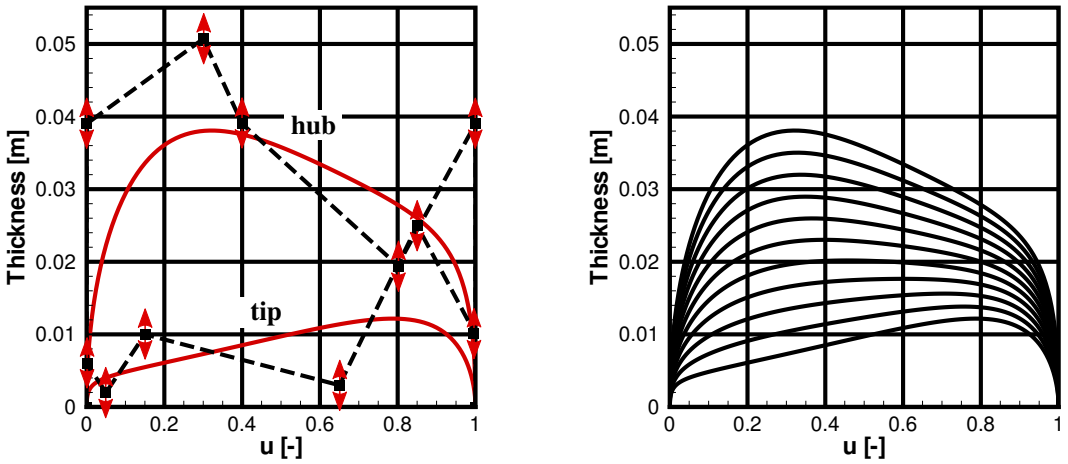
The parametrised cross section of the dovetail root is shown in Fig. 7.7. The dovetail width is variable and adapts to the blade hub thickness as described in detail in chapter 4. Apart from that, the root cross sectional shape is kept unchanged throughout the optimisation presented in this chapter. The example in Fig. 7.7 shows the cross section adapted to the baseline blade. Figure 7.8 gives an overview of the resulting fully constructed solid geometry of the baseline blade.

7.1.6 Number of blades

The number of blades is added as a last optimisation parameter with a lower bound of 17 and an upper bound of 19 (baseline).

7.1.7 Summary of optimisation parameters

The parametrisation results in a total of 26 optimisation parameters. Table 7.1 gives a global overview of all defined optimisation parameters and their range.



(a) **(b)**

Figure 7.6: Parametrisation of the blade thickness distributions at hub and tip (a) and spanwise interpolated distributions (b).

7.2 Aerodynamic performance evaluation

The three-dimensional steady flow through the fan is computed using the commercially available 3D Reynolds-Averaged Navier-Stokes solver FINE^{TM}/Turbo. The theoretical background of the fluid solver was given in Chapter 5. This section will give an explanation of the case-specific setup.

7.2.1 Fluid mesh generation

The fluid domain comprising one periodic section of the full annulus is discretised with a multi-block structured mesh consisting of about 2 million grid points. In total 14 structured mesh blocks are used for the entire domain. The average non-dimensional height of the wall-adjacent cell (y^+) is approximately equal to three. 156 mesh points are used to discretise the geometry in the radial direction, including 17 points in the tip gap. A global overview of the surface mesh of one periodic section of the baseline fan is given in Fig. 7.9.

A template-based semi-automated meshing process is employed for the generation of fluid meshes for designs generated during the optimisation process. This meshing method uses the structured baseline multi-block (template) mesh and maps the blocking and grid point distributions to the fluid domain of a different geometry. The process works well for geometries that do not differ strongly, and it was found

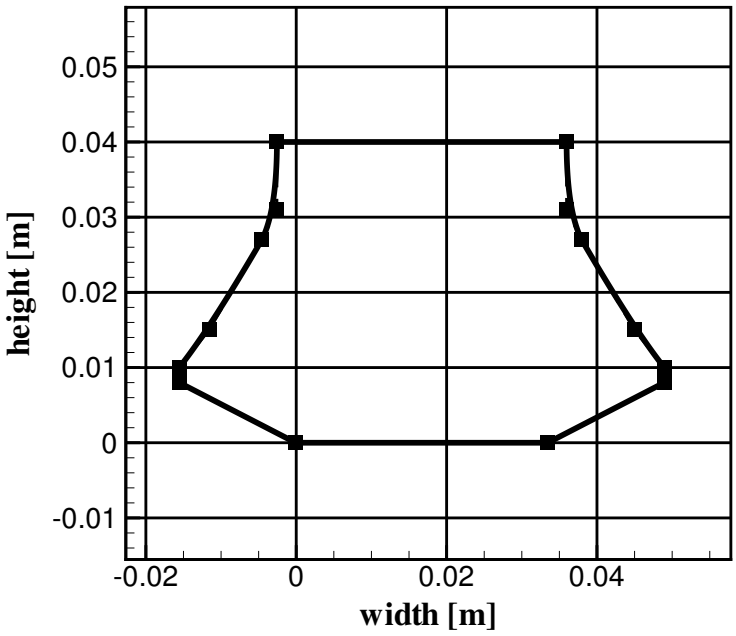


Figure 7.7: Parametrisation of the dovetail root cross section. The width of the shape shown is adapted to the thickness of the baseline blade.

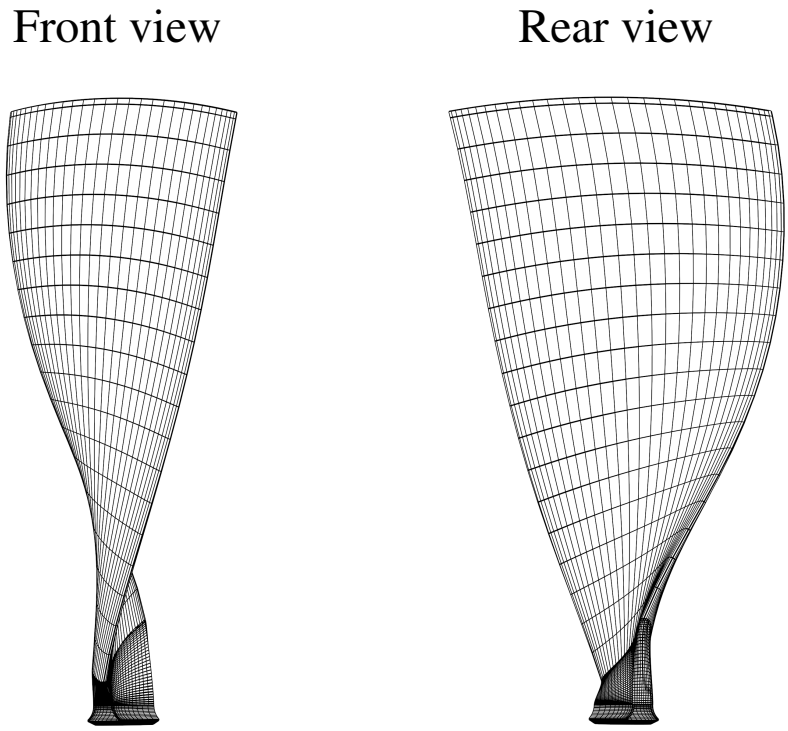


Figure 7.8: Geometry of the parametrised baseline blade with fully constructed root geometry.

Table 7.1: Summary of optimisation parameters and their range.

#	Name	Unit	Lower bound	Upper bound
1	Beta.LeadingEdge.Hub	deg	-38.6	-18.6
2	Beta.LeadingEdge.Hub.Angle	deg	27.9	37.9
3	Beta.LeadingEdge.Tip.Angle	deg	0.0	6.0
4	Beta.LeadingEdge.Tip	deg	-66.6	-60.6
5	Beta.Mid.Hub	[-]	0.2	0.8
6	Beta.Mid.NearHub	[-]	0.2	0.8
7	Beta.Mid.MidSpan	[-]	0.1	0.5
8	Beta.Mid.Tip	[-]	-0.3	0.5
9	CamberAngle.Hub	[deg]	37.0	47.0
10	CamberAngle.NearHub	[deg]	28.0	38.0
11	CamberAngle.MidSpan	[deg]	12.0	22.0
12	CamberAngle.Tip	[deg]	0.0	4.0
13	ChordScalingFactor.Hub	[-]	1.0	1.3
14	ChordScalingFactor.NearHub	[-]	1.0	1.3
15	ChordScalingFactor.MidSpan	[-]	1.0	1.25
16	ChordScalingFactor.Tip	[-]	1.0	1.2
17	NearHub.Cntrl.Pnt.Span	[-]	0.15	0.25
18	Lean.MidSpan	[-]	-0.05	0.05
19	Lean.Tip	[-]	-0.03	0.03
20	SweepAngle.Hub	[deg]	-45.0	45.0
21	SweepAngle.NearTip	[deg]	-80.0	80.0
22	SweepAngle.Tip	[deg]	-50.0	50.0
23	Sweep.Tip	[-]	-0.004	0.004
24	Hub.ThicknessScalingFactor	[-]	1.0	2.5
25	Tip.ThicknessScaling.Factor	[-]	0.5	1.5
26	NumberOfBlades	[-]	17	19

that it provides a robust and efficient method for the shape optimisations performed in the framework of this thesis.

A mesh independence study was performed a-priori to the optimisation process for the estimation of the discretisation error. Three systematically refined meshes were used with respectively 41 515 (coarse), 283 606 (medium) and 2 084 902 (fine) mesh points. The results obtained for the baseline design operating at the cruise design point are shown in Fig. 7.10 for the obtained total outlet mass flow, absolute total pressure ratio and isentropic total-to-total efficiency. It can be seen, that the results show a monotonous convergence towards an exact solution. An approximation of the exact solution was computed here using Richardson extrapolation (e.g. [48]). The

resulting estimated absolute values of the respective discretisation errors are indicated in the plots. The maximum error magnitude for the fine mesh (occurring for the efficiency) is 0.16%, or 0.00143 in absolute terms, which is deemed sufficiently accurate for the optimisation process. The fine mesh is hence used for the optimisation.

7.2.2 Fluid and turbulence modelling

The air is modelled as an ideal gas. Turbulence effects are resolved using the one-equation turbulence model of Spalart and Allmaras [154].

7.2.3 Operating points and boundary conditions

To meet the aerodynamic design and off-design targets originating from the engine cycle design given in Table 1.1, the fan's performance is evaluated at two key flight conditions, namely top-of-climb (ToC) and cruise. Top-of-climb presents the aerodynamically highest load condition for the fan, having a comparable pressure ratio requirement to take-off, but at a reduced rotational speed (c.f. Table 1.1). For that reason and for the sake of computational cost, take-off is not explicitly evaluated for aerodynamic performance.

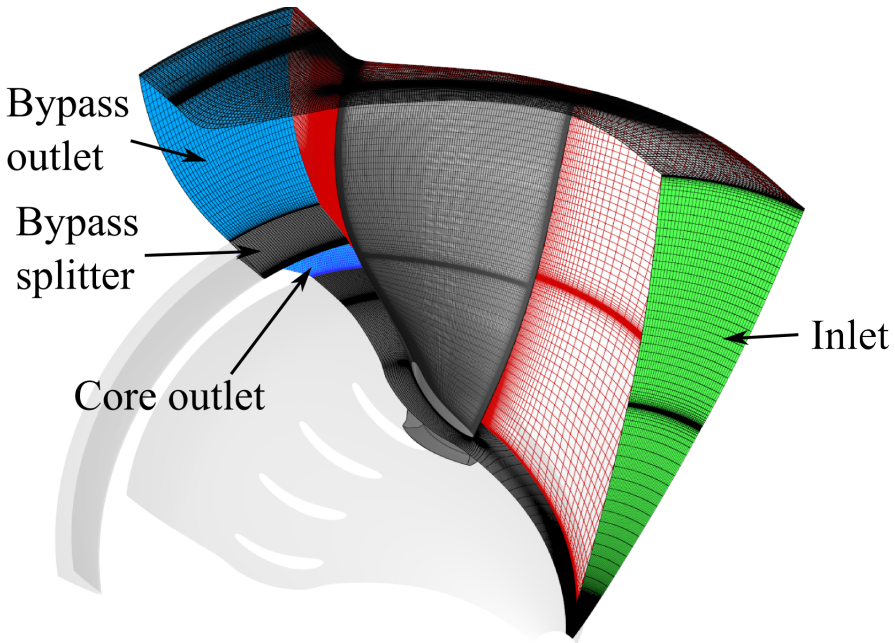


Figure 7.9: Fluid domain and surface mesh point distributions of the baseline design.

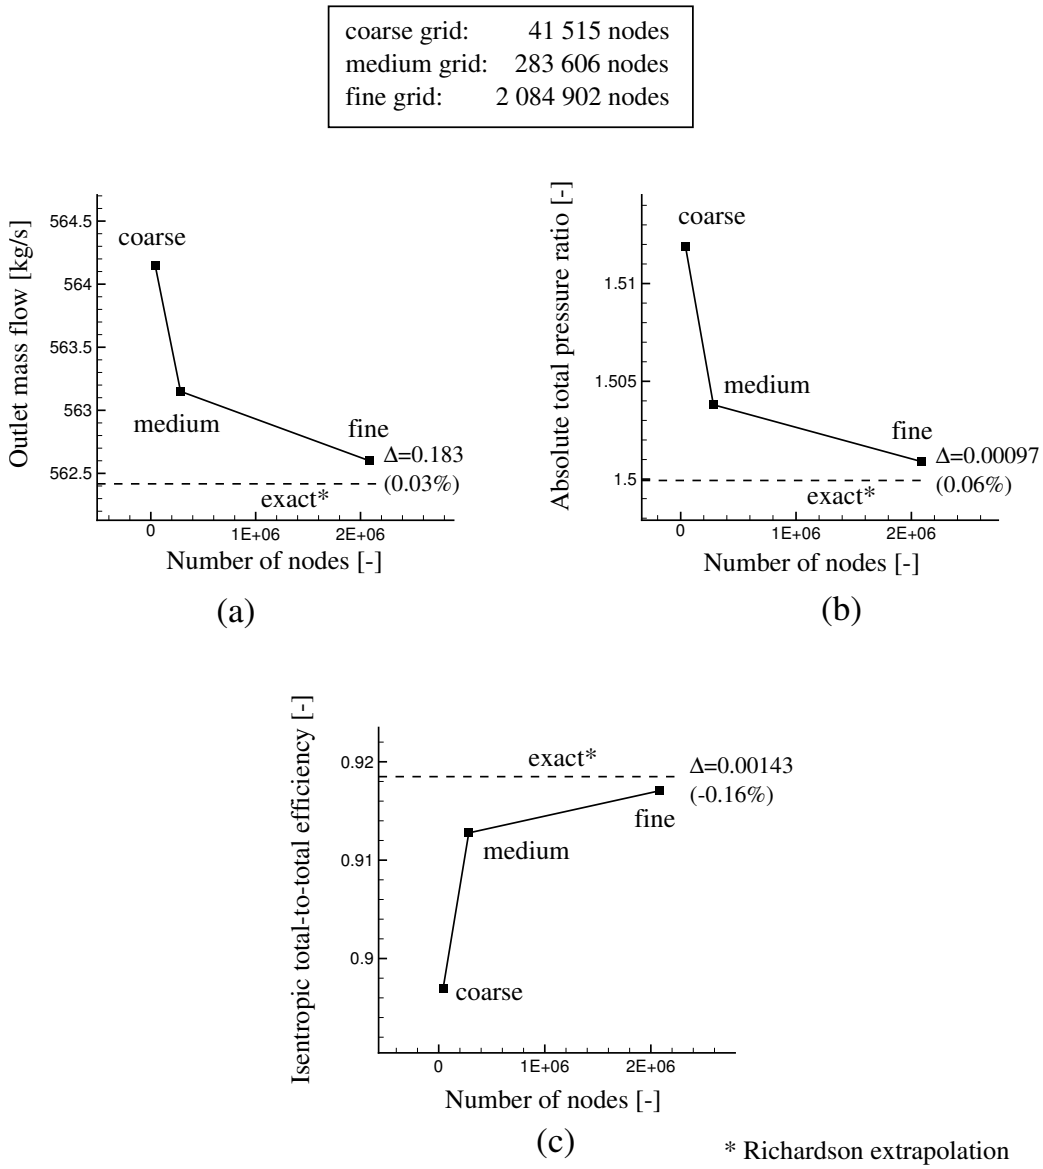


Figure 7.10: Mesh independence study for the baseline blade at the cruise design point.

In total nine operating points are computed for each design, including four points on the ToC speedline and five points on the cruise speedline. Total pressure, total temperature, the absolute inlet flow angle and the turbulent kinematic viscosity are imposed as boundary conditions at the inlet. The mass flow is imposed as boundary condition at the core outlet, while the static pressure using the radial equilibrium law is defined as boundary condition at the bypass outlet. Performance curves are computed via adaptation of the bypass outlet static pressure. Periodic boundary conditions are imposed on the periodic surfaces of the domain. Walls are treated as

adiabatic.

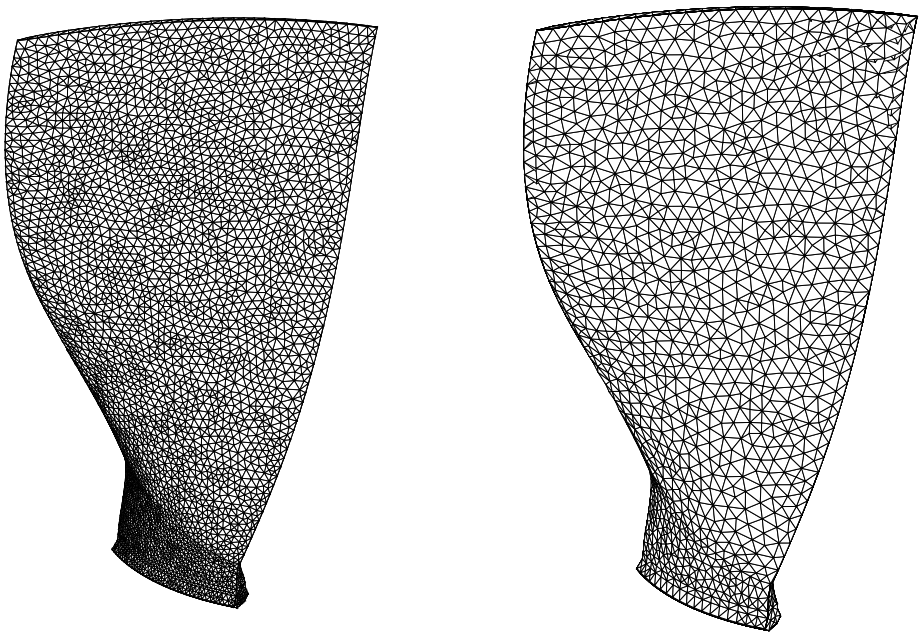
An automatic convergence check is performed after each CFD computation to assess the mass flow error between domain inlet and outlet as well as the iteration errors of efficiency and pressure ratio. Only cases with sufficient convergence are subsequently assessed by an automatic postprocessing step, which extracts the required performance parameters needed by the optimiser. Non-converged cases are considered as failed and are automatically excluded from the optimisation process.

7.3 Structural performance evaluation

The structural response of the fan blade including static stresses and vibrations is evaluated using the open-source finite element solver CalculiX [41]. Details of the finite element solution procedure were provided in Chapter 6.

7.3.1 Solid mesh generation

The solid fan blade is discretised using quadratic tetrahedral elements. An automated unstructured meshing process is applied, as described in Section 6.4. For the sake of



(a) Static analysis mesh.

(b) Frequency analysis mesh.

Figure 7.11: Solid meshes of the baseline blade.

Table 7.2: Properties of Ti-6Al-4V (Grade 5) titanium alloy.

Young's modulus	Poisson's ratio	Mass density	Yield strength
E	ν	ρ	X
GPa	-	kg/m^3	MPa
113.8	0.342	4430	880

computational efficiency, two different mesh resolutions are used for the static stress and vibration analyses. A fine mesh is used for the static analyses to enable an accurate resolution of stress concentrations. In contrast, a coarser mesh is sufficient to obtain an accurate frequency response. The exact number of nodes and elements varies between geometries due to the automated unstructured mesh generation process used. To provide a reference, the baseline static analysis mesh consists of 60833 nodes. The baseline frequency mesh consists of 16077 nodes. An overview of both meshes is provided in Fig. 7.11.

7.3.2 Displacement boundary conditions and material model

In the present optimisation, the fan blade is modelled using material properties of a high-strength isotropic titanium alloy. An overview of the corresponding material properties is given in Table 7.2.

Displacement boundary conditions are defined for all nodes being located on the dovetail flanks. Displacements in all three translational degrees of freedom of the element nodes are set to zero in order to simulate the disk attachment, as shown in Fig. 7.12 by example of the baseline static analysis mesh.

7.3.3 Static stress analysis

A static finite element computation including geometric non-linear effects is performed to evaluate the static stresses in the fan blade. The computation is performed at take-off conditions, where both aerodynamic and centrifugal loads are at their peak values throughout the mission. The take-off rotational speed is listed in Table 1.1. The blade is subjected to centrifugal and gas loads whereas the gas loads are extracted from the converged take-off CFD computation of the baseline geometry and interpolated onto the non-matching FEM surface grid using an inverse distance weighting approach. Due to the negligence of the contact problem between the blade root and the disk, the resulting stresses on the dovetail flanks can be considered to be inaccurate. Those

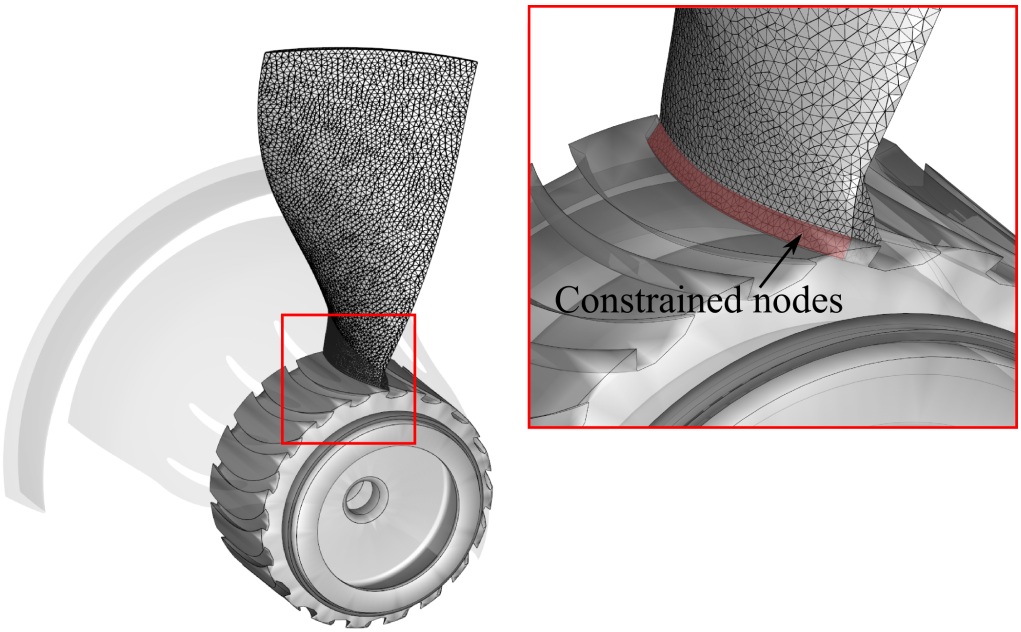


Figure 7.12: Solid mesh of the baseline fan blade and its restraint in a typical disk. The nodes on the dovetail flank surfaces shown in red are fixed in all degrees of freedom to simulate the disk attachment (opposite constrained surface not visible).

stresses are therefore disregarded in the postprocessing step for the assessment of the peak stress values in the blade.

7.3.4 Vibration analysis

The free vibration response of the fan blade is computed at all previously mentioned key operating points (cruise, ToC and T/O) in order to assess the risk of high cycle fatigue failure. Natural frequencies of the fan blade are determined by means of finite element modal analyses. A geometric non-linear static pre-stressing computation is performed prior to each modal analysis to take the non-linear effect of the centrifugal forces on the blade stiffness at different rotational speeds into account. The Campbell diagram is used to compute the margins between first and second engine order excitation frequencies and the blade natural frequencies at the rotational speeds associated with the aforementioned operating points.

7.4 Objectives and constraints

The optimisation problem is composed of two objectives and a total of twelve aerodynamic and structural constraints. For the long range aircraft mission considered in this thesis, low specific fuel consumption at cruise operation is usually of highest concern for a reduction of direct operating cost, because the aircraft operates longest in this condition. Fan blades have a very favourable exchange rate between efficiency and SFC, with a 1% improvement in fan efficiency typically translating to an 0.7% improvement of engine SFC [151], making the achievement of high aerodynamic fan efficiency of prime concern in the design process. The first objective is therefore defined to maximise the peak efficiency at cruise rotational speed. The second objective is defined to maximise the fan's stall margin, defined with respect to the mass flow. Both objectives are translated to an equivalent minimisation problem, formally stated as

$$\text{mimimise} \begin{cases} -\eta_{cruise,peak} \\ -\left(\frac{\dot{m}_{cruise,design}-\dot{m}_{cruise,stall}}{\dot{m}_{cruise,stall}}\right) \end{cases} \tag{7.3}$$

where $\eta_{cruise,peak}$ is the isentropic total-to-total peak efficiency of the blade at cruise rotational speed, $\dot{m}_{cruise,design}$ is the cruise design point mass flow as required by the engine cycle design (Table 1.1), and $\dot{m}_{cruise,stall}$ is the mass flow of the last converged CFD computation towards the stall limit of the blade. It is hence assumed that numerical stall approximates the physical stall limit. The objectives are visualised in the performance map of the baseline design shown in Fig. 7.13.
The constraints are defined as follows

$$\dot{m}_{cruise,stall} < \dot{m}_{cruise,design} < \dot{m}_{cruise,choke} \tag{7.4}$$

$$0.97 \cdot \dot{m}_{toc,design} < \dot{m}_{eval,toc} < 1.03 \cdot \dot{m}_{toc,design} \tag{7.5}$$

$$\Pi_{t-t,bypass,eval,toc} > 1.7 \tag{7.6}$$

$$\Pi_{t-t,core,eval,toc} > 1.4 \tag{7.7}$$

$$Ma_{max,bypass,eval,toc} < 0.9 \tag{7.8}$$

$$Ma_{max,core,eval,toc} < 0.9 \tag{7.9}$$

$$\alpha_{max,bypass,eval,toc} < 70 \text{ deg} \tag{7.10}$$

$$\alpha_{max,core,eval,toc} < 70 \text{ deg} \tag{7.11}$$

$$\sigma_{vM} < 800 \text{ MPa} \tag{7.12}$$

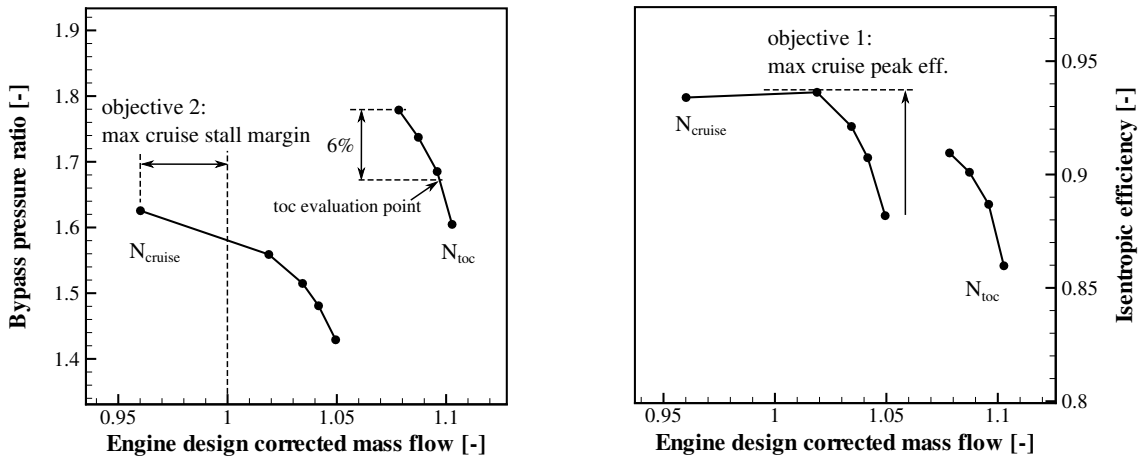


Figure 7.13: Definition of objectives and the top-of-climb evaluation point shown in the performance map of the baseline design.

$$\Delta\hat{F}_{cruise} + \Delta\hat{F}_{ToC} + \Delta\hat{F}_{T/O} > \Delta\hat{F}_{total,req} \tag{7.13}$$

The first two constraints, Eq. (7.4), are defined to ensure that the first and last converged CFD operating points on the cruise speedline, the "stall" and "choke" mass flows, are smaller, and respectively higher than the cruise design point mass flow. In other words, the constraints ensure that the design point mass flow is within the operating range of the cruise speedline. The requirement for those constraints arises from the relatively wide design variable ranges defined in this optimisation, which may lead to blade designs with significantly shifted operating ranges.

Following a similar reasoning, the third and fourth constraint, defined in Eq. (7.5), are formulated to ensure that the mass flow of a pre-defined evaluation point on the top-of-climb speedline is located within a six percent range around the top-of-climb design point mass flow. The top-of-climb evaluation point is determined as the interpolated point on the speedline having a six percent stall margin with respect to the total pressure ratio, as shown in Fig. 7.13.

Two constraints are defined to ensure that the pressure ratios for the core and bypass flows at the top-of-climb evaluation point are not lower than the design requirements, Eq. (7.6) and (7.7). Additionally, Eqs. (7.8) to (7.11) limit the maximum absolute inlet Mach numbers and the maximum absolute flow angles at the bypass and core inlet to ensure that the stators located further downstream (Outlet Guide Vane in the bypass and Engine Section Stator in the core) receive a healthy inlet flow and performance improvements of the fan blade are not obtained at the cost of overall

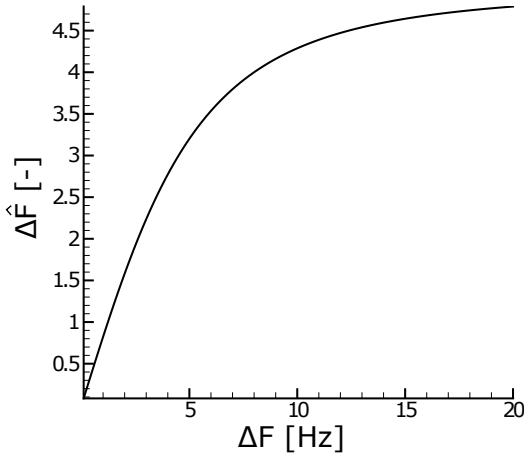


Figure 7.14: Definition of the non-dimensional frequency margin.

stage performance. The absolute flow angle is measured with respect to the meridional plane. Equation (7.12) defines a constraint on the static maximum von Mises stress in the fan blade, which needs to be lower than 800 MPa, leaving an 80 MPa dynamic stress margin towards the yield strength of the titanium alloy. Finally, a last constraint is defined in Eq. (7.13) to assess blade vibration at cruise, top-of-climb and take-off. The constraint imposes the requirement of a minimum margin between the eigenfrequency of the first bending mode of the blade and the first and second engine order excitation frequencies at the three aforementioned operating points. The constraint definition utilises a non-dimensional frequency margin at each operating point, defined as

$$\Delta\hat{F} = \xi_f \cdot \frac{\Delta F}{\sqrt{(\zeta_f + \Delta F^2)}} \tag{7.14}$$

where ΔF is the dimensional minimal margin between the excitation and natural frequency at the respective operating point, for example top-of-climb. The coefficient values are set to $\xi_f = 5$ and $\zeta_f = 36$. Equation (7.14) in combination with the constraint definition, Eq. (7.13), enforces that all operating points must simultaneously exhibit a frequency margin above a threshold value in order for the design to be feasible. More specifically, Eq. (7.14) limits the impact of the dimensional margin at a single-operating point on the overall non-dimensional constraint function value, once a sufficiently large dimensional margin is existing. Hence, for example, if a design has extensive margins at take-off and top-of-climb, but a critically small margin at cruise, it will remain infeasible until the margin at all operating points has been raised above

a minimum value. Equation (7.14) is plotted in Fig. 7.14. In the present case the total required non-dimensional frequency margin (Eq. (7.13)) is defined as $\Delta\hat{F}_{total,req} = 12$. This value, as well as the coefficient values ξ_f and ζ_f used in the relation here have been chosen to yield overall sensible results. The constants can be adapted to satisfy any required minimum dimensional margin for different design problems.

7.5 Optimisation settings

The optimisation was performed with the differential evolution constants set to $C = 0.8$ and $F = 0.6$, respectively. A population size of 100 individuals was chosen. 1000 differential evolution generations were computed per iteration. As the optimisation is performed using the Kriging metamodel as performance predictor, this step is computationally cheap, taking only a few seconds to complete. After each iteration, eight candidate solutions were chosen for high-fidelity re-evaluation from the resulting metamodel-predicted Pareto front.

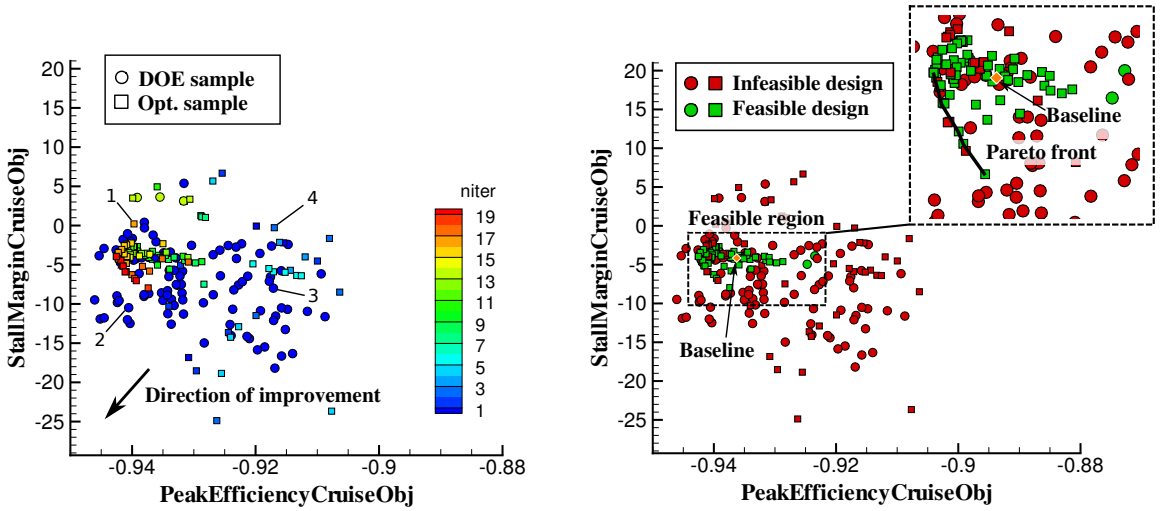
7.6 DOE database

The initial database used for the generation of the first metamodel was generated using a fractional factorial design of experiments (DOE) technique containing a total of 128 samples. The DOE samples the design space at a subset of the full factorial design at 20 and 80% of the design range. In addition to the DOE samples covering the boundaries of the design range, the baseline design and the "central" design (generated with all design variables set to the centre value of their range) were added to the initial database. 103 out of the resulting 130 samples were successfully evaluated by the high-fidelity evaluation chain. Only three samples in the DOE database satisfied all imposed constraints.

7.7 Optimisation results

7.7.1 Survey of optimisation process

The optimisation was run for a total number of 20 iterations. In total 122 out of 160 submitted individuals were successfully evaluated by the high-fidelity analysis tools. With the initial database samples included, 2025 steady CFD computations, 225 non-linear static finite element computations and 675 finite element frequency analyses (each including a non-linear static pre-stressing computation) were performed.



(a) Global view.

(b) Detail view on feasible individuals.

Figure 7.15: Objective space after 20 iterations. Circles indicate DOE samples, while squares indicate designs generated during the optimisation.

A global view on the resulting objective space is shown in Fig. 7.15 (a). Each symbol in the plot represents a design that was analysed by the high fidelity evaluation chain. A geometry overview of four designs that have been successfully evaluated is shown in Fig. 7.16. Improved performance in the objective space is obtained towards the lower left corner of the objective space, corresponding to the minimisation of both objectives. The symbols are coloured according to the iteration number in which they have been generated, with the initial database samples are represented by iteration 0, hence the rainbow colour evolution from blue to red allows to trace the optimisation process.

As many individuals initially showed constraint violations, the first task of the optimiser was the identification of the feasible region of the design space, while objective improvement is not considered at this stage of the optimisation process. The amount of individuals showing constraint violations is more clearly visible in Fig.. 7.15 (b) in which the individuals in the objective space are coloured according to their constraint satisfaction. Infeasible designs are indicated in red colour, whereas feasible ones are shown in green. It can be seen, that a large fraction of all analysed designs showed constraint violations. Feasible designs are seen to be clustered within a smaller part of the entire objective space. A close-up view on the feasible region is shown in the upper right hand corner of Fig. 7.15 (b). The comparison of both plots allows to see the evolution of the optimiser towards the region in the objective space, where the

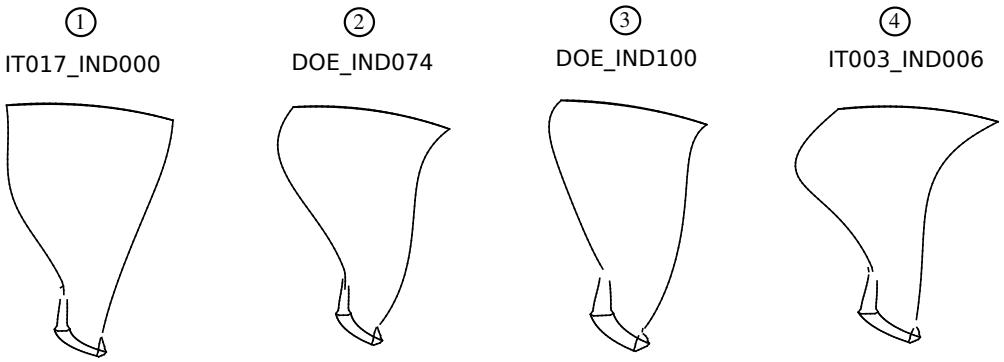


Figure 7.16: Geometries of selected individuals that have been successfully evaluated by the high-fidelity evaluation chain. Numbers correspond to Fig. 7.15 (a).

feasible designs are located.

A better understanding of the difficulty to satisfy individual constraints can be obtained by consideration of the constraint sub-spaces shown in Fig. 7.17 for all twelve constraints imposed in the present optimisation problem. The symbols shown represent designs that have been analysed by the high-fidelity evaluation chain throughout the optimisation process. Only negative values of the constraint functions indicate constraint satisfaction, hence individuals that are located in the first quadrant violate both considered constraints, while the second and fourth quadrants indicate the violation of one of the two respective constraints. Individuals that are located in the third quadrant satisfy both constraints. In conjunction to these plots, Fig. 7.18 shows the metamodel prediction errors for each performance parameter, averaged over the eight candidate solutions per iteration.

In Fig. 7.17 it can be seen that most constraints shown on the right hand side plots (Figures (b), (d) and (f)) are satisfied by most individuals, including many or all of the initial database samples, shown as circles. The fact that most individuals are located in the feasible area of those constraint sub-spaces indicates that the constraints are easy enough to satisfy that even early iterations being characterised by large metamodel prediction errors do not lead to constraint violations. This is perhaps most clearly visible for the maximum Mach number constraints shown in Fig. 7.17 (b) where all individuals are satisfying the constraints, even with metamodel prediction errors as large as -100% seen for the bypass maximum Mach number constraint in Fig. 7.18. Further, also the maximum flow angle constraints for the core and bypass ducts, shown in Fig. 7.17 (d), as well as the cruise stall mass flow constraint shown in Fig. 7.17 (f) seem to be relatively easy to satisfy. However, the choke mass flow con-

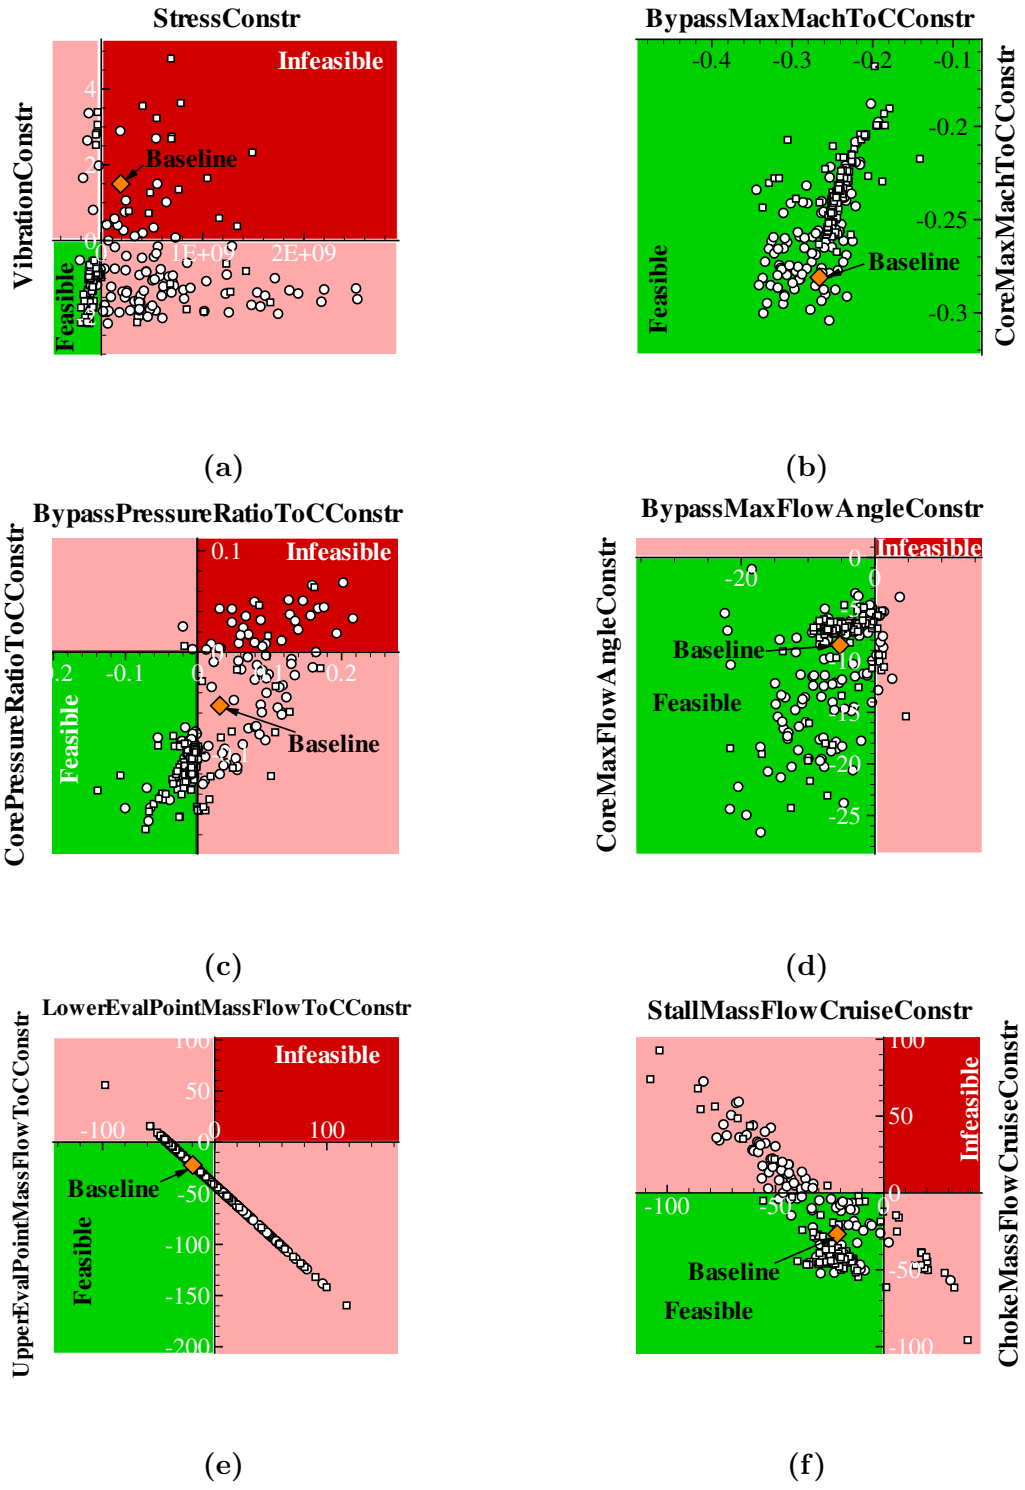


Figure 7.17: Overview of constraint spaces after 20 optimisation iterations. Every symbol shown represents a design evaluated by the high-fidelity evaluation chain. Symbols in red quadrants are violating the respective constraint. Circles indicate DOE samples, while squares indicate designs generated during the optimisation.

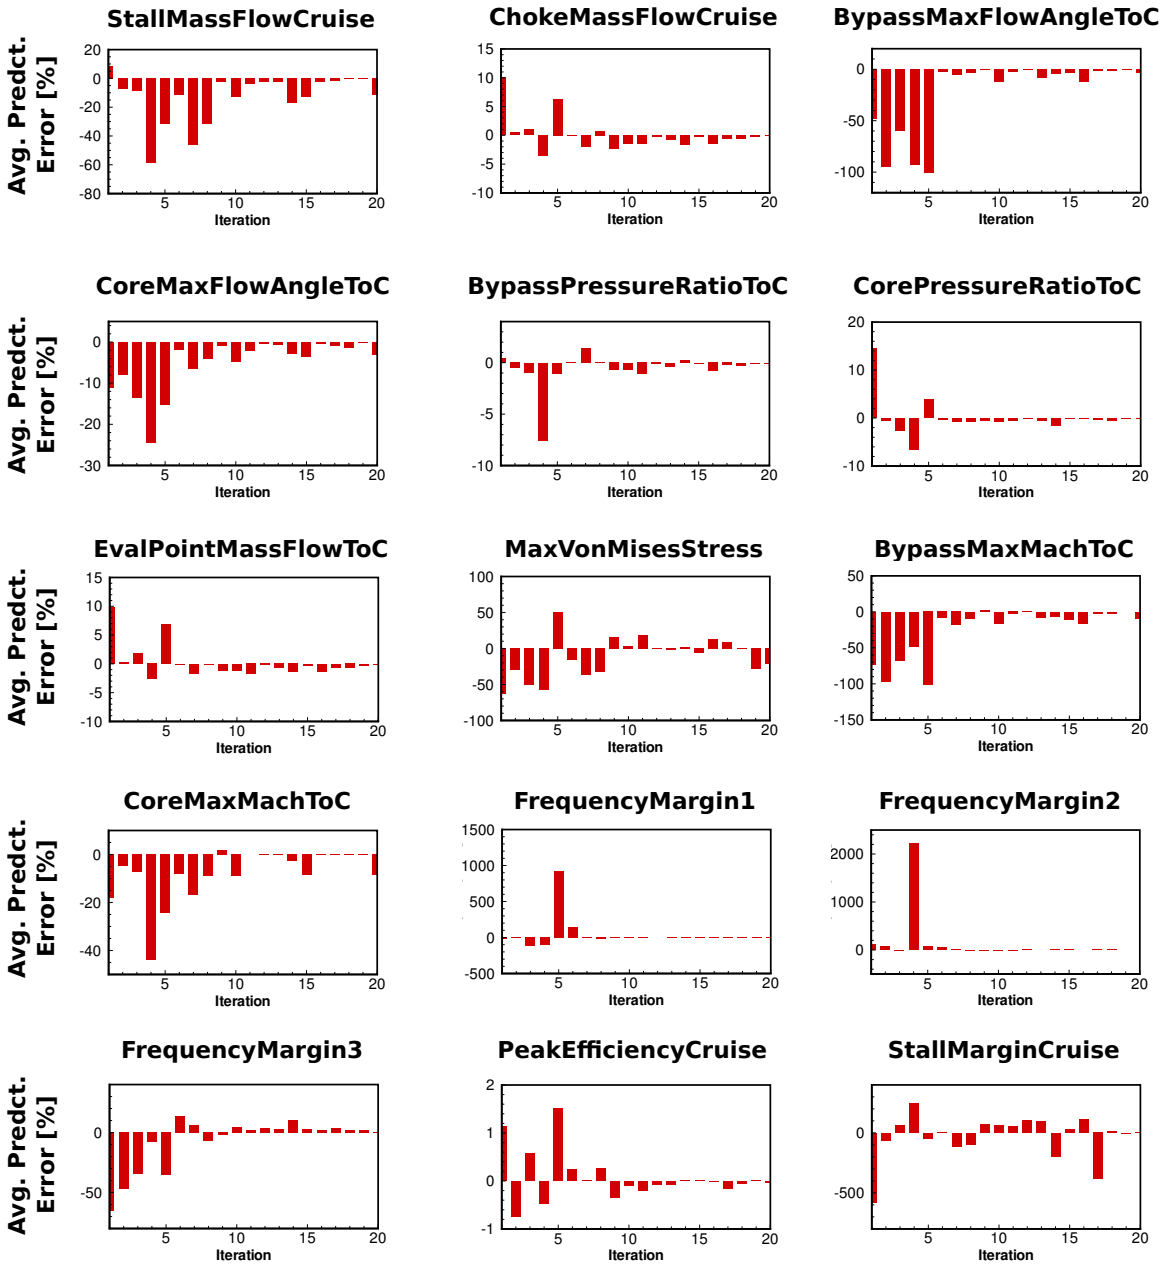


Figure 7.18: Overview of the metamodel prediction errors for all performance parameters in dependency of the optimisation progress. The error values are computed as the average prediction error of all candidate individuals that are re-evaluated by the high-fidelity evaluation per iteration. Errors are given in percent.

straint shown in Fig. 7.17 (f) is more restrictive, indicating that many designs show a choking mass flow that is lower than the design point mass flow. The plots shown on the left hand side of Fig. 7.17 present generally more restrictive constraints. As can be seen in plot (a), the stress constraint appears to be particularly challenging, while quite a number of designs are shown to satisfy the vibration constraint. The bypass pressure ratio constraint (Fig. 7.17 (c)) and the evaluation point mass flow constraints (Fig. 7.17 (e)) similarly seem to be harder to satisfy. The core pressure ratio constraint is well satisfied by many designs, as shown in Fig. 7.17 (c). Figure 7.17 also provides a good overview of the extent of constraint satisfaction and violations of the baseline design. It can be seen that the baseline does satisfy most of the aerodynamic constraints, with the exception of the bypass pressure ratio constraint shown in Fig. 7.17 (c). The baseline does violate both the stress and vibration constraint - which is not surprising considering that it was generated purely based on aerodynamic considerations.

For the constraints that are hard to satisfy, the accuracy of the metamodel predictions becomes of increasing importance. It can be seen in Fig. 7.18, that the accuracy of the metamodel predictions improves considerably in the first 5-10 iterations for most performance parameters. Two notable exceptions are the stall margin and the maximum von Mises stress constraint, for which comparably large prediction errors are observed throughout major parts of the optimisation process. The difficulties observed with the stall margin prediction are likely associated with the relatively coarse resolution of the cruise speedline in the CFD computations, which inevitably may lead to relatively large variations of predicted stall margin values.

Figure 7.19 shows the evolution of the design variables over the optimisation progress. As before, each symbol in the plot represents a design evaluated by the high-fidelity evaluation tools. Individuals with constraint violations are shown in red, while a green symbol represents a feasible design. Two circles shown at iteration 0 indicate the values of the DOE samples, while the baseline value is shown as orange diamond. Figure 7.19 gives a very good indication of the search progress of the optimiser. It can be seen that the optimiser initially explores large parts of the design space in the attempt to satisfy the constraints. The first feasible individuals are generated in the ninth iteration, thus being closely linked to the improvement of the metamodel prediction accuracy discussed above and shown in Fig. 7.18. A few distinct design changes can be identified to reach feasibility with regard to the baseline design. Those are

- A reduction of the number of blades

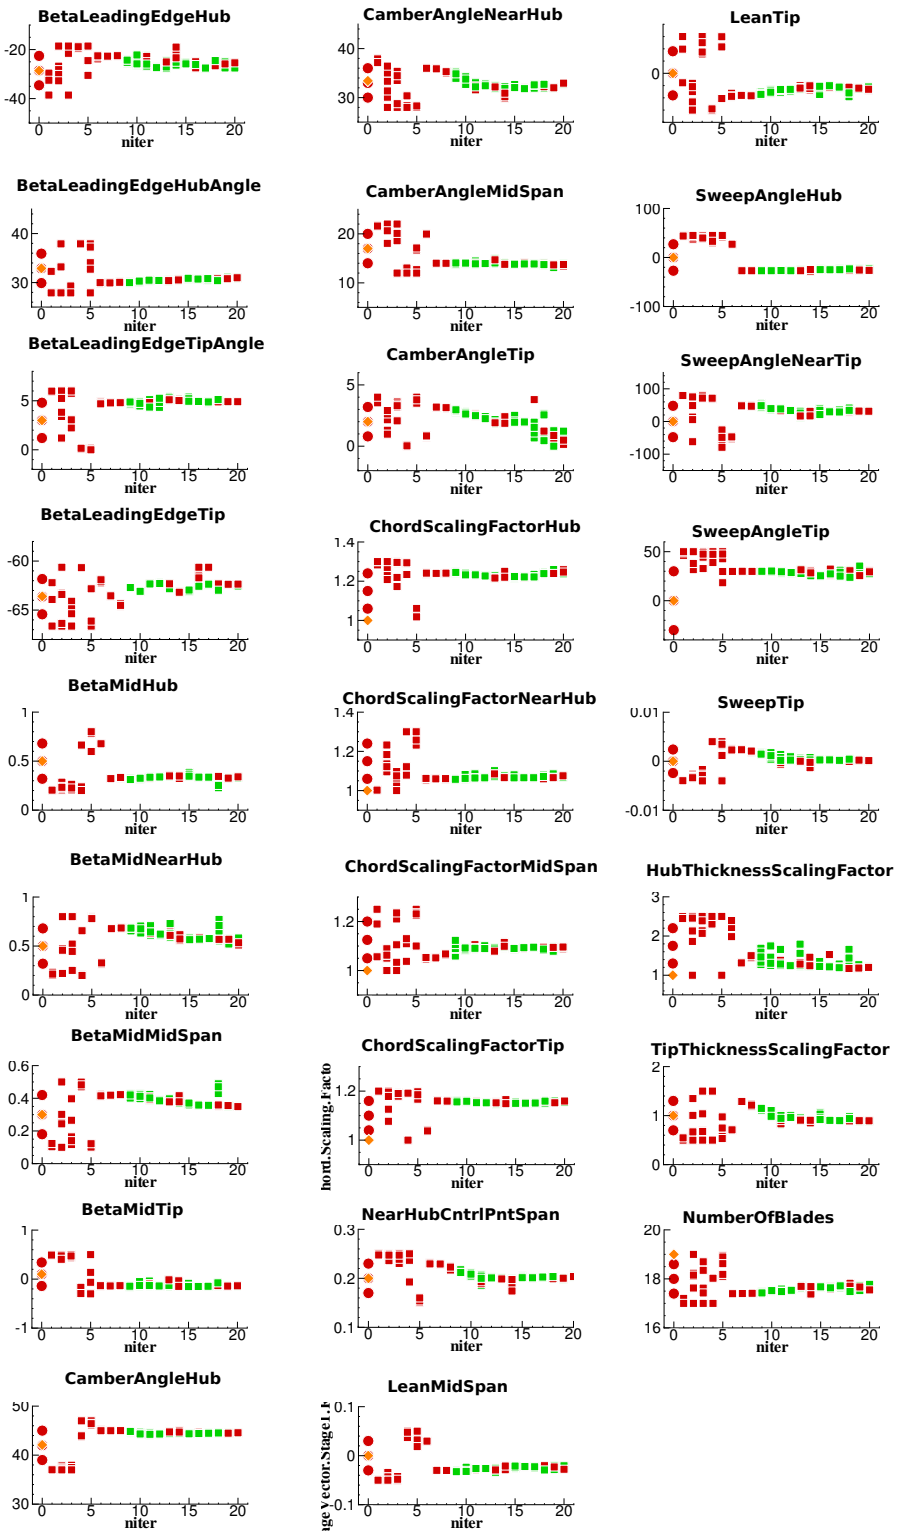
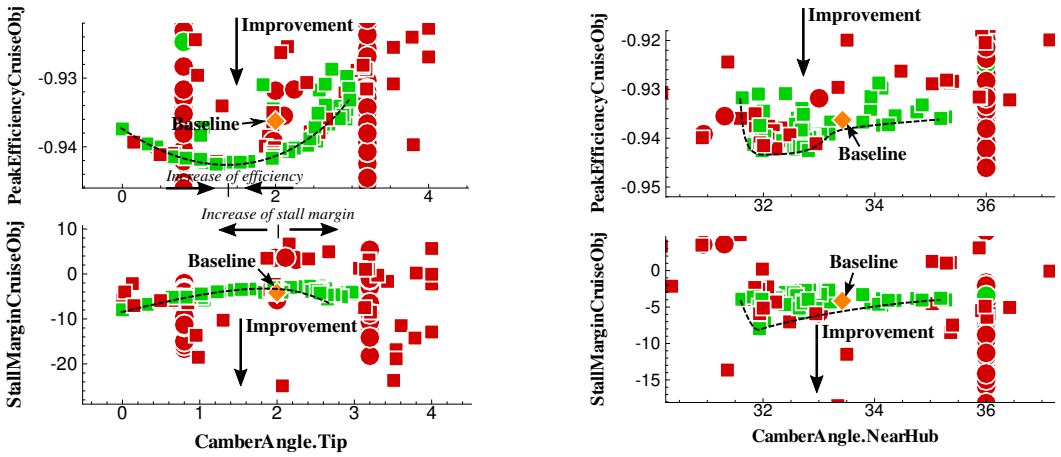


Figure 7.19: Evolution of design parameters. Green symbols denote feasible designs, while red symbols indicate infeasible ones. The baseline design is shown as orange diamond.



(a) (b)

Figure 7.20: Impact of the tip and near-hub camber angles on the cruise peak efficiency objective and the stall margin objective. Green symbols denote feasible designs, while red symbols indicate infeasible ones.

- An increased chord length from hub to tip
- An increase of the hub camber angle
- A backward swept tip

In the further optimisation progress, feasibility is overall maintained for most designs, as the optimiser attempts to improve the objectives. The feasible region of the design space can be seen to be rather restricted for many design parameters, being visible by small parameter diversities as the optimisation progresses, combined with immediate constraint violations with small parameters deviations. However, some parameters continue to evolve in the feasible design space, indicating their contribution towards the improvement of the objectives. Three parameters can be identified that appear to be particularly influential. Those are

- The camber angle at the tip
- The camber angle near the hub
- The hub thickness scaling factor

In Fig. 7.20 (a), the impact of the camber angle at the tip on the two objectives is shown in further detail. This parameter can be seen to exhibit a pronounced trade-off between peak efficiency and stall margin improvement. Maximum peak efficiency is

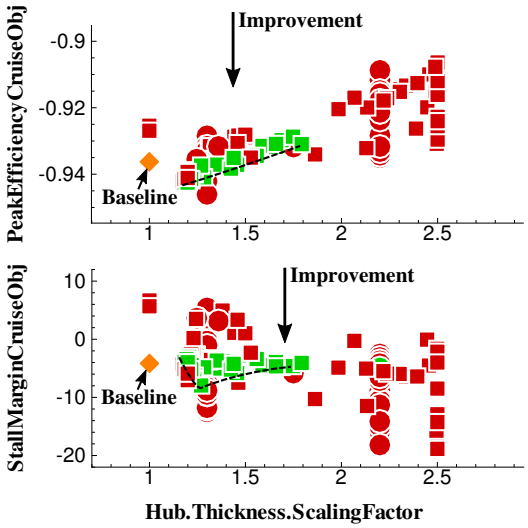


Figure 7.21: Impact of the hub thickness scaling factor on the cruise peak efficiency objective and the stall margin objective. Green symbols denote feasible designs, while red symbols indicate infeasible ones.

observed to be achieved with a tip camber angle value of 1.5 degrees, while smaller and larger values of the camber angle deteriorate the efficiency. The stall margin is almost exactly opposed to this trend, with the worst value obtained for a camber angle of about 2 degrees, and improvements seen for lower and higher camber angle values. In contrast, Fig. 7.20 (b) shows that both objectives are simultaneously maximised by the reduction of the near-hub camber angle towards a value of 32 degrees.

The impact of the hub thickness scaling factor is shown in Fig. 7.21. The peak efficiency can be seen to be overall strongly correlated to the design parameter, with hub section thickness reductions leading to improved efficiencies. The stall margin can be seen to be maximised for a scaling factor of 1.3.

7.7.2 Comparison of optimised and baseline performance

After 20 iterations, the optimiser made considerable progress towards the improvement of both objectives. Figure 7.22, provides a close-up on the feasible region in the objective space. A Pareto front can be identified, showing a trade-off between peak efficiency and stall margin improvement. One design from the Pareto front was chosen for further examination and comparison to the baseline design. The chosen design, generated in the 20th iteration, is labelled as IT020_IND002 and is shown as red triangle in Fig. 7.22. The design was chosen with a higher weight on the efficiency

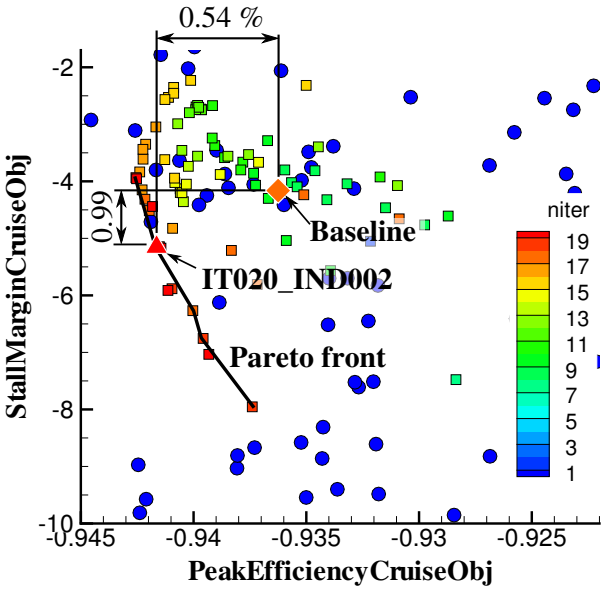


Figure 7.22: Detail view on the feasible region in the objective space and the chosen individual from the Pareto front (IT020_IND002).

improvement, at the cost of a less pronounced stall margin improvement. In comparison to the (infeasible) baseline design, the optimised design shows an efficiency improvement of 0.54% and a slight stall margin improvement while satisfying all imposed constraints. The aerodynamic and structural performances of the optimised and baseline designs will be compared in further detail in the following subsections.

7.7.2.1 Aerodynamic performance

Figure 7.23 depicts the meridional contours of both designs. By comparison of the leading and trailing edge curves, a general design change in the optimised design becomes immediately apparent, which is the overall increase of the axial chord length. The leading edge can be seen to exhibit a slight forward sweep with a backward swept tip section. A comparison of the pitch averaged absolute total pressure in the bypass inlet at the top-of-climb evaluation point is shown in the top right corner plot for both designs. In comparison to the baseline distribution, the optimised design can be seen to have raised the absolute total pressure over the entire span, thus leading to the satisfaction of the bypass pressure ratio constraint (c.f. Fig. 7.17 (c)).

Figure 7.24 shows the isentropic Mach number distributions and blade profile shapes for both designs at 10, 50 and 90% span. The blade loading can be seen to have increased in the optimised design for all blade sections shown, as indicated by

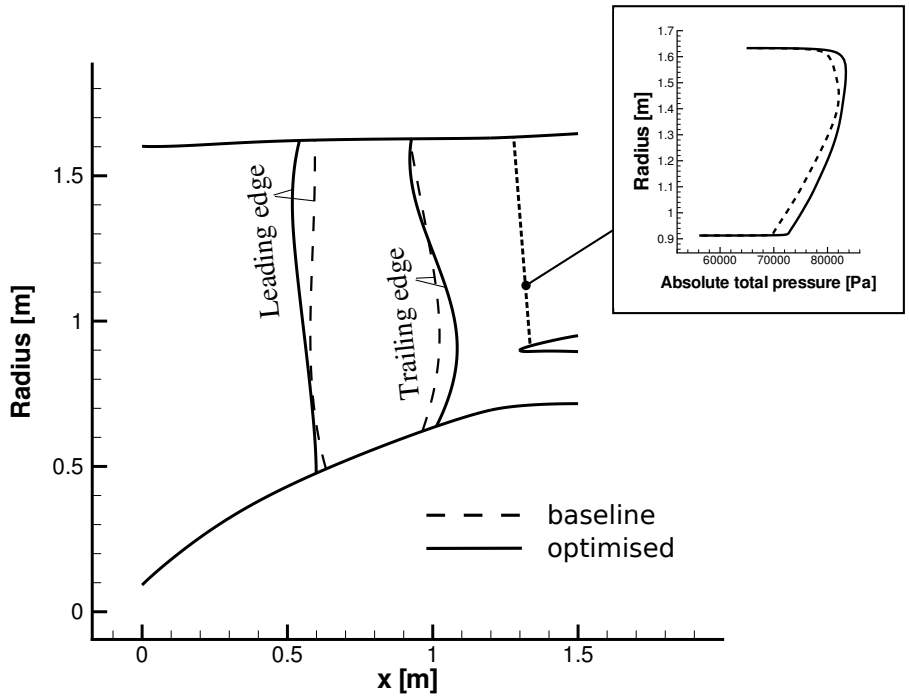


Figure 7.23: Comparison of the meridional geometries of baseline and optimised designs and pitch averaged absolute total pressure distributions in the bypass inlet.

an increase of the enclosed area of the isentropic Mach number distributions. The loading increase can be seen to be primarily attributed to the overall extension of the chord length. While the blade count of the optimised design has dropped to 18 (from a blade count of 19 for the baseline blade), a sufficient net loading increase was achieved in order to meet the required bypass pressure rise.

The leading edge blade angle has been reduced consistently for all blade sections, leading to an overall increased incidence with a consequent removal of negative loading in the 50% and 90% span sections. The camber angle of the tip section profile was reduced from two degrees in the baseline design to 0.8 degrees in the optimised design. The optimised tip profile shape has a very slight amount of pre-compression.

It was shown in Fig. 7.22 that the optimised design has an 0.54% improved cruise peak efficiency with respect to the baseline. To evaluate where this efficiency improvement originates from, Fig. 7.25 shows an overview of the entropy generation at a plane downstream of the blade as well as at 2, 50 and 90% span. A comparison of the downstream contour shows that the baseline design exhibits increased losses in the tip region. The blade to blade plot at 90% span reveals that those losses originate

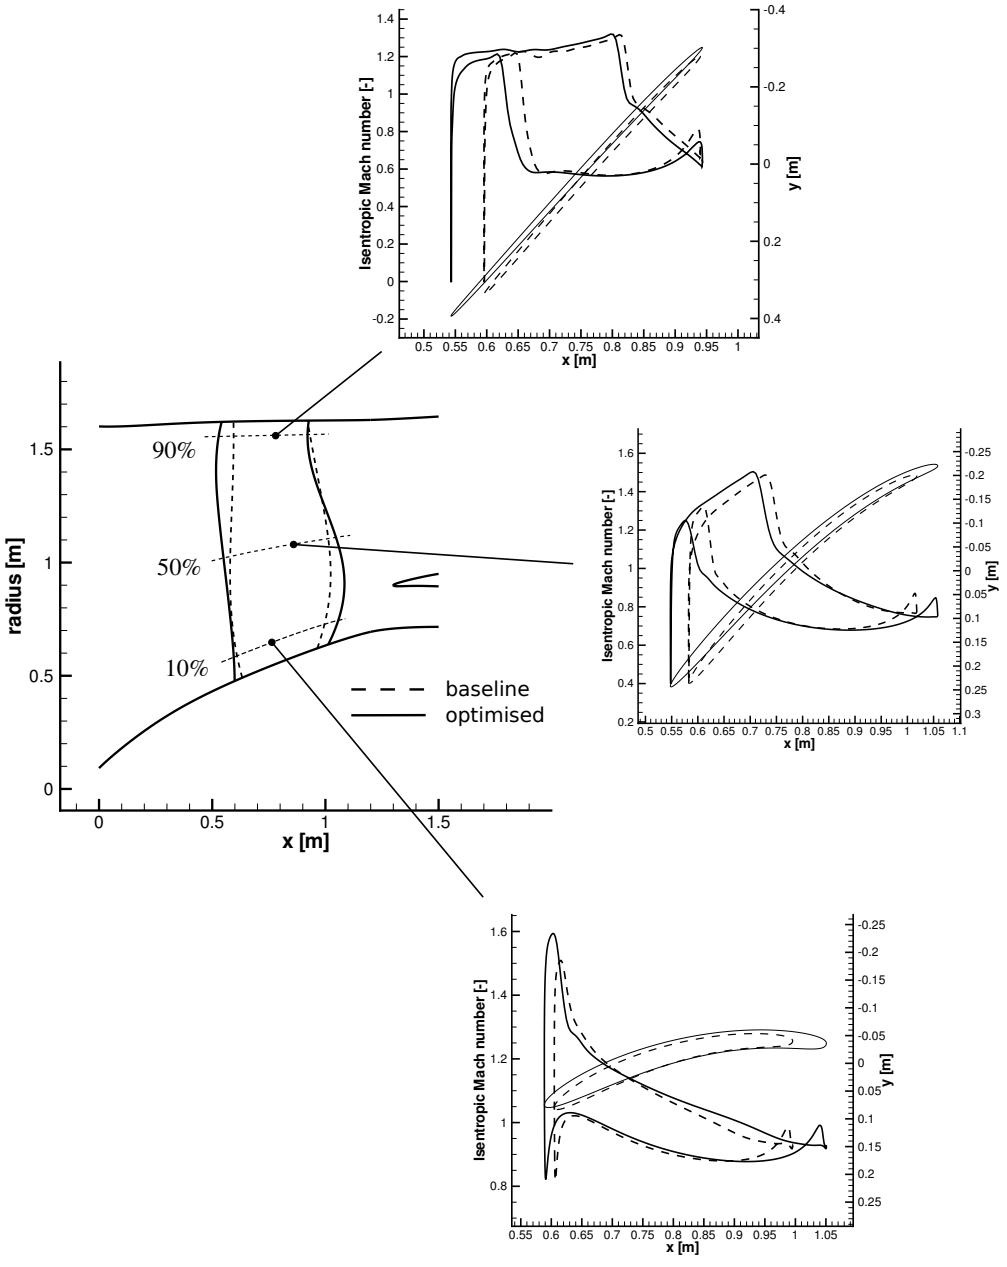


Figure 7.24: Comparison of the isentropic Mach number distributions at 10, 50 and 90% span of baseline and optimised design.

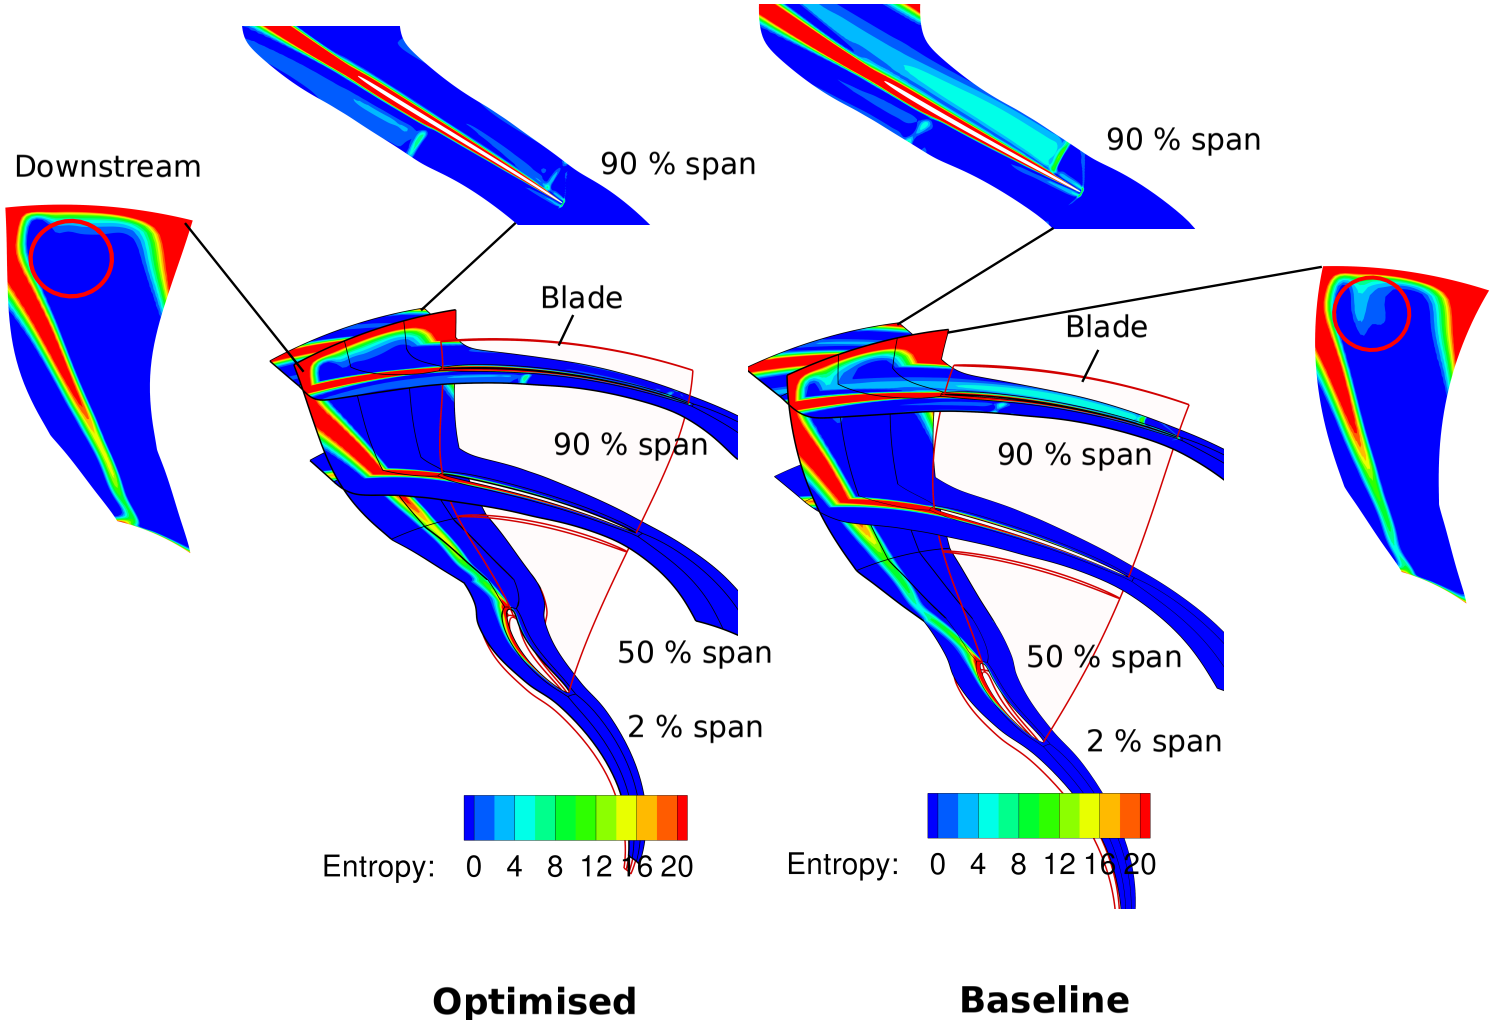


Figure 7.25: Comparison of entropy contours for baseline and optimised design at their respective cruise peak efficiency point.

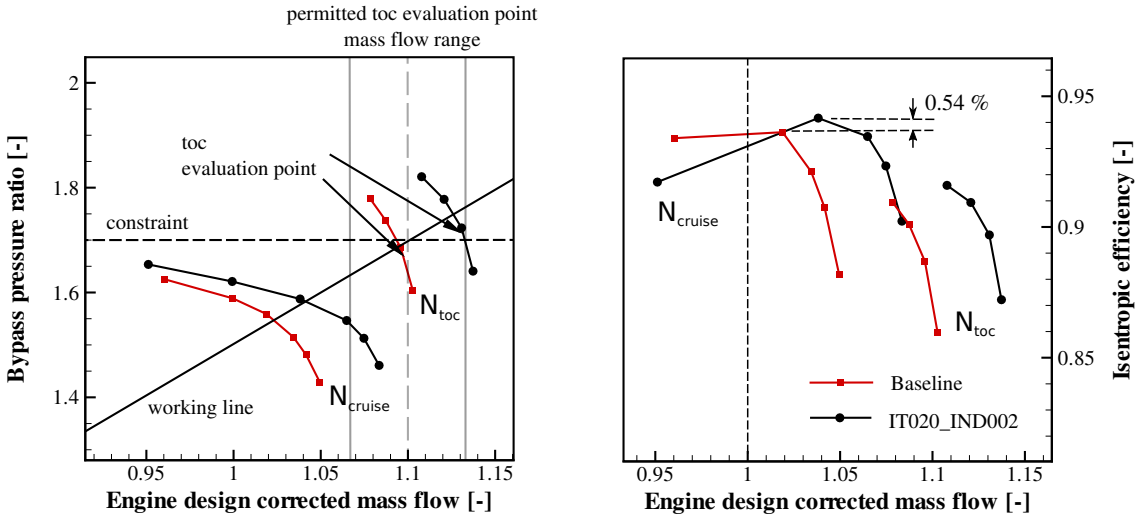


Figure 7.26: Performance map comparison of baseline and optimised designs.

from a strong passage shock impinging on the pressure side surface of the blade. In the optimised design, the shock strength has been reduced with the result of lower losses.

The fan performance maps of both designs are shown in Fig. 7.26, where the total-to-total bypass pressure ratio and the isentropic fan efficiency are plotted over the total corrected engine mass flow. The evaluation point on the top-of-climb speedlines is specifically indicated. As it was previously seen, the corresponding pressure rise of the baseline fan is below the required bypass pressure ratio constraint indicated by the dashed horizontal line in the plot on the left hand side. An overall elevation of the bypass pressure ratio can be observed for both speedlines in the optimised design, leading to the satisfaction of the constraint value at top-of-climb. However, primarily due to the reduced blade count, an overall shift of the performance curves towards increased mass flows can be observed. While the top-of-climb evaluation point lies within the permitted 6% mass flow range of the top-of-climb design point mass flow, the shifted operating range can be seen to result in an unfavourable crossing with the fan working line, leading to a reduced stall margin. A reduced permitted top-of-climb mass flow range would be needed to circumvent this issue.

The efficiency curves shown on the right hand side of Fig. 7.26 show the increase of the cruise peak efficiency, as well as the slight increase of the cruise stall margin of the optimised design. However, similarly to the observation made for the top-of-climb evaluation point, due to a missing constraint on the peak efficiency mass flow, a significant mass flow offset between the cruise design point and the peak

efficiency mass flow of the optimised design is observed. As a result, the efficiency of the optimised design is actually lower than the efficiency of the baseline design at the cruise design point mass flow. In order to circumvent this issue, additional constraints on the peak efficiency mass flow would be needed.

The shortcomings of the mass flow constraints that have been discussed here will be addressed in the optimisation presented in the next chapter.

7.7.2.2 Structural performance

A comparison of the von Mises stress distributions of the baseline and optimised designs occurring at take-off conditions is shown in Fig. 7.27. Two stress concentrations exceeding the constraint value set for the titanium alloy can be identified for the baseline design on the suction side surface, which are located at the blade leading edge at the connection of the aerofoil section to the shank and close to the trailing edge at about 80 percent chord. As shown on the right hand side of Fig. 7.27, the stresses in these regions were successfully reduced to satisfy the allowed levels in the optimised design. It should be recalled here that the high stress levels located on the dovetail flanks are disregarded for the computation of the maximum von Mises stresses, as the contact problem between dovetail flanks and disk is not resolved in the finite element computations. The stresses on the pressure side surfaces were found to be uncritical.

The vibration characteristics of both designs are shown in the Campbell diagram in Fig. 7.28. The natural frequencies of the first four eigenmodes (1F, 2F, 1T and 3F) are plotted for both designs. Note, that although the first five engine order lines are shown, only the first two have been used in the optimisation to compute the frequency margins. The rotational speeds corresponding to cruise, top-of-climb (ToC) and take-off (T/O) operation, at which the frequency margins were computed, are indicated. The Campbell diagram shows that a dangerously small frequency margin exists between the eigenfrequency of the first bending mode (1F) of the baseline design and the first engine order at take-off rotational speed, resulting in a violation of the frequency constraint. To raise this margin above the constraint value, the optimiser increased the stiffness of the blade resulting in globally higher eigenfrequencies for the optimised geometry. The second engine order is seen to have a sufficient margin to the closest blade modes 1F and 2F of both designs and is thus seen to be non-critical at all three key operating conditions. Figure 7.29 provides an overview of the first four vibration blade modes at take-off operation.

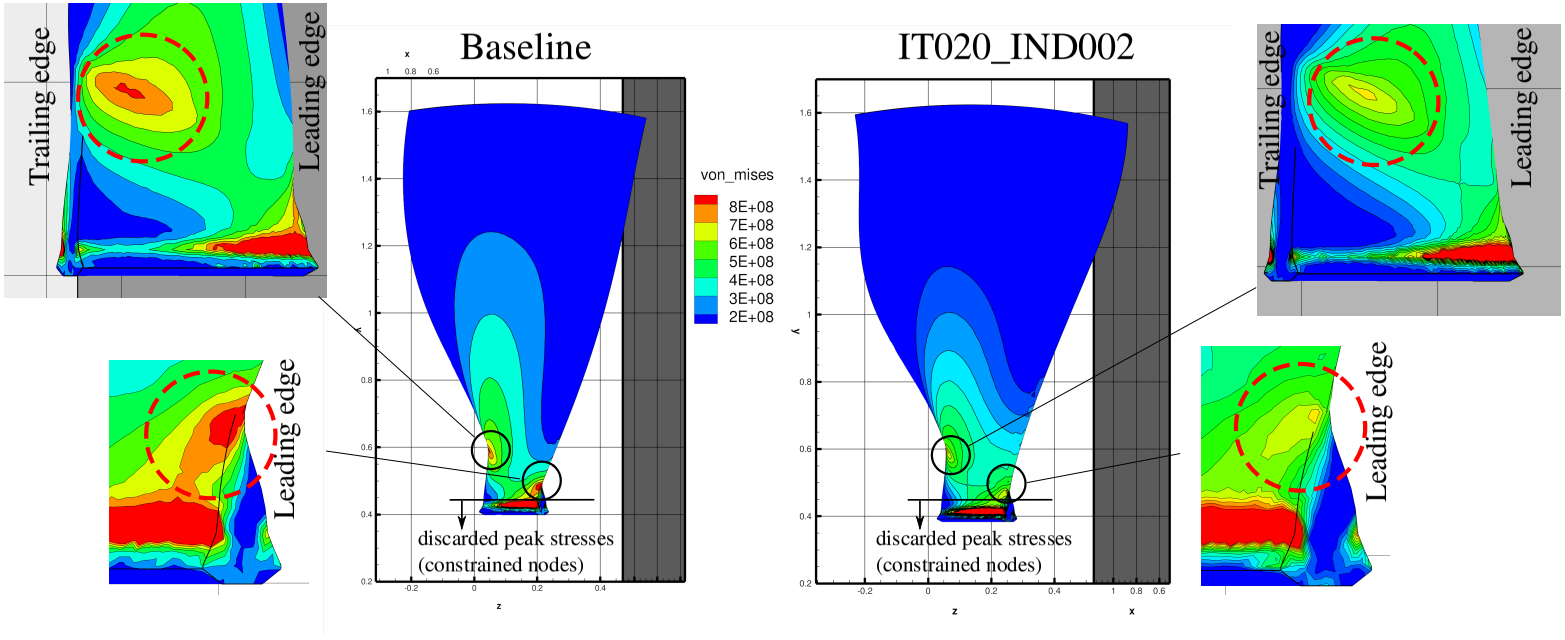


Figure 7.27: Comparison of the von Mises stress distributions in the baseline and optimised design.

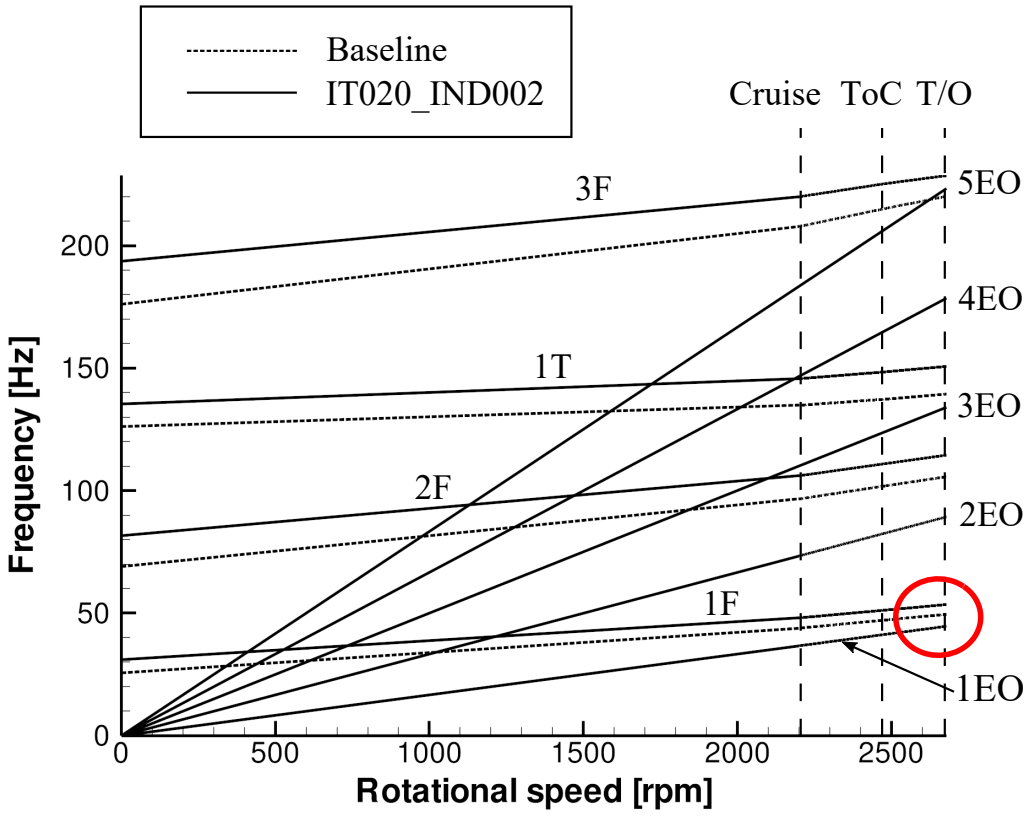


Figure 7.28: Campbell diagram of baseline and optimised designs.

7.7.2.3 Further geometry and performance data

A global comparison of the blade geometries and full fan assemblies of the baseline and optimised design is shown in Fig. 7.30 and 7.31. The stall margins shown correspond to the margins computed with the CFD setup used in the present optimisation. The actual stall margin is likely higher than that. Figure 7.32 shows the relative Mach number contours at 10, 50 and 90 percent span. Fig. 7.33 shows the pitchwise averaged spanwise distribution of the absolute total pressure at the fan exit, upstream of the bypass splitter.

7.8 Summary

The application of a multidisciplinary and multiobjective optimisation system to the design of a full titanium aero-engine fan blade for a high-bypass ratio turbofan engine was presented. Particular novelties introduced in this chapter include the consideration of the bypass splitter in the fluid domain as well as the integration of a shape

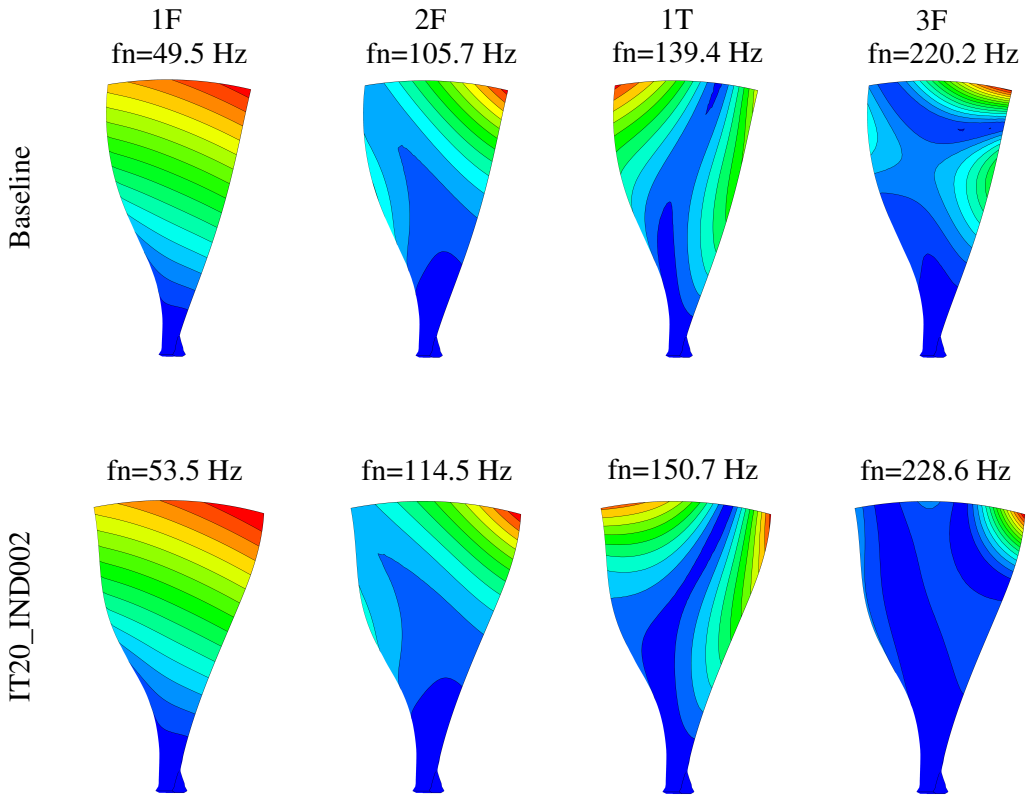


Figure 7.29: Comparison of blade mode shapes and eigenfrequencies at take-off.

adaptive dovetail root geometry in the solid domain. The bypass splitter configuration allows the core and bypass flow requirements from the engine cycle design to be directly included in the optimisation process. Aerodynamic performance was evaluated at cruise and top-of-climb conditions, while static blade stresses were evaluated at take-off operation. Vibration was furthermore assessed over the entire operating range. The baseline design in this work was generated solely based on aerodynamic considerations. However, it was not attempted to satisfy all aerodynamic requirements in the manual aerodynamic design step. This task was purposefully left to the optimisation system, which was fulfilled flawlessly. The optimiser successfully identified the feasible region in the design space after 9 iterations and subsequently improved both objectives until the optimisation was stopped after 20 iterations. The chosen design from the Pareto front has a peak efficiency improvement of 0.54% and a slight stall margin improvement with respect to the baseline design while it satisfied all imposed aerodynamic and structural constraints.

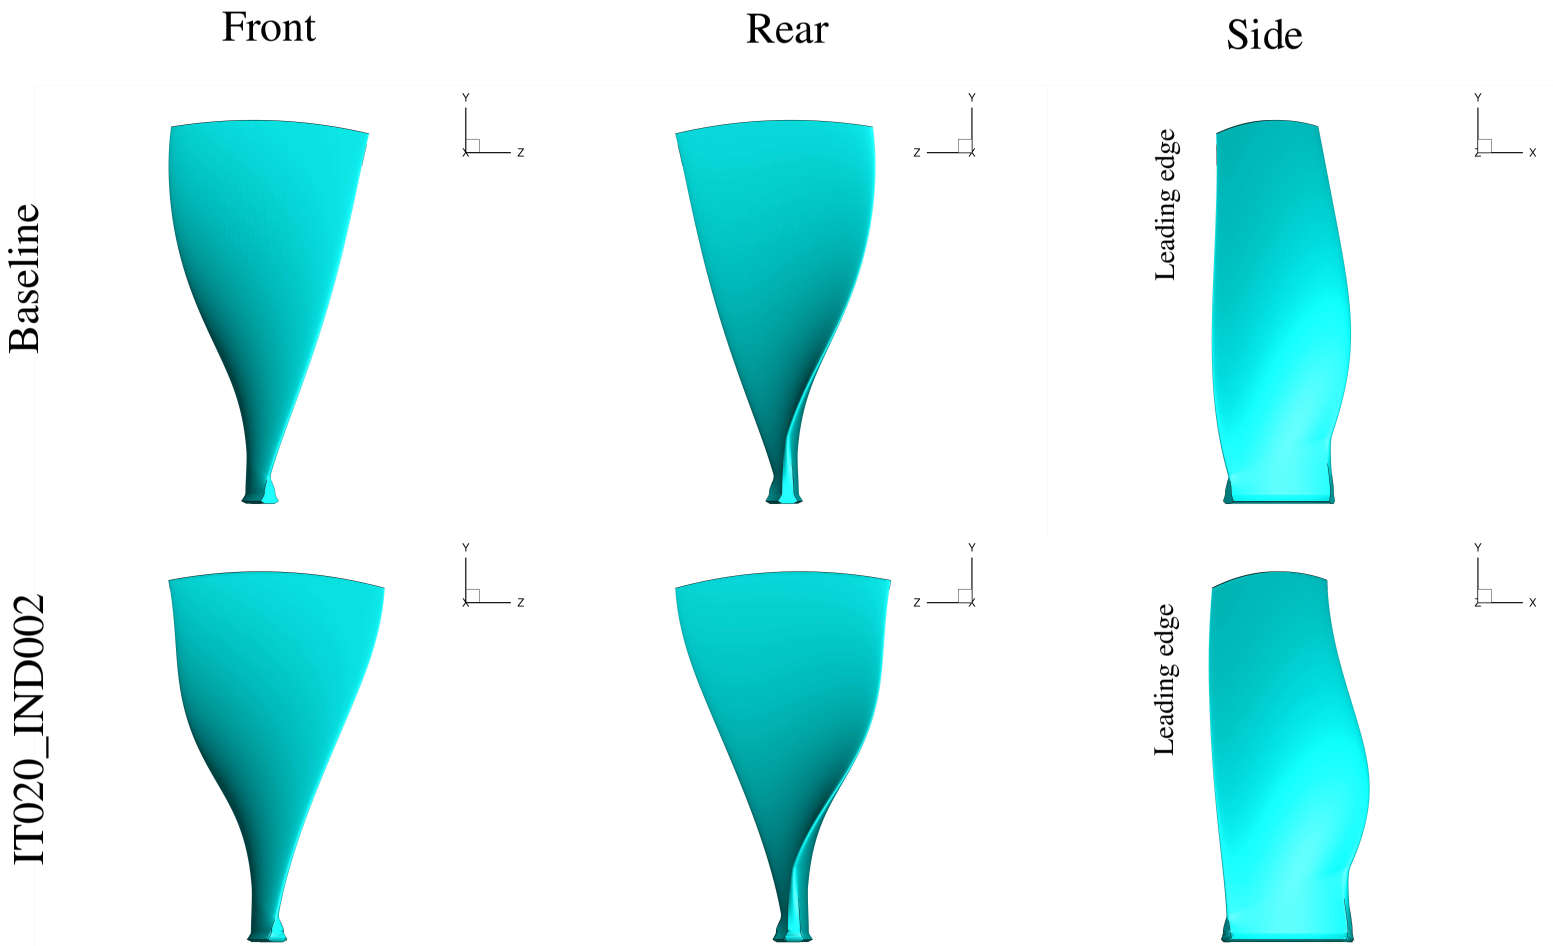


Figure 7.30: Blade geometries of baseline and optimised design.

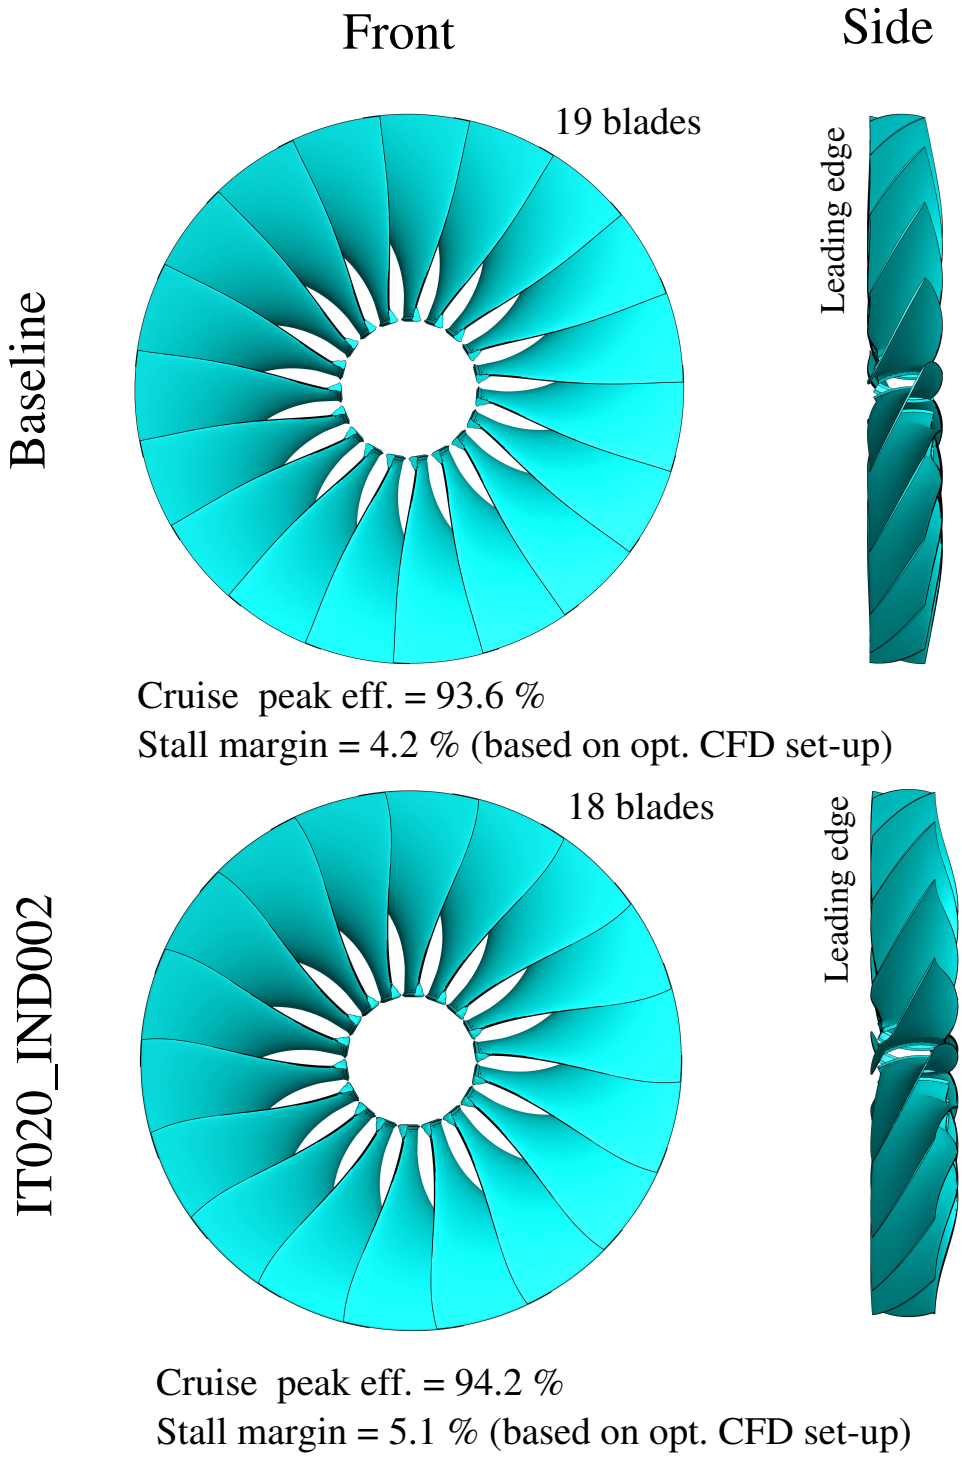


Figure 7.31: Full fan assemblies of baseline and optimised design.

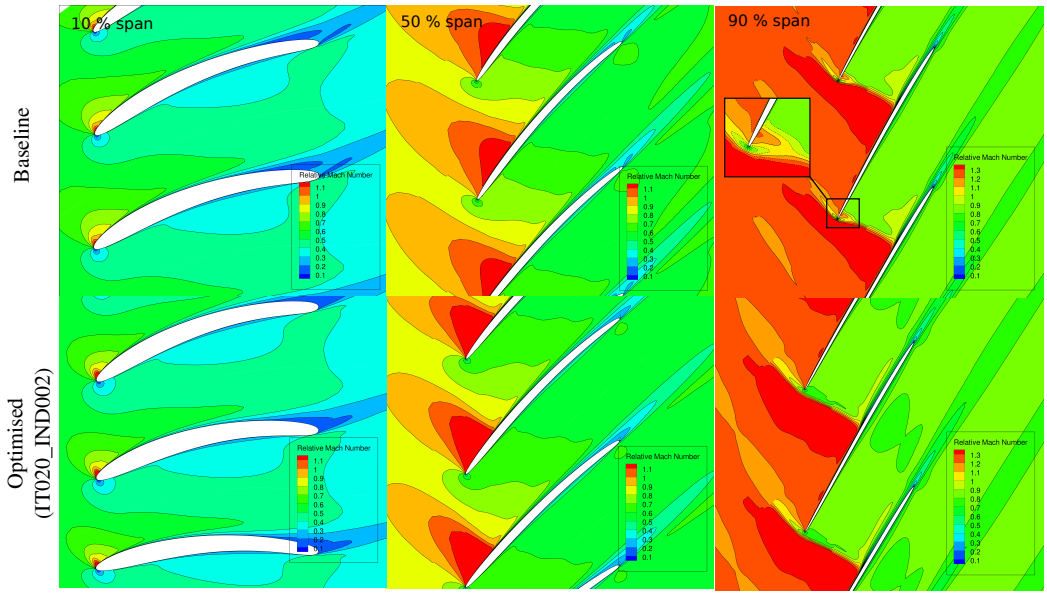


Figure 7.32: Relative Mach number contours at 10, 50 and 90 percent span for baseline and optimised design at the cruise design point.

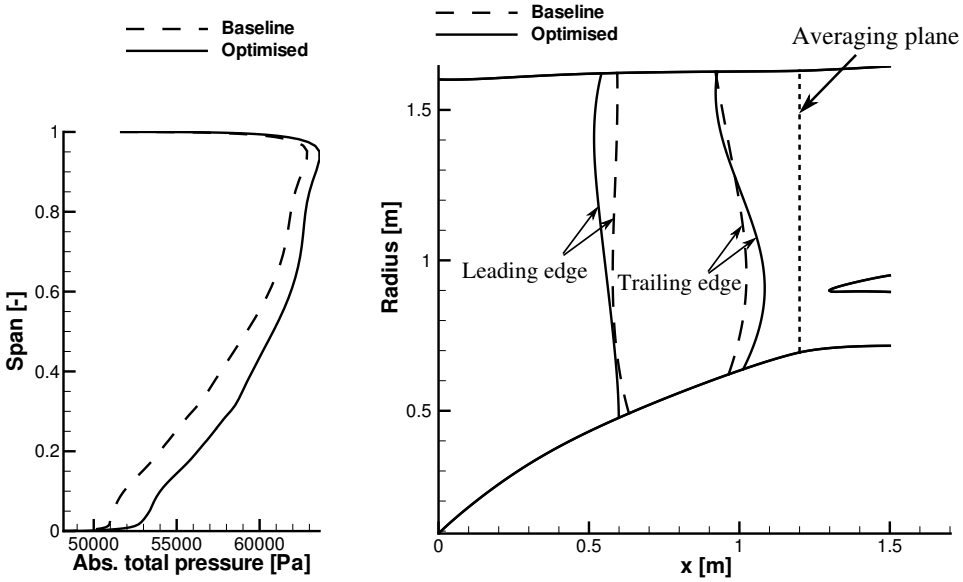


Figure 7.33: Pitch averaged spanwise distribution of absolute total pressure at the fan exit, upstream of the bypass splitter, at the cruise design point.

8 Multidisciplinary Design Optimisation of a Composite Fan Blade

In this chapter a second aeromechanical fan blade optimisation is presented. A fan blade made from a carbon-fibre reinforced plastic composite material is now considered in addition to a modified blade parametrisation and a refined optimisation problem formulation according to the observations made in the previous chapter.

8.1 Fan blade parametrisation

8.1.1 Meridional contour

The parametrised meridional contour of the baseline design defining the hub, casing and bypass splitter end-walls was presented in Section 7.1.1. The meridional contour is not altered throughout the optimisation process presented in this chapter.

8.1.2 Camber surface

The parametric curves defining the spanwise blade angle and chord length distributions for the definition of the blade's camber surface are shown in Fig. 8.1. Arrows in the plots indicate how the control points are allowed to change throughout the optimisation process. Black arrows indicate direct degrees of freedom, i.e. parameters that can be directly modified by the optimiser, while red arrows indicate indirect degrees of freedom, i.e. parameters that are defined in dependency to other optimisation parameters.

The spanwise curves defining the blade angle distributions are defined by four control points at 0, 20, 50 and 100% span. In order to allow additional control of the camber surface close to the bypass splitter, the spanwise coordinate of the control

point initially defined at 20% span is allowed to change. The value of the optimisation parameter simultaneously affects all distributions in order to retain consistency of the control points.

All four control point values defining the leading edge blade angle as well as those defining the axial position of the mid-point blade angle, $u_{\beta_{mid}}$ are allowed to be directly modified by the optimiser.

Two pseudo camber angles are defined which determine the control point values of the mid and trailing edge blade angle distributions. The front pseudo camber angle is defined as

$$\lambda'_{front} = \beta_{mid} - \beta_{LE} \tag{8.1}$$

where β_{mid} denotes the blade angle coordinate value of the control point defining the mid point blade angle, and β_{LE} denotes the respective value of the control point defining the leading edge blade angle. Similarly, a pseudo camber angle at the rear of the camber line is defined as

$$\lambda'_{rear} = \beta_{TE} - \beta_{mid} \tag{8.2}$$

where β_{TE} is the blade angle coordinate value of the control point defining the trailing edge blade angle.

The front and rear pseudo camber angles at the four dedicated spanwise positions are defined as optimisation parameters. The camber angles have a direct physical

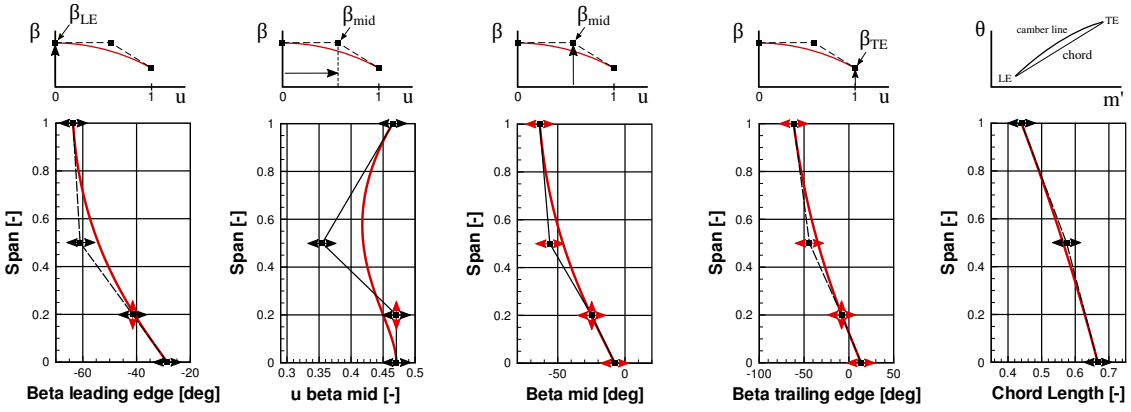


Figure 8.1: Camber surface parametrisation. Black arrows indicate direct degrees of freedom of control points (optimisation parameters), while red arrows indicate indirect degrees of freedom (quantities that are defined in dependency to other optimisation parameters).

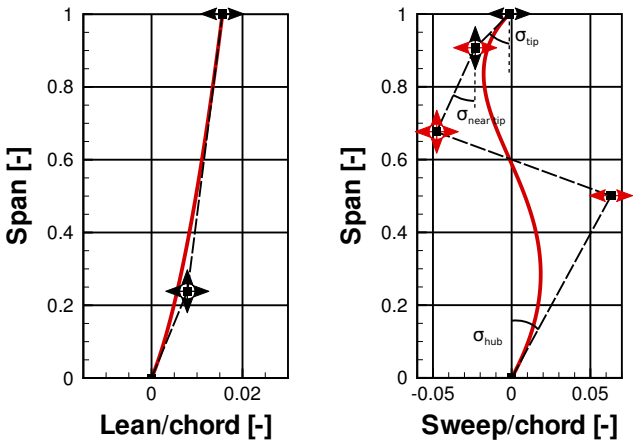


Figure 8.2: Parametrisation of lean and sweep. Note that the baseline values of all control points defining lean and sweep are zero.

implication on the blade loading distribution, which makes it easier for the optimiser to identify the impact of optimisation parameter changes on the blade performance.

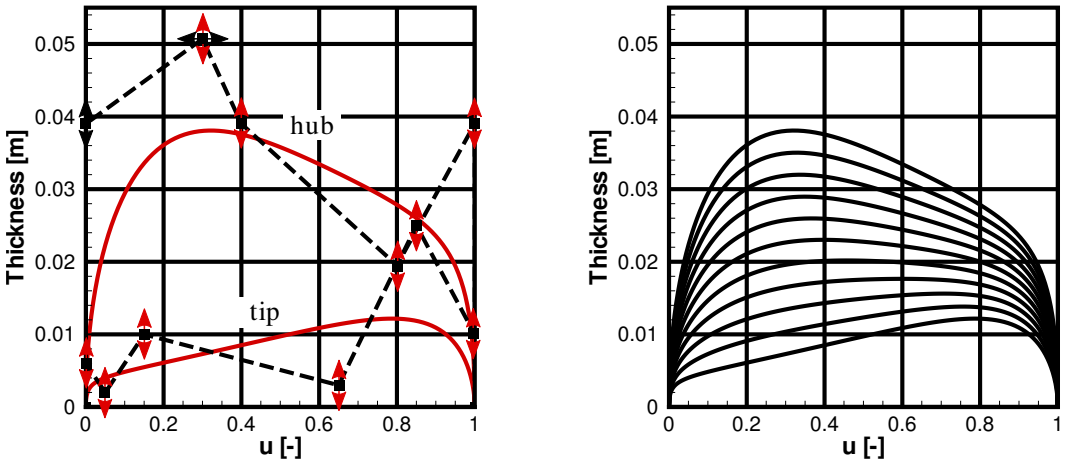
The spanwise chord length distribution is defined by three control points at 0, 50 and 100% span. Each control point value can be changed individually by the optimiser by the application of a chord length scaling factor, which is defined for each of the respective blade sections (i.e. hub, mid-span and tip).

8.1.3 Lean and sweep

The parametrisation of the spanwise lean-to-chord and sweep-to-chord ratio distributions is shown in Fig. 8.2. The distribution of the lean-to-chord ratio is defined with a second order Bézier curve with control points being initially specified at 0, 50 and 100% span. The control points at 50 and 100% span are allowed to be directly modified by the optimiser, whereas the mid-height point is also allowed to move in the spanwise direction (note that its spanwise position in Fig. 8.2 has been altered from its initial position to highlight this point).

A fourth order Bézier curve is used for the definition of the sweep-to-chord distribution. Control points are initially defined at 0, 50, 70, 80 and 100% span, enabling an increased control of sweep in the higher span region, where the relative inflow is supersonic.

The impact of blade sweep on the flow field mainly stems from the spanwise variation of sweep and not from the absolute value of sweep at any one blade section. Three sweep angles are therefore defined as optimisation parameters, indicated as



(a) **(b)**

Figure 8.3: Parametrisation of the blade thickness distributions at hub and tip (a) and spanwise interpolated distributions (b).

σ_{hub}, $\sigma_{neartip}$ and σ_{tip} in Fig. 8.2. Solely the control point defining the tip value of sweep is allowed to be varied independently. The spanwise position of the two near tip control points is equally allowed to change.

8.1.4 Thickness

Figure 8.3 (a) shows the parametrisation of the blade thickness distributions at hub and tip. The thickness distributions are allowed to be uniformly scaled with an independent thickness scaling factor for the hub and tip section, respectively. An example of the spanwise interpolation of the parametrised thickness distributions as resulting from the baseline distributions shown here is presented in Fig. 8.3 (b).

In addition to the uniform scaling, two optimisation parameters are defined to allow the leading edge hub thickness profile to change, as indicated with the black arrows in Fig. 8.3 (a), enabling improved control of the front hub profile.

8.1.5 Blade root parametrisation

The parametrised cross section of the dovetail root is shown in Fig. 8.4. The dovetail width is variable and adapts to the blade hub thickness as described in detail in Chapter 4. Apart from that, the root cross sectional shape is kept unchanged throughout

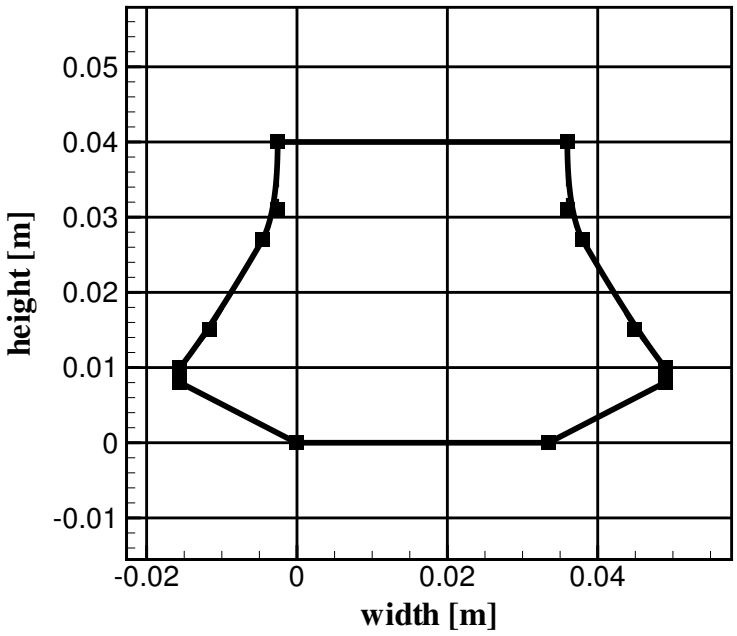


Figure 8.4: Parametrisation of the dovetail root cross section. The width of the shape shown is adapted to the thickness of the baseline blade.

the optimisation presented in this chapter. The example in Fig. 8.4 shows the cross section adapted to the baseline blade.

In addition, the root leading edge offset controlling the amount of front root tilting, as explained in Section 4.3, is added as an optimisation parameter.

8.1.6 Number of blades

The blade count is added as a last optimisation parameter and is allowed to change between 17 and 19 (baseline).

8.1.7 Summary of optimisation parameters

The previously described fan parametrisation amounts to a total of 35 optimisation parameters. Table 8.1 gives a global overview of all defined optimisation parameters and their range.

Table 8.1: Summary of optimisation parameters and their range.

#	Name	Unit	Lower bound	Upper bound
1	Beta.LeadingEdge.Hub	deg	-38.6	-18.6
2	Beta.LeadingEdge.NearHub	deg	-51.5	-31.5
3	Beta.LeadingEdge.MidSpan	deg	-66	-56
4	Beta.LeadingEdge.Tip	deg	-66.6	-60.6
5	U.BetaMid.Hub	[-]	0.2	0.6
6	U.BetaMid.NearHub	[-]	0.2	0.6
7	U.BetaMid.MidSpan	[-]	0.2	0.5
8	U.BetaMid.Tip	[-]	0.2	0.5
9	CamberAngle.Hub.Front	[deg]	18	30
10	CamberAngle.NearHub.Front	[deg]	14	25
11	CamberAngle.MidSpan.Front	[deg]	4.3	13
12	CamberAngle.Tip.Front	[deg]	0	10
13	CamberAngle.Hub.Rear	[deg]	10	24
14	CamberAngle.NearHub.Rear	[deg]	10	19.3
15	CamberAngle.MidSpan.Rear	[deg]	4	10
16	CamberAngle.Tip.Rear	[deg]	0	-10
17	ChordScalingFactor.Hub	[-]	1.0	1.3
18	ChordScalingFactor.MidSpan	[-]	1.0	1.25
19	ChordScalingFactor.Tip	[-]	1.0	1.2
20	NearHub.CntrlPnt.Span	[-]	0.15	0.25
21	Lean.Intermediate	[-]	-0.07	0.07
22	Lean.Intermediate.Span	[-]	0.2	0.9
23	Lean.Tip	[-]	-0.05	0.05
24	SweepAngle.Tip	[deg]	-45.0	45.0
25	SweepAngle.NearTip	[deg]	-80.0	80.0
26	SweepAngle.Hub	[deg]	-40.0	20.0
27	Sweep.Tip	[-]	-0.004	0.004
28	Sweep.NearTip.Span	[-]	0.8	0.95
29	Sweep.Span.Distance.NearTipPnts	[-]	0.1	0.25
30	Hub.ThicknessScalingFactor	[-]	1.0	2.5
31	Tip.ThicknessScalingFactor	[-]	0.5	1.5
32	U.Hub.Thickness.Pnt.3	[-]	0.1	0.3
33	Hub.Thickness.Dist.Pnt2And3	[m]	-0.01	0.0117
34	Dovetail.Front.LeadingEdge.Pnt.Offset	[m]	0.0	0.03
35	NumberOfBlades	[-]	17	19

8.2 Aerodynamic performance evaluation

The steady aerodynamic performance of the fan at multiple operating points is evaluated using a commercially available three-dimensional RANS CFD solver. Details of the modelling and solution process were given in Chapter 5. An explanation of the case specific setup will be provided below.

8.2.1 Fluid mesh generation

The template mesh used for the optimisation in this chapter is identical to the one used in Chapter 7, see Fig. 7.9.

8.2.2 Operating points and boundary conditions

In total ten operating points are computed for each design, including four points on the ToC speedline and six points on the cruise speedline. Total pressure, total temperature, the absolute inlet flow angle and the turbulent kinematic viscosity are imposed as boundary conditions at the inlet. The mass flow is imposed as boundary condition at the core outlet, while the static pressure using the radial equilibrium law is defined as boundary condition at the bypass outlet. Performance curves are computed via adaptation of the bypass outlet static pressure. Periodic boundary conditions are imposed on the periodic surfaces of the domain. Walls are treated as adiabatic.

An automatic convergence check is performed after each CFD computation to assess the mass flow error between domain inlet and outlet as well as the iteration errors of efficiency and pressure ratio. Only cases with sufficient convergence are subsequently assessed by an automatic postprocessing step, which extracts the required performance parameters needed by the optimiser. Non-converged cases are considered as failed and are automatically excluded from the optimisation process.

8.3 Structural performance evaluation

The structural static and free vibration response of the fan blade is evaluated using a three-dimensional finite element solver. Details of the solution procedure were provided in Chapter 6. The case specific setup will be described below.

8.3.1 Solid mesh generation

The solid mesh generation process is similar to the one described in Chapter 7.

Table 8.2: Elastic properties and mass density of the chosen woven fabric composite material.

Young's moduli			Shear moduli			Poisson's ratios			Mass density
E_1	E_2	E_3	G_{12}	G_{23}	G_{13}	ν_{12}	ν_{13}	ν_{23}	ρ
GPa	GPa	GPa	GPa	GPa	GPa	-	-	-	kg/m^3
67.5	67.5	13.1	6.8	5.1	5.1	0.04	0.44	0.33	1569.5

Table 8.3: Engineering strengths of the chosen woven fabric composite material.

X_t	X_c	Y_t	Y_c	Z_t	Z_c	Q	R	S
MPa	MPa	MPa	MPa	MPa	MPa	MPa	MPa	MPa
897	862	897	862	59.8	812.8	74.8	74.8	71.7

8.3.2 Material properties

Material properties of a laminate made from two-dimensional woven fabric carbon-epoxy composite plies are applied, which are obtained from measurements reported in [32]. The nine elastic engineering constants along with the overall mass density are listed in Table 8.2. The density was determined based on the rule of mixtures (e.g. [80]), with an overall fibre volume fraction of the weave, V_f, of 58% [64]. Table 8.3 lists the measured engineering strengths. The symbols X_t, X_c, Y_t, Y_c, Z_t, Z_c denote tension and compression uniaxial strengths in the material directions 1, 2 and 3 respectively. Q, R and S denote the shear strengths in the 2-3, 3-1 and 1-2 planes. The element coordinate systems required for the transformation of the orthotropic material properties are generated using the method explained in Section 6.7.

8.3.3 Boundary conditions

Displacement boundary conditions are defined for all element nodes being located on the dovetail flank surfaces, as shown in Fig. 8.5. Nodal displacements are not permitted in all three translational degrees of freedom in order to simulate the disk attachment.

8.3.4 Static stress analysis

A static finite element computation including geometric non-linear effects is performed to evaluate the static stresses in the fan blade. The computation is performed at

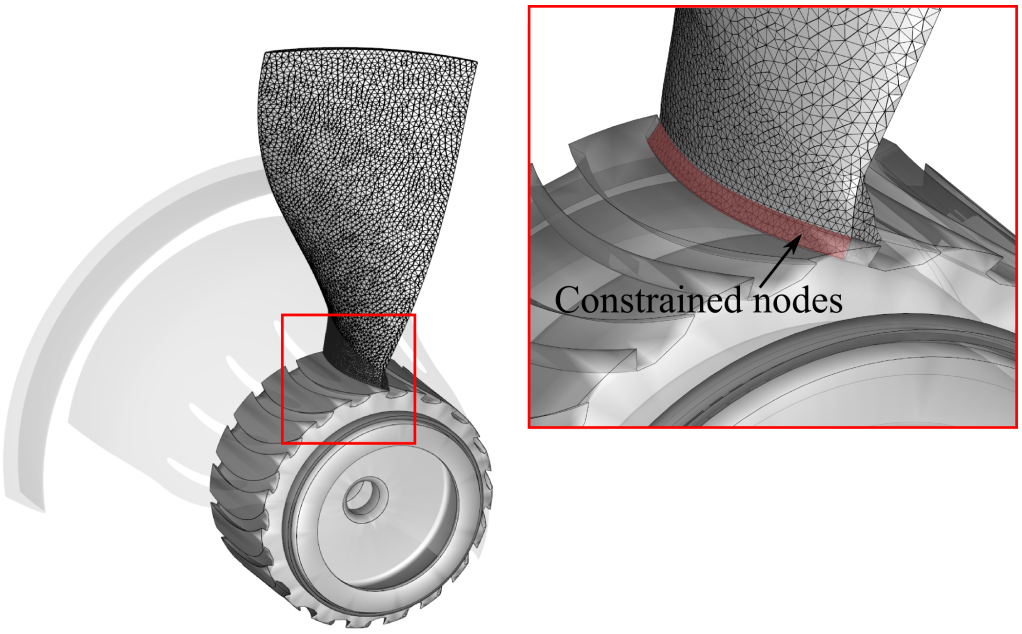


Figure 8.5: Solid mesh of the baseline fan blade and its restraint in a typical disk. The nodes on the dovetail flank surfaces shown in red are fixed in all degrees of freedom to simulate the disk attachment (opposite constrained surface not visible).

take-off conditions, where both aerodynamic and centrifugal loads are at their peak values throughout the mission. The blade is subjected to centrifugal and gas loads whereas the gas loads are extracted from the converged take-off CFD computation of the baseline geometry and interpolated onto the non-matching FEM surface grid using an inverse distance weighting approach. Due to the negligence of the contact problem between the blade root and the disk, the resulting stresses on the dovetail flanks can be considered to be inaccurate. Those stresses are therefore disregarded in the postprocessing step for the assessment of the peak stress values in the blade. Failure of the composite material under three-dimensional stress states is evaluated using the Tsai-Wu failure criterion (see Section 6.8).

8.3.5 Vibration analysis

Fan blade vibration is assessed at all previously mentioned key operating points in order to assess the risk of high cycle fatigue failure. Natural frequencies of the fan blade are determined by means of finite element modal analyses. A geometric non-linear static pre-stressing computation is performed prior to each frequency computation to take the non-linear effect of centrifugal forces on the blade stiffness at different

rotational speeds into account. The Campbell diagram is used to compute the margins between excitation frequencies and blade natural frequencies at the rotational speeds associated with the aforementioned operating points. Excitations from one-per-revolution and two-per-revolution disturbances are considered covering possible sources like unbalance and inlet flow distortions due to cross-wind.

8.4 Objectives and constraints

Two objectives and a total of 13 constraints are imposed in the present optimisation case. The first objective is defined to maximise the peak efficiency at the cruise rotational speed, while the second objective is defined to minimise the total fan mass, computed as the mass of an individual blade multiplied by the number of blades

$$\text{mimimise} \begin{cases} -\eta_{cruise,peak} \\ m_{fan,total} \end{cases} \tag{8.3}$$

A permitted mass flow range is imposed to ensure that the cruise peak efficiency is located within a small margin around the design point mass flow. The corresponding constraints are defined as

$$0.98 \cdot \dot{m}_{cruise,design} < \dot{m}_{cruise,peakeff} < 1.02 \cdot \dot{m}_{cruise,design} \tag{8.4}$$

Hence, peak efficiency needs to occur at a mass flow which differs no more than 2% from the design point mass flow.

In addition, a minimum cruise stall margin of five percent is defined as a constraint based on the last converged CFD solution towards small mass flows, assuming that numerical stall approximates the occurrence of physical stall

$$\text{SM} = \left(\frac{\dot{m}_{cruise,design} - \dot{m}_{cruise,stall}}{\dot{m}_{cruise,stall}} \right) > 5\% \tag{8.5}$$

Top-of-climb performance is evaluated at an operating point defined to have a six percent total pressure margin with the last converged operating point on the ToC speedline.

Bypass and fan root pressure ratios, maximum absolute Mach numbers and maximum absolute flow angles in the bypass and core inlets are respectively imposed as constraints for the ToC evaluation point. The required fan root pressure ratio was elevated in comparison to the optimisation case presented in the previous chapter in order to increase the difficulty of the optimisation problem.

$$\Pi_{t-t,bypass,eval,toc} > 1.7 \tag{8.6}$$

$$\Pi_{t-t,core,eval,toc} > 1.5 \tag{8.7}$$

To ensure that the stages downstream of the fan receive a healthy inlet flow, the permitted maximum absolute Mach number in the core and bypass inlets is defined as

$$Ma_{max,bypass,eval,toc} < 0.9 \tag{8.8}$$

$$Ma_{max,core,eval,toc} < 0.9 \tag{8.9}$$

For a similar reason, the permitted maximum absolute flow angles in the bypass and core inlet are restricted to 65 degrees measured from the meridional plane, which represents a 5 degree reduction with respect to the previous optimisation

$$\alpha_{max,bypass,eval,toc} < 65 \text{ deg} \tag{8.10}$$

$$\alpha_{max,core,eval,toc} < 65 \text{ deg} \tag{8.11}$$

The ToC evaluation point mass flow is restricted to a 4% range around the ToC design point mass flow

$$0.98 \cdot \dot{m}_{toc,design} < \dot{m}_{eval,toc} < 1.02 \cdot \dot{m}_{toc,design} \tag{8.12}$$

Two constraints are imposed to ensure the structural integrity of the fan blade. The first structural constraint imposes a limitation of the maximum Tsai-Wu failure criterion and is defined as

$$\zeta_{TW,max} < 1 \tag{8.13}$$

The second structural constraint is defined to ensure a minimum margin between the eigenfrequency of the first bending mode of the blade and the first and second engine order excitation frequencies at cruise, top-of-climb and take-off. The constraint formulation is identical to the previous optimisation and is given as

$$\Delta\hat{F}_{cruise} + \Delta\hat{F}_{ToC} + \Delta\hat{F}_{T/O} > 12 \tag{8.14}$$

8.5 Optimisation settings

The optimisation was performed with the Differential Evolution constants set to $C = 0.8$ and $F = 0.6$. The population size was set to 50. 10000 generations were computed per optimisation iteration. Eight candidate solutions were chosen from the metamodel-predicted Pareto front after each complete DE optimisation for re-evaluation using the high-fidelity evaluation tools.

8.6 DOE database

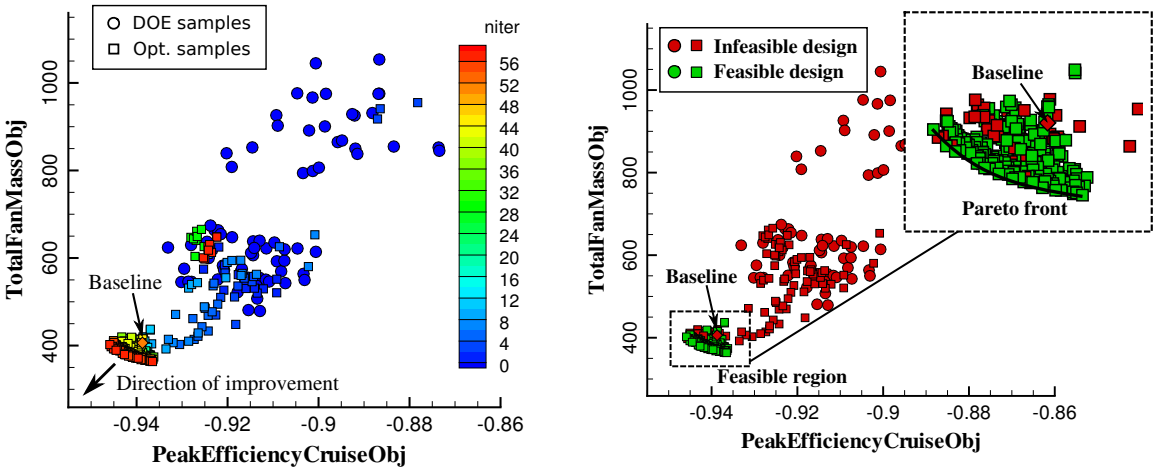
A fractional factorial design of experiments (DOE) technique containing 128 samples was used for the generation of the initial database. The DOE samples the design space at a subset of the full factorial design at 20 and 80% of the design range. In addition to the DOE samples covering the boundaries of the design variable ranges, the baseline design and the "central" design (generated with all design variables set to the centre value of their respective range) were added to the initial database. 72 samples were successfully evaluated by the high-fidelity evaluation chain. With the baseline design included, the DOE database therefore consists of 73 designs. Of the 57 individuals that failed, 36 were due to non-convergence of the CFD solver. None of the designs contained in the initial database simultaneously satisfied all imposed constraints.

8.7 Optimisation results

8.7.1 Survey of optimisation process

The objective space obtained after a total of 59 optimisation iterations is shown in Fig. 8.6. Improved performance is obtained towards the lower left corner of the objective space, corresponding to the minimisation of both objectives. In total 524 individuals were successfully evaluated by the high-fidelity evaluation chain, which amounts to 5240 CFD evaluations, 524 non-linear static finite element computations and 1572 finite element frequency analyses (each including a non-linear static pre-stressing computation). The overall computational time was 24 days on 128 AMD EPYC 7551 2GHz CPU cores.

Eleven iterations were required until the optimiser successfully identified the feasible region in the design space. As indicated previously, none of the designs in the initial database simultaneously satisfied all of the imposed constraints, presenting a



(a) Global view.

(b) Detail view on feasible individuals.

Figure 8.6: Objective space after 59 iterations. Circles indicate DOE samples, while squares indicate designs generated during the optimisation.

significant challenge for the optimisation system. Nevertheless, each constraint was individually satisfied by some of the initial designs, hence allowing the optimiser to detect and combine design features that each lead to the satisfaction of a particular constraint. This is illustrated in the constraint spaces plotted in Fig. 8.7 and Fig. 8.8. Negative values indicate constraint satisfaction. Due to an odd total number of constraints, the stall margin constraint is plotted against itself in Fig. 8.8 for the sole reason of enabling a consistent visualisation of the constraint spaces.

It can be seen that some constraints are harder to satisfy than others. For instance, all designs in the database are found to satisfy the bypass and core maximum Mach number constraints, Fig. 8.7 (b), while for example only a few designs in the database satisfy the maximum Tsai-Wu failure criterion constraint seen in Fig. 8.7 (f).

The constraint spaces also give a direct overview of the baseline constraint violations. The baseline design is shown as orange diamond. With regard to the aerodynamic constraints, it can be seen that the design satisfies all but the pressure ratio constraints for both, bypass and core flow (Fig. 8.7 (c)). In terms of structural performance, the maximum Tsai-Wu failure criterion constraint is satisfied, while the vibration constraint is violated (Fig. 8.7(f)).

Note, that the appearance of all individuals on a straight line for the four mass flow constraints, i.e. Fig. 8.7 (d) and (e) is due to the constraint definition (Eqs. (8.4) and (8.12)). For a tighter mass flow constraint, this line would move closer to the

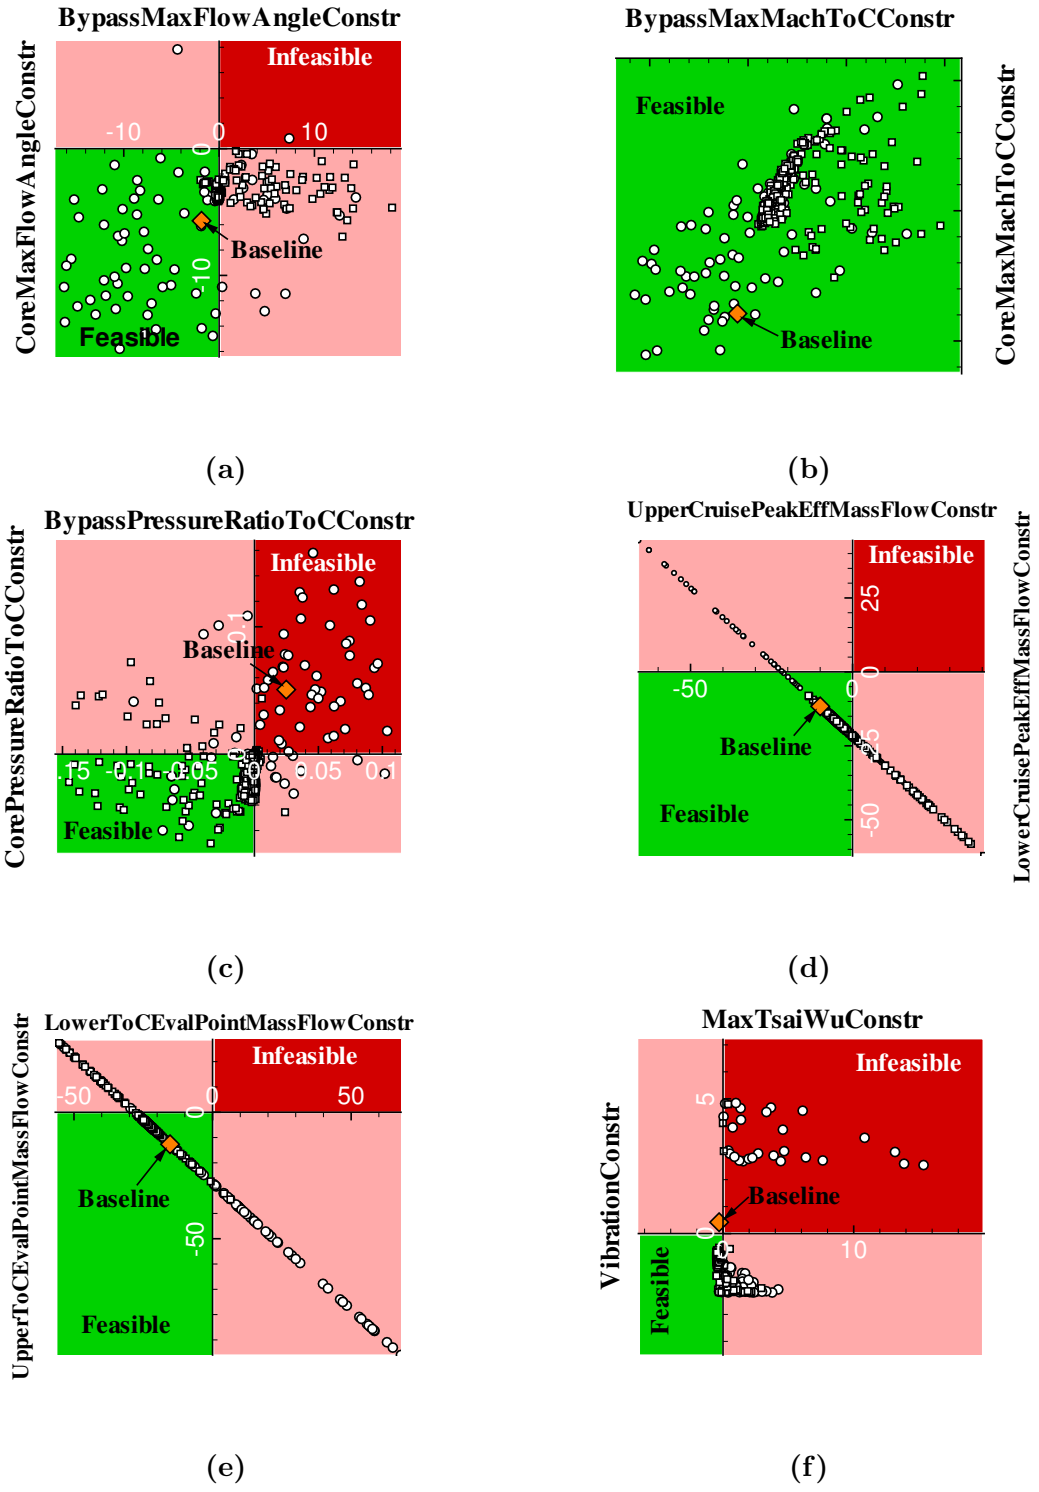


Figure 8.7: Overview of constraint spaces. Every symbol shown represents a design evaluated by the high-fidelity evaluation chain. Circles indicate DOE samples, while squares indicate designs generated during the optimisation. Negative values indicate constraint satisfaction. The baseline design is shown as orange diamond.

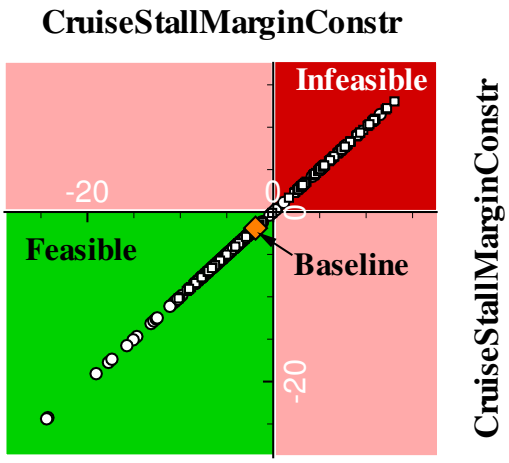


Figure 8.8: Overview of the cruise stall margin constraint. The symbol convention follows that of Fig. 8.7. Negative values indicate constraint satisfaction.

origin and would intersect it at the limit of an equality constraint.

Figure 8.9 provides a view on the overall learning progress of the Kriging metamodel throughout the optimisation progress. The figure shows the percentage error between the metamodel predictions and the high-fidelity performance evaluations, averaged over the candidate designs in each iteration. An overall improvement of the prediction accuracy with the optimisation progress can be observed. However, it is equally visible that some performance parameters appear to be harder to predict than others. This includes the cruise stall margin, which shows by far the highest overall error magnitudes throughout the optimisation process, reaching up to 6000%. Nevertheless, an improvement of the prediction accuracy with the optimisation progress is observable. Large error fluctuations are further visible for the Tsai-Wu failure criterion, indicating that this is another quantity that is hard to predict, similarly to the maximum von Mises stress constraint applied in the optimisation presented in the previous chapter. In contrast, very high prediction accuracies can be observed for the core and bypass pressure ratios.

The evolution of the design parameters over the optimisation process is depicted in Fig. 8.10 and 8.11. As before, red symbols represent infeasible designs, while green symbols show feasible ones. It can be observed that the optimiser initially explores a wide range of the design space in the search for the feasible region. This goal is successfully achieved in the 11th iteration. In subsequent iterations, feasibility of the designs can be seen to be overall maintained. However, as the search progresses

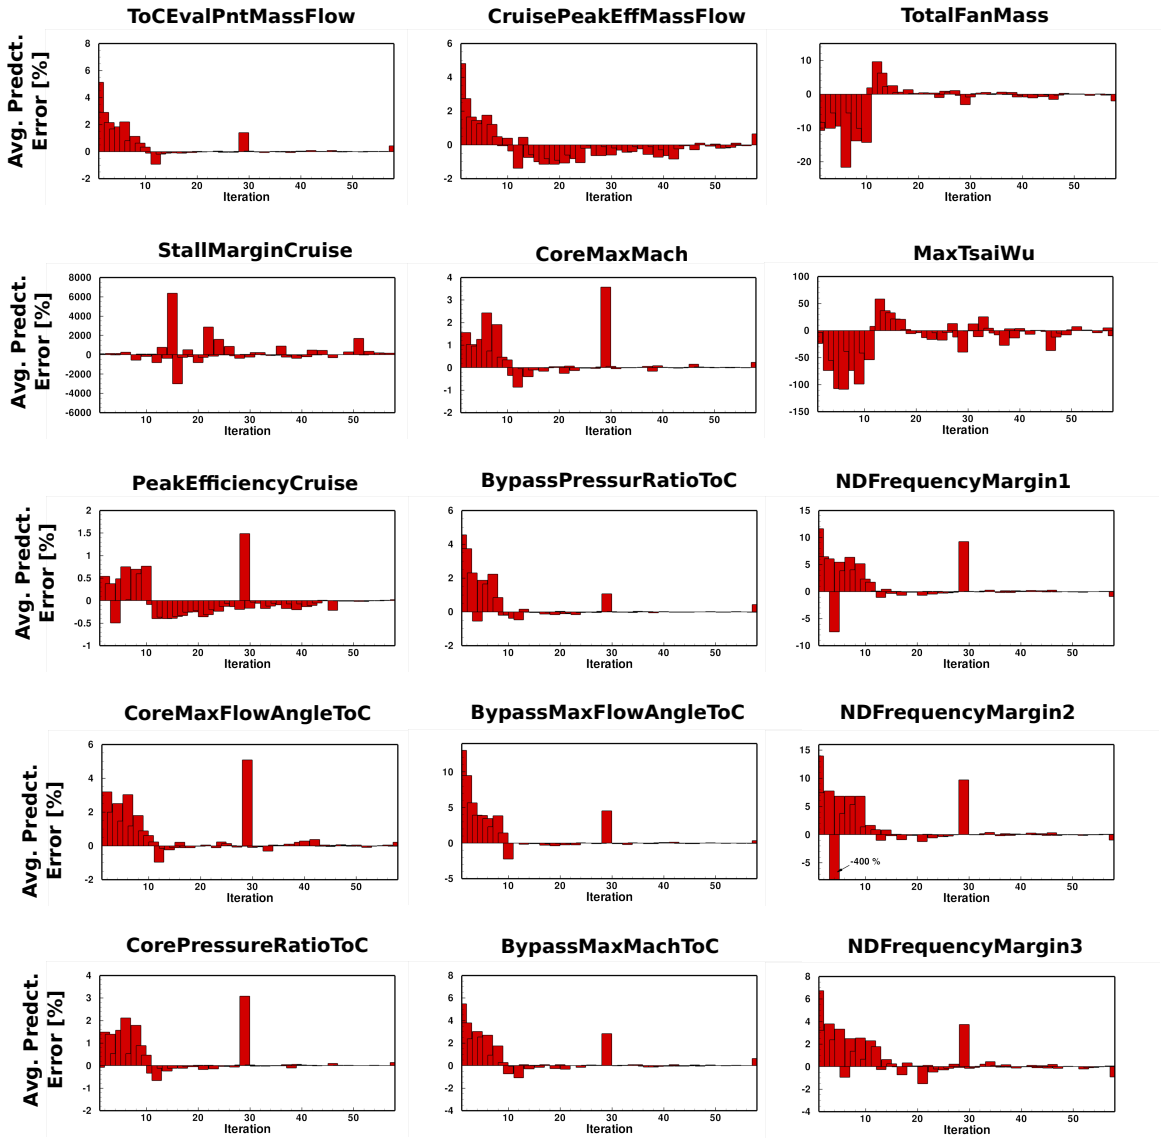


Figure 8.9: Overview of the metamodel prediction errors for all performance parameters in dependency of the optimisation progress. The error values represent the average error of all candidate designs per iteration.

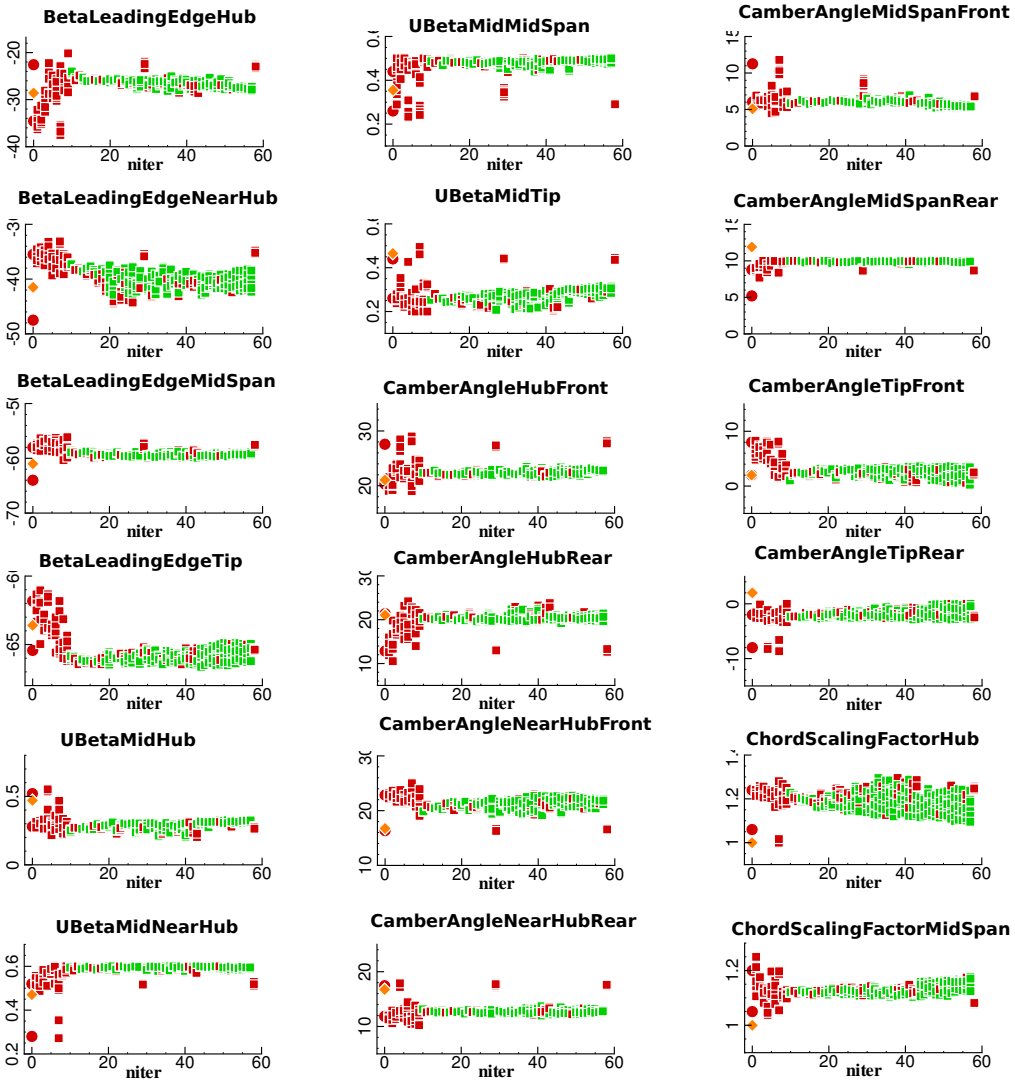


Figure 8.10: Evolution of design parameters - part 1. Green symbols denote feasible designs, while red symbols indicate infeasible ones. The baseline design is shown as orange diamond.

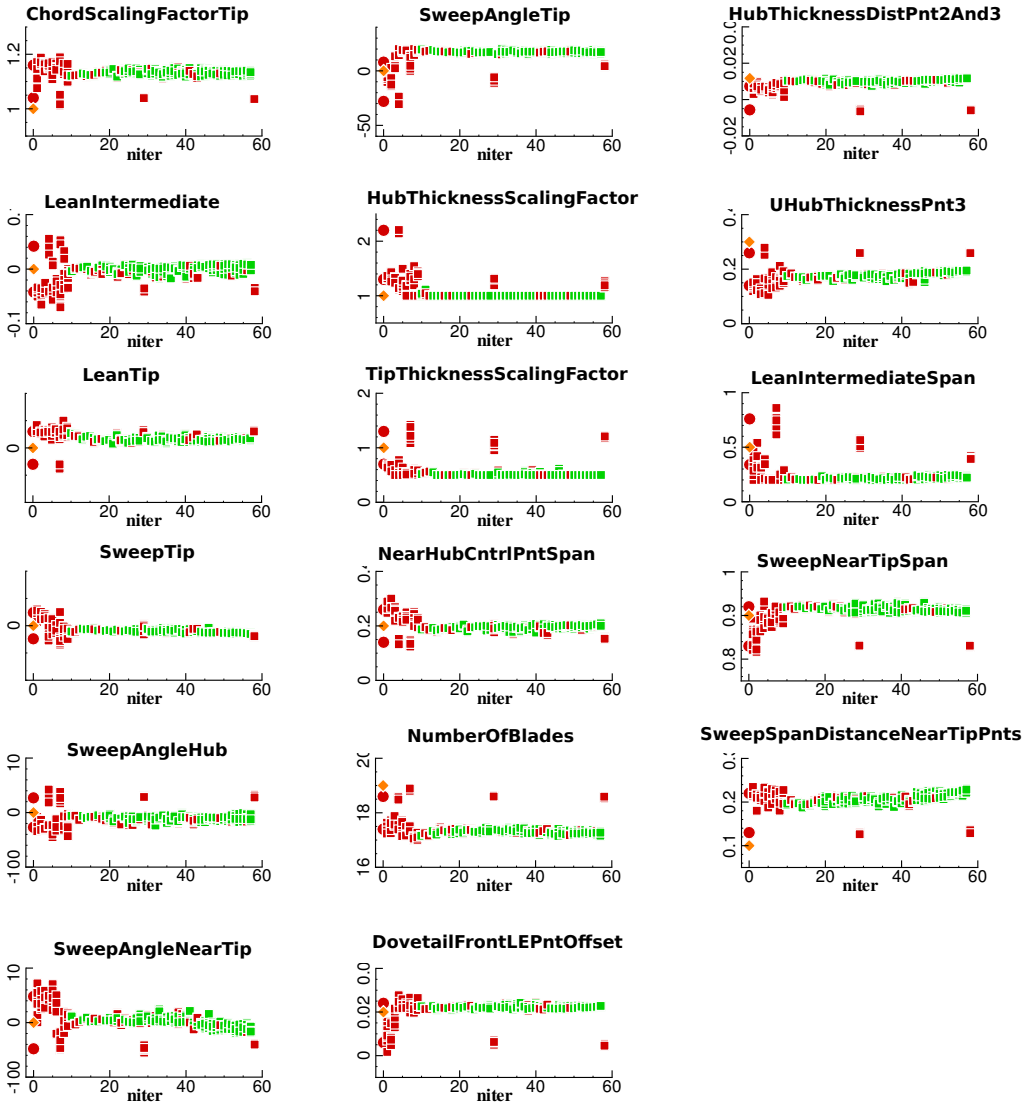


Figure 8.11: Evolution of design parameters - part 2. Green symbols denote feasible designs, while red symbols indicate infeasible ones. The baseline design is shown as orange diamond.

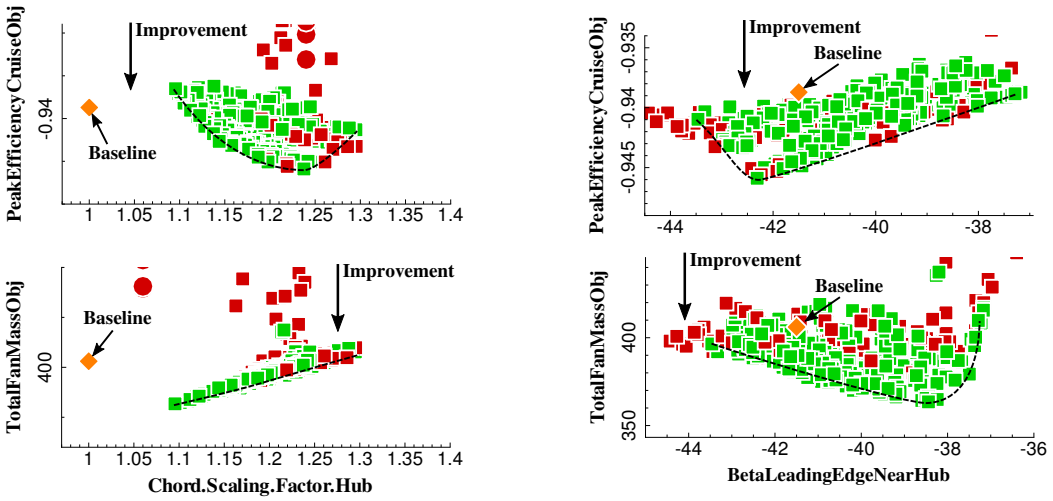


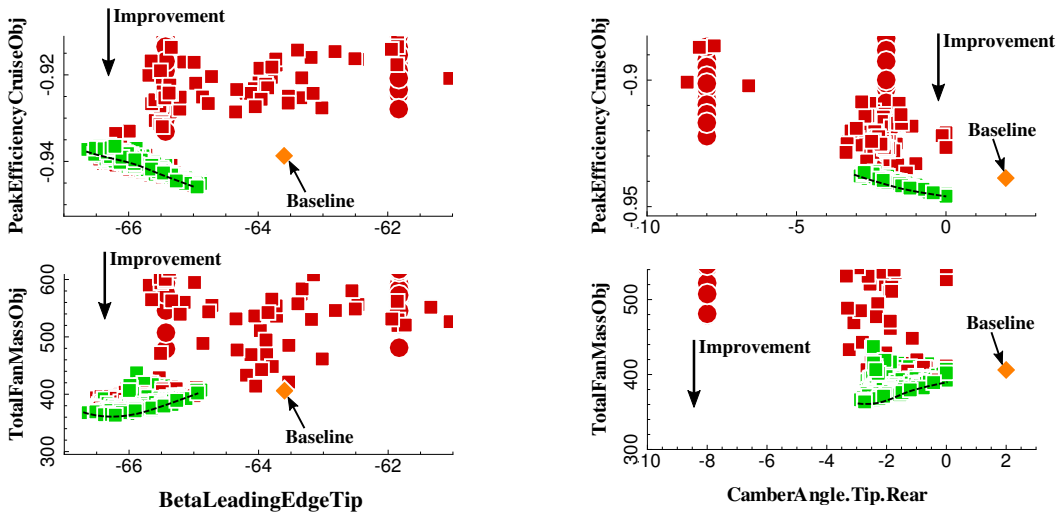
Figure 8.12: Impact of hub chord scaling factor and the near hub leading edge blade angle on the cruise peak efficiency objective and the total fan mass objective. Green symbols denote feasible designs, while red symbols indicate infeasible ones.

towards improved objective performance, a number of infeasible designs can still be seen to be generated, which can be primarily related to inaccuracies of the metamodel predictions.

To achieve feasibility, the following main design changes can be observed with respect to the baseline:

- A reduction of the number of blades
- An increased chord length from hub to tip
- Parameters defining the blade angle distributions have changed significantly over the blade height
- The tip section thickness has been reduced

Some parameters are seen to present very little diversity as the optimiser attempts to improve the objectives. This indicates a tightly constrained feasible design space for that particular parameter. An example is the sweep angle at the tip. The hub and tip thickness scaling factors are equally kept almost constant. Both are furthermore pushed towards their lower design space limits, which has been found to be due to the interaction of several constraints, with the mass flow and stress constraints being the most dominant for the blade thickness.



(a) (b)

Figure 8.13: Impact of leading edge blade angle and the rear camber angle at the tip on the cruise peak efficiency objective and the total fan mass objective. Green symbols denote feasible designs, while red symbols indicate infeasible ones.

In contrast, considerable diversity can be observed in other parameters. Particularly pronounced parameter variations are existing in the evolution of the hub chord scaling factor, the beta leading edge values at near-hub and tip, and the rear camber angle at the tip. It can be seen in Figs. 8.12 and 8.13 that all mentioned parameters present a trade-off between peak efficiency and fan mass objectives. While the hub chord scaling factor and the blade mass, shown in Fig. 8.12 (a), are directly related over blade volume and material density for the integral blade construction considered here, blade angle values are generally not strongly correlated to the blade mass. The observed dependencies between the fan mass objective and the blade angle values in Figs. 8.12 (b), 8.13 (a) and (b) are thus more likely the result of coupling effects to other design and performance parameters.

A detail view on the Pareto front is shown in Fig. 8.14. As can be seen, the optimiser achieved a considerable performance improvement in both objectives in comparison to the baseline design. A trade-off between the total fan mass and the peak efficiency can be observed, indicating that designs with lower weight can be generated at the cost of a lower efficiency. A comparison of the designs at the extremities of the Pareto front (IT059_IND001 and IT059_IND015) shows that a weight reduction of 30.3 kg comes at the cost of 0.75 % in efficiency. These two designs, in addition to a compromise design located at the centre of the Pareto front (IT057_IND038), were chosen to investigate the trade-off between the objectives in further detail.

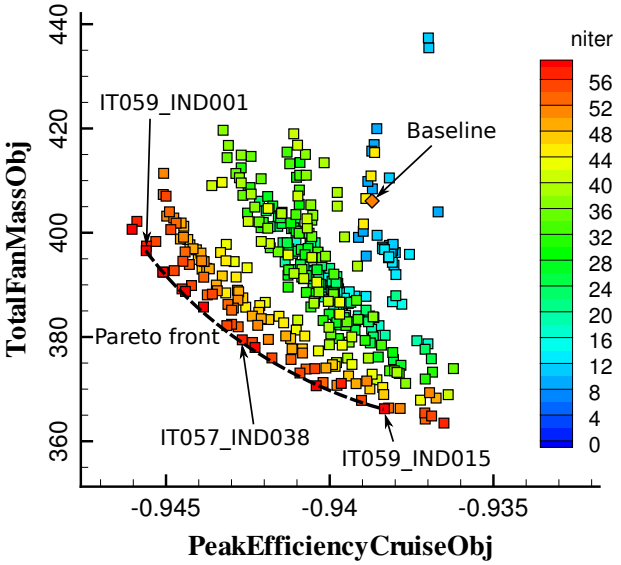


Figure 8.14: Detail view on the Pareto front obtained after 59 optimisation iterations and indication of individuals chosen for further examination.

8.7.2 Comparison of Pareto optimal solutions

The three chosen Pareto optimal solutions shown in Fig. 8.14 will now be compared to investigate the observed trade-off between total fan weight and cruise peak efficiency in further detail and to assess how the chosen design parameters influence the aeromechanical fan performance. The baseline design is added to the comparison to provide a reference. All optimised designs have a blade count equal to 17, which represents a reduction of two blades in comparison to the baseline fan and equals the specified lower bound of the design variable.

8.7.2.1 Aerodynamic performance

10% span A comparison of the blade profiles at 10% span (equal to about mid-height of the core inlet) is depicted in Fig. 8.15. It can be observed that the hub chord length has been significantly increased in all optimised designs with respect to the baseline profile, which in effect counters the reduction of solidity originating from the reduced blade count. The comparison shows that, among the three Pareto optimal designs, the lowest weight design is characterised by the shortest chord length, followed by the compromise design, while the highest efficiency design has the longest chord. All blade sections have an identical maximum thickness, corresponding to a hub thickness scaling factor of one, which equals the minimum permitted value. As

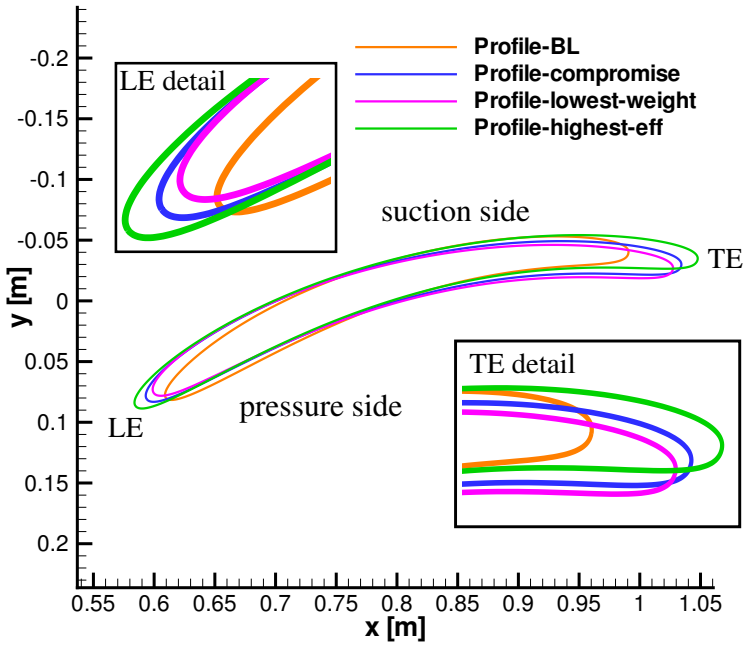


Figure 8.15: Comparison of baseline and optimised blade profiles at 10% span.

shown previously, the hub chord length strongly affects the overall blade mass as well as the efficiency, see Fig. 8.12 (a).

Aerodynamically, all optimised blade profiles have to deliver the same pressure rise in order to meet the top-of-climb constraint for the core flow. Indeed, the optimiser did an excellent job in ensuring that the constraint is almost active in all optimised designs, implying that the design just delivers the required pressure rise, and not more. This is also due to the excellent metamodel prediction accuracy for this performance parameter, see Fig. 8.9. The fan root pressure ratios are 1.506 for the highest efficiency design, 1.505 for the lowest weight design and 1.503 for the compromise design, which can be considered as practically identical to the required pressure rise of 1.5. For reference, the baseline design delivers a root pressure rise of 1.45.

The above implies that the lightweight, low-chord design needs to operate at a higher aerodynamic loading than the longer chord designs, which has adverse consequences for the efficiency. The isentropic Mach number distributions shown in Fig. 8.16 provide an indication of the blade loading distributions of the four chosen designs at the two considered flight conditions, cruise and top-of-climb.

The most significant difference to be observed in the isentropic Mach number distributions is a successive increase of front loading from the highest efficiency design

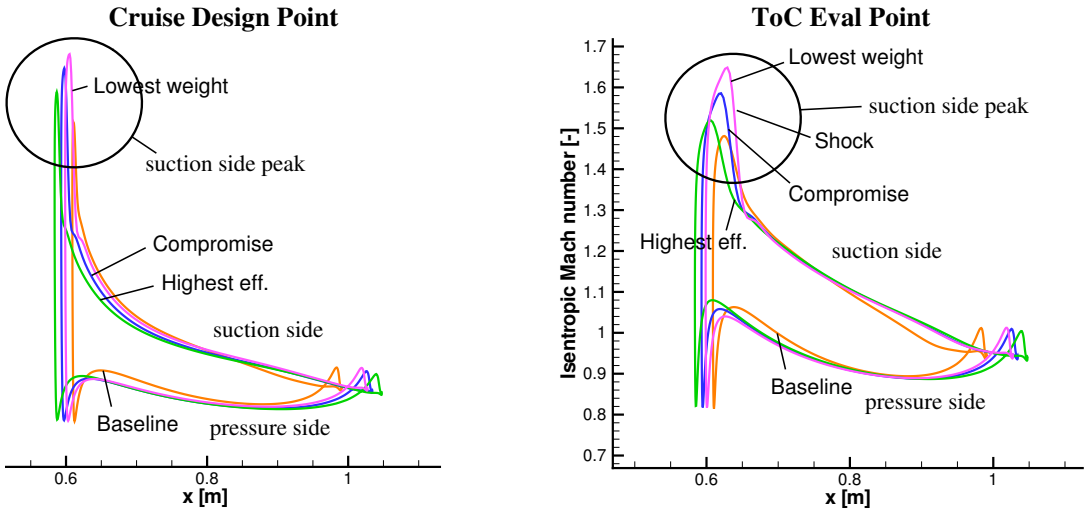


Figure 8.16: Comparison of isentropic Mach number distributions at 10% span at the cruise design point and the top-of-climb evaluation point.

towards the compromise and lightweight design. This is most clearly indicated by the visible increase of the suction side peak Mach number for both operating points. Loading differences are overall far less pronounced between about 20 and 80% chord. While the most marked difference at the trailing edge is seen to occur due to the different axial profile extensions, the trailing edge Mach numbers are overall comparable.

Small differences in the leading edge blade angles are existing. The low weight design has a two degree leading edge angle reduction with respect to the highest efficiency design (-33 vs -35 degrees), while the compromise design has a value of -34 degrees. The reduced leading edge blade angle causes the incidence angle to rise, which is the reason for the front loading increase. At top-of-climb, the operating point at which the profile is subjected to the highest loading, the increased incidence raises the suction side peak Mach number above sonic speeds, such that a shock occurs. The shock significantly increases the pressure rise, leading to the satisfaction of the required fan root pressure ratio, but comes at the cost of a considerable increase in losses.

Figure 8.17 depicts the entropy production contours for the lowest weight and highest efficiency blade profiles at 10% span at the cruise design point. An increased entropy production for the lowest weight design can be identified downstream of about 50% chord on the suction side, pointing to higher diffusion losses and a consequently

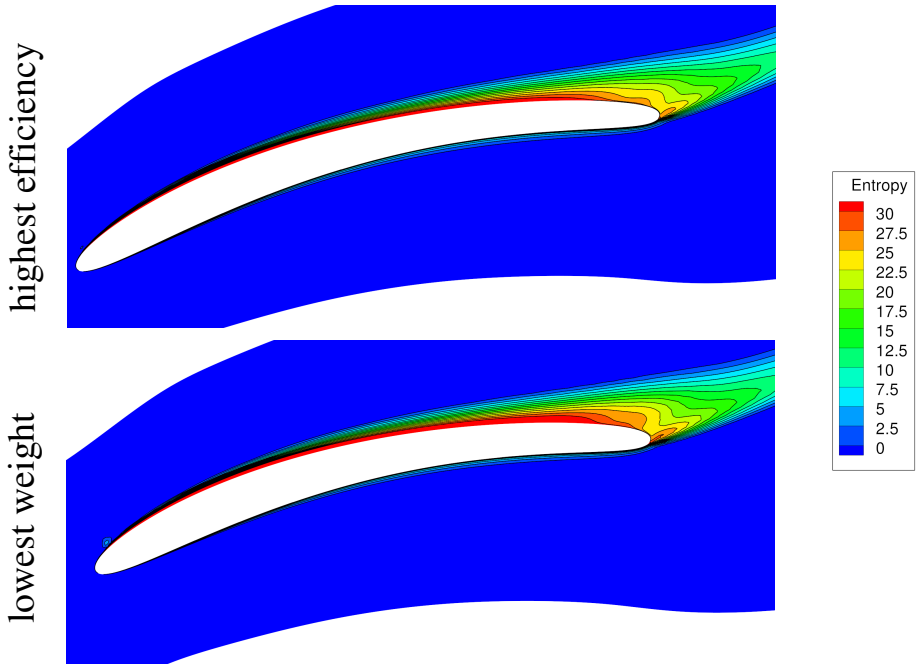


Figure 8.17: Comparison of the entropy contours at 10% span between the highest efficiency and lowest weight design at the cruise design point.

lower efficiency.

90% span The optimised blade profiles at 90% span are shown in Fig. 8.18. They can be seen to have significantly changed with respect to the baseline design. The optimised designs show a markedly different leading edge angle, having increased by about 2 degrees. More specifically, the optimised blade leading edge angles are -64.6, -64, and -63.7 degrees, for the lowest weight, compromise and highest efficiency blades, respectively. The baseline leading edge blade angle is -62.7 degrees. In addition, the blade thickness was considerably reduced in all optimised designs, corresponding to the lowest permitted tip thickness scaling factor of 0.5. Similarly to the blade sections at lower span, the chord lengths have been increased with respect to the baseline design.

The loading differences between the optimised tip shapes are identifiable in Fig. 8.19 for both cruise and top-of-climb operation. However, the overall differences between the optimised geometries are less pronounced than for the lower span sections discussed previously.

An overview of the tip solidities of all four designs is given in Table 8.4. As can be seen, the increased tip chord lengths of the optimised blades have resulted

Table 8.4: Overview of tip solidities.

Designation	Number of blades	Tip solidity (90% span)
Baseline	19	1.38
Lowest weight	17	1.38
Highest efficiency	17	1.40
Compromise	17	1.40

in a close match of the baseline solidity, thus compensating for the lost solidity due to the blade count decrease. A commonly applied rule of thumb for the design of transonic compressor blades is to match the tip solidity to the tip inflow relative Mach number [175]. In the present case, the inflow relative Mach number at 90% span is approximately equal to 1.3 at the cruise design point, which is shown to be in good proximity to the optimised solidity values.

Three-dimensional design features - blade stacking A comparison of the stacking axes of the three optimised designs in the meridional plane is shown in Fig. 8.20. Recall, that the baseline geometry has a straight radial stacking axis

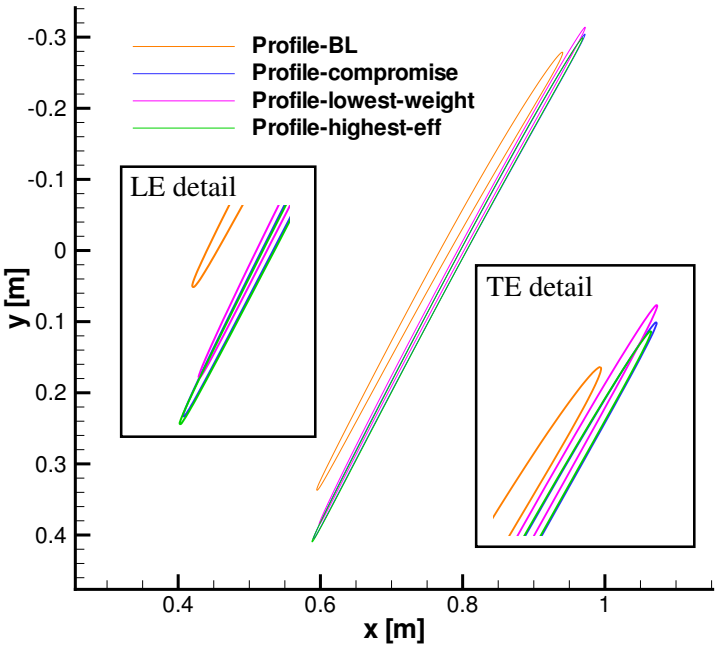


Figure 8.18: Comparison of baseline and optimised blade profiles at 90% span.

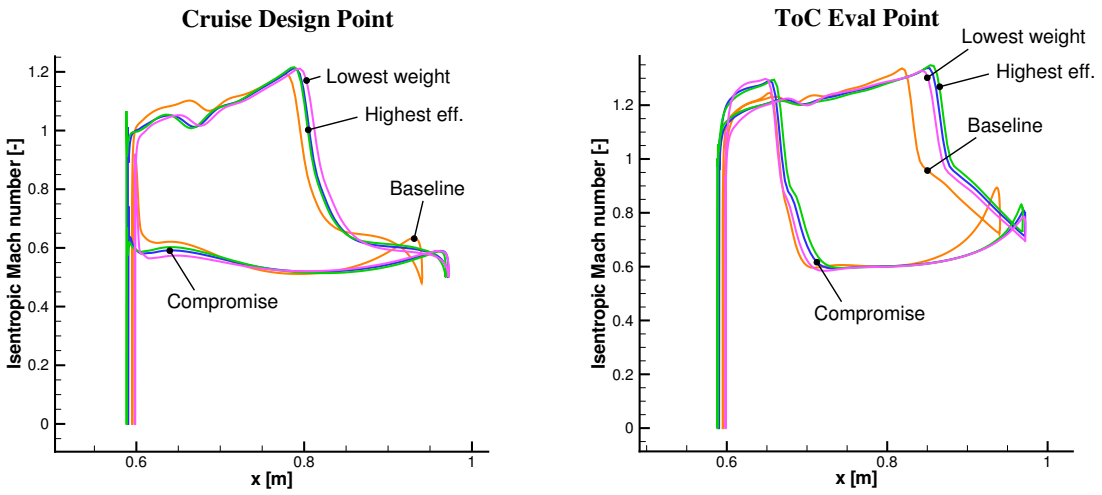


Figure 8.19: Comparison of isentropic Mach number distributions at 10% span at the cruise design point and the top-of-climb evaluation point.

through the centre of gravity of all blade sections and is therefore not shown. It can be seen that the blade sections of all three optimised designs only show a minor offset from the radial stacking axis. All designs have a similar, slight S-shape, trend of the stacking axis, with a small amount of backward sweep below about 80% span and a backward swept tip.

The backward sweep angle at the tip can be seen to be least pronounced for the lowest weight design, while it is most distinct for the highest efficiency design. Indeed, it can be seen in Fig. 8.21 (a) that all optimised designs have a negatively valued tip sweep angle, resulting in the observed backward swept tip. Moreover, increasing values of the backward sweep angle at the tip can be seen to increase the cruise peak efficiency up to a value of about -15 degrees, while further increased levels of backward sweep are observed to have an adverse effect on the efficiency. The tip sweep angle has a considerable influence on the stresses in the blade. As can be seen in Fig. 8.21 (b), where the Tsai-Wu stress constraint is plotted in dependency of the tip sweep angle, designs satisfying the stress constraint were found exclusively for blades with a backward swept tip; the vast majority of forward swept designs introduced an excessive amount of stresses. Hence, the optimiser did not continue to search for designs with forward swept tip sections.

The aerodynamic mechanisms leading to the improvement of efficiency with a backward swept tip of transonic fans have been discussed in detail in the literature, and the effects found here are pretty much in line with previous observations (see

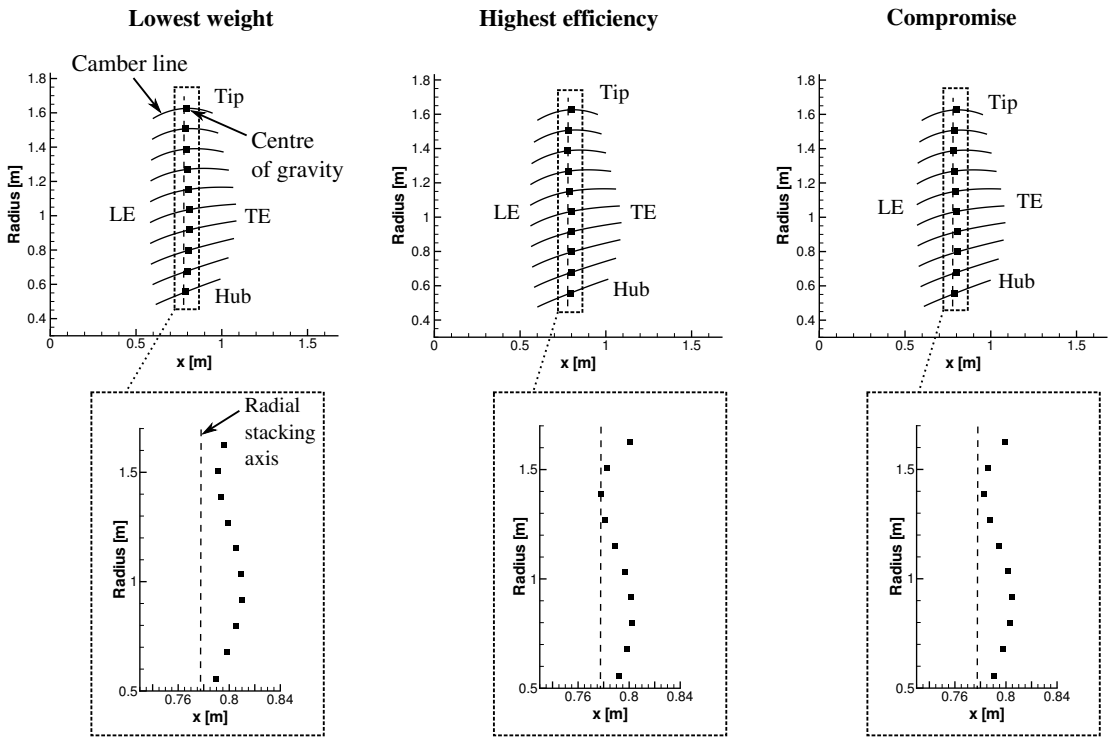


Figure 8.20: Comparison of the stacking axes of the optimised geometries in the meridional plane. Squares indicate the centre of gravities of the blade sections. The straight dashed lines represent a radial stacking axis.

e.g. [175]). Backward sweep has been observed to be beneficial for the efficiency up to a certain value, above which the efficiency and stall margin degrade. Using the CFD solutions at the cruise design point for the baseline design, the highest efficiency design and a design with a more pronounced value of backward sweep shown as IT009_IND047 in Fig. 8.21, the flow mechanisms of increasing levels of backward sweep will now be further investigated.

Figure 8.22 shows the iso-surfaces for a relative Mach number of one for the three previously indicated designs with increasing levels of backward sweep. It is well known, that the operating conditions in civil aero-engine fans cause the shock surface to intersect the casing perpendicularly, irrespective of the amount of blade sweep [39], [175]. For the baseline design with zero sweep, it can be seen that the passage shock at the tip is attached to the leading edge and sits slightly inwards of the covered passage. It is important to observe that the shock is tilted in the spanwise direction despite the absence of any geometrical blade sweep, as best observable by consideration of the shock impingement line on the suction surface. In the highest efficiency design, which features a sweep angle of -14.5 degrees at the tip, it can be seen that the passage shock

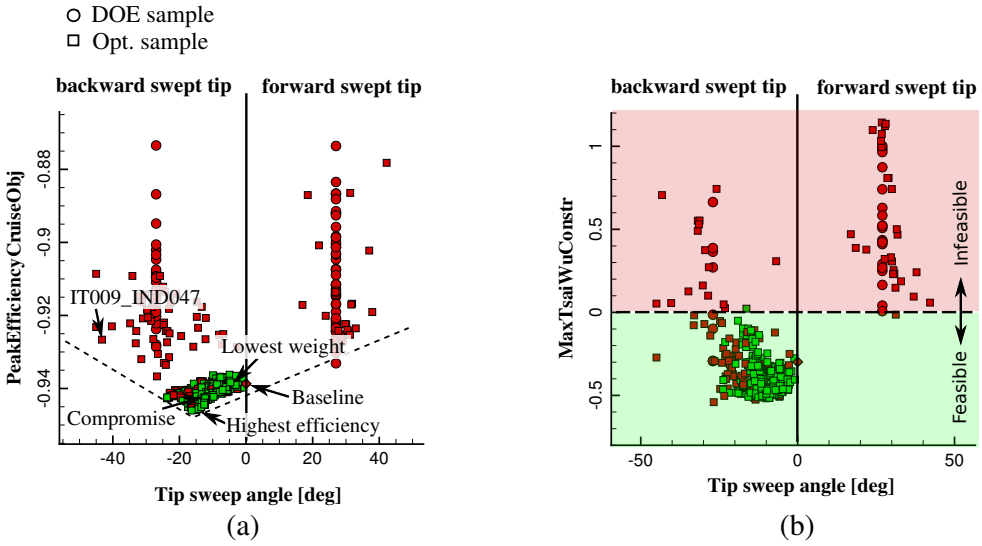


Figure 8.21: Influence of the tip sweep angle on the cruise peak efficiency objective (a) and the Tsai-Wu stress constraint (b). Green symbols denote design that are satisfying all constraints, while red symbols indicate designs with at least one constraint violation.

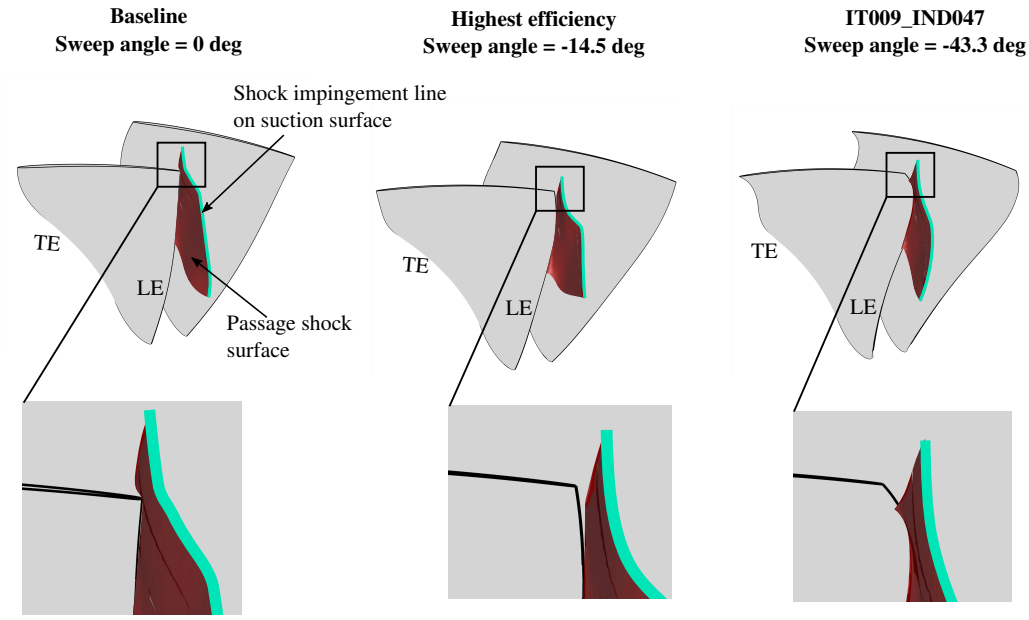


Figure 8.22: Impact of different levels of backward sweep at the tip on the shock surface at the cruise design point (red surface shows the iso-surface of unity relative Mach number).

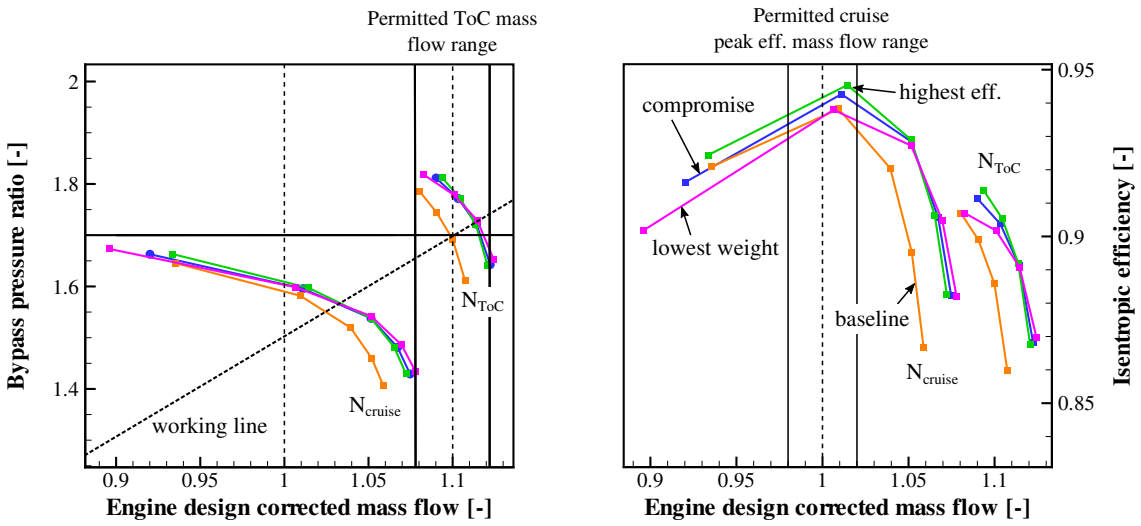


Figure 8.23: Performance curves of the chosen designs.

is located just in front of the covered passage, which is typically the configuration at which the highest efficiency occurs in supersonic blade sections e.g. [26]. In the design with the highest amount of backward sweep corresponding to a sweep angle of -43.3 degrees, it can be seen that the tip section is pulled further backwards, while the passage shock inevitably stays normal to the casing, as it transitions upward in span from the leading edge attachment at around 90% span. This causes the tip section to operate closer to stall, and consequently degrades the blade tip section efficiency and stall margin.

Performance curves The performance curves for the three optimised designs and the baseline design are shown in Fig. 8.23. The bypass pressure ratio versus total engine mass flow characteristics for both operating speeds can be seen to be overall very similar for the three optimised designs. All optimised designs can further be seen to deliver an overall increased bypass pressure rise at the top-of-climb rotational speed in comparison to the baseline fan, thus leading to the satisfaction of the bypass pressure ratio constraint. A slight shift of the speedlines towards increased mass flows can be observed, being likely due to the reduced blade count with respect to the baseline design. However, the operating range is still well within the permitted mass flow range.

All designs satisfy the required minimum cruise stall margin of 5%. It is worthwhile to point out that the last CFD operating point towards reduced mass flows failed to converge for all shown designs, therefore indicating that the computed range

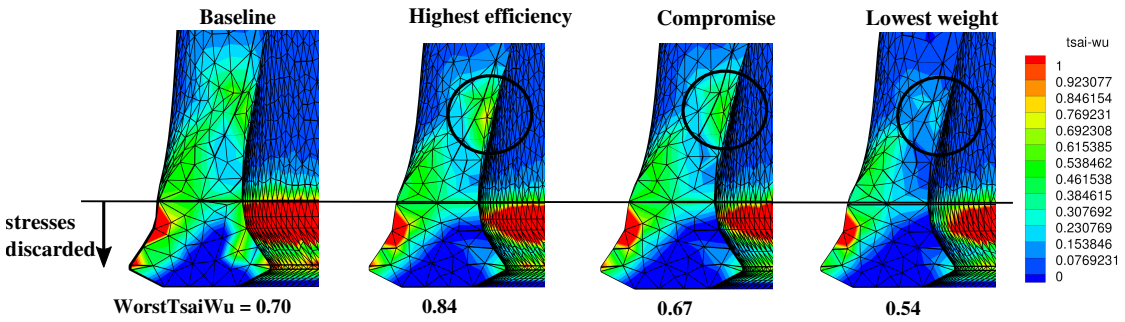


Figure 8.24: Comparison of the Tsai-Wu failure criterion distribution in the front root of the three optimised and the baseline design.

should not be too far off from the numerical stall limit. The peak efficiency of all optimised geometries can be seen to be located well within the cruise peak efficiency mass flow constraint, which was set to a 4% range around the design point mass flow.

8.7.2.2 Structural performance

The distributions of the Tsai-Wu failure criterion in the front root, computed at take-off operation, is shown in Fig. 8.24. Disregarding the stresses in the dovetail flanks due to reasons discussed earlier, the peak failure criterion value in all designs is smaller than one and therefore satisfies the imposed constraint. However, significant differences in the worst (maximum) Tsai-Wu failure criterion values exist between the designs. As shown in Fig. 8.24, the lowest weight design has the lowest value of the maximum failure criterion, and therefore would have the highest safety factor towards blade failure due to static loading. Reduced safety margins are observable for the compromise solution and the highest efficiency designs.

All designs shown in Fig. 8.24 have a tilted front root section, and therefore a non-zero value of the dovetail front leading edge point offset (see Section 4.3). The overall influence of this parameter on the blade stresses in terms of the maximum Tsai-Wu failure constraint is shown in Fig. 8.25. A parameter value of zero translates to a straight front root section, while a non-zero value introduces a tilted front root section. Designs satisfying the stress constraint can be seen to be clustered between parameter values of about 0.02 and 0.026 m, while no design was found to be feasible for smaller values of the offset parameter (with the exception of one initial database sample which just satisfies the constraint).

To explain this matter, Fig. 8.26 shows the Tsai-Wu failure criterion distribution in the frontal root section for two different values of the leading edge offset parameter

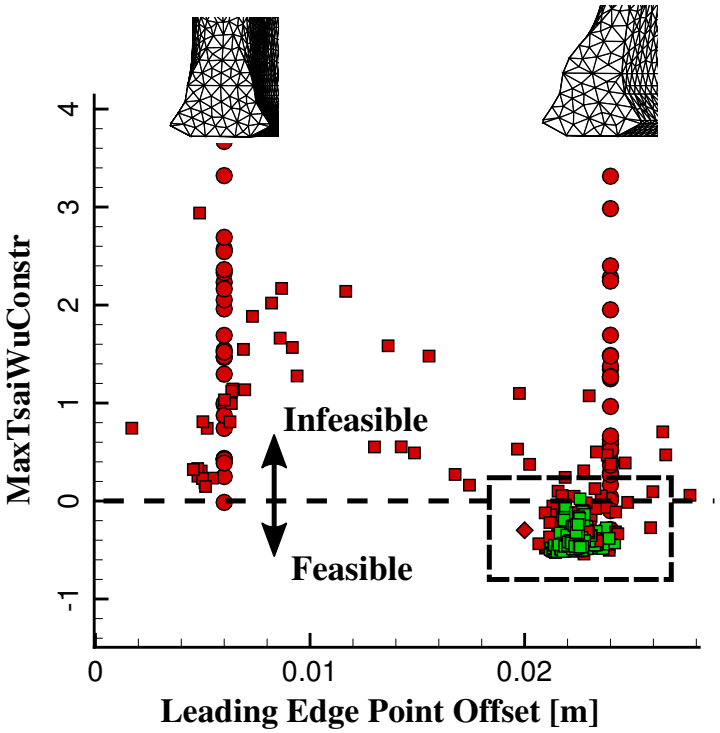


Figure 8.25: Influence of the leading edge point offset value (front root tilting) on the Tsai-Wu stress constraint.

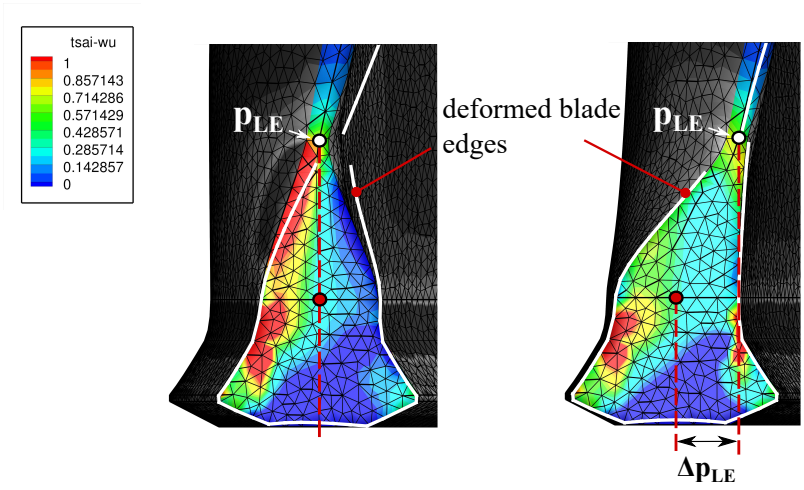


Figure 8.26: Stress distributions in a straight vs. tilted front root section in terms of the Tsai-Wu failure criterion. Values greater than one indicate material failure. White edges indicate the deformed shape with tenfold magnification.

for an otherwise identical blade geometry exposed to centrifugal loading at take-off conditions. The red portions of the contour plot indicate a failure criterion value greater than one and therefore material failure. It can be seen that the straight root construction shown on the left hand side exhibits a significant region of excessive stresses spanning over almost the entire height of the shank surface. The reason for this becomes clearer when considering the deformed blade geometry, illustrated with tenfold magnification with the white blade edges in the figure. It can be seen that as the centrifugal force pulls the blade radially outwards, the leading edge curve is moving to the right, therefore inducing significant bending moments in the transition region with corresponding out-of-plane stresses exceeding the out-of-plane strength of the composite laminate. In contrast, the tilted root results in an almost pure tension stress state in the frontal root region with a corresponding alleviation of the harmful out-of-plane stresses.

The Campbell diagrams for all four designs are shown in Figs. 8.27-8.31. The baseline design shown in Fig. 8.27 has a dangerously small frequency margin between the first engine order and the first bending mode (1F) at take-off (T/O) operation, resulting in the violation of the imposed frequency constraint. In contrast, the second engine order can be seen to be uncritical. By increasing the overall blade stiffness, the eigenfrequency of the first bending mode was successfully raised by the optimiser, thus leading to the satisfaction of the frequency constraint. The corresponding Campbell diagrams of the three chosen optimised designs are shown in Figs. 8.28-8.31. A comparison of the figures shows that the optimised designs have an overall very similar vibration response.

8.7.2.3 Further geometry and performance data

A global comparison of the blade geometries and full fan assemblies of baseline and optimised designs is shown in Figs. 8.32 to 8.35. Figure 8.36 and Fig. 8.37 show the relative Mach number contours for all blades at 10, 50 and 90 percent span. Fig. 8.38 shows the pitchwise averaged spanwise distribution of the absolute total pressure at the fan exit, upstream of the bypass splitter at the cruise design point.

8.8 Summary

A second multidisciplinary fan blade design optimisation was presented in this chapter. The optimisation problem complexity was raised in comparison to the case presented in the previous chapter. The fan blade geometry was parametrised with

35 design variables. A woven fabric composite material was applied. The isentropic total-to-total peak efficiency at cruise operation and the total fan mass were defined as objectives in the optimisation problem. Further, in total 13 structural and aerodynamic constraints were imposed.

A fractional factorial design of experiments (DOE) technique was used to initially sample the design space in order to initialise the optimisation process. The DOE contained 128 samples, of which 72 were successfully evaluated by the high-fidelity evaluation chain. No individual in the DOE database simultaneously satisfied all imposed constraints.

The optimisation was run for a total of 59 optimisation iterations, equalling a total of 524 individuals that were successfully evaluated by the high-fidelity analysis tools. The overall computational time was 24 days on 128 CPU cores. The highly constrained optimisation problem presented a challenge for the optimiser. However, after 11 optimisation iterations, representing 88 high-fidelity performance evaluations, the feasible region in the design space was successfully identified. In subsequent iterations a considerable improvement of both objectives was achieved. The optimisation process resulted in a distinct Pareto-front, revealing the existence of a trade-off between the cruise peak efficiency and total fan mass objectives.

A detailed analysis of the influence of single optimisation parameters on the aerodynamic and structural performances allowed to identify the most important parameters that contributed to design improvements. It was found that the hub chord length was a key parameter in the observed trade-off between efficiency and blade mass; while designs with shorter chord lengths were seen to have a reduced blade mass, they need to operate with higher loading in order to meet the imposed core pressure ratio constraint, which consequently penalised the efficiency.

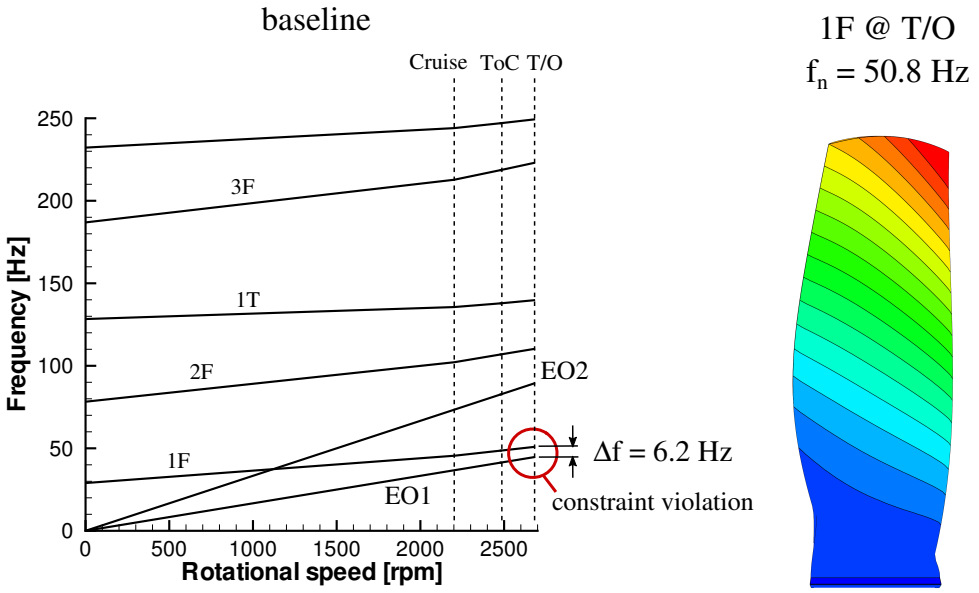


Figure 8.27: Campbell diagram of the baseline design.

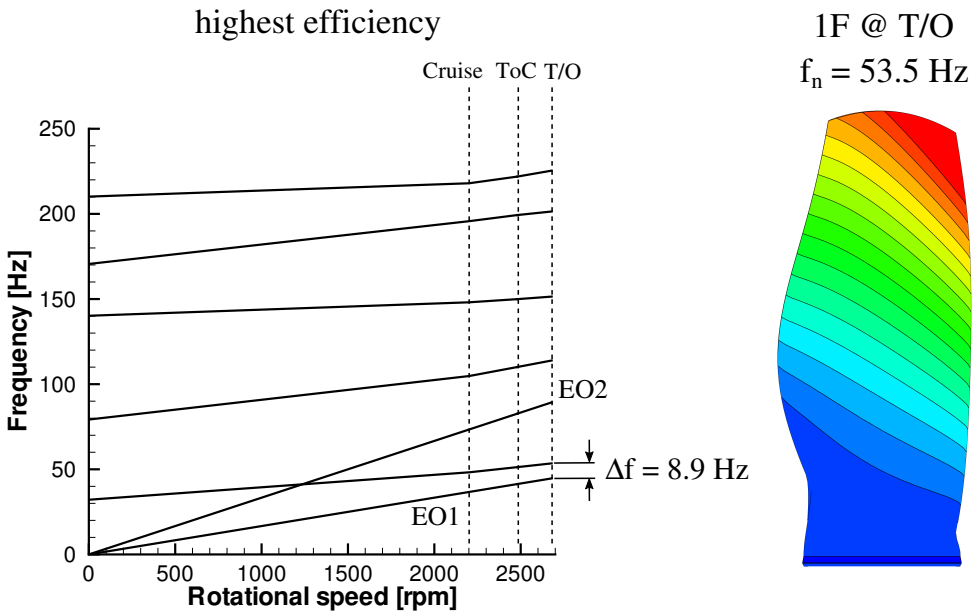


Figure 8.28: Campbell diagram of the optimised design with highest efficiency.

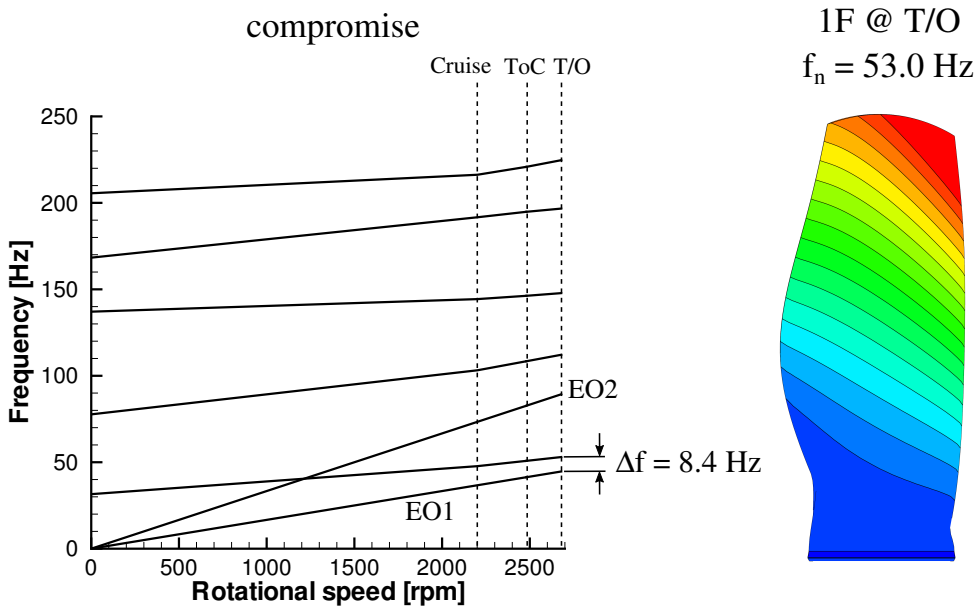


Figure 8.29: Campbell diagram of the optimised design with compromise between efficiency and weight.

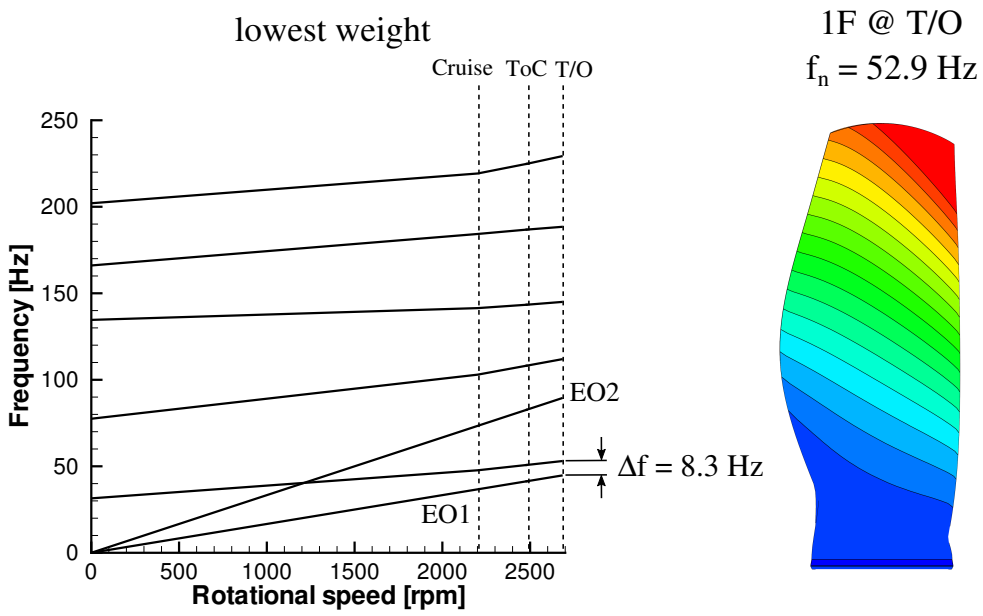


Figure 8.30: fig:campbell-diagram-he

Figure 8.31: Campbell diagram of the optimised design with lowest weight.

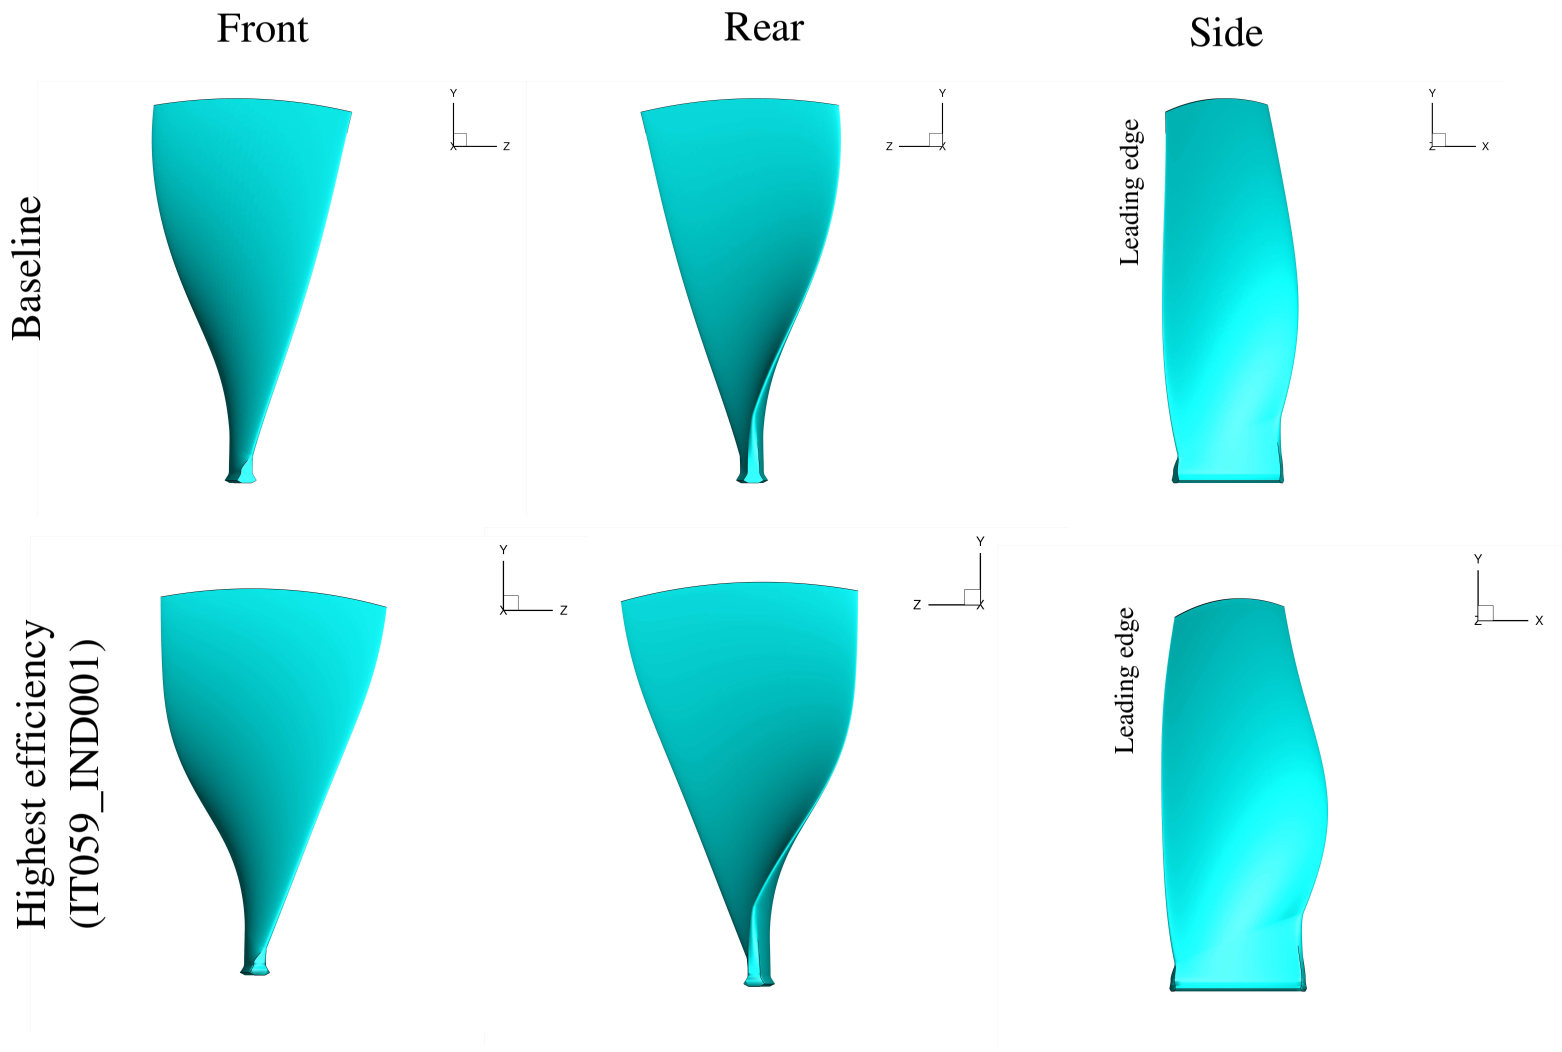


Figure 8.32: Blade geometries of baseline and optimised design with highest efficiency.

Figure 8.33: Blade geometries of the compromise and lowest weight design.

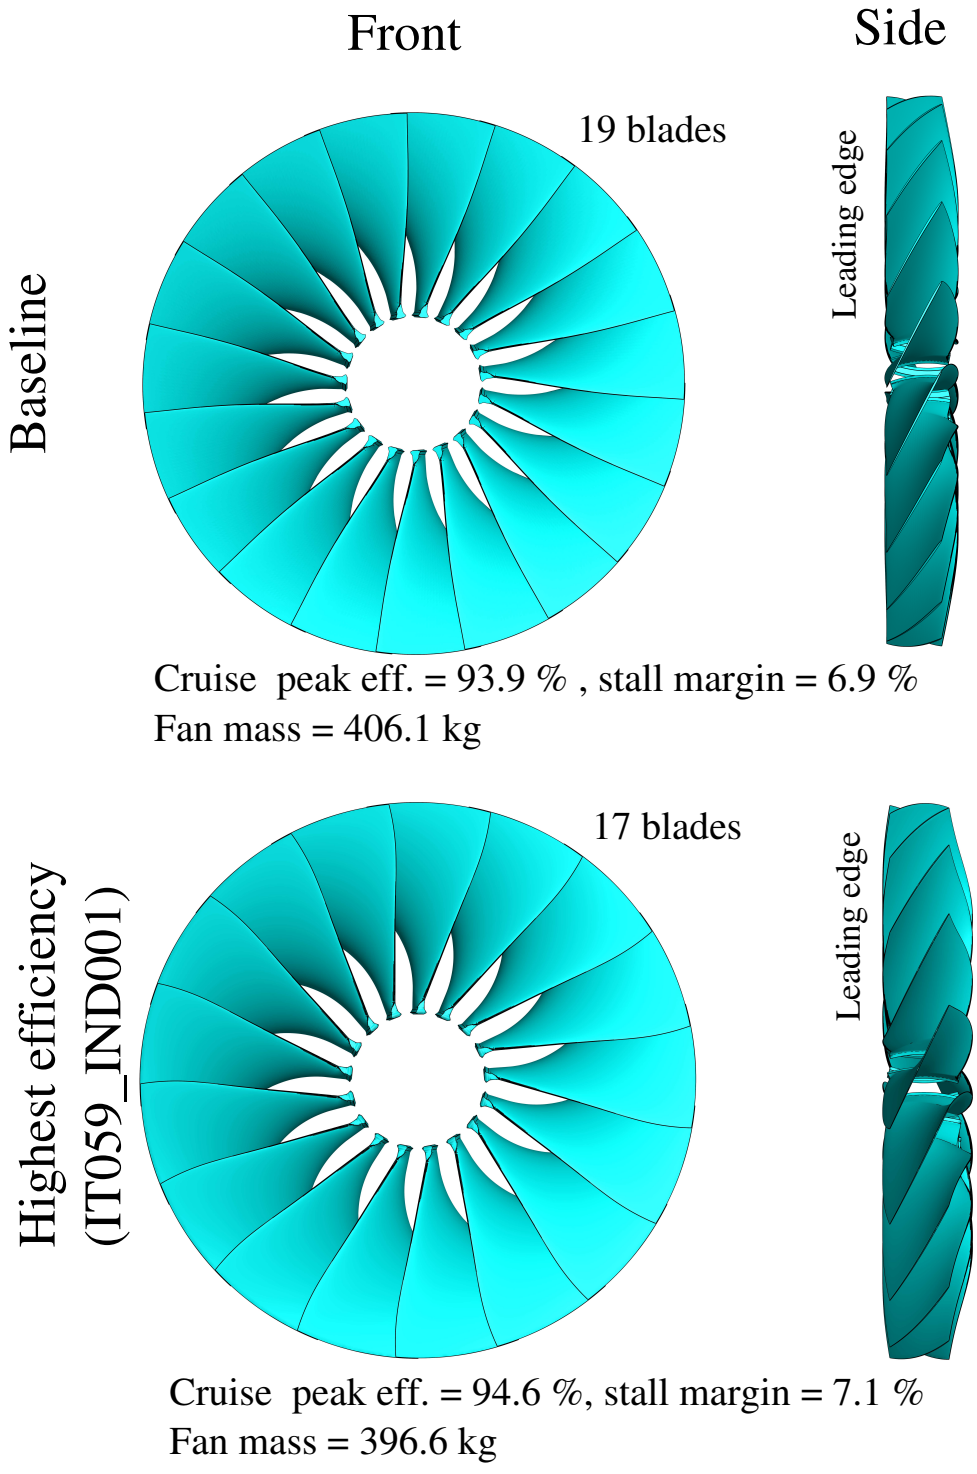


Figure 8.34: Full fan assemblies of the baseline and highest efficiency optimised design.

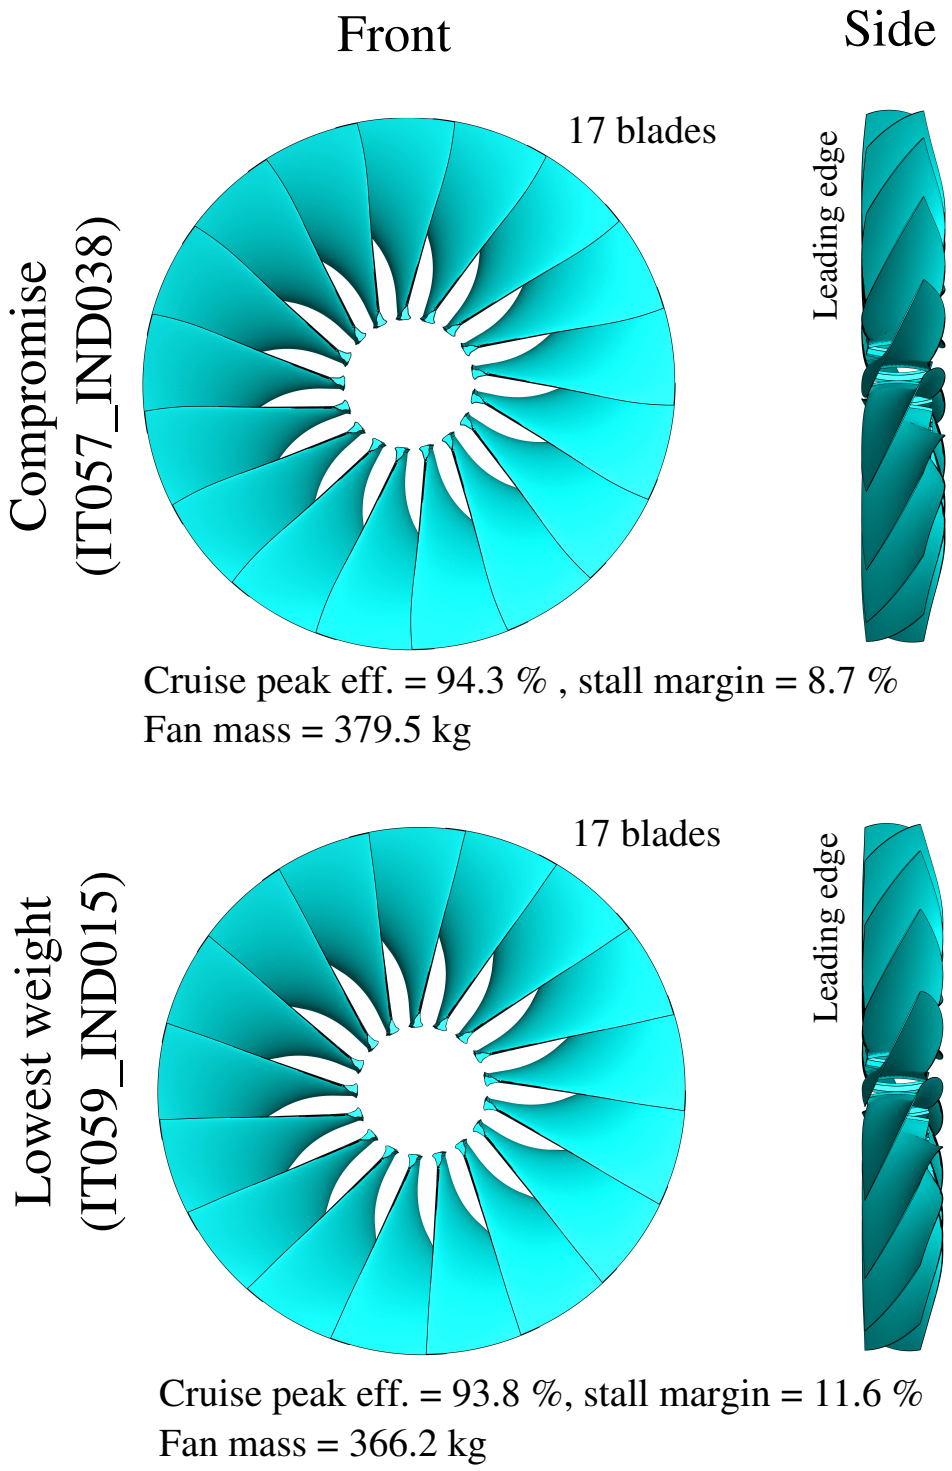


Figure 8.35: Full fan assemblies of the compromise and lowest weight design.

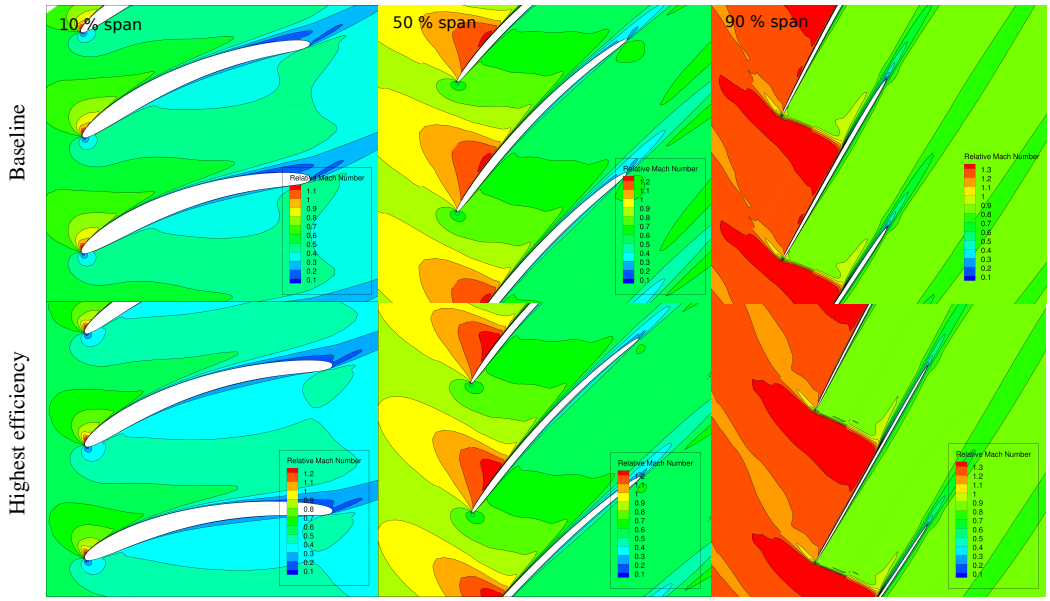


Figure 8.36: Relative Mach number contours at 10, 50 and 90 percent span for baseline and optimised design with highest efficiency at the cruise design point.

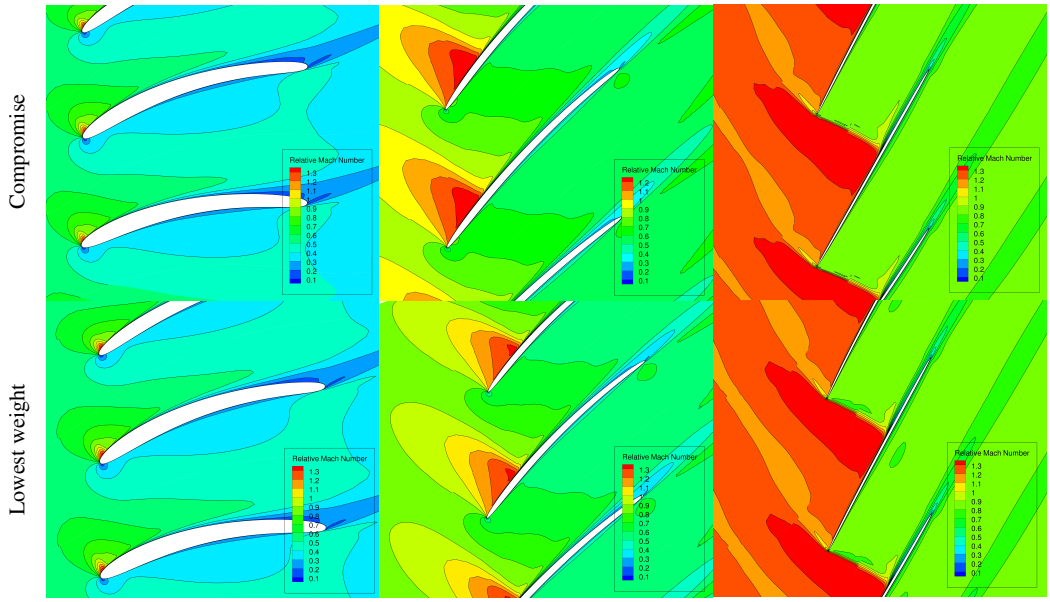


Figure 8.37: Relative Mach number contours at 10, 50 and 90 percent span for the compromise and lowest weight optimised design at the cruise design point.

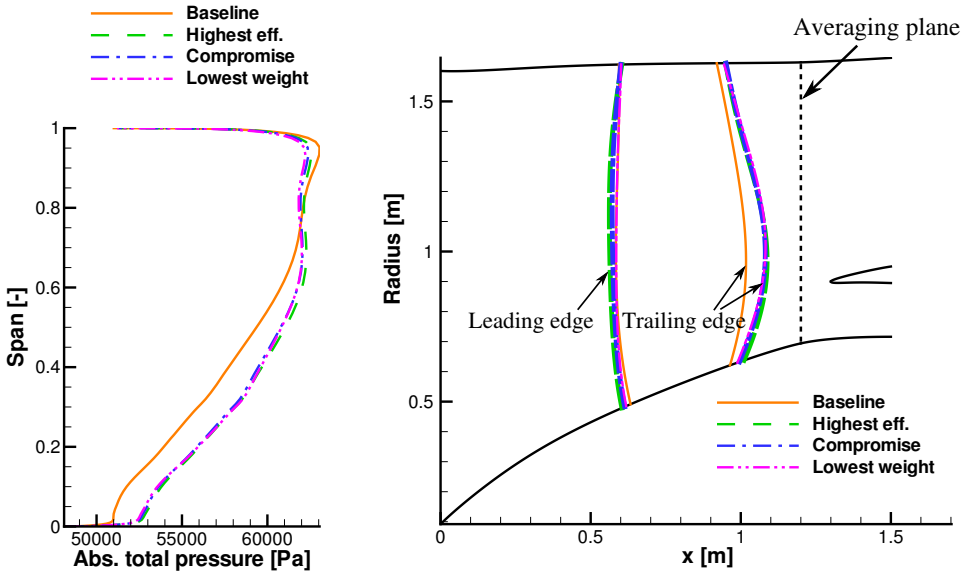


Figure 8.38: Pitch averaged spanwise distribution of absolute total pressure at the fan exit, upstream of the bypass splitter, at the cruise design point.

9

Analysis of Aerodynamics-Structural Dynamics Coupling for Blade Flutter Prediction

Self-excited aeroelastic instabilities, known as flutter, are a dangerous phenomenon for fan blades of modern civil aero-engines. Flutter occurs when small blade perturbations in certain structural vibrational modes become amplified by the unsteady flow field surrounding the blades. If the so introduced excess vibrational energy cannot be dissipated by some form of structural damping, the vibration amplitude grows rapidly to reach harmful values leading to the destruction of the component. The severity of a flutter incident mainly stems from the fact that it can appear suddenly and its consequences can be disastrous. In aero-engines, flutter mostly occurs in fan blades, which are characterised by highly three-dimensional, large slender structures, but also low pressure turbine blades can suffer from this instability.

While flutter has long been known to be of importance for fan blades, the trend of modern engine design towards increasing fan diameters to enable so-called ultra-high bypass ratios, as well as the simultaneously growing application of low-density composite materials leads to an increased susceptibility of the blades to flutter. But not only does its likelihood increase; a corresponding decrease of the structure-to-air mass ratio resulting from current design trends also influences in which form flutter may occur and the prediction methods that can be used to accurately predict its onset.

To date, turbomachinery flutter has been treated mainly as a decoupled phenomenon, based on an early proposed method by Carta [28], most commonly known as the *energy method*. In this method, aeroelastic coupling effects are assumed to be negligible, which implies that unsteady aerodynamic forces acting on the blade surface are small enough to not alter the vibrational behaviour of the blade. This means, in the aeroelastic system with the blade subjected to an airflow, that the blade

is assumed to vibrate purely in one of its structural modes. The implication for the prediction of fan flutter is that the feedback from the fluid to the structure does not need to be taken into account, and the flutter behaviour can be predicted by imposing the structural free vibration in terms of bladed disk modeshape and eigenfrequency in an unsteady flow computation, hence the term *decoupled*. To be more precise, the energy method is based on two major assumptions: First, the aeroelastic vibration mode is identical to one of the structural blade modes (i.e. modal coupling does not occur). And second, the aeroelastic response frequency is identical to the blade's natural frequency.

While those assumptions are valid for turbomachinery blades with high mass ratios and stiffness, their validity for large, low density and low stiffness fan blades is less clear. Despite the existence of a large body of research work on the development and application of various flutter prediction methods, publications targeting the validity of the energy method for lightweight blade designs are relatively scarce and are often limited to two-dimensional cases or subsonic flows. This is not least because of the significant computational cost associated with fully coupled fluid-structure interaction analyses for three-dimensional transonic cases, that are needed as a reference for the comparison with the decoupled approach. This cost is further increased by the need to compute a large number of vibration cycles for an accurate determination of the aeroelastic response frequency.

This chapter presents a systematic comparative study contrasting the results of the decoupled energy method with those of a strongly coupled fluid-structure interaction approach for a three-dimensional transonic fan blade. The blade mass ratio and blade stiffness are parametrically varied in order to assess their respective impact on the strength of aeroelastic coupling in terms of the aeroelastic frequency shift, i.e. the difference between the aeroelastic response frequency and the natural frequency of the blade. This strategy was chosen over the usage of a more specific material system, such as for instance a woven fabric composite material, in order to avoid an over specification of the case with a consequent loss of generality. This work focuses on single mode fan flutter, hence it neglects the modal coupling effects. For reasons of computational cost, the results shown are limited to operation at an interblade phase angle of zero.

The test case is NASA's Rotor 67 [158], which was been widely used in the framework of computational aeroelasticicy (e.g. [106], [135], [155], [96], [97], [42], [60], [166], and [30]), despite the lack of experimental flutter data in the open literature. The rotor has 22 blades with an aspect ratio of 1.56 (average span over root axial chord).

Tip solidity (chord to pitch ratio) is 1.29 and the tip relative inlet Mach number at the aerodynamic design point is 1.38.

In the next section the aeroelastic system and the derivation of the structural modal model that is used for both flutter prediction methods applied in this work will be presented. The structural finite element model, which forms the basis for the modal model is presented in the subsequent section. In this section also the parametric material changes are explained, that are used to alter the blade mass ratio and stiffness. A short explanation of the applied CFD solver, mesh and an introduction to the investigated operating points follows in section 9.3. The decoupled and coupled flutter prediction methods applied are detailed in section 9.4. This is followed by the results and the conclusions of the chapter.

9.1 Aeroelastic system and modal equation

A discrete approximation of the continuous aeroelastic system comprising an elastic blade and surrounding air flow can be given by the equation of motion

$$\underbrace{[\mathbf{M}]\ddot{\boldsymbol{x}} + [\mathbf{C}]\dot{\boldsymbol{x}} + [\mathbf{K}]\boldsymbol{x}}_{\text{structure}} = \underbrace{\boldsymbol{F}_d \ (+\boldsymbol{F}_f)}_{\text{fluid}} \tag{9.1}$$

where $[\mathbf{M}]$, $[\mathbf{C}]$ and $[\mathbf{K}]$ are respectively the mass, structural damping and stiffness matrices, while $\boldsymbol{x}$ represents the vector of blade displacements and $\boldsymbol{F}_d$ the vector of blade motion dependent aerodynamic damping forces on the blade surface. For the purposes of this work, externally forced excitations originating from blade movement independent forcing, denoted above as $\boldsymbol{F}_f$, are omitted. Those may be due to, for example, inlet flow distortions and constitute the problem of forced response.

The equation of motion, Eq. (9.1), represents a coupled system of (in the most general form non-linear) equations of size equal to the number of degrees of freedom chosen for the system discretisation. In the current state-of-the-art of CFD and FEM modelling this number is typically of the order of several millions. A solution of the full system of equations is therefore often prohibitively expensive for flutter analyses in an industrial setting.

However, it is well known that the structural response of turbomachinery bladings can be sufficiently well approximated by only a few significant low order modes. By making use of the natural orthogonality of the free vibration modeshapes $[\mathbf{\Phi}] = [\boldsymbol{\phi_1}, \boldsymbol{\phi_2}, ..., \boldsymbol{\phi_N}]$ with respect to the mass and stiffness matrices, the set of N coupled equations in Eq. (9.1) in physical space can be transformed to a set of N decoupled

equations in modal space, where each equation describes the system state for one blade mode in separation. Defining a modal mass $[\mathbf{m}] = [\mathbf{\Phi}]^T[\mathbf{M}][\mathbf{\Phi}]$, modal damping $[\mathbf{c}] = [\mathbf{\Phi}]^T[\mathbf{C}][\mathbf{\Phi}]$ and modal stiffness $[\mathbf{k}] = [\mathbf{\Phi}]^T[\mathbf{K}][\mathbf{\Phi}]$, and using the coordinate transformation $\boldsymbol{x} = [\mathbf{\Phi}]\boldsymbol{q}$ with $\boldsymbol{q}$ being the vector of modal displacements, the single degree of freedom modal equation for the first blade mode is given by

$$m\ddot{q} + c\dot{q} + kq = f_d \tag{9.2}$$

The modal aerodynamic damping force can be obtained by pre-multiplication of the physical aerodynamic damping forces with the transpose of the modeshape vector $f_d = \boldsymbol{\phi}^T \boldsymbol{F_d}$. For simplicity, modeshapes are typically scaled to unity modal mass $m = 1$, which results in the modal stiffness becoming the square of the natural frequency $k = \omega_n^2$. The structural damping matrix can be similarly decoupled in case of proportional damping. However, in this work, structural damping is omitted in all computations, as is a customary practice to yield a conservative estimate of flutter stability. In practice, fluid and structure discretisations are normally not coincident, and hence the transfer of quantities between solid and fluid domains require an interpolation in the non-matching mesh interface on the blade surface. A thin plate spline radial basis function is applied in this work for that purpose. The modal equation, Eq. (9.2), is the basis for all flutter computations presented in this work.

9.2 Structural model

The free vibration modeshapes and eigenfrequencies of the blade required for the construction of the modal equation, Eq. (9.2), were computed using a three-dimensional finite element code ([41]). The blade was discretised in space using 17512 quadratic tetrahedral elements. Centrifugal stiffening was taken into account by augmentation of the stiffness matrix with the stiffness contribution from a preceding linear static analysis subject to centrifugal loading. The blade was set to rotate at design rotational speed of 16043 rpm. Nodes on the hub surface were constrained in all degrees of freedom in order to model a clamped blade. The finite element mesh as well as the contour of the first bending mode, that is used for all flutter computations in this work, is shown in Fig. 9.1.

The rotor 67 geometry used here corresponds to the running blade geometry reported in [158]. Hence we did not perform a hot-to-cold transformation to recover the blade manufacturing geometry for the structural analyses. Consequently, this geometry is subsequently assumed to represent the cold blade geometry.

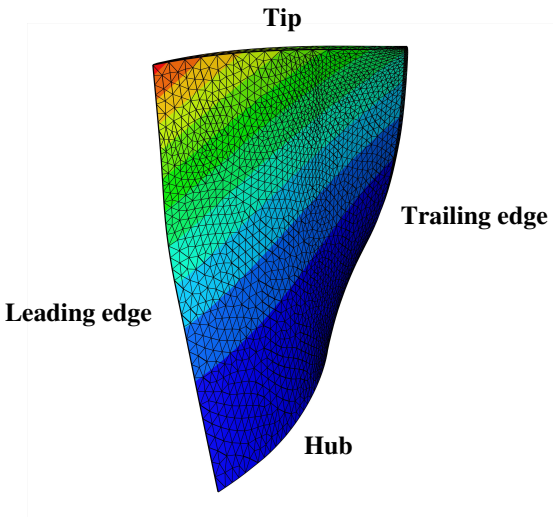


Figure 9.1: Solid mesh and mode shape contour of the first bending mode at design speed, calculated with the baseline titanium alloy.

Mass ratio and blade stiffness were altered by changing respectively the blade density and Young's modulus of an isotropic material model. A baseline set of properties was chosen to represent a typical titanium alloy. The blade density and Young's modulus were altered separately while keeping the other respective quantity constant. This resulted in five blade configurations, which are listed in Table 9.1 along with the corresponding material properties used.

The mass ratio is a commonly used indicator for the amount of aeroelastic coupling in a system and is defined as the ratio of blade mass m_b over the mass of air of a blade enclosing cylinder

$$\mu = \frac{m_b}{\pi(c_{tip}/2)^2 \rho_\infty h} \tag{9.3}$$

Table 9.1: Blade material properties used for the change of mass ratio and blade stiffness. Poisson's ratio = 0.27.

Designation	**Mass density** $[kg/m^3]$	**Young's modulus** $[GPa]$	**Mass ratio**
Baseline (bl)	4110	125	178
Low mass ratio (lm)	797	125	34
High mass ratio (hm)	23040	125	1000
Low stiffness (ls)	4110	60	178
High Stiffness (hs)	4110	205	178

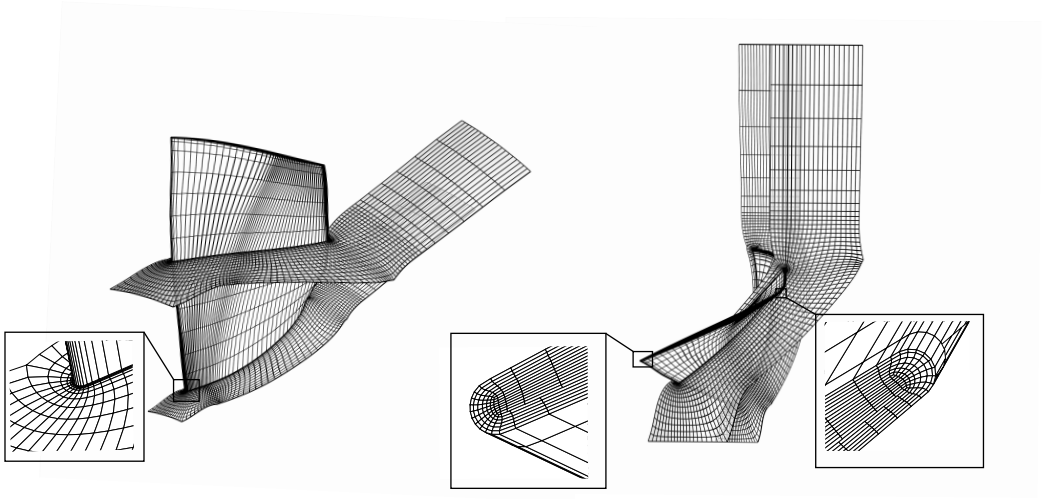


Figure 9.2: Overview of CFD mesh. Every second mesh point is shown.

where c_{tip} is the tip chord length, ρ_∞ is the undisturbed air density and h is the blade height. High mass ratios typically translate to low levels of aeroelastic coupling, hence representing a situation in which the aeroelastic response is mainly influenced by the structural dynamics while the impact of the unsteady aerodynamic forces is small. The mass ratio range shown in Table 9.1 was chosen to include a low mass ratio configuration to approximate values of modern lightweight aero-engine fan blades (e.g. a composite sandwich construction), as well as one unrealistically high mass ratio to investigate the agreement of both flutter prediction methods towards a high mass ratio limit.

9.3 CFD solver, mesh and operating points

For small perturbations, the unsteady flow is often assumed to behave linearly with respect to the vibration amplitude, which has motivated the usage of time-linearised methods for its solution. The time-linearised flow equations can be solved in the frequency domain, which is computationally much more efficient than a corresponding solution of the unsteady flow equations in a time marching fashion (time domain). However, in transonic flows, even small oscillations can lead to non-linear flow effects due to the presence of strong shocks and/or boundary layer separation (e.g. [70], [27]). Non-linearities manifest themselves in the form of higher flow harmonics which alter the time averaged unsteady flow from its steady counterpart. The non-linear harmonic method ([61], [115]), has been developed as an extension to linear frequency domain

methods, allowing for the interaction of higher harmonics to be taken into account. In order to assess the general capabilities of the utilised flow solver for the prediction of unsteady flows subject to oscillating blades using the non-linear harmonic method, a validation was performed, as detailed in Appendix A. However, in order to avoid that any differences in flow modelling might affect the comparison of decoupled and coupled flutter results, full non-linear unsteady time-domain computations are used for all flutter computations shown in this chapter.

The unsteady RANS equations are solved in time using a dual-time stepping approach. Turbulence closure is obtained using the one-equation turbulence model of Spalart and Allmaras [154]. The details of the flow solver were given in Chapter 5.

The fluid domain is discretised with a multi-block structured mesh consisting of 0.3 million grid points with an average non-dimensional height of the wall-adjacent cell of $y^+ = 3.6$. The tip gap is discretised with five grid points in the radial direction. Figure 9.2 gives an overview of the mesh with details of the leading edge and blade tip discretisations. The computations were performed for two operating points on the design speedline - one at peak efficiency and one at a higher loading, off-design condition at 96% of the design point mass flow. The experimental measurements of the steady flow performance reported in [158] also include a near stall operating point at 93% of the design point mass flow. However, this operating point was not considered in this work due to convergence reasons.

Table 9.2 summarises the values of reduced frequencies for all blade configurations at both operating points. The reduced frequency is given as

$$\nu = \frac{\omega_n C_{tip}}{V_{tip,inl,rel}} \tag{9.4}$$

where C_{tip} is the tip chord length and $V_{tip,inl,rel}$ denotes the relative tip inlet velocity. Note here the change of natural frequency due to mass ratio changes, which is a result of the constant elasticity modulus. The impact of the operating point change on the reduced frequencies can be seen to be negligible.

9.4 Description of flutter prediction methods

9.4.1 Decoupled method

The decoupled energy method introduced by [28], sometimes also referred to as work-per-cycle method, is based on the assumption that the unsteady fluid forces are small compared to the blade's inertia and stiffness forces and therefore do not alter the

Table 9.2: Reduced frequencies for all blade configurations at peak efficiency and higher load condition.

Designation	Peak efficiency	Higher load condition
Baseline	0.708	0.711
Low mass ratio	1.262	1.268
High mass ratio	0.503	0.505
Low stiffness	0.595	0.597
High Stiffness	0.821	0.825

vibrational mode in terms of modeshape and frequency. Under this assumption, the unsteady flow field and in turn the flutter stability of a blade assembly can be computed by imposing a forced low-amplitude vibrational blade motion in the unsteady fluid computation, which corresponds to one of the vibrational modeshapes and eigenfrequencies of the structure either in vacuum, or at times with the steady state gas loads imposed. Flutter stability is inferred from the work that the fluid performs on the blade over one vibration period. A positive net work input to the structure indicates flutter. Structural damping is omitted, leading to a conservative estimate of the flutter stability.

Given a forced harmonic blade motion with period T and modal amplitude $\hat{q}$, the aerodynamic work W_A is computed as the time integral of the total aerodynamic power on the blade suction side and pressure side surfaces over one vibration cycle as

$$W_A = \int_T \int_{Surface} (p\boldsymbol{n} + \tau\boldsymbol{n}) \cdot \boldsymbol{u} \; dA \; dt \tag{9.5}$$

where p is the unsteady pressure, τ is the unsteady wall shear stress, $\boldsymbol{n}$ is the unit surface normal and $\boldsymbol{u}$ is the local wall vibratory velocity. While the viscous force contribution due to shear stress is typically small in comparison to the pressure forces (and is therefore often neglected), its contribution is required for an accurate prediction of stall flutter [167] and is therefore taken into consideration.

Based on the assumption of linear flow behaviour, aerodynamic damping can be cast in the form of an aerodynamic damping ratio as

$$\zeta_A = \frac{-W_A}{2\pi\omega_n^2\hat{q}^2} \tag{9.6}$$

Flutter is consequently indicated by negative values of the damping ratio. The corresponding logarithmic decrement (log dec) is given as

$$\text{log dec} = 2\pi\zeta_A. \tag{9.7}$$

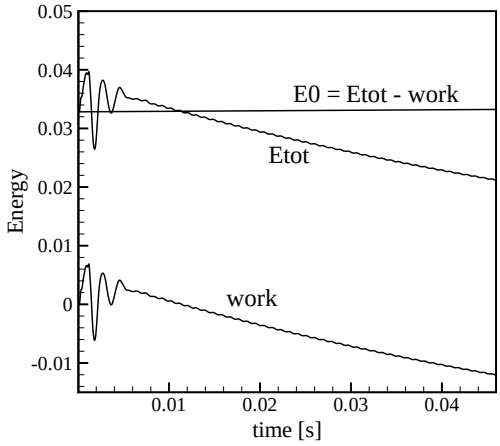


Figure 9.3: Energy conservation test of coupled computation for baseline blade at peak efficiency.

While the energy method neglects any frequency shift resulting from aeroelastic coupling effects, it can be used to predict the frequency shift. How this is achieved is shown later in the chapter.

9.4.2 Coupled method

In contrast to the decoupled method introduced above, the coupled method enables the consideration of the two-directional coupling between fluid and structure by allowing for the transfer of blade displacements as well as fluid forces between the domains. Therefore, the blade response frequency is no longer predetermined, but results from the solution of the aeroelastic system. The coupled fluid-structure interaction method applied in this work uses a strongly coupled staggered partitioned approach. Strong coupling is achieved via full convergence sub-iterations at every physical timestep [165], [47]. The scheme allows to obtain second-order time accuracy, and ensures energy conservation at the fluid-structure interface [47].

A complementary function and particular integral method is used for the solution of the modal equation, which was found to be more accurate than the commonly used Newmark algorithm [36]. Accurate energy conservation of the scheme was verified as shown in Fig. 9.3 by comparison of total structural energy Etot and aerodynamic work for rotor 67 based on the baseline full titanium structural model. A constant initial energy E0 in time indicates that the energy transfer across the fluid structure interface is accurately modelled in the computation.

Each vibration period is resolved with 50 physical timesteps. It was found that a higher number of physical timesteps per period had no significant impact on the aeroelastic response in terms of stability and aeroelastic response frequency. Aerodynamic damping is quantified in terms of the logarithmic decrement which is directly obtained from the resulting time evolution of the generalised displacements. The logarithmic decrement can be computed from the amplitudes of successive vibration cycles as

$$\log \text{ dec} = \frac{1}{n} log \frac{\hat{q}(t_0)}{\hat{q}(t_n)} \tag{9.8}$$

where n denotes the number of cycles considered (here typically the last six).

The aeroelastic response frequency is computed by means of a Fast Fourier Transform of the overall generalised displacement signal. Zero padding was applied in order to obtain a frequency resolution of approximately 0.025 Hz. The aeroelastic frequency shift, i.e. the percentage difference of the aeroelastic response frequency f from the natural vibration frequency f_{eig} then follows from

$$\text{freq. shift } [\%] = \frac{f - f_{eig}}{f_{eig}} \cdot 100 \tag{9.9}$$

In the peak efficiency computations, the blade vibrational motion is initiated by a pressure impulse, while at the higher load condition the unsteady aerodynamic forces were found to be sufficient to initiate the blade movement. Similarly to the decoupled computations, structural damping was neglected in all coupled computations.

9.4.3 Static aeroelastic blade untwist

Blade stagger changes due to static aeroelastic blade untwist were shown to significantly affect the aerodynamic damping characteristics of transonic fans (e.g. [155]). While the blade automatically assumes its aeroelastic equilibrium position (if existing) in a coupled computation, it has to be imposed in the decoupled computations in order to avoid that differences in blade untwist affect the comparison of the aerodynamic damping results. For a single mode modal model this translates to the determination of the first mode's generalised displacement at which static aeroelastic equilibrium occurs. To correctly account for this effect, the average generalised displacement to be imposed in the decoupled computations was predetermined for each blade configuration and operating point by a coupled untwist computation.

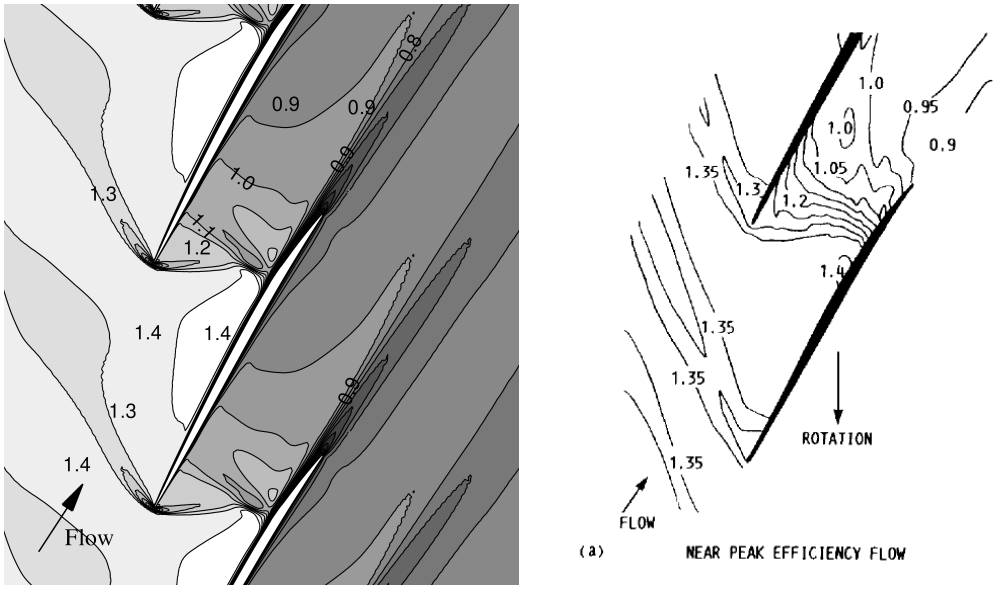


Figure 9.4: Relative Mach number contours at 90 percent blade span for peak efficiency operation. Left: CFD, right: Experiments [158].

9.5 Comparative results and discussion

9.5.1 Global aero-damping comparisons and frequency shift results

9.5.1.1 Peak efficiency

The steady flow relative Mach number contours at 90% blade span for the peak efficiency operating point are shown in Fig. 9.4, along with the corresponding experimental results from [158]. The CFD results are shown for the running design point geometry of rotor 67 as obtained from the publication, hence no additional blade untwist was applied here for the comparison to the experimental results. The flow field served as the initial solution for the coupled untwist computations which were in turn used as starting solutions for the flutter computations.

It can be observed that a reasonably good agreement is obtained between CFD and experiment. The flow field is characterised by a weak oblique shock wave attached at the leading edge of the blade which extends slightly downstream into the covered blade passage where the shock tilts to become normal on the suction side surface.

In Fig. 9.5 the tip leading edge displacements (max blade displacements) resulting from the coupled computations for all blade configurations are shown in terms of percentage tip chord. In total 30 vibration cycles were computed for each configuration

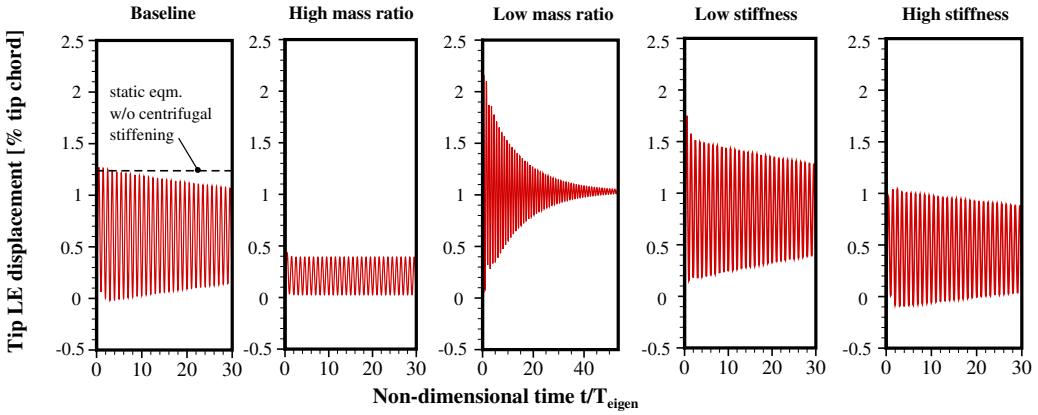


Figure 9.5: Coupled time evolution of the tip leading edge displacements at peak efficiency.

with the exception of the low mass ratio case for which the computations were continued for a total of 52 vibration cycles to obtain an accurate frequency response. It can be seen that all blade configurations except for the high mass ratio configuration are stable, indicated by a decaying vibration amplitude with time. Neutral stability is observed for the high mass ratio case, indicated by a constant vibration amplitude.

The global comparison of the coupled and decoupled aerodynamic damping results for mass ratio and stiffness variations, along with the corresponding coupled frequency shifts is shown in Fig. 9.6. The plots reveal straight away that coupled and decoupled log dec predictions show an excellent agreement for all mass ratios and stiffness values under consideration. In Fig. 9.6 b) the aeroelastic frequency shift is seen to follow an exponential increase towards low mass ratios. As would be expected, the frequency shift is zero at the highest mass ratio. However, even at a mass ratio of 34 the frequency shift remains relatively small with a value of about -1.25%. Interestingly, aeroelastic coupling effects thus do not appear to be significant at this condition.

Under stiffness variations, Fig. 9.6 c), the blade can be seen to be overall stable, with a reduction of stability towards increasing blade stiffness. This result might appear to be somewhat counterintuitive, as stiffness and associated reduced frequency increases are usually understood to increase blade stability (with the corresponding reduced frequency being a frequently used design parameter). However, for a stable blade, frequency increases can lead to decreasing blade stability, as was similarly observed in [168]. In terms of aeroelastic coupling, the effects of stiffness variations appear to be negligible for the stiffness range considered, leading to a flat curve for the frequency shift as shown in Fig. 9.6 d).

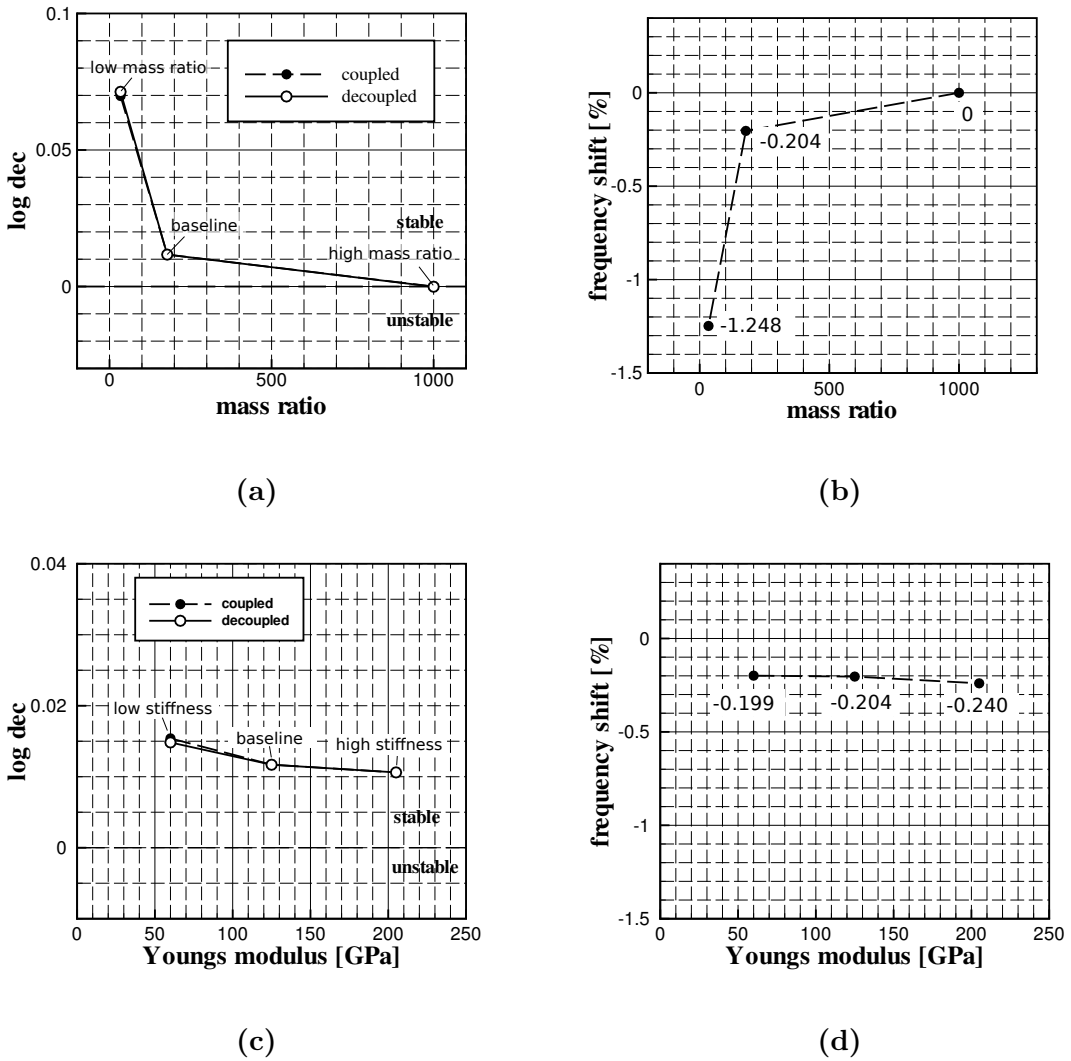


Figure 9.6: Comparison of decoupled and coupled predictions of aeroelastic stability (left) and frequency shift (right) over mass ratio and stiffness changes at peak efficiency.

9.5.1.2 Higher load condition

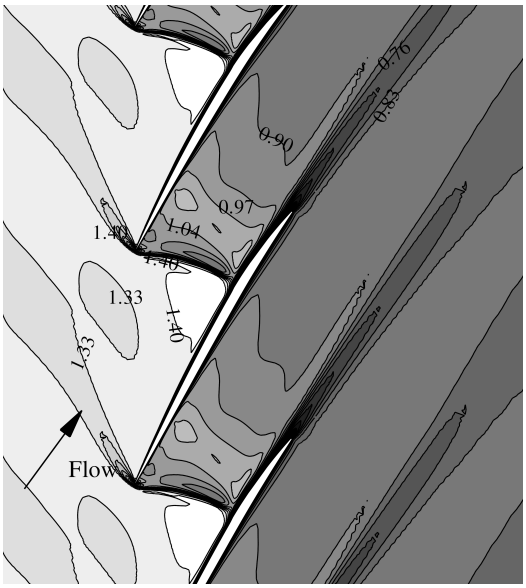


Figure 9.7: Relative Mach number contours at 90 percent blade span for operation at the higher load condition.

The steady relative Mach number contours for operation at the higher load condition are shown in Fig. 9.7. Experimental measurements at this operating condition are not available. A comparison to the steady flow field at peak efficiency, Fig. 9.4, shows that increased back pressure has now caused the passage shock to move upstream towards the leading edge resulting in a strong and nearly normal shock.

The corresponding coupled time evolutions of the tip leading edge displacements for all blade configurations are depicted in Fig. 9.8. A comparison of the plots with the peak efficiency results, Fig. 9.5, shows that the operating point change had an overall destabilising effect. Increasing vibration amplitude with time shows that the high mass ratio and low stiffness configurations are now unstable, while the baseline configuration has become neutrally stable.

The global comparison of decoupled and coupled values of aerodynamic damping and the associated frequency shifts for mass ratio and blade stiffness variations are shown in Fig. 9.9. The agreement between both flutter prediction methods can be seen to be of a similarly good level as previously observed for the peak efficiency operation. The overall system destabilisation with respect to the peak efficiency results is more clearly visible in these plots. However, unstable and stable behaviour, as well as the neutrally stable baseline behaviour are accurately predicted by the

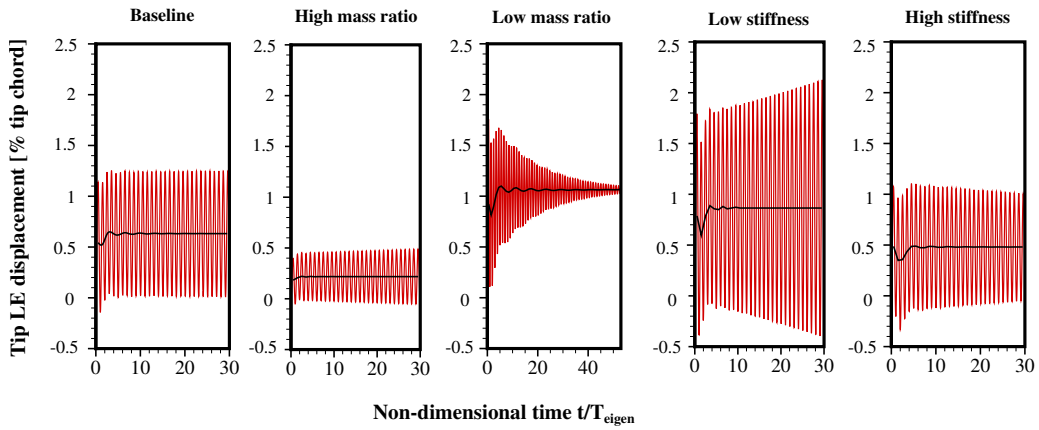


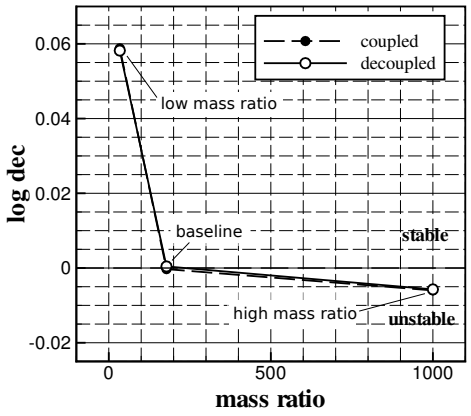
Figure 9.8: Coupled time evolution of the tip leading edge displacements at the higher load condition.

energy method. The flutter mode observable in Fig. 9.9 c) for the low stiffness design can be characterised as part speed stall flutter, which typically occurs for fan blades in the first bending mode at a raised working line (increased loading) and part speed operation (i.e. reduced stiffness).

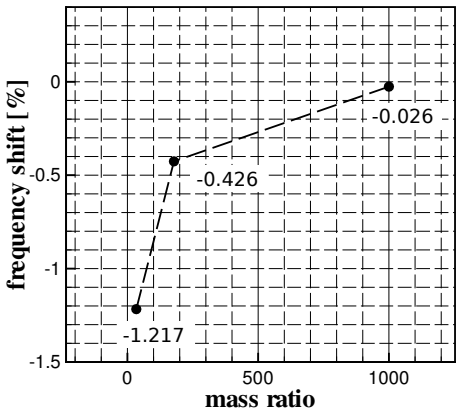
A comparison of Fig. 9.9 b) and Fig. 9.6 b) shows that the overall frequency shift curve has not changed drastically due to the change of the operating point. The biggest effect can be seen to occur for the baseline configuration, where the frequency shift roughly doubled. However, the impact on the low mass ratio case is minor, with an actual small frequency shift reduction with respect to the corresponding peak efficiency result. A small frequency shift can now also be seen to occur for the high mass ratio configuration. The sensitivity of the frequency shift towards stiffness variations (Fig. 9.9 d) can be observed to have slightly increased in comparison to the peak efficiency trend. However, it still remains small on a global level.

9.5.2 Discussion of frequency shift

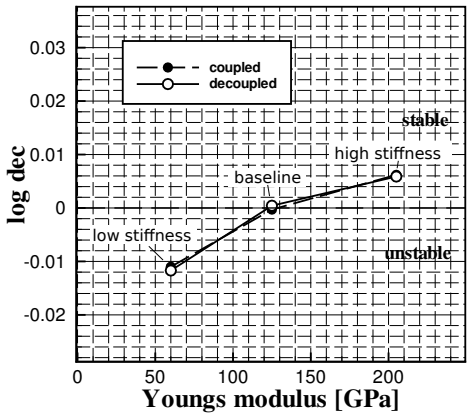
It was seen in the results presented above that despite small values of mass ratio, the aeroelastic coupling strength in the form of a frequency shift did not appear to be significant. As a result, a very good agreement was obtained between flutter stability predictions by decoupled and coupled methods. In view of the frequency shift results reported elsewhere in the literature for comparably low mass ratios (e.g. [134], [109]) this result was slightly surprising and motivated a more detailed investigation into the mechanisms governing the frequency shift, which will be discussed next.



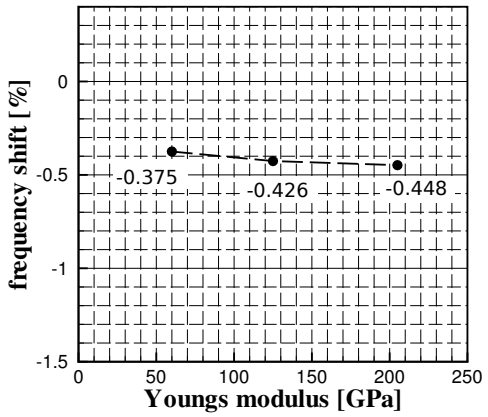
(a)



(b)



(c)



(d)

Figure 9.9: Comparison of decoupled and coupled predictions of aeroelastic stability (left) and coupled frequency shift (right) over mass ratio and stiffness changes at the higher load condition.

Frequency variations in aeroelastic systems can generally be attributed to two effects, from which the first is viscous damping (or an equivalent damping force component), and the second is an alteration of the dynamic mass of the system resulting from the presence of the fluid. The theoretical basis of the former effect is well known and is given by the relation

$$\omega_d = \omega_n \sqrt{1 - \zeta^2} \tag{9.10}$$

representing the natural frequency of an underdamped system in dependency of the natural frequency and the damping ratio. In the context of turbomachinery flutter, the aerodynamic damping ratio is typically small, and as a result the frequency shift due to damping is normally negligible.

To explain the second contribution, consider the modal equation given in Eq. (9.2). If we assume a harmonic blade motion $q = \hat{q}e^{i\omega t}$, and a harmonic modal aerodynamic damping force at the same frequency $f_d = \hat{f}_d^{\,i(\omega t+\phi)}$, the force amplitudes in the modal equation can be represented in the Argand diagram, as shown in Fig. 9.10. Note, that

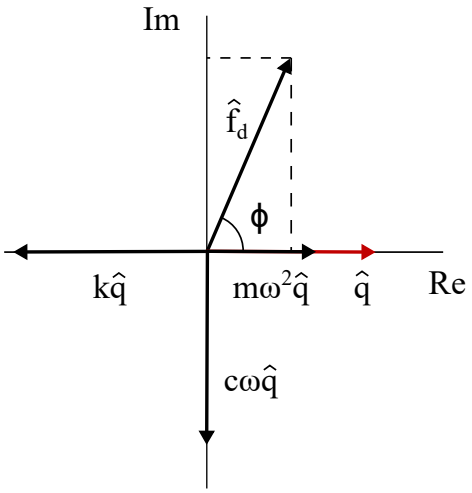


Figure 9.10: Argand diagram of the forces in the aeroelastic system.

the blade displacement is here considered to be purely real pointing in the positive direction. The force components are the modal force amplitudes of inertia force, $m\omega^2\hat{q}$, stiffness force, $k\hat{q}$, (structural) damping force, $c\omega\hat{q}$, and aerodynamic damping force, $\hat{f}_d$. For the sake of simplicity and clarity, we will drop the denotation 'modal amplitude' for the forces in the following discussion, and we further neglect structural damping.

It is well known from vibration theory, that only a force component that is out-of-phase with the blade displacement, is able to do work to the system, and therefore

can contribute to the aerodynamic damping, be it positive or negative. In case of a purely real blade displacement as shown here, this force component is the imaginary component of the aerodynamic damping force $\mathrm{Im}(\hat{f}_d)$, c.f. Fig. 9.10, which contributes to the alteration of the vibration frequency via the same mechanism as for viscous damping, as shown above in Eq. (9.10). On the other hand, the aerodynamic damping force component that is in-phase with the blade displacement, $\mathrm{Re}(\hat{f}_d)$, can be seen to be similarly in-phase with the inertia and stiffness forces. It therefore effectively resembles the force contribution of an *added mass* in the aeroelastic system, which in turn affects the system's response frequency.

Real and imaginary components of the aerodynamic damping force can be easily determined from its phase shift and amplitude obtained either from the combined time evolution of the aerodynamic damping force and the generalised displacement or from a frequency domain solution, from which the real and imaginary components of pressure and viscous forces are directly available.

For a further explanation of this matter, consider the force equilibrium of the real components of the forces in Fig. 9.10

$$m\omega^2\hat{q} - k\hat{q} + \mathrm{Re}(\hat{f}_d) = 0. \tag{9.11}$$

As can be seen here, the only mechanism in place to obtain a force equilibrium in the case of a non-zero real component of $\hat{f}_d$, is the adaptation of the inertia force $(m\omega^2\hat{q})$ via a change of the response frequency ω, because the modal mass m and modal stiffness k are constant and the generalised displacement amplitude $\hat{q}$ linearly scales both the stiffness and inertia forces.

Such a change of frequency may alter the unsteady flow field and in turn the aerodynamic damping force $\hat{f}_d$. An iteration in the form of a coupled computation would be required to find a converged solution of the frequency shift. However, in a first approximation, it can be assumed that the aerodynamic damping force is invariant with ω, which enables to give a prediction of the dynamic response frequency using a decoupled approach.

It follows directly from Eq. (9.11)

$$\omega_{dyn} = \sqrt{\frac{k}{m} - \frac{\mathrm{Re}(\hat{f}_d)}{m\hat{q}}} \tag{9.12}$$

where $\mathrm{Re}(\hat{f}_d)$ and $\hat{q}$ can be substituted for their corresponding pendants from a decoupled computation to yield a first order approximation of the dynamic response frequency ω_{dyn}. Note here again, that $k = \omega_n^2$, $m = 1$, so that the common relation

$k/m = \omega_n^2$ holds. Very similar routes to the ones outlined here have been taken by [106] for the investigation of aeroelastic coupling effects in the framework of turbomachinery forced response computations.

It was mentioned above that the apparent increase or decrease of the dynamic mass of the system due to the in-phase component of the aerodynamic damping force can be suitably described as an added mass effect. However, it is important to note that the effect discussed here should be clearly discerned from the more commonly debated added mass effect that results from the acceleration of an attached potion of fluid in the framework of mainly incompressible hydroelastic applications (e.g. [24]). In contrast, the presently discussed apparent change of the dynamic mass of the system is due to the phase shift between the unsteady aerodynamic forces and blade displacements and results from the complex aeroelastic relations in a compressible flow. To more clearly differentiate the two physical mechanisms that are both found in the scope of fluid-structure interaction, the authors of present chapter believe that the term *aerodynamics added mass* or *aero added mass*, in short, seems to be better suited to describe the phasing induced change of the dynamic mass of the system that is the subject of the discussions presented here.

Before we use the force equilibrium in Eq. (9.11) to further investigate the factors influencing the magnitude of the frequency shift, it is sensible to validate the assumptions made above about the physical effects which are governing it. To this end, Eq. (9.12) can be used in combination with the decoupled computational results presented earlier, to yield a prediction of the frequency shift. Those can be compared to the actual values of frequency shift as obtained from the coupled computations. Such a comparison is shown in Table 9.3 for operation at the higher load condition. In addition to the decoupled frequency shift originating from the aero added mass, the frequency shift due to the viscous (out-of-phase) force component of the aerodynamic damping force is shown, resulting from Eq. (9.10). It can be seen that the decoupled frequency shift due to the aero added mass closely matches the values of the coupled computations, therefore reassuring that the frequency shift model, as well as the assumption of a frequency invariant aerodynamic damping force, are sensible for the cases considered. In addition, it can be seen that the frequency shift due to the viscous aero damping force component is several orders of magnitude smaller than the frequency shift due to the aero added mass effect, hence confirming that its contribution to the overall frequency shift can be considered as negligible.

We can now investigate the force equilibrium in further detail to understand how the force interactions are governing the actual frequency shift magnitude originating

Table 9.3: Comparison of coupled and decoupled frequency shift values at the higher load condition. Decoupled values are shown for both aero added mass and viscous damping.

	bl	**hm**	**lm**	**hs**	**ls**
Δf_c [%]	-0.426	-0.026	-1.217	-0.448	-0.375
$\Delta f_{d,aam}$ [%] (Eq.(9.12))	-0.49	-0.022	-1.36	-0.46	-0.39
$\Delta f_{d,visc}$ [%] (Eq.(9.10))	-1.1e-7	-4.6e-5	-4.4e-3	-4.9e-5	-1.5e-4

from the aero added mass effect. It is apparent at this point, that the magnitude of the real component of the aerodynamic damping force plays a key role for the frequency shift, which depends on the phase angle and the amplitude of the aerodynamic damping force. The larger the magnitude of the real component of the aerodynamic damping force, the larger will be the induced frequency shift. However, the frequency shift magnitude is established by the combined effect of all in-phase forces, and as we will show next, it can be attributed to the ratio between the real component of the aerodynamic damping force and the vibrational inertia force of the system, which we express as the force ratio

$$Fr = \frac{\mathrm{Re}(\hat{f}_d)}{m\omega_n^2 \hat{q}} \tag{9.13}$$

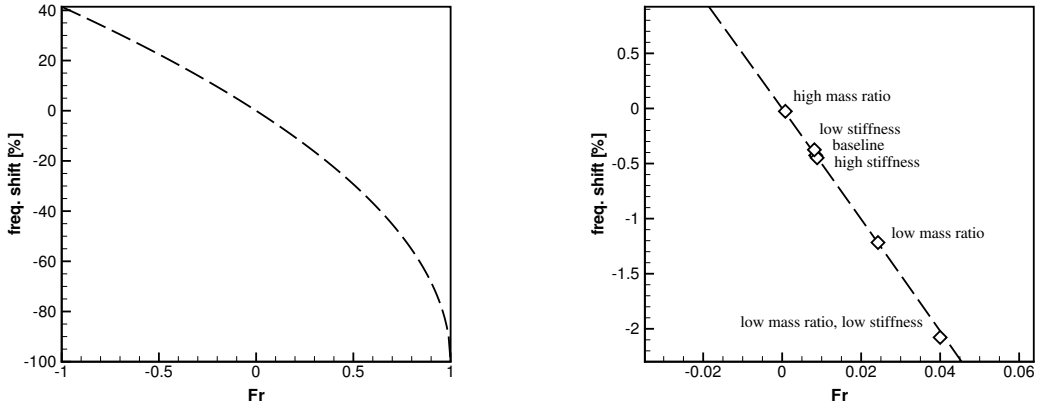
Note that the inertia force in the denominator equals the inertia force at the blade's natural frequency. Inserting Eq. (9.13) into Eq. (9.11) and solving for the frequency yields again an expression for the dynamic response frequency, but this time as a function of Fr

$$\omega_{dyn} = \sqrt{\frac{k}{m} - Fr \cdot \omega_n^2} = \omega_n \sqrt{(1 - Fr)} \tag{9.14}$$

Equation (9.14) can now be expressed in terms of a frequency shift in percent as

$$\Delta\omega \ [\%] = (\frac{\omega_{dyn}}{\omega_n} - 1) \cdot 100 = (\sqrt{(1 - Fr)} - 1) \cdot 100 \tag{9.15}$$

The variation of frequency shift with the force ratio is plotted in Fig. 9.11 a). In this figure, we can see that, in the theoretical case of the force ratio increase towards a value of one, i.e. the real component of the aerodynamic damping force would reach the same magnitude as the inertia force, the frequency shift reaches -100%. However, due to the quadratic dependency of the inertia force from the vibration frequency,



(a) **(b)**

Figure 9.11: a) Decoupled frequency shift prediction due to aero added mass as function of the force ratio. b) the same plot with coupled results for the higher load condition included.

only an approximately 41% increase of the frequency would be needed to compensate for the case of $Fr = -1$. Again, those are theoretical limits.

Values of force ratio and frequency shift obtained from the coupled rotor 67 computations can now be included in the plot as shown in Fig. 9.11 b) in which the higher load condition results are depicted. As can be seen, the predicted dependency of the frequency shift on the force ratio is indeed very well matched by the coupled results. Results for the low stiffness and high stiffness configurations are clustered closely around the baseline design, while the mass ratio changes have a higher impact on the force ratio and consequently the frequency shift. However, the overall values of the force ratio remain rather small for all blade configurations considered thus far in this chapter, where either mass ratio or blade stiffness were altered in isolation. This strategy was chosen in order to allow an investigation of the respective impact of each parameter in isolation. However, in reality, blade mass and stiffness are typically not decoupled from one another. The notable exception might be a composite material in which the directional stiffness of the material can be tailored without necessarily having a significant effect on the blade mass. Nevertheless, a low mass ratio blade, e.g. a blade made from carbon fibre reinforced plastic, would normally have a lower stiffness than a corresponding integral titanium blade, which should be taken into consideration in the assessment of potential aeroelastic coupling effects. Based on the derivations presented above, c.f. Eq.(9.13) to Eq. (9.15), it can be expected that

a corresponding low mass ratio and low stiffness blade would potentially show an increased frequency shift magnitude. An additional case was therefore added to the investigations, as will be shown in the next subsection.

9.5.3 Low mass ratio, low stiffness case

An additional blade configuration with the same mass ratio as the previously discussed low mass ratio configuration, but with a decreased stiffness to match the eigenfrequency of the baseline design is presented in this subsection. The material properties for this low mass ratio, low stiffness configuration are listed in Table 9.4.

The coupled results obtained for this blade were added to the plot shown in Fig. 9.11 b). Also this new case matches the predicted curve well, however, it can be seen that the force ratio and the corresponding magnitude of the frequency shift have increased considerably in comparison to the case where the mass ratio was altered in isolation. The frequency shift magnitude originating from the coupled computations has increased to about 2% for this new case. It can also be seen that the point is deviating slightly from the decoupled prediction, which might be an indicator that the frequency shift begins to affect the aerodynamic damping force.

The corresponding logarithmic decrement values resulting from the decoupled and coupled computations for the low mass ratio, low stiffness configuration are shown in Table 9.5 and to allow for a straightforward comparison, the values for the low mass ratio case (mass ratio change in isolation) were added to the table. It can be seen that the reduction of the blade stiffness leads to a destabilisation with respect to the isolated mass ratio change. In addition, the stability prediction mismatch between the coupled and decoupled methods can now be seen to be of significance for the low mass ratio and low stiffness blade. The decoupled method is seen to overpredict the stability by 27.3%. A comparison of the coupled and decoupled frequency shift values included in Table 9.5 for the two cases shows here more clearly the capability of the decoupled method to predict the frequency shift with good accuracy, therefore

Table 9.4: Material properties used for the low mass ratio, low stiffness blade configuration.

Mass density $[kg/m^3]$	**Young's modulus** $[GPa]$	**Poisson's ratio**	**Mass ratio**
797.15	24.2	0.27	34

Table 9.5: Comparison of decoupled and coupled logarithmic decrement predictions and frequency shift for the low mass ratio, low stiffness configuration as well as the low mass ratio configuration at the higher load condition. The decoupled frequency shift shown is due to the aero added mass effect.

	log dec			**freq. shift [%]**	
	coupled	decoupled	difference [%]	coupled	decoupled
lmls	0.022	0.028	27.3	-2.078	-2.1
lm	0.059	0.058	-1.7	-1.217	-1.36

providing a useful means to estimate the applicability of the decoupled approach for stability predictions.

This result illustrates that, while the mass ratio can be an indicator for the potential strength of aeroelastic coupling effects, the manifestation of aeroelastic coupling in the form of a frequency shift is mainly resulting from the ratio of the in-phase component of the aerodynamic damping force and the vibrational inertia force of the blade. While the former is dependent on the operating conditions and the vibration mode (blade and disk modeshape and frequency), the vibrational inertia force depends solely on the blade mass and the blade stiffness, which are purely structural parameters. The consequence is that the frequency shift experienced by low mass ratio blades with high stiffness might be negligible, while blades with comparable mass ratios and low stiffness (large, lightweight fan blades) can be expected to show a higher susceptibility to aeroelastic coupling effects. As a result, stability prediction differences between decoupled and coupled methods may become of further significance as well, which should be of consideration in the analysis and design process.

9.6 Summary

In this chapter comparisons of decoupled and coupled flutter analyses of NASA's transonic fan rotor 67 subject to parametric variations of mass ratio and blade stiffness are presented. The aeroelastic coupling strength is evaluated in terms of the shift of the response frequency from the natural frequency of the blade and its effect on the agreement between the stability predictions of the two methods is shown.

From the results obtained, the following main conclusions can be drawn:

1. Overall, decoupled and coupled stability predictions show a very good agreement for all cases in which isolated variations of mass ratio and stiffness from

the baseline, full titanium blade are performed. The corresponding frequency shift values reach maximum values of about 1.2 percent for the lowest mass ratio considered ($\mu = 34$) without negatively affecting the prediction quality of the energy method. For a combined low mass ratio and low stiffness blade configuration, which more closely resembles a realistic lightweight fan blade design, the stability prediction accuracy of the energy method is seen to decrease with a corresponding frequency shift increase to about 2 percent. For this case the energy method overpredicts the stability by about 27 percent.

2. It was shown that the aeroelastic frequency shift mainly originates from an aerodynamics added mass effect which results from the component of the aerodynamic damping force that is in-phase with the vibrational blade displacement. Consequently, the magnitude of the frequency shift is shown to depend on the ratio between this in-phase damping force component and the blade's vibrational inertia force, which is affected by blade mass and blade stiffness. The mass ratio in isolation can therefore only be an indicative measure for the expected aeroelastic coupling strength in transonic fans and other turbomachinery blades operating in high-speed compressible flows. Fan designers should evaluate carefully whether good experience with the energy method with previous blade designs can be transferred to new designs with comparable mass ratios but modified stiffness. This aspect might become of increasing importance if advanced composite blade design and manufacturing techniques are applied, which enable the targeted tailoring of blade stiffness, without having a major impact on the blade mass.

3. The decoupled method can be used to predict the aeroelastic frequency shift rapidly and with acceptable accuracy, using the technique shown in this chapter. While the prediction accuracy can be expected to decline for larger values (e.g. in excess of 2%), it provides a good indication of whether the aeroelastic coupling strength is negligible, and therefore presents a useful means to check the potential applicability of a decoupled (energy method) approach without resorting to a costly coupled method.

10 Conclusions and Future Work

10.1 Conclusions and Recommendations

This thesis aimed to introduce a multidisciplinary design optimisation framework for the efficient aerostructural design of modern high-bypass ratio aero-engine fan blades. Further, the applicability of common decoupled flutter prediction methods for modern lightweight fan blades was to be investigated. This section contains the main conclusions of the research work, which are subdivided into two parts.

10.1.1 Multidisciplinary Design Optimisation

The starting point of this work was an existing general purpose optimisation method that has been developed at the von Karman Institute. The optimisation strategy is based on a two-level approach, coupling a multi-objective Differential Evolution algorithm to an ordinary Kriging metamodel and to accurate, but computationally expensive, high-fidelity evaluation tools, which serve to continuously update the metamodel throughout the optimisation process. A strategy for the direct handling of constraints is implemented based on an adaptation of the Differential Evolution selection operator. A Pareto-raking process proposed by Deb et al. [35] is used to handle multi-objective optimisation problems.

In order to enable a multidisciplinary fan design optimisation process, a fully automated high-fidelity evaluation chain was developed in this work. The evaluation chain consists of a parametric fan blade model, a fully automated geometry generation process, automated solid and fluid mesh generations, automated CFD and CSM performance analyses and postprocessing for the extraction of relevant performance parameters for the optimiser.

A CAD-based fan blade parametrisation method was developed and implemented. The method applies spanwise variations of geometric quantities for the definition of the blade aerofoil, thus providing a straightforward link to two-dimensional through-flow design approaches commonly used in the preliminary design stage in industry. Fan blades contain complex root geometries which serve to secure the blade to the fan disk. Those represent a significant part of the blade geometry, and need to be included in the structural blade model if an accurate representation of stresses and vibrational mode shapes and natural frequencies is desired. For this purpose, the aerofoil blade generation process was augmented with a dovetail root geometry generation. A fully parametrised bypass splitter configuration was further included in the fluid domain description, in order to enable a realistic consideration of the aerodynamic design problem in accordance with the engine cycle design requirements.

To enable an accurate prediction of the relevant flow physics, a three-dimensional Reynolds-Averaged Navier-Stokes CFD solver was applied in this thesis. A commercial, structured, density-based solver was applied which uses a finite volume method. Turbulence effects were modelled using the one-equation model of Spalart and Allmaras [154]. A template-based semi-automatic multi-block structured mesh generation process was applied for the discretisation of the fluid domain.

The finite element method was applied for the evaluation of the structural fan performance in terms of its static stress and vibration response. The solid domain was discretised using an unstructured mesh consisting of quadratic tetrahedral elements. A fully automated meshing process was used for that purpose. Material models for an isotropic titanium alloy and an orthotropic composite material were applied. Geometric non-linearities were taken into consideration in the static finite element computations in order to account for large fan blade deformations. The blade's vibration response was further predicted using a finite element modal analysis, including the effect of centrifugal stiffening.

The successful application of the optimisation framework was demonstrated in two test cases in this thesis. The main conclusions that can be drawn from the work are given below.

- The implemented geometry generation process was found to be highly robust and able to model a wide range of different fan blade geometries. Among all candidate designs that were generated in the optimisations carried out in the framework of this thesis, none of the designs was found to fail due to an error in the geometry generation process.

- The optimisation system was shown to be capable to successfully handle highly constrained optimisation problems in high-dimensional search spaces. The considered cases were defined by 12 and 13 aerodynamic and structural constraints, respectively. In the first optimisation case, only three individuals in the initial database were satisfying all imposed constraints, while in the second case, not a single design was satisfying all constraints. Nevertheless, in both cases, the optimiser successfully identified the feasible region in the design space after only a few iterations; 9, and 11 iterations were respectively needed.

- The sequentially updated ordinary Kriging metamodel used in this work was generally found to show a very good learning behaviour even in the highly dimensional design spaces considered (26 and 35 design variables, respectively). Some quantities proved however difficult to predict in both optimisation cases. In particular the stress constraint (von Mises and Tsai-Wu failure criteria) and the stall margin predictions were characterised by reoccurring large prediction errors throughout the entire optimisation process. The accuracy of the metamodel predictions considerably influences the ability of the optimiser to identify the feasible region in the design space. This is particularly true for constraints that are difficult to satisfy. On the other hand, if a constraint is not overly restrictive, a larger metamodel error may still result in the generation of feasible designs.

- Despite different material models and different blade parametrisations used, both optimisations resulted in designs with a backward swept tip. It was consistently found that a forward swept tip resulted in stress constraint violations. Interestingly, no significant efficiency advantage was found in either of the two optimisation cases for designs with a forward swept tip.

- Imposing proper mass flow constraints to ensure that the designs satisfy the required range is not a trivial task. A 3% mass flow tolerance as defined in the first optimisation for the top-of-climb evaluation point mass flow was found to be too relaxed. It led eventually to a significant shift of the performance curve of the optimised design towards increased mass flows, which resulted in an unfavourable crossing with the engine working line, and consequently in a reduced stall margin. The situation was improved in the second optimisation case where the permitted top-of-climb mass flow tolerance was reduced to 2%.

- The limited out-of-plane tension strength of the applied composite material in the second optimisation led to the requirement to introduce a tilted front shank geometry in order to limit the out-of-plane tension stresses. It was ultimately observed that a tilting angle may be found which enables a dominantly in-plane stress state in the front root section. Comparably low out-of-plane tension strength is a general property of laminated composite materials due to the risk of layer debonding. A tilted front shank geometry was not found to be required for isotropic metallic materials.

10.1.2 Flutter prediction

The applicability of frequently used decoupled flutter prediction techniques for the prediction of transonic fan flutter was investigated in this thesis. A systematic comparative study was presented, contrasting the results of a decoupled method with those of a strongly coupled fluid-structure interaction approach. The blade (structure-to-air) mass ratio and stiffness were parametrically varied in order to assess their respective impact on the strength of aeroelastic coupling in terms of the aeroelastic frequency shift, i.e. the difference between the aeroelastic response frequency and the natural frequency of the blade.

The results showed that the decoupled method provides accurate predictions of fan flutter stability over wide ranges of mass ratios and blade stiffnesses. However, for a design with simultaneously low stiffness and low weight, which more closely resembles realistic lightweight fan blade designs, aeroelastic coupling effects were observed to become of increasing importance. For this case, a frequency shift of 2% was observed in conjunction with a considerable stability overprediction by the decoupled method.

In a detailed investigation into the mechanisms that govern the aeroelastic frequency shift, it was found that its magnitude mainly depends on the ratio between the in-phase aerodynamic damping force component and the blade's vibrational inertia force, which is affected by both blade mass and blade stiffness. This implied that the mass ratio in isolation can only be an indicative measure for the expected aeroelastic coupling strength in transonic fans.

A methodology was presented with which the decoupled method can be used to predict the aeroelastic frequency shift without the need to resort to a costly coupled method. This should prove useful in an a-priori assessment of the applicability of the method.

10.2 Future work

The conclusions above have shown that the stated objectives of this thesis were successfully achieved. However, several open questions and many avenues for future work are remaining. Some suggestions are given below.

- Evolutionary gradient-free optimisation methods, as used in this thesis, represent a class of powerful global optimisation algorithms. The main strengths of those methods for engineering design are their tolerance to noisy objective and constraint functions and their ability to handle multiple objectives as well as constraints in a direct fashion. Those features make evolutionary algorithms, and many other metaheuristic optimisation methods, attractive tools for the exploration of complex design spaces. However, in the close vicinity to an optimal design, where only small design changes might be needed for further performance improvements, gradient-free methods become increasingly inefficient. Gradient-based methods, on the other hand, are powerful in local design space exploitation, but perform poorly for design space exploration. The strengths and weaknesses of both methods thus appear to be almost diametrically opposed. It appears to be a promising idea to combine the two into a unified optimisation framework. Such hybrid optimisation strategies require, however, the solution of many challenges, including the definition of criteria to specify when and where to start local gradient-based searches, a feasible approach for a unified design parametrisation and a strategy for the handling constraints and multiple objectives.

- In this thesis, a fixed blade parametrisation was applied in each respective optimisation problem. The initial choice of blade parameters and their range, however, intrinsically pre-defines the optimal shape to be found. This poses a substantial problem for design problems where the optimal shape is by the very nature of the problem unknown. The fact of the matter is that, the more knowledge exists about the design problem, the better the parametrisation can be designed. This dependency oftentimes defends the purpose of optimisation in situations where very little knowledge of the problem is existing a-priori to its solution. This paradox can be approached by a concept known as adaptive parametrisation. In such an approach, the design parameters are themselves subject to change. The idea is that a high number of parameters should be

existing in regions of the geometry which are most influential, while parameters should be removed where they have little or not impact on the desired performance. The realisation of such an approach would require a geometry representation which enables parameter addition and removal without geometry changes. Additionally, sensitivity information is necessary in order to decide on where to add parameters and where to remove them. The adjoint method for the efficient computation of gradients would pose a good starting point for that. However, in complex design spaces, one has to consider that parameter influences are not constant. In some parts of the design space, a parameter might show little relevance, but might become of increasing importance in the further optimisation process.

- The consideration of flutter in the design optimisation process presents an important requirement for the design of novel thin and lightweight fan blades. It was shown in this work, that for some low-density and low-stiffness blade designs the magnitude of the unsteady aerodynamic forces can become of significance in comparison to the structural inertia and stiffness forces. It was observed that this can result in a significant alteration of the blade's vibration frequency, with a consequent effect on the flow field and the aerodynamic damping. In addition to frequency alterations, the aerodynamic forces may lead to coupling of structural modes which may further affect the aerodynamic damping - an effect that has not been considered in this thesis. Nevertheless, the results presented in this work suggest that coupled fluid-structure interaction methods seem to be required to enable an accurate prediction of fan flutter stability. Interestingly, the implementation of coupled methods into gradient-free optimisation systems is conceptually rather straightforward. In a coupled computation, the occurrence of flutter can be assessed by monitoring the time evolution of the total structural energy in the system over several periods of vibration, for all structural modes under consideration. Flutter in one or several modes is indicated by an increase of the total structural energy in those modes - effectively representing a net work input to the blade vibration from the flow field. Such an approach lends itself to automation and can easily be translated into a performance parameter that can be fed to the optimiser. The computational cost associated with the integration of coupled flutter into an optimisation system might be offset to some degree by the application of a global sequentially adaptive metamodel, similarly to the approach used in this work for the prediction

of the CFD and CSM response. However, further work is needed to verify such an approach.

- In addition to flutter, static aeroelastic blade untwist is a relevant problem in fan blade design. Fan blades are subjected to considerable static aerodynamic and centrifugal forces in operation, causing them to undergo significant elastic deformations. In transonic flows, even seemingly small geometry changes can have a substantial effect on the shock structure, with consequences for important performance parameters such as efficiency and pressure ratio but also stall and choke limits. In this thesis, the effects of static aeroelastic deformations have been neglected in the design process. Two strategies can generally be followed for their consideration. The first is to initiate the design from the so-called hot or design point geometry. Then a hot-to-cold geometry transformation (also known as blade un-running) is performed to recover the cold or manufacturing geometry of the blade, which, under the known aerodynamic and centrifugal forces at the design point, elastically deforms to yield the design point geometry. From this manufacturing geometry, the blade shape at different operating points can be determined in a so-called cold-to-hot transformation, in which the fluid-structure interaction problem is solved iteratively until the shape is found at which aeroelastic equilibrium is obtained, thus closely resembling the deformation process as it occurs in reality. In this approach, the parametrised geometry would resemble the design point geometry. The second strategy is to start from the cold or manufacturing geometry and perform a cold-to-hot transformation for each operating point under consideration. In this approach, the cold geometry would be the parametrised one. The need for the solution of several fluid structure interaction problems for each individual design would introduce a considerable computational cost using either of those approaches. Hence, similarly to the prediction of flutter, computational cost poses an issue. It can be anticipated that the prediction of the static aeroelastic response using global metamodels would likely be associated with considerable difficulties due to potentially strongly non-linear relations between design variables and aero-structural blade performance at different states of deformation. Finding a valid and computationally efficient approach for the integration of static aeroelasticity into fan design optimisation processes represents an important issue to be solved.

- Recent advances in composite material manufacturing techniques are now enabling the design of highly tailored materials. This opens up an enormous additional design freedom for fan blades. It can be envisioned that the material properties might be tailored to e.g. optimise the blade stiffness distribution to enable thinner blades, or to increase the flutter margin. In addition, the static aeroelastic deformation of blades can be influenced by the adaptation of fibre directions, which may be used to enable improved off-design performance or an increased operating range. In any of those cases, the design problem of material and shape is a coupled one. Numerical optimisation techniques can be expected to play a key enabling role in such a design scenario.

A

Validation Study of Unsteady Aerodynamics for Blade Dynamic Aeroelasticity (Flutter)

For the prediction of flutter stability, the capability of the flow solver to accurately predict the unsteady flow resulting from an oscillating blade is of central importance. In low-aspect ratio aero-engine fans the flow field is inherently three-dimensional. However, most test data available in the open literature suitable for code validation are two-dimensional in nature, e.g. [19], [25]. Validation of the commercial CFD solver used in this work was performed for oscillating two-dimensional cases, including the 11th International Standard Configuration proposed by [52], as described in [37].

In the open literature, only two cases are existing that present detailed data for 3D oscillating cascades, which are presented in Huang et al. (2009) [69] for a turbine and Yang and He (2004) [180] for a compressor. The turbine test case has been chosen for the validation of the commercial flow solver. The results are presented in this Appendix.

A.1 Test case description

The test case is a low-speed linear turbine cascade. An overview of the test facility is given in Fig. A.1. The cascade consists of seven blades, whereas the central blade is hinged at the hub and driven at the tip end to oscillate in a simulated bending mode by an external DC motor. The hinge at the hub is located outside of the test section, thus introducing a small displacement at the blade hub, which is equal to one-tenth of the

The results obtained in this chapter were included in the company internal quality control database of Numeca International.

tip displacement. The hub wall is adequately sealed to prevent flow leakage through the hub gap that is required to accommodate the vibrational blade motion. The top and bottom walls of the cascade are profiled to represent the suction/pressure side surfaces of adjacent blades, hence resulting in a total of eight blade passages. Table A.1 provides an overview of the geometric cascade configuration and the operating conditions.

The reduced frequency is defined as

$$\nu = \frac{\omega_n C}{V_{ref}} \tag{A.1}$$

where ω_n is the natural frequency in rad/s, C is the chord length and V_{ref} is the reference, isentropic exit velocity given as

$$V_{ref} = \sqrt{\frac{2(p_{01} - p_2)}{\rho}} \tag{A.2}$$

with p_{01} being the inlet stagnation pressure, p_2 the exit static pressure and ρ the density.

Steady and unsteady pressure measurements were taken using pressure tappings on the oscillating and one non-oscillating blade located at six spanwise positions (at 10, 30, 50, 70, 80, and 95% span). At each of the blade sections, 14 pressure tappings are located on the suction surface and 10 on the pressure surface, hence resulting in 144 tappings per blade. An overview of the location of the pressure tappings for one blade section is shown in Fig. A.2. For a detailed explanation of the data acquisition and reduction procedure, the reader is referred to the publication of Huang et al. [69].

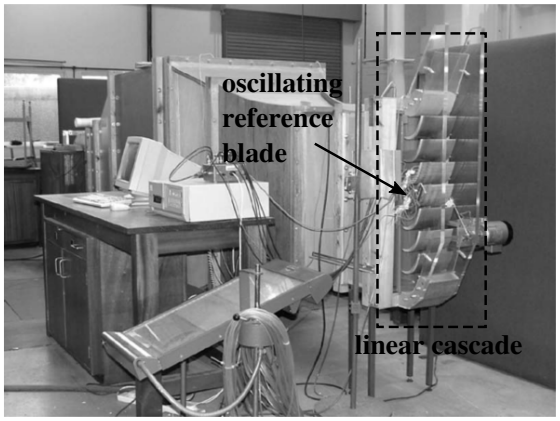


Figure A.1: Low-speed linear turbine cascade test rig. Adapted from [69].

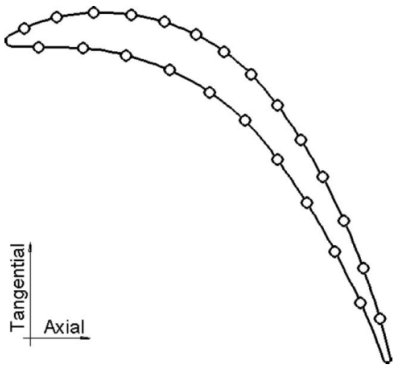


Figure A.2: Locations of the pressure tappings on one blade section [69].

Table A.1: Cascade data and operating conditions.

Blade chord, C	0.143 m
Blade height, h	0.180 m
Pitch, S	0.100 m
Stagger angle	40 deg
Inlet flow angle	0.0 deg
Vibration mode	First bending normal to absolute chord
Exp. bending amplitude	3%C at tip, 0.3%C at hub
Reynolds number, Re	$2.2 \cdot 10^5$
Reduced frequency, ν	0.4
Nominal frequency, f	10.3 Hz

A.2 Computational method description

The unsteady flow subject to the oscillating blades is computed using the non-linear harmonic (NLH) method [61], [115] with an extension to moving meshes as reported in [37]. A single flow harmonic is computed assuming the influence of higher flow harmonics to be negligible in the present case.

The computational domain containing one periodic sector of the linear cascade is shown in Fig. A.3. The domain is discretised with a multi-block structured mesh consisting of 616683 points. The radial direction is resolved with 57 layers. The tip clearance is neglected in the present configuration. Inlet and outlet blocks are discretised respectively with an H-block topology while an O-block around the blade and four H-blocks are used for the discretisation of the blade passage. The block boundaries are equally shown in Fig. A.3. Translational periodic boundary conditions

are applied on each of the periodic surfaces of the domain. A detail view of the blade-to-blade mesh at 50% span is shown in Fig. A.4.

Total temperature, pressure and flow angle are specified at the inlet. A fixed static pressure is imposed at the outlet. Merkle preconditioning is applied in order to improve convergence of the density based solver at the low Mach number encountered in this case. Further, the flow velocity is increased to match an exit Mach number of 0.3, for which compressibility effects are still negligible. The fluid viscosity has been adapted in order to recover the experimental Reynolds number. On the endwall surfaces the tangential flow is allowed to slip while no-slip boundary conditions are applied on the blade surface. Turbulence is modelled using the one-equation turbulence model of Spalart and Allmaras [154]. The oscillating tip displacement amplitude is reset to 1% of chord in the computations in order to improve the convergence.

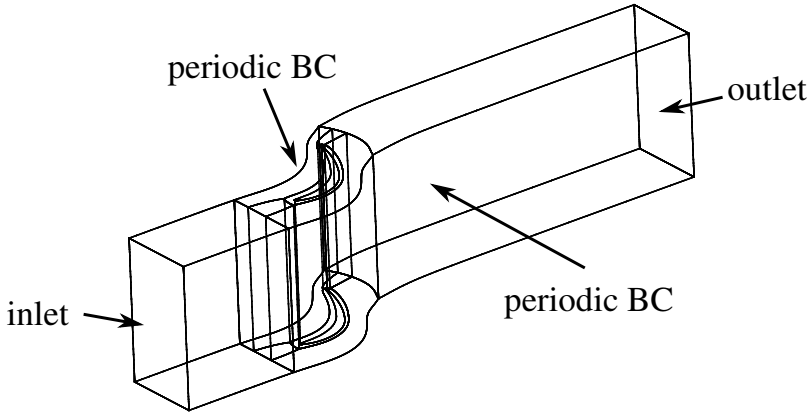


Figure A.3: Computational domain.

A.3 Comparison of numerical and experimental results

A comparison of the predicted and measured steady pressure distributions in terms of the pressure coefficient at 50% span is shown in Fig. A.2. Overall, the steady results

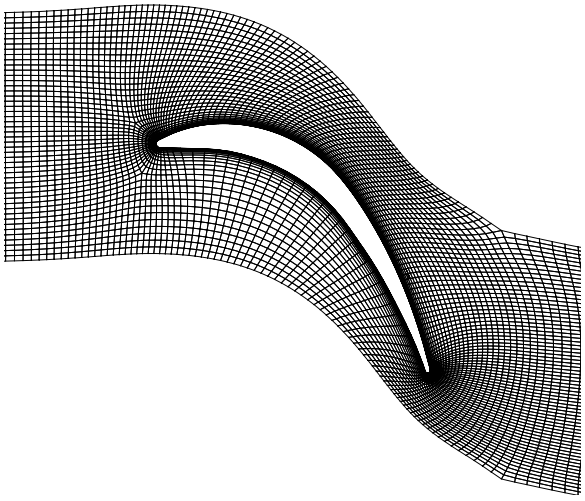


Figure A.4: Blade detail of the blade-to-blade mesh at 50% span.

can be seen to be in good agreement.

The measured and predicted first harmonic pressure obtained at a reduced frequency of $\nu = 0.4$ is shown in Fig. A.6 for the least stable and in Fig. A.7 for the most stable interblade phase angle (IBPA), respectively. The results are shown at three spanwise sections at 10, 50 and 95% span. The left hand plot shows the amplitude in terms of a pressure coefficient, defined as

$$\mathrm{cp}_1 = \frac{\hat{p}_1}{(p_{01} - p_2)\hat{x}_{tip}} \tag{A.3}$$

where $\hat{p}_1$ is the amplitude of the first harmonic pressure and $\hat{x}_{tip}$ is the blade displacement amplitude at the tip in percentage of the chord length. The right hand plots in Figs. A.6 and A.7 show the respective phase angle of the first harmonic pressure.

Overall a very good agreement is obtained between the pressure amplitudes for both interblade phase angles Fig. A.6 and Fig. A.7. Slight deviations can be observed to occur at 95% span starting from about 80% axial chord for both IBPAs, which is likely attributable to tip clearance effects, which have been neglected in the computations. The predicted phase angle shows a very good agreement with the experimental results for the most stable IBPA (Fig. A.7), while the deviations appear slightly more pronounced at the least stable IBPA in Fig. A.6. However, in a global sense, the predictions are in good agreement with the measurements.

A comparison of the predicted and measured global aerodynamic damping coefficient is shown in Fig. A.8. The aerodynamic damping coefficient is computed as

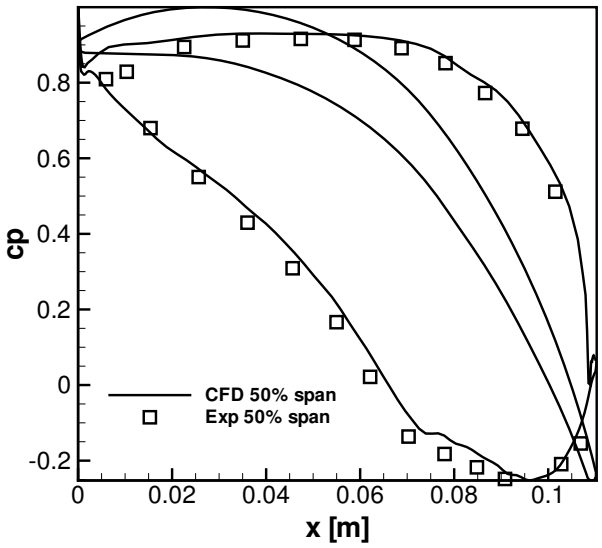


Figure A.5: Comparison of the steady pressure distributions at 50% span.

$$\xi_A = \frac{1}{h \cdot \int_0^h \xi_{A,l} dr} \tag{A.4}$$

with the local aerodynamic damping coefficient at a specified radius defined as

$$\xi_{A,l} = \int_C \frac{-\pi \hat{x}_l \ \mathrm{cp}_1 sin(\phi_1)}{(C \cdot \hat{x}_{tip})} ds \tag{A.5}$$

where $\hat{x}_l$ is the local displacement amplitude in percentage chord and ϕ is the phase angle. Computations were performed at -180, -150, -120, -90, -60, -30, 30, 60, 90, 120 and 150 degrees IBPA. The aerodynamic damping curve shows a typical S-shape, with computations and experiments being overall in good agreement.

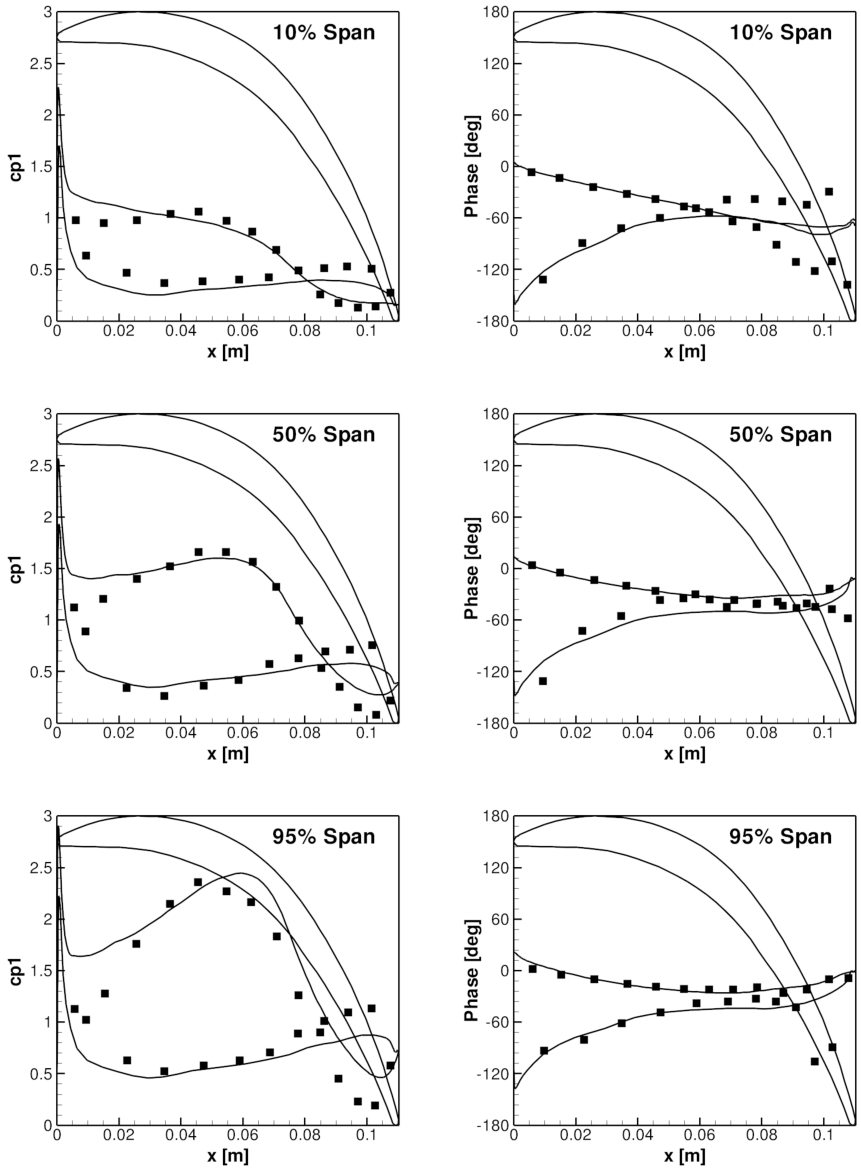


Figure A.6: Amplitude and phase of the unsteady pressure at a reduced frequency $\nu = 0.4$, and the least stable IBPA (-60 deg). Lines indicate CFD results, while squares indicate experiments.

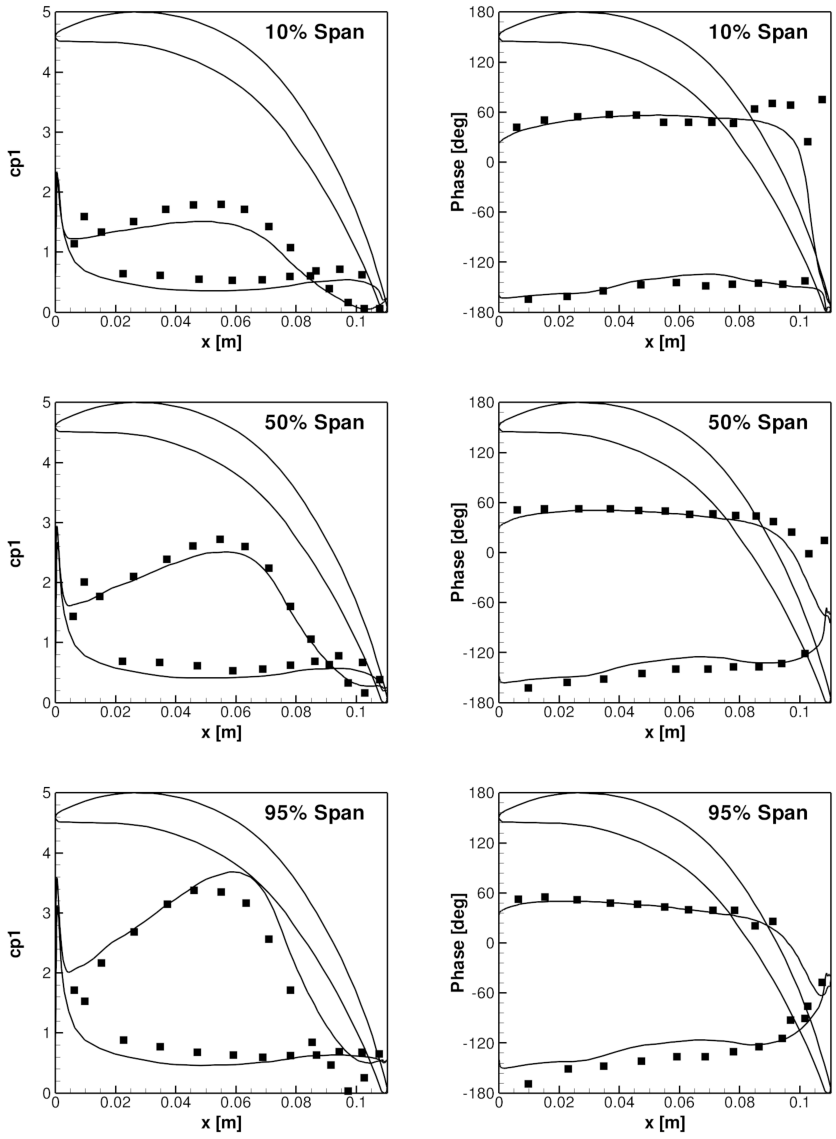


Figure A.7: Amplitude and phase of the unsteady pressure at a reduced frequency $\nu = 0.4$, and the least stable IBPA (120 deg). Lines indicate CFD results, while squares indicate experiments.

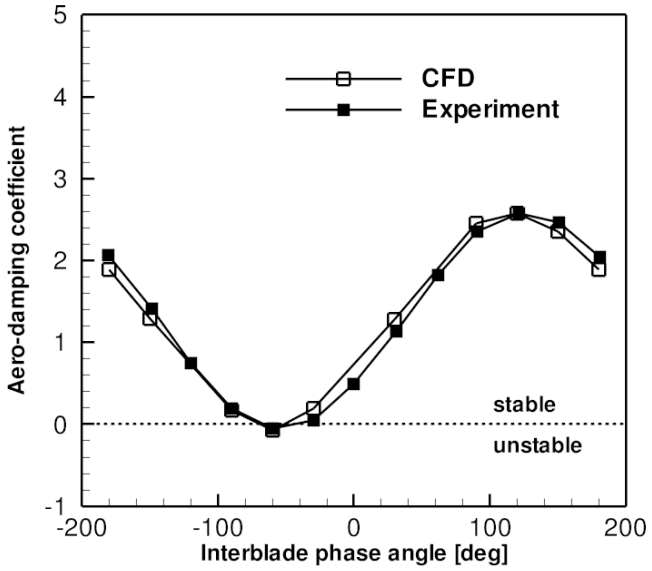


Figure A.8: Global comparison of aero-damping versus the interblade-phase angle (IBPA) for a reduced frequency of $\nu = 0.4$.

Bibliography

[1] http://gosculptor.com.

[2] http://www.integrityware.com.

[3] H.A. Abbass, R. Sarker, and C. Newton. PDE: a Pareto-frontier differential evolution approach for multi-objective optimization problems. In *Proceedings of the 2001 Congress on Evolutionary Computation*, pages 971–978, 2001.

[4] A. P. Adamson. Composite Airfoil Construction. *US Patent 4 051 289*, 1977.

[5] C.-S. Ahn and K.-Y. Kim. Aerodynamic design optimization of an axial flow compressor rotor. In *Proceedings of ASME Turbo Expo 2002*, Amsterdam, The Netherlands, 2002.

[6] C.-S. Ahn and K.-Y. Kim. Aerodynamic design optimization of a compressor rotor with Navier-Stokes analysis. *Journal of Power and Energy*, 217:179–183, 2003.

[7] B. Alatas, E. Akin, and A. Karci. MODENAR: Multi-objective differential evolution algorithm for mining numeric association rules. *Applied Soft Computing*, 8(1):646–656, 2008.

[8] P. Alexander and R. A. Naik. Composite Fan Blade. *US Patent US2011/0176927A1*, 2011.

[9] N. Anand, S. Vitale, P. Colonna, and M. Pini. Assessment of FFD and CAD-based shape parametrization methods for adjoint-based turbomachinery shape

optimization. In *Proceedings of the Global Power and Propulsion Forum*, Montreal, Canada, 2018.

[10] G. R. Anderson, M. J. Aftosmis, and M. Nemec. Parametric deformation of discrete geometry for aerodynamic shape design. In *50th AIAA Aerospace Sciences Meeting*, Nashville, Tennessee, USA, 2012.

[11] C. Audet, J. Denni, D. Moore, A. Booker, and P. Frank. A surrogate-model-based method for constrained optimization. In *8th Symposium on Multidisciplinary Analysis and Optimization*, Long Beach, CA, USA, 2000.

[12] B. V. Babu and M. M. L. Jehan. Differential evolution for multi-objective optimization. In *The 2003 Congress on Evolutionary Computation*, volume 4, pages 2696–2703, Canberra, Australia, 2003.

[13] R. C. Backhouse. Article Formed From a Composite Material. *US Patent US2009/0246446A1*, 2009.

[14] E. J. Barbero. *Finite Element Analysis of Composite Materials*. CRC Press, 2008.

[15] J.-F. M. Barthelemy and R. T. Haftka. Approximation concepts for optimum structural design - a review. *Structural Optimization*, 5:129–144, 1993.

[16] G. Becker, M. Schäfer, and A. Jameson. An advanced NURBS fitting procedure for post-processing of grid-based shape optimizations. In *49th AIAA Aerospace Sciences Meeting*, Orlando, Florida, USA, 2011.

[17] E. Benini. Three-dimensional multi-objective design optimization of a transonic compressor rotor. *Journal of Propulsion and Power*, 20(3):559–565, 2004.

[18] E. Benini and A. Toffolo. Development of high-performance airfoils for axial flow compressors using evolutionary computation. *Journal of Propulsion and Power*, 18(3):544–554, 2002.

[19] A. Bölcs and T. H. Fransson. Aeroelasticity in Turbomachines. Comparison of Theoretical and Experimental Cascade Results. Technical report, Ecole Polytechnique Federale de Lausanne, Lab de Thermique Appliquee, 1986.

[20] N. M. Bond and M. C. Hemsworth. Assembly of a Composite Blade Root and a Rotor. *US Patent 5573377*, 1996.

[21] A. J. Booker, J. E. Dennis, P. D. Frank, D. B. Serafini, V. Torczon, and M. W. Trosset. A rigorous framework for optimization of expensive functions by surrogates. *Structural Optimization*, 17(1):1–13, 1999.

[22] G. E. P. Box. Evolutionary operation: a method of increasing industrial productivity. *Applied Statistics*, 6:81–101, 1957.

[23] G. E. P. Box and N. R. Draper. *Empirical Model Bulding and Response Surfaces.* John Wiley & Sons, 1987.

[24] C. E. Brennen. A review of added mass and fluid inertial forces. Technical report, Naval Civil Engineering Laboratory, Port Hueneme, CA, USA, 1982.

[25] D. H. Buffum, V. R. Capece, A. J. King, and Y. M. El-Aini. Oscillating cascade aerodynamics at large mean incidence. *Journal of Turbomachinery*, 120:122–130, 1998.

[26] W. J. Calvert and R. B. Ginder. Transonic fan and compressor design. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 213(5):419–436, 1999.

[27] V. Carstens and J. Belz. Numerical investigation of nonlinear fluid-structure interaction in vibrating compressor blades. *Journal of Turbomachinery*, 123:402–408, 2001.

[28] F. O. Carta. Coupled blade-disk-shroud flutter instabilities in turbojet engine rotors. *Journal of Engineering for Power*, 89(3):419–426, 1967.

[29] C. J. R Champenois, L. J. P. David, and G. F. Inizan. Turbomachine Blade Made of Composite Material. *US Patent 5672417*, 1997.

[30] H. A. Chuang and J. M. Verdon. A nonlinear numerical simulator for three-dimensional flows through vibrating blade rows. *Journal of Turbomachinery*, 121(2):348–357, 1999.

[31] N. Cumpsty. *Jet propulsion.* Cambridge University Press, 2nd edition, 2003.

[32] I. M. Daniel, J. J. Luo, and P. M. Schubel. Three-dimensional characterization of textile composites. *Composites Part B: Engineering*, 39(1):13–19, 2008.

[33] F. Daoud, M. Firl, and K.-U. Bletzinger. Filter techniques in shape optimization with CAD-free parametrization. In *6th World Congress of Structural and Multidisciplinary Optimization*, Rio de Janeiro, Brazil, 2005.

[34] K. Deb. Evolutionary algorithms for multi-criterion optimization in engineering design. *Evolutionary Algorithms in Engineering and Computer Science*, pages 135–161, 1999.

[35] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.

[36] F. Debrabandere, B. Tartinville, C. Hirsch, and G. Coussement. Fluid-structure interaction using a modal approach. *Journal of Turbomachinery*, 134(5), 2012.

[37] F. Debrabandere, B. Tartinville, C. Hirsch, and G. Coussement. Extension of the non-linear harmonic method to flow computations around moving and deforming structures. In *International Forum on Aeroelasticity & Structural Dynamics 2013*, Bristol, UK, 2013.

[38] X. Deng, F. Guo, Y. Liu, and P. Han. Aero-mechanical optimization design of a transonic fan blade. In *Proceedings of ASME Turbo Expo 2013*, San Antonio, Texas, USA, 2013.

[39] J. D. Denton. The effects of lean and sweep on transonic fan performance: a computational study. In *Proceedings of ASME Turbo Expo 2002*, Amsterdam, The Netherlands, 2002.

[40] J. D. Denton and L. Xu. The exploitation of three-dimensional flow in turbomachinery design. *Proc. I. Mech. E., Part C, Journal of Mechanical Engineering Science*, 213:125–137, 1999.

[41] G. Dhondt. CalculiX CrunchiX USER' S MANUAL version 2.5, 2012.

[42] H. Doi and J. J. Alonso. Fluid/structure coupled aeroelastic computations for transonic flows in turbomachinery. In *Proceedings of ASME Turbo Expo 2002*, Amsterdam, The Netherlands, 2002.

[43] M. C. Duta, S. Shahpar, and M. B. Giles. Turbomachinery design optimization using automatic differentiated adjoint code. In *Proceedings of ASME Turbo Expo 2007*, Montreal, Canada, 2007.

[44] European Commission. Flightpath 2050: Europe's Vision for Aviation, 2011.

[45] European Commission. Annual Analyses of the EU Air Transport Market 2016. Technical report, 2017.

[46] K. Fang, R. Z. Li, and A. Sudjianto. *Design and modeling for computer experiments*. Computer Science and Data Analysis Series, 2006.

[47] C. Farhat and M. Lesoinne. Two efficient staggered algorithms for the serial and parallel solution of three-dimensional nonlinear transient aeroelastic problems. *Computer Methods in Applied Mechanics and Engineering*, 182(3-4):499–515, 2000.

[48] J. H. Ferziger and M. Perić. *Computational Methods for Fluid Dynamics*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002.

[49] S. R. Finn, G. Rengarajan, D. W. Crall, and I. F. Prentice. Reinforced Fan Blade. *US Patent US7575417B2*, 2009.

[50] L. J. Fogel, A. J. Owens, and M. J. Walsh. *Artificial Intelligence Through Simulated Evolution*. Wiley, 1966.

[51] J. M. Forrester, J. T. Stevenson, J. C. Potter, and S. J. Corrigan. Circular Arc Multi-Bore Fan Disk Assembly. *US Patent US6634863*, 2003.

[52] T. H. Fransson, M. Jöcker, A. Bölcs, and P. Ott. Viscous and inviscid linear/nonlinear calculations versus quasi-three-dimensional experimental cascade data for a new aeroelastic turbine standard configuration. *Journal of Turbomachinery*, 121(4), 1999.

[53] A. S. Fraser. Simulation of genetic systems by automatic digital computers. *Australian Journal of Biological Science*, 10:484–491, 1957.

[54] D. F. Gray, C. E. Doughty, and F. B. Pinney. Composite Blade. *US Patent 3600103*, 1971.

[55] Z. Guerdal, R. T. Haftka, and P. Hajela. *Design and Optimization of Laminated Composite Materials*. John Wiley & Sons, 1999.

[56] R. T. Haftka and Z. Gurdal. *Elements of Structural Optimization*. Kluwer Academic Publishers, 1992.

[57] C. Hah, S. L. Puterbaugh, and W. W. Copenhaver. Unsteady aerodynamic flow phenomena in a transonic compressor stage. *Journal of Propulsion and Power*, 13(3):329–333, 1997.

[58] P. Hajela and L. Berke. Neural networks in structural analysis and design: An overview. *Computing Systems in Engineering*, 3(1-4):525–538, 1992.

[59] D. J. Harris. Composite Blade for Turbofan Engine Fan. *US Patent 4343593*, 1982.

[60] L. He and J. D. Denton. Three dimensional time-marching inviscid and viscous solutions for unsteady flows around vibrating blades. *Journal of Turbomachinery*, 116(3):469–476, 1993.

[61] L. He and W. Ning. Efficient approach for analysis of unsteady viscous flows in turbomachines. *AIAA Journal*, 36(11):2005–2012, 1998.

[62] L. He and D. X. Wang. Concurrent blade aerodynamic-aero-elastic design optimization using adjoint method. *Journal of Turbomachinery*, 133(1), 2011.

[63] L. He and S. Zhou. A check on the energy method of predicting blade transonic stall flutter. In *Beijing International Gas Turbine Symposium and Exposition*, Beijing, China, 1985.

[64] Hexcel. 3501-6 Epoxy Matrix. *Material Datasheet*, 2016.

[65] J. E. Hicken and D. W. Zingg. Aerodynamic optimization algorithm with integrated geometry parameterization and mesh movement. *AIAA Journal*, 48(2):400–413, 2010.

[66] C. Hirsch. *Numerical Computation of Internal and External Flows*. Elsevier, 2007.

[67] J. H. Holland. *Adaptation in Natural and Artificial Systems*. The University of Michigan Press, 1975.

[68] X. Hu and R. Eberhart. Solving constrained nonlinear optimization problems with particle swarm optimization. In *Proceedings of the sixth world multiconference on systemics, cybernetics and informatics*, pages 203–206, Orlando, FL, USA, 2002.

[69] X. Q. Huang, L. He, and D. L. Bell. Experimental and computational study of oscillating turbine cascade and influence of part-span shrouds. *Journal of Fluids Engineering*, 131(5), 2009.

[70] D. L. Huff, T. W. Swafford, and T. S. R. Reddy. Euler flow predictions for an oscillating cascade using a high resolution wave-split scheme. In *Proceedings of International Gas Turbine and Aeroengine Congress and Exposition*, Orlando, Florida, USA, 1991.

[71] T. Ishikawa and T.-W. Chou. Nonlinear behavior of woven fabric composites. *Journal of Composite Materials*, 17(5):399–413, 1983.

[72] D. I. James and N. M. Merriman. Composite Turbomachine Blade. *US Patent US2012/0003100A1*, 2012.

[73] A. Jameson, N. A. Pierce, and L. Martinelli. Optimum aerodynamic design using the Navier-Stokes equations. In *35th Aerospace Sciences Meeting and Exhibit, AIAA Paper 97-0101*, 1997.

[74] C.-M. Jang, P. Li, and K.-Y. Kim. Optimization of blade sweep in a transonic axial compressor rotor. *JSME International Journal Series B*, 48(4):793–801, 2005.

[75] C.-M. Jang, A. Samad, and K.-Y. Kim. Optimal design of swept, leaned and skewed blades in a transonic axial compressor. In *Proceedings of ASME Turbo Expo 2006*, Barcelona, Spain, 2006.

[76] I. K. Jennions and M. G. Turner. Three-dimensional navier-stokes computations of transonic fan flow using an explicit flow solver and an implicit K-E solver. In *Proceedings of International Gas Turbine and Aeroengine Congress and Exposition*, Cologne, Germany, 1992.

[77] A. John, S. Shahpar, and N. Qin. Alleviation of shock-wave effects on a highly loaded axial compressor through novel blade shaping. In *Proceedings of ASME Turbo Expo 2016*, Seoul, South Korea, 2016.

[78] M. M. Joly, T. Verstraete, and G. Paniagua. Full design of a highly loaded and compact contra-rotating fan using multidisciplinary evolutionary optimization. GT2013-94433. In *Proceedings of ASME Turbo Expo 2013*, San Antonio, Texas, USA, 2013.

[79] M. M. Joly, T. Verstraete, and G. Paniagua. Multidisciplinary design optimization of a compact highly loaded fan. *Structural and Multidisciplinary Optimization*, 49(3):471–483, 2014.

[80] R. M. Jones. *Mechanics of Composite Materials*. Tailor and Francis, Inc., 2nd edition, 1999.

[81] J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proceedings of ICNN'95 - International Conference on Neural Networks*, pages 1942–1948, Perth, WA, Australia, 1995.

[82] J. H. Kim, D. W. Kim, and K. Y. Kim. Aerodynamic optimization of a transonic axial compressor with a casing groove combined with tip injection. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 227(8):869–884, 2013.

[83] G. Kiss, C. Giannelli, U. Zore, B. Jüttler, D. Großmann, and J. Barner. Adaptive CAD model (re-)construction with THB-splines. *Graphical Models*, 76(5):273–288, 2014.

[84] J. P. C. Kleijnen. Kriging metamodeling in simulation: A review. *European Journal of Operational Research*, 192(3):707–716, 2009.

[85] N. J. Kray, T. Davis, C. L. McAfee, M. J. Franks, K. L. Kirkeng, and D. Crall. Method and System for Manufacturing a Blade. *US Patent US8123463B2*, 2012.

[86] D. G. Krige. A statistical approach to some basic mine valuation problems on the Witwatersrand. *Journal of the Chemical, Metal and Mining Society of South Africa*, 52(6):119–139, 1951.

[87] L. Lapworth. Challenges and methodologies in the design of axial flow fans for high-bypass-ratio, gas turbine engines using steady and unsteady CFD. In *Advances of CFD in Fluid Machinery Design*. John Wiley and Sons Ltd., 2003.

[88] G. Lauzeille. Turbomachine Blade Made of Composite Material. *US Patent US 5292231*, 1994.

[89] C. Le, T. Bruns, and D. Tortorelli. A gradient-based, parameter-free approach to shape optimization. *Computer Methods in Applied Mechanics and Engineering*, 200(9-12):985–996, 2011.

[90] Y. Lian and M.-S. Liou. Multiobjective optimization using coupled response surface model and evolutionary algorithm. In *10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Albany, NY, USA, 2004.

[91] Y. Lian and M.-S. Liou. Multi-objective optimization of transonic compressor blade using evolutionary algorithm. *Journal of Propulsion and Power*, 21(6), 2005.

[92] Y. Lian and M.-S. Liou. Aerostructural optimization of a transonic compressor rotor. *Journal of Propulsion and Power*, 22(4):880–888, 2006.

[93] C. Luo, L. Song, J. Li, and Z. Feng. A study on multidisciplinary optimization of an axial compressor blade based on evolutionary algorithms. *Journal of Turbomachinery*, 134(5):054501, 2012.

[94] J. Luo, C. Zhou, and F. Liu. Multipoint design optimization of a transonic compressor blade by using an adjoint method. *Journal of Turbomachinery*, 136(5):051005, 2013.

[95] N. K. Madavan. Multiobjective optimization using a Pareto differential evolution approach. In *Proceedings of the Congress on Evolutionary Computation (CEC 2002)*, pages 1145–1150, Honolulu, HI, USA, 2002.

[96] J. G. Marshall and M. Imregun. An analysis of the aeroelastic behaviour of a typical fan-blade with emphasis on the flutter mechanism. In *Proceedings of the International Gas Turbine and Aeroengine Congress and Exhibition*, Birmingham, United Kingdom, 1996.

[97] J.G. Marshall. *Prediction of turbomachinery aeroelasticity effects using a 3D non-linear integrated method*. Doctoral dissertation, Imperial College of Science, Technology and Medicine London, 1996.

[98] A. Massardo, A. Satta, and M. Marini. Axial flow compressor design optimization: Part II - throughflow analysis. *Journal of Turbomachinery*, 112(3), 1990.

[99] G. Matheron. Principles of geostatistics. *Economic Geology*, 58(8):1246–1266, 1963.

[100] J. D. Mattingly. *Elements of Gas Turbine Propulsion*. McGraw-Hill, New York, 1996.

[101] A. J. McMillan. Composite Blade. *US Patent US20060216154A1*, 2006.

[102] A. J. McMillan. Composite Blade. *US Patent US7837446B2*, 2010.

[103] A. J. McMillan. Composite Aerofoil. *US Patent US8333565B2*, 2012.

[104] A. J. McMillan. A Composite Blade. *EU Patent EP1681439B1*, 2014.

[105] A. J. McMillan and A. D. Jackson. Composite Blade. *US Patent US7384240B2*, 2008.

[106] S. Moffatt and L. He. On decoupled and fully-coupled methods for blade forced response prediction. *Journal of Fluids and Structures*, 20(2):217–234, 2005.

[107] D. C. Montgomery. *Design and Analysis of Experiments*. John Wiley & Sons, 2017.

[108] J. G. Morley. Composite Blade. *US Patent 3602608*, 1969.

[109] F. Moyroud, G. Jacquet-Richardet, and T. H. Fransson. Aeroelasticity in turbomachines: some aspects of the effect of coupling modeling and blade material changes. *International Journal of Rotating Machinery*, 6(4):265–273, 2000.

[110] G. C. Murphy and B. J. Duhrmann. Wide Chord Fan Blade. *EU Patent EP0496550B1*, 1997.

[111] G. C. Murphy and B. J. Fuhrman. Wide Chord Fan Blade. *US Patent 5141400*, 1992.

[112] R. A. Naik. Three-Dimensionally Woven Composite Blade with Spanwise Weft Yarns. *US Patent US2011/0182743A1*, 2011.

[113] R. A. Naik. Three-Dimensionally Woven Composite Blade With Spanwise Weft Yarns. *US Patent US2013/0243603A1*, 2013.

[114] R. A. Naik and S. R. Clarke. Integrally Woven Composite Fan Blade Using Progressively Larger Weft Yarns. *US Patent US8662855B2*, 2014.

[115] W. Ning and L. He. Computation of unsteady flows around oscillating blades using linear and nonlinear harmonic Euler methods. *ASME Journal of Turbomachinery*, 120(3):508–514, 1998.

[116] G. C. Oates. *The Aerothermodynamics of Gas Turbine and Rocket Propulsion.* AIAA Education Series, Reston, Virigina, 3rd edition, 1997.

[117] H. Okui, T. Verstraete, R.A. Van den Braembussche, and Z. Alsalihi. Three-dimensional design and optimization of a transonic rotor in axial flow compressors. *Journal of Turbomachinery*, 135(3), 2013.

[118] A. Oyama, M.-S. Liou, and S. Obayashi. Transonic axial-flow blade shape optimization using evolutionary algorithm and three-dimensional Navier-Stokes solver. *Journal of Propulsion and Power*, 20(4):612–619, 2004.

[119] V. Pareto. *Manual of Political Economy.* Augustus M. Kelley Publishers, New York, 1971.

[120] M. Parkin, P. Alexander, B. P. Huth, and C. B. Klinetob. Method of Making Composite Fan Blade. *EU Patent EP2540486A2*, 2013.

[121] M. Parkin, B. P. Huth, and W. M. Bedard. Method of Manufacturing a Composite Fan Blade with Co-Cured Sheath, and Corresponding Fan Blade. *EU Patent EP2353830A2*, 2011.

[122] K. K. E. Parsopoulos and M. N. Vrahatis. Particle swarm optimization method for constrained optimization problems. *Intelligent Technologies-Theory and Application: New Trends in Intelligent Technologies*, 76:214–220, 2002.

[123] S. Pierret, R. Filomeno Coelho, and H. Kato. Multidisciplinary and multiple operating points shape optimization of three-dimensional compressor blades. *Structural and Multidisciplinary Optimization*, 33(1):61–70, 2007.

[124] C. E. Platt. Composite Rotor Blade. *US Patent 4108572*, 1978.

[125] K. Price, R. Storn, and J. A. Lampinen. *Differential evolution: a practical approach to global optimization.* Springer-Verlag New York, 2005.

[126] S. S. Rao. *The Finite Element Method in Engineering.* Butterworth Heinemann, 1999.

[127] T. Ray and K. M. Liew. A swarm with an effective information sharing mechanism for unconstrained and constrained single objective optimisation problems. In *Procceedings of the IEEE Congress on Evolutionary Computation*, pages 75–80, Seoul, South Korea, 2001.

[128] I. Rechenberg. *Evolutionsstrategie: Optimierung Technischer Systeme nach Prinzipien der Biologischen Evolution*. Frommann-Holzboog, 1973.

[129] L. Reid and R. D. Moore. Design and overall performance of four highly loaded, high-speed inlet stages for an advanced high-pressure-ratio core compressor. *NASA Technical Paper 1337*, 1978.

[130] M. Reyes-Sierra and C. A. Coello Coello. Multi-objective particle swarm optimizers: a survey of the state-of-the-art. *International Journal of Computational Intelligence Research*, 2(3):287–308, 2006.

[131] C. M. Rhie, R. M. Zacharias, D. E. Hobbs, K. P. Sarathy, B. P. Biederman, C. R. Lejambre, and D. A. Spear. Advanced transonic fan design procedure based on a Navier-Stokes method. *Journal of Turbomachinery*, 116:291–297, 1994.

[132] E. A. Rothman. Composite Fan Blade and Method of Construction. *US Patent 3903578*, 1975.

[133] J. Sacks, W. J. Welch, T. J. Mitchell, and H. P. Wynn. Design and analysis of computer experiments. *Statistical Science*, 4(4):409–435, 1989.

[134] M. Sadeghi and F. Liu. Computation of cascade flutter by uncoupled and coupled methods. *International Journal of Computational Fluid Dynamics*, 19(8):559–569, 2005.

[135] M. Sadeghi and F. Liu. Coupled fluid-structure simulation for turbomachinery blade rows. In *43rd AIAA Aerospace Sciences Meeting and Exhibit*, pages 10–13, Reston, Virigina, USA, 2005.

[136] A. Samad and K. Y. Kim. Multi-objective optimization of an axial compressor blade. *Journal of Mechanical Science and Technology*, 22(5):999–1007, 2008.

[137] J. A. Samareh. Multidisciplinary aerodynamic-structural shape optimization using deformation (MASSOUD). In *8th AIAA/NASA/USAF/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, Long Beach, CA, USA, 2000.

[138] J. A. Samareh. Aerodynamic shape optimization based on free-form deformation. In *10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Reston, Virigina, USA, 2004.

[139] M. Sasena, P. Papalambros, and P. Goovaerts. Global optimization of problems with disconnected feasible regions via surrogate modeling. In *9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, Atlanta, Georgia, USA, 2002.

[140] L.A. Schmit. Structural design by systematic synthesis. In *2nd Conference on Electronic Computation*, pages 105–122, New York, USA, 1960.

[141] L.A. Schmit and B. Farshi. Some approximation concepts for structural synthesis. *AIAA Journal*, 12(5):692–699, 1973.

[142] K. Schreiber. Fan Blade for a Gas-Turbine Engine. *US Patent US8251664B2*, 2012.

[143] T. W. Sederberg and S. R. Parry. Free-form deformation of solid geometric models. *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques - SIGGRAPH '86*, pages 151–160, 1986.

[144] P. Seshadri, G. T. Parks, and S. Shahpar. An aerodynamic analysis of a robustly redesigned modern aero-engine fan. *arXiv preprint*, (arXiv:1604.02345vl), 2016.

[145] S. Shahpar, A. Polynkin, and V. Toropov. Large scale optimization of transonic axial compressor rotor blades. In *49th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Schaumburg, IL, USA, 2008.

[146] U. Siller and M. Aulich. Multidisciplinary 3D-optimization of a fan stage performance map with consideration of the static and dynamic rotor mechanics. In *Proceedings of ASME Turbo Expo 2010*, Glasgow, UK, 2010.

[147] U. Siller, C. Voss, and E. Nicke. Automated multidisciplinary optimization of a transonic axial compressor. In *Proceedings of the 47th AIAA Aerospace Sciences Meeting*, Florida, USA, 2009.

[148] T. W. Simpson, A. J. Booker, D. Ghosh, A. A. Giunta, P. N. Koch, and R.-J. Yang. Approximation methods in multidisciplinary analysis and optimization: a panel discussion. *Structural and Multidisciplinary Optimization*, 27(5):302–313, 2004.

[149] T. W. Simpson, J. D. Peplinski, P. N. Koch, and J. K. Allen. Metamodels for computer-based engineering design: survey and recommendations. *Engineering with Computers*, 17:129–150, 2001.

[150] G. B. Sinclair, N. G. Cormier, J. H. Griffin, and G. Meda. Contact Stresses in Dovetail Attachments: Finite Element Modeling. *Journal of Engineering for Gas Turbines and Power*, 124(1):182, 2002.

[151] N. H. S. Smith. The aerodynamic design of the LPC system. In *VKI LS 2013-04: Aeroengine design: From state of the art of turbofans towards innovative architectures*. Von Karman Institute for Fluid Dynamics, Rhode-Saint-Genese, 2013.

[152] J. Sobieszczanski-Sobieski. Multidisciplinary design optimization: An emerging new engineering discipline. In *Advances in Structural Optimization*, volume 25, pages 483–496. 1995.

[153] P. Song, J. Sun, and K. Wang. Axial flow compressor blade optimization through flexible shape tuning by means of cooperative co-evolution algorithm and adaptive surrogate model. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 228(7):782–798, 2014.

[154] P. R. Spalart and S. R. Allmaras. A one-equation turbulence model for aerodynamic flows. *Recherche Aerospatiale*, (1):5–21, 1994.

[155] R. Srivastava, M. A. Bakhle, T. G. Keith Jr, and G. L. Stefko. Flutter analysis of a transonic fan. *NASA Technical Memorandum No. 211818*, 2002.

[156] R. Srivastava and T. S. R. Reddy. Comparative study of coupled-mode flutter-analysis methods for fan configurations. *Journal of Propulsion and Power*, 15(3):447–453, 1999.

[157] R. Storn and K. Price. Differential evolution - a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11(4):341–359, 1997.

[158] A. J. Strazisar, R. Wood, M. D. Hathaway, and K. L. Suder. Laser anemometer measurements in a transonic axial-flow fan rotor. *NASA Technical Paper 2879*, 1989.

[159] S. Teichel, T. Verstraete, and J. Seume. Optimized preliminary design of compact axial compressors, a comparison of two design tools. In *Proceedings of the 31st AIAA Applied Aerodynamics Conference*, San Diego, California, 2013.

[160] A. Toffolo, V. Venezia, and E. Benini. Genetic diversity as an objective in multi-objective evolutionary algorithms. *Evolutionary Computation*, 11(2):151–167, 2003.

[161] V. V. Toropov, L. F. Alvarez, and H. Ravaii. Structural damage recognition using optimization techniques and genetic programming methodology. In *3rd World Congress on Structural and Multidisciplinary Optimization*, Amherst, NY, USA, 1999.

[162] V. V. Toropov and H. Carlsen. Optimization of stirling engine performance based on multipoint approximation technique. In *Advances in Design Automation, Vol 2*, 1994.

[163] S. W. Tsai and E. M. Wu. A general theory of strength for anisotropic materials. *Journal of Composite Materials*, 5(1):58–80, 1971.

[164] T. Tušar and B. Filipič. Differential evolution versus genetic algorithms in multi-objective optimization. In *Evolutionary Multi-Criterion Optimization*. Springer, 2007.

[165] J. M. Vaassen, P. De Vincenzo, C. Hirsch, and B. Leonard. Strong coupling algorithm to solve fluid-structure-interaction problems with a staggered approach. In *Proceedings of the 7th European Symposium on Aerothermodynamics*, Brugge, Belgium, 2011.

[166] M. Vahdati and M. Imregun. A non-linear aeroelasticity analysis of a fan blade using unstructured dynamic meshes. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 210(6):549–564, 1996.

[167] M. Vahdati, A. I. Sayma, J. G. Marshall, and M. Imregun. Mechanisms and prediction methods for fan blade stall flutter. *Journal of Propulsion and Power*, 17(5):1100–1108, 2001.

[168] M. Vahdati, G. Simpson, and M. Imregun. Mechanisms for wide-chord fan blade flutter. *Journal of Turbomachinery*, 133(4), 2011.

[169] G. N. Vanderplaats. *Numerical Optimization Techniques for Engineering Design*. McGraw-Hill, 1984.

[170] T. Verstraete. CADO : a computer aided design and optimization tool for turbomachinery applications. In *2nd International Conference on Engineering Optimization*, Lisbon, Portugal, 2010.

[171] T. Verstraete. Multidisciplinary optimization of turbomachinery components using differential evolution. In *Introduction to Optimization and Multidisciplinary Design in Aeronautics and Turbomachinery, VKI LS 2010-07*. Von Karman Institute for Fluid Dynamics, Rhode-Saint-Genese, Belgium, 2010.

[172] H.-P. Wang and M. K. Lee. Method for Forming a Composite Airfoil Structure. *US Patent 5939006*, 1999.

[173] X. Wang, S. Wang, and W. Han. Multi-objective aerodynamic design optimization based on camber line and thickness distribution for a transonic compressor rotor. In *Proceedings of the IMECE 2008*, Boston, Massachusetts, USA, 2008.

[174] K. C. Warren, R. A. Lopez-Anido, and J. Goering. Experimental investigation of three-dimensional woven composites. *Composites Part A: Applied Science and Manufacturing*, 73:242–259, 2015.

[175] A. J. Wennerstrom. *Design of Highly Loaded Axial Flow Fans and Compressors*. Concepts ETI, 2001.

[176] A. Wilson, R. Stieger, J. Coupland, N. Smith, and N. Humphreys. Multidisciplinary optimisation of a transonic fan for low tone noise. In *17th AIAA/CEAS Aeroacoustics Conference (32nd AIAA Aeroacoustics Conference)*, Portland, USA, 2011.

[177] C. H. Wu. A general through flow theory of fluid flow with subsonic or supersonic velocities in turbomachines of arbitrary hub and casing shapes. Technical report, NACA, 1951.

[178] S. Xu, W. Jahn, and J.-D. Müller. CAD-based shape optimisation with CFD using a discrete adjoint. *International Journal for Numerical Methods in Fluids*, 74:153–168, 2013.

[179] B. Yang, Q. Xu, L. He, L. H. Zhao, and P. Ren. A novel global optimization algorithm and its application to airfoil optimization. *Journal of Turbomachinery*, 137, 2015.

[180] H. Yang and L. He. Experimental study on linear compressor cascade with three-dimensional blade oscillation. *Journal of Propulsion and Power*, 20(1):180–188, 2004.

[181] A. A. Yavuz, B. Glas, M. Ihle, H. Hacioglu, and K. Wehefritz. Composite Fan Annulus Filler with Nano-Coating. *US Patent US 2016/0047261 A1*, 2016.

[182] W. Yi, H. Huang, and W. Han. Design optimization of transonic compressor rotor using CFD and genetic algorithm. In *Proceedings of ASME Turbo Expo 2006*, Barcelona, Spain, 2006.

[183] R. Zheng, J. Xiang, and J. Sun. Blade geometry optimization for axial flow compressor. In *Proceedings of ASME Turbo Expo 2010*, Glasgow, UK, 2010.

[184] O. C. Zienkiewicz and R. L. Taylor. *The Finite Element Method for Solid and Structural Mechanics*. Elsevier, 6th edition, 2005.

[185] O. C. Zienkiewicz, R. L. Taylor, and J. Z. Zhu. *The Finite Element Method - Its Basis and Fundamentals*. Elsevier, 6th edition, 2005.